

The Coq Proof Assistant

Reference Manual

December 16, 2017

Version 8.7.1¹

The Coq Development Team

πr^2 Project (formerly LogiCal, then TypiCal)

¹This research was partly supported by IST working group “Types”

V8.7.1, December 16, 2017

©INRIA 1999-2004 (CoQ versions 7.x)

©INRIA 2004-2017 (CoQ versions 8.x)

This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, v1.0 or later (the latest version is presently available at

<http://www.opencontent.org/openpub>). Options A and B of the licence are *not* elected.

Introduction

This document is the Reference Manual of version 8.7.1 of the COQ proof assistant. A companion volume, the COQ Tutorial, is provided for the beginners. It is advised to read the Tutorial first. A book [14] on practical uses of the COQ system was published in 2004 and is a good support for both the beginner and the advanced user.

The COQ system is designed to develop mathematical proofs, and especially to write formal specifications, programs and to verify that programs are correct with respect to their specification. It provides a specification language named GALLINA. Terms of GALLINA can represent programs as well as properties of these programs and proofs of these properties. Using the so-called *Curry-Howard isomorphism*, programs, properties and proofs are formalized in the same language called *Calculus of Inductive Constructions*, that is a λ -calculus with a rich type system. All logical judgments in COQ are typing judgments. The very heart of the Coq system is the type-checking algorithm that checks the correctness of proofs, in other words that checks that a program complies to its specification. COQ also provides an interactive proof assistant to build proofs using specific programs called *tactics*.

All services of the COQ proof assistant are accessible by interpretation of a command language called *the vernacular*.

COQ has an interactive mode in which commands are interpreted as the user types them in from the keyboard and a compiled mode where commands are processed from a file.

- The interactive mode may be used as a debugging mode in which the user can develop his theories and proofs step by step, backtracking if needed and so on. The interactive mode is run with the `coqtop` command from the operating system (which we shall assume to be some variety of UNIX in the rest of this document).
- The compiled mode acts as a proof checker taking a file containing a whole development in order to ensure its correctness. Moreover, COQ's compiler provides an output file containing a compact representation of its input. The compiled mode is run with the `coqc` command from the operating system.

These two modes are documented in Chapter 14.

Other modes of interaction with COQ are possible: through an emacs shell window, an emacs generic user-interface for proof assistant (PROOF GENERAL [1]) or through a customized interface (PCoq [138]). These facilities are not documented here. There is also a COQ Integrated Development Environment described in Chapter 16.

How to read this book

This is a Reference Manual, not a User Manual, so it is not made for a continuous reading. However, it has some structure that is explained below.

- The first part describes the specification language, Gallina. Chapters 1 and 2 describe the concrete syntax as well as the meaning of programs, theorems and proofs in the Calculus of Inductive Constructions. Chapter 3 describes the standard library of COQ. Chapter 4 is a mathematical description of the formalism. Chapter 5 describes the module system.
- The second part describes the proof engine. It is divided in five chapters. Chapter 6 presents all commands (we call them *vernacular commands*) that are not directly related to interactive proving: requests to the environment, complete or partial evaluation, loading and compiling files. How to start and stop proofs, do multiple proofs in parallel is explained in Chapter 7. In Chapter 8, all commands that realize one or more steps of the proof are presented: we call them *tactics*. The language to combine these tactics into complex proof strategies is given in Chapter 9. Examples of tactics are described in Chapter 10.
- The third part describes how to extend the syntax of COQ. It corresponds to the Chapter 12.
- In the fourth part more practical tools are documented. First in Chapter 14, the usage of `coq` (batch mode) and `coqtop` (interactive mode) with their options is described. Then, in Chapter 15, various utilities that come with the COQ distribution are presented. Finally, Chapter 16 describes the COQ integrated development environment.
- The fifth part documents a number of advanced features, including coercions, canonical structures, typeclasses, program extraction, and specialized solvers and tactics. See the table of contents for a complete list.

At the end of the document, after the global index, the user can find specific indexes for tactics, vernacular commands, and error messages.

List of additional documentation

This manual does not contain all the documentation the user may need about COQ. Various informations can be found in the following documents:

Tutorial A companion volume to this reference manual, the COQ Tutorial, is aimed at gently introducing new users to developing proofs in COQ without assuming prior knowledge of type theory. In a second step, the user can read also the tutorial on recursive types (document `RecTutorial.ps`).

Installation A text file `INSTALL` that comes with the sources explains how to install COQ.

The COQ standard library A commented version of sources of the COQ standard library (including only the specifications, the proofs are removed) is given in the additional document `Library.ps`.

Credits

COQ is a proof assistant for higher-order logic, allowing the development of computer programs consistent with their formal specification. It is the result of about ten years of research of the Coq project. We shall briefly survey here three main aspects: the *logical language* in which we write our axiomatizations and specifications, the *proof assistant* which allows the development of verified mathematical proofs, and the *program extractor* which synthesizes computer programs obeying their formal specifications, written as logical assertions in the language.

The logical language used by COQ is a variety of type theory, called the *Calculus of Inductive Constructions*. Without going back to Leibniz and Boole, we can date the creation of what is now called mathematical logic to the work of Frege and Peano at the turn of the century. The discovery of antinomies in the free use of predicates or comprehension principles prompted Russell to restrict predicate calculus with a stratification of *types*. This effort culminated with *Principia Mathematica*, the first systematic attempt at a formal foundation of mathematics. A simplification of this system along the lines of simply typed λ -calculus occurred with Church's *Simple Theory of Types*. The λ -calculus notation, originally used for expressing functionality, could also be used as an encoding of natural deduction proofs. This Curry-Howard isomorphism was used by N. de Bruijn in the *Automath* project, the first full-scale attempt to develop and mechanically verify mathematical proofs. This effort culminated with Jutting's verification of Landau's *Grundlagen* in the 1970's. Exploiting this Curry-Howard isomorphism, notable achievements in proof theory saw the emergence of two type-theoretic frameworks; the first one, Martin-Löf's *Intuitionistic Theory of Types*, attempts a new foundation of mathematics on constructive principles. The second one, Girard's polymorphic λ -calculus F_ω , is a very strong functional system in which we may represent higher-order logic proof structures. Combining both systems in a higher-order extension of the Automath languages, T. Coquand presented in 1985 the first version of the *Calculus of Constructions*, CoC. This strong logical system allowed powerful axiomatizations, but direct inductive definitions were not possible, and inductive notions had to be defined indirectly through functional encodings, which introduced inefficiencies and awkwardness. The formalism was extended in 1989 by T. Coquand and C. Paulin with primitive inductive definitions, leading to the current *Calculus of Inductive Constructions*. This extended formalism is not rigorously defined here. Rather, numerous concrete examples are discussed. We refer the interested reader to relevant research papers for more information about the formalism, its meta-theoretic properties, and semantics. However, it should not be necessary to understand this theoretical material in order to write specifications. It is possible to understand the Calculus of Inductive Constructions at a higher level, as a mixture of predicate calculus, inductive predicate definitions presented as typed PROLOG, and recursive function definitions close to the language ML.

Automated theorem-proving was pioneered in the 1960's by Davis and Putnam in propositional calculus. A complete mechanization (in the sense of a semi-decision procedure) of classical first-order logic was proposed in 1965 by J.A. Robinson, with a single uniform inference rule called *resolution*. Resolution relies on solving equations in free algebras (i.e. term structures), using the *unification algorithm*.

Many refinements of resolution were studied in the 1970's, but few convincing implementations were realized, except of course that PROLOG is in some sense issued from this effort. A less ambitious approach to proof development is computer-aided proof-checking. The most notable proof-checkers developed in the 1970's were LCF, designed by R. Milner and his colleagues at U. Edinburgh, specialized in proving properties about denotational semantics recursion equations, and the Boyer and Moore theorem-prover, an automation of primitive recursion over inductive data types. While the Boyer-Moore theorem-prover attempted to synthesize proofs by a combination of automated methods, LCF constructed its proofs through the programming of *tactics*, written in a high-level functional meta-language, ML.

The salient feature which clearly distinguishes our proof assistant from say LCF or Boyer and Moore's, is its possibility to extract programs from the constructive contents of proofs. This computational interpretation of proof objects, in the tradition of Bishop's constructive mathematics, is based on a realizability interpretation, in the sense of Kleene, due to C. Paulin. The user must just mark his intention by separating in the logical statements the assertions stating the existence of a computational object from the logical assertions which specify its properties, but which may be considered as just comments in the corresponding program. Given this information, the system automatically extracts a functional term from a consistency proof of its specifications. This functional term may be in turn compiled into an actual computer program. This methodology of extracting programs from proofs is a revolutionary paradigm for software engineering. Program synthesis has long been a theme of research in artificial intelligence, pioneered by R. Waldinger. The Tablog system of Z. Manna and R. Waldinger allows the deductive synthesis of functional programs from proofs in tableau form of their specifications, written in a variety of first-order logic. Development of a systematic *programming logic*, based on extensions of Martin-Löf's type theory, was undertaken at Cornell U. by the Nuprl team, headed by R. Constable. The first actual program extractor, PX, was designed and implemented around 1985 by S. Hayashi from Kyoto University. It allows the extraction of a LISP program from a proof in a logical system inspired by the logical formalisms of S. Feferman. Interest in this methodology is growing in the theoretical computer science community. We can foresee the day when actual computer systems used in applications will contain certified modules, automatically generated from a consistency proof of their formal specifications. We are however still far from being able to use this methodology in a smooth interaction with the standard tools from software engineering, i.e. compilers, linkers, run-time systems taking advantage of special hardware, debuggers, and the like. We hope that COQ can be of use to researchers interested in experimenting with this new methodology.

A first implementation of CoC was started in 1984 by G. Huet and T. Coquand. Its implementation language was CAML, a functional programming language from the ML family designed at INRIA in Rocquencourt. The core of this system was a proof-checker for CoC seen as a typed λ -calculus, called the *Constructive Engine*. This engine was operated through a high-level notation permitting the declaration of axioms and parameters, the definition of mathematical types and objects, and the explicit construction of proof objects encoded as λ -terms. A section mechanism, designed and implemented by G. Dowek, allowed hierarchical developments of mathematical theories. This high-level language was called the *Mathematical Vernacular*. Furthermore, an interactive *Theorem Prover* permitted the incremental construction of proof trees in a top-down manner, subgoalng recursively and backtracking from dead-alleys. The theorem prover executed tactics written in CAML, in the LCF fashion. A basic set of tactics was predefined, which the user could extend by his own specific tactics. This system (Version 4.10) was released in 1989. Then, the system was extended to deal with the new calculus with inductive types by C. Paulin, with corresponding new tactics for proofs by induction. A new standard set of tactics was streamlined, and the vernacular extended for tactics execution. A package to compile programs extracted from proofs to actual computer programs in CAML or some other functional language was designed and implemented by B. Werner. A new user-interface, relying on a CAML-X interface by D.

de Rauglaudre, was designed and implemented by A. Felty. It allowed operation of the theorem-prover through the manipulation of windows, menus, mouse-sensitive buttons, and other widgets. This system (Version 5.6) was released in 1991.

COQ was ported to the new implementation Caml-light of X. Leroy and D. Doligez by D. de Rauglaudre (Version 5.7) in 1992. A new version of COQ was then coordinated by C. Murthy, with new tools designed by C. Parent to prove properties of ML programs (this methodology is dual to program extraction) and a new user-interaction loop. This system (Version 5.8) was released in May 1993. A Centaur interface CTCOQ was then developed by Y. Bertot from the Croap project from INRIA-Sophia-Antipolis.

In parallel, G. Dowek and H. Herbelin developed a new proof engine, allowing the general manipulation of existential variables consistently with dependent types in an experimental version of COQ (V5.9).

The version V5.10 of COQ is based on a generic system for manipulating terms with binding operators due to Chet Murthy. A new proof engine allows the parallel development of partial proofs for independent subgoals. The structure of these proof trees is a mixed representation of derivation trees for the Calculus of Inductive Constructions with abstract syntax trees for the tactics scripts, allowing the navigation in a proof at various levels of details. The proof engine allows generic environment items managed in an object-oriented way. This new architecture, due to C. Murthy, supports several new facilities which make the system easier to extend and to scale up:

- User-programmable tactics are allowed
- It is possible to separately verify development modules, and to load their compiled images without verifying them again - a quick relocation process allows their fast loading
- A generic parsing scheme allows user-definable notations, with a symmetric table-driven pretty-printer
- Syntactic definitions allow convenient abbreviations
- A limited facility of meta-variables allows the automatic synthesis of certain type expressions, allowing generic notations for e.g. equality, pairing, and existential quantification.

In the Fall of 1994, C. Paulin-Mohring replaced the structure of inductively defined types and families by a new structure, allowing the mutually recursive definitions. P. Manoury implemented a translation of recursive definitions into the primitive recursive style imposed by the internal recursion operators, in the style of the ProPre system. C. Muñoz implemented a decision procedure for intuitionistic propositional logic, based on results of R. Dyckhoff. J.C. Filliâtre implemented a decision procedure for first-order logic without contraction, based on results of J. Ketonen and R. Weyhrauch. Finally C. Murthy implemented a library of inversion tactics, relieving the user from tedious definitions of “inversion predicates”.

Rocquencourt, Feb. 1st 1995
G rard Huet

Credits: addendum for version 6.1

The present version 6.1 of COQ is based on the V5.10 architecture. It was ported to the new language OCAML by Bruno Barras. The underlying framework has slightly changed and allows more conversions between sorts.

The new version provides powerful tools for easier developments.

Cristina Cornes designed an extension of the COQ syntax to allow definition of terms using a powerful pattern-matching analysis in the style of ML programs.

Amokrane Saïbi wrote a mechanism to simulate inheritance between types families extending a proposal by Peter Aczel. He also developed a mechanism to automatically compute which arguments of a constant may be inferred by the system and consequently do not need to be explicitly written.

Yann Coscoy designed a command which explains a proof term using natural language. Pierre Crégut built a new tactic which solves problems in quantifier-free Presburger Arithmetic. Both functionalities have been integrated to the COQ system by Hugo Herbelin.

Samuel Boutin designed a tactic for simplification of commutative rings using a canonical set of rewriting rules and equality modulo associativity and commutativity.

Finally the organisation of the COQ distribution has been supervised by Jean-Christophe Filliâtre with the help of Judicaël Courant and Bruno Barras.

Lyon, Nov. 18th 1996
Christine Paulin

Credits: addendum for version 6.2

In version 6.2 of COQ, the parsing is done using `camlp4`, a preprocessor and pretty-printer for CAML designed by Daniel de Rauglaudre at INRIA. Daniel de Rauglaudre made the first adaptation of COQ for `camlp4`, this work was continued by Bruno Barras who also changed the structure of COQ abstract syntax trees and the primitives to manipulate them. The result of these changes is a faster parsing procedure with greatly improved syntax-error messages. The user-interface to introduce grammar or pretty-printing rules has also changed.

Eduardo Giménez redesigned the internal tactic libraries, giving uniform names to Caml functions corresponding to COQ tactic names.

Bruno Barras wrote new more efficient reductions functions.

Hugo Herbelin introduced more uniform notations in the COQ specification language: the definitions by fixpoints and pattern-matching have a more readable syntax. Patrick Loiseleur introduced user-friendly notations for arithmetic expressions.

New tactics were introduced: Eduardo Giménez improved a mechanism to introduce macros for tactics, and designed special tactics for (co)inductive definitions; Patrick Loiseleur designed a tactic to simplify polynomial expressions in an arbitrary commutative ring which generalizes the previous tactic implemented by Samuel Boutin. Jean-Christophe Filliâtre introduced a tactic for refining a goal, using a proof term with holes as a proof scheme.

David Delahaye designed the `SearchIsos` tool to search an object in the library given its type (up to isomorphism).

Henri Laulhère produced the COQ distribution for the Windows environment.

Finally, Hugo Herbelin was the main coordinator of the COQ documentation with principal contributions by Bruno Barras, David Delahaye, Jean-Christophe Filliâtre, Eduardo Giménez, Hugo Herbelin and Patrick Loiseleur.

Orsay, May 4th 1998
Christine Paulin

Credits: addendum for version 6.3

The main changes in version V6.3 was the introduction of a few new tactics and the extension of the guard condition for fixpoint definitions.

B. Barras extended the unification algorithm to complete partial terms and solved various tricky bugs related to universes.

D. Delahaye developed the `AutoRewrite` tactic. He also designed the new behavior of `Intro` and provided the tacticals `First` and `Solve`.

J.-C. Filliâtre developed the `Correctness` tactic.

E. Giménez extended the guard condition in fixpoints.

H. Herbelin designed the new syntax for definitions and extended the `Induction` tactic.

P. Loiseleur developed the `Quote` tactic and the new design of the `Auto` tactic, he also introduced the index of errors in the documentation.

C. Paulin wrote the `Focus` command and introduced the reduction functions in definitions, this last feature was proposed by J.-F. Monin from CNET Lannion.

Orsay, Dec. 1999
Christine Paulin

Credits: versions 7

The version V7 is a new implementation started in September 1999 by Jean-Christophe Filliâtre. This is a major revision with respect to the internal architecture of the system. The COQ version 7.0 was distributed in March 2001, version 7.1 in September 2001, version 7.2 in January 2002, version 7.3 in May 2002 and version 7.4 in February 2003.

Jean-Christophe Filliâtre designed the architecture of the new system, he introduced a new representation for environments and wrote a new kernel for type-checking terms. His approach was to use functional data-structures in order to get more sharing, to prepare the addition of modules and also to get closer to a certified kernel.

Hugo Herbelin introduced a new structure of terms with local definitions. He introduced “qualified” names, wrote a new pattern-matching compilation algorithm and designed a more compact algorithm for checking the logical consistency of universes. He contributed to the simplification of COQ internal structures and the optimisation of the system. He added basic tactics for forward reasoning and coercions in patterns.

David Delahaye introduced a new language for tactics. General tactics using pattern-matching on goals and context can directly be written from the COQ toplevel. He also provided primitives for the design of user-defined tactics in CAML.

Micaela Mayero contributed the library on real numbers. Olivier Desmettre extended this library with axiomatic trigonometric functions, square, square roots, finite sums, Chasles property and basic plane geometry.

Jean-Christophe Filliâtre and Pierre Letouzey redesigned a new extraction procedure from COQ terms to CAML or HASKELL programs. This new extraction procedure, unlike the one implemented in previous version of COQ is able to handle all terms in the Calculus of Inductive Constructions, even involving universes and strong elimination. P. Letouzey adapted user contributions to extract ML programs when it was sensible. Jean-Christophe Filliâtre wrote `coqdoc`, a documentation tool for COQ libraries usable from version 7.2.

Bruno Barras improved the reduction algorithms efficiency and the confidence level in the correctness of COQ critical type-checking algorithm.

Yves Bertot designed the `SearchPattern` and `SearchRewrite` tools and the support for the PCOQ interface (<http://www-sop.inria.fr/lemme/pcoq/>).

Micaela Mayero and David Delahaye introduced `Field`, a decision tactic for commutative fields.

Christine Paulin changed the elimination rules for empty and singleton propositional inductive types.

Loïc Pottier developed `Fourier`, a tactic solving linear inequalities on real numbers.

Pierre Crégut developed a new version based on reflexion of the `Omega` decision tactic.

Claudio Sacerdoti Coen designed an XML output for the COQ modules to be used in the Hypertextual Electronic Library of Mathematics (HELM cf <http://www.cs.unibo.it/helm>).

A library for efficient representation of finite maps using binary trees contributed by Jean Goubault was integrated in the basic theories.

Pierre Courtieu developed a command and a tactic to reason on the inductive structure of recursively defined functions.

Jacek Chrząszcz designed and implemented the module system of COQ whose foundations are in Judicaël Courant's PhD thesis.

The development was coordinated by C. Paulin.

Many discussions within the *Démons* team and the *LogiCal* project influenced significantly the design of COQ especially with J. Courant, J. Duprat, J. Goubault, A. Miquel, C. Marché, B. Monate and B. Werner.

Intensive users suggested improvements of the system : Y. Bertot, L. Pottier, L. Théry, P. Zimmerman from INRIA, C. Alvarado, P. Crégut, J.-F. Monin from France Telecom R & D.

Orsay, May. 2002

Hugo Herbelin & Christine Paulin

Credits: version 8.0

COQ version 8 is a major revision of the COQ proof assistant. First, the underlying logic is slightly different. The so-called *impredicativity* of the sort `Set` has been dropped. The main reason is that it is inconsistent with the principle of description which is quite a useful principle for formalizing mathematics within classical logic. Moreover, even in an constructive setting, the impredicativity of `Set` does not add so much in practice and is even subject of criticism from a large part of the intuitionistic mathematician community. Nevertheless, the impredicativity of `Set` remains optional for users interested in investigating mathematical developments which rely on it.

Secondly, the concrete syntax of terms has been completely revised. The main motivations were

- a more uniform, purified style: all constructions are now lowercase, with a functional programming perfume (e.g. abstraction is now written `fun`), and more directly accessible to the novice (e.g. dependent product is now written `forall` and allows omission of types). Also, parentheses and are no longer mandatory for function application.
- extensibility: some standard notations (e.g. “<” and “>”) were incompatible with the previous syntax. Now all standard arithmetic notations (`=`, `+`, `*`, `/`, `<`, `<=`, ... and more) are directly part of the syntax.

Together with the revision of the concrete syntax, a new mechanism of *interpretation scopes* permits to reuse the same symbols (typically $+$, $-$, $*$, $/$, $<$, $<=$) in various mathematical theories without any ambiguities for COQ, leading to a largely improved readability of COQ scripts. New commands to easily add new symbols are also provided.

Coming with the new syntax of terms, a slight reform of the tactic language and of the language of commands has been carried out. The purpose here is a better uniformity making the tactics and commands easier to use and to remember.

Thirdly, a restructuration and uniformisation of the standard library of COQ has been performed. There is now just one Leibniz' equality usable for all the different kinds of COQ objects. Also, the set of real numbers now lies at the same level as the sets of natural and integer numbers. Finally, the names of the standard properties of numbers now follow a standard pattern and the symbolic notations for the standard definitions as well.

The fourth point is the release of COQIDE, a new graphical gtk2-based interface fully integrated to COQ. Close in style from the Proof General Emacs interface, it is faster and its integration with COQ makes interactive developments more friendly. All mathematical Unicode symbols are usable within COQIDE.

Finally, the module system of COQ completes the picture of COQ version 8.0. Though released with an experimental status in the previous version 7.4, it should be considered as a salient feature of the new version.

Besides, COQ comes with its load of novelties and improvements: new or improved tactics (including a new tactic for solving first-order statements), new management commands, extended libraries.

Bruno Barras and Hugo Herbelin have been the main contributors of the reflexion and the implementation of the new syntax. The smart automatic translator from old to new syntax released with COQ is also their work with contributions by Olivier Desmettre.

Hugo Herbelin is the main designer and implementor of the notion of interpretation scopes and of the commands for easily adding new notations.

Hugo Herbelin is the main implementor of the restructuration of the standard library.

Pierre Corbineau is the main designer and implementor of the new tactic for solving first-order statements in presence of inductive types. He is also the maintainer of the non-domain specific automation tactics.

Benjamin Monate is the developer of the COQIDE graphical interface with contributions by Jean-Christophe Filliâtre, Pierre Letouzey, Claude Marché and Bruno Barras.

Claude Marché coordinated the edition of the Reference Manual for COQ V8.0.

Pierre Letouzey and Jacek Chrząszcz respectively maintained the extraction tool and module system of COQ.

Jean-Christophe Filliâtre, Pierre Letouzey, Hugo Herbelin and contributors from Sophia-Antipolis and Nijmegen participated to the extension of the library.

Julien Narboux built a NSIS-based automatic COQ installation tool for the Windows platform.

Hugo Herbelin and Christine Paulin coordinated the development which was under the responsibility of Christine Paulin.

Palaiseau & Orsay, Apr. 2004
Hugo Herbelin & Christine Paulin
(updated Apr. 2006)

Credits: version 8.1

COQ version 8.1 adds various new functionalities.

Benjamin Grégoire implemented an alternative algorithm to check the convertibility of terms in the COQ type-checker. This alternative algorithm works by compilation to an efficient bytecode that is interpreted in an abstract machine similar to Xavier Leroy's ZINC machine. Convertibility is performed by comparing the normal forms. This alternative algorithm is specifically interesting for proofs by reflection. More generally, it is convenient in case of intensive computations.

Christine Paulin implemented an extension of inductive types allowing recursively non uniform parameters. Hugo Herbelin implemented sort-polymorphism for inductive types (now called template polymorphism).

Claudio Sacerdoti Coen improved the tactics for rewriting on arbitrary compatible equivalence relations. He also generalized rewriting to arbitrary transition systems.

Claudio Sacerdoti Coen added new features to the module system.

Benjamin Grégoire, Assia Mahboubi and Bruno Barras developed a new more efficient and more general simplification algorithm on rings and semi-rings.

Laurent Théry and Bruno Barras developed a new significantly more efficient simplification algorithm on fields.

Hugo Herbelin, Pierre Letouzey, Julien Forest, Julien Narboux and Claudio Sacerdoti Coen added new tactic features.

Hugo Herbelin implemented matching on disjunctive patterns.

New mechanisms made easier the communication between COQ and external provers. Nicolas Ayache and Jean-Christophe Filliâtre implemented connections with the provers CVCL, SIMPLIFY and ZENON. Hugo Herbelin implemented an experimental protocol for calling external tools from the tactic language.

Matthieu Sozeau developed RUSSELL, an experimental language to specify the behavior of programs with subtypes.

A mechanism to automatically use some specific tactic to solve unresolved implicit has been implemented by Hugo Herbelin.

Laurent Théry's contribution on strings and Pierre Letouzey and Jean-Christophe Filliâtre's contribution on finite maps have been integrated to the COQ standard library. Pierre Letouzey developed a library about finite sets "à la OCAML". With Jean-Marc Notin, he extended the library on lists. Pierre Letouzey's contribution on rational numbers has been integrated and extended..

Pierre Corbineau extended his tactic for solving first-order statements. He wrote a reflection-based intuitionistic tautology solver.

Pierre Courtieu, Julien Forest and Yves Bertot added extra support to reason on the inductive structure of recursively defined functions.

Jean-Marc Notin significantly contributed to the general maintenance of the system. He also took care of `coqdoc`.

Pierre Castéran contributed to the documentation of (co-)inductive types and suggested improvements to the libraries.

Pierre Corbineau implemented a declarative mathematical proof language, usable in combination with the tactic-based style of proof.

Finally, many users suggested improvements of the system through the Coq-Club mailing list and bug-tracker systems, especially user groups from INRIA Rocquencourt, Radboud University, University of Pennsylvania and Yale University.

Palaiseau, July 2006

Hugo Herbelin

Credits: version 8.2

COQ version 8.2 adds new features, new libraries and improves on many various aspects.

Regarding the language of Coq, the main novelty is the introduction by Matthieu Sozeau of a package of commands providing Haskell-style type classes. Type classes, that come with a few convenient features such as type-based resolution of implicit arguments, plays a new role of landmark in the architecture of Coq with respect to automatization. For instance, thanks to type classes support, Matthieu Sozeau could implement a new resolution-based version of the tactics dedicated to rewriting on arbitrary transitive relations.

Another major improvement of Coq 8.2 is the evolution of the arithmetic libraries and of the tools associated to them. Benjamin Grégoire and Laurent Théry contributed a modular library for building arbitrarily large integers from bounded integers while Evgeny Makarov contributed a modular library of abstract natural and integer arithmetics together with a few convenient tactics. On his side, Pierre Letouzey made numerous extensions to the arithmetic libraries on \mathbb{Z} and \mathbb{Q} , including extra support for automatization in presence of various number-theory concepts.

Frédéric Besson contributed a reflexive tactic based on Krivine-Stengle Positivstellensatz (the easy way) for validating provability of systems of inequalities. The platform is flexible enough to support the validation of any algorithm able to produce a “certificate” for the Positivstellensatz and this covers the case of Fourier-Motzkin (for linear systems in \mathbb{Q} and \mathbb{R}), Fourier-Motzkin with cutting planes (for linear systems in \mathbb{Z}) and sum-of-squares (for non-linear systems). Evgeny Makarov made the platform generic over arbitrary ordered rings.

Arnaud Spiwack developed a library of 31-bits machine integers and, relying on Benjamin Grégoire and Laurent Théry’s library, delivered a library of unbounded integers in base 2^{31} . As importantly, he developed a notion of “retro-knowledge” so as to safely extend the kernel-located bytecode-based efficient evaluation algorithm of Coq version 8.1 to use 31-bits machine arithmetics for efficiently computing with the library of integers he developed.

Beside the libraries, various improvements contributed to provide a more comfortable end-user language and more expressive tactic language. Hugo Herbelin and Matthieu Sozeau improved the pattern-matching compilation algorithm (detection of impossible clauses in pattern-matching, automatic inference of the return type). Hugo Herbelin, Pierre Letouzey and Matthieu Sozeau contributed various new convenient syntactic constructs and new tactics or tactic features: more inference of redundant information, better unification, better support for proof or definition by fixpoint, more expressive rewriting tactics, better support for meta-variables, more convenient notations, ...

Élie Soubiran improved the module system, adding new features (such as an “include” command) and making it more flexible and more general. He and Pierre Letouzey improved the support for modules in the extraction mechanism.

Matthieu Sozeau extended the RUSSELL language, ending in a convenient way to write programs of given specifications, Pierre Corbineau extended the Mathematical Proof Language and the automatization tools that accompany it, Pierre Letouzey supervised and extended various parts the standard library, Stéphane Glondou contributed a few tactics and improvements, Jean-Marc Notin provided help in debugging, general maintenance and `coqdoc` support, Vincent Siles contributed extensions of the `Scheme` command and of `injection`.

Bruno Barras implemented the `coqchk` tool: this is a stand-alone type-checker that can be used to certify `.vo` files. Especially, as this verifier runs in a separate process, it is granted not to be “hijacked” by virtually malicious extensions added to COQ.

Yves Bertot, Jean-Christophe Filliâtre, Pierre Courtieu and Julien Forest acted as maintainers of features they implemented in previous versions of Coq.

Julien Narboux contributed to COQIDE. Nicolas Tabareau made the adaptation of the interface of the old “setoid rewrite” tactic to the new version. Lionel Mamane worked on the interaction between Coq and its external interfaces. With Samuel Mimram, he also helped making Coq compatible with recent software tools. Russell O’Connor, Cezary Kaliszyk, Milad Niqui contributed to improved the libraries of integers, rational, and real numbers. We also thank many users and partners for suggestions and feedback, in particular Pierre Castéran and Arthur Charguéraud, the INRIA Marelle team, Georges Gonthier and the INRIA-Microsoft Mathematical Components team, the Foundations group at Radboud university in Nijmegen, reporters of bugs and participants to the Coq-Club mailing list.

Palaiseau, June 2008
Hugo Herbelin

Credits: version 8.3

COQ version 8.3 is before all a transition version with refinements or extensions of the existing features and libraries and a new tactic `nsatz` based on Hilbert’s Nullstellensatz for deciding systems of equations over rings.

With respect to libraries, the main evolutions are due to Pierre Letouzey with a rewriting of the library of finite sets `FSets` and a new round of evolutions in the modular development of arithmetic (library `Numbers`). The reason for making `FSets` evolve is that the computational and logical contents were quite intertwined in the original implementation, leading in some cases to longer computations than expected and this problem is solved in the new `MSets` implementation. As for the modular arithmetic library, it was only dealing with the basic arithmetic operators in the former version and its current extension adds the standard theory of the division, min and max functions, all made available for free to any implementation of \mathbb{N} , \mathbb{Z} or $\mathbb{Z}/n\mathbb{Z}$.

The main other evolutions of the library are due to Hugo Herbelin who made a revision of the sorting library (including a certified merge-sort) and to Guillaume Melquiond who slightly revised and cleaned up the library of reals.

The module system evolved significantly. Besides the resolution of some efficiency issues and a more flexible construction of module types, Élie Soubiran brought a new model of name equivalence, the Δ -equivalence, which respects as much as possible the names given by the users. He also designed with Pierre Letouzey a new convenient operator `<+` for nesting functor application, what provides a light notation for inheriting the properties of cascading modules.

The new tactic `nsatz` is due to Loïc Pottier. It works by computing Gröbner bases. Regarding the existing tactics, various improvements have been done by Matthieu Sozeau, Hugo Herbelin and Pierre Letouzey.

Matthieu Sozeau extended and refined the type classes and `Program` features (the `RUSSELL` language). Pierre Letouzey maintained and improved the extraction mechanism. Bruno Barras and Élie Soubiran maintained the Coq checker, Julien Forest maintained the `Function` mechanism for reasoning over recursively defined functions. Matthieu Sozeau, Hugo Herbelin and Jean-Marc Notin maintained `coqdoc`. Frédéric Besson maintained the `MICROMEGA` platform for deciding systems of inequalities. Pierre Courtieu maintained the support for the Proof General Emacs interface. Claude Marché maintained the plugin for calling external provers (`dp`). Yves Bertot made some improvements to the libraries of lists and integers. Matthias Puech improved the search functions. Guillaume Melquiond usefully contributed here and there. Yann Régis-Gianas grounded the support for Unicode on a more standard and more robust basis.

Though invisible from outside, Arnaud Spiwack improved the general process of management of existential variables. Pierre Letouzey and Stéphane Glondou improved the compilation scheme of the Coq archive. Vincent Gross provided support to COQIDE. Jean-Marc Notin provided support for benchmarking and archiving.

Many users helped by reporting problems, providing patches, suggesting improvements or making useful comments, either on the bug tracker or on the Coq-club mailing list. This includes but not exhaustively Cédric Auger, Arthur Charguéraud, François Garillot, Georges Gonthier, Robin Green, Stéphane Lescuyer, Eelis van der Weegen, ...

Though not directly related to the implementation, special thanks are going to Yves Bertot, Pierre Castéran, Adam Chlipala, and Benjamin Pierce for the excellent teaching materials they provided.

Paris, April 2010
Hugo Herbelin

Credits: version 8.4

COQ version 8.4 contains the result of three long-term projects: a new modular library of arithmetic by Pierre Letouzey, a new proof engine by Arnaud Spiwack and a new communication protocol for COQIDE by Vincent Gross.

The new modular library of arithmetic extends, generalizes and unifies the existing libraries on Peano arithmetic (types `nat`, `N` and `BigN`), positive arithmetic (type `positive`), integer arithmetic (`Z` and `BigZ`) and machine word arithmetic (type `Int31`). It provides with unified notations (e.g. systematic use of `add` and `mul` for denoting the addition and multiplication operators), systematic and generic development of operators and properties of these operators for all the types mentioned above, including `gcd`, `pcm`, `power`, square root, base 2 logarithm, division, modulo, bitwise operations, logical shifts, comparisons, iterators, ...

The most visible feature of the new proof engine is the support for structured scripts (bullets and proof brackets) but, even if yet not user-available, the new engine also provides the basis for refining existential variables using tactics, for applying tactics to several goals simultaneously, for reordering goals, all features which are planned for the next release. The new proof engine forced to reimplement `info` and `Show Script` differently, what was done by Pierre Letouzey.

Before version 8.4, COQIDE was linked to COQ with the graphical interface living in a separate thread. From version 8.4, COQIDE is a separate process communicating with COQ through a textual channel. This allows for a more robust interfacing, the ability to interrupt COQ without interrupting the interface, and the ability to manage several sessions in parallel. Relying on the infrastructure work made by Vincent Gross, Pierre Letouzey, Pierre Boutillier and Pierre-Marie Pédrot contributed many various refinements of COQIDE.

COQ 8.4 also comes with a bunch of many various smaller-scale changes and improvements regarding the different components of the system.

The underlying logic has been extended with η -conversion thanks to Hugo Herbelin, Stéphane Glondou and Benjamin Grégoire. The addition of η -conversion is justified by the confidence that the formulation of the Calculus of Inductive Constructions based on typed equality (such as the one considered in Lee and Werner to build a set-theoretic model of CIC [97]) is applicable to the concrete implementation of COQ.

The underlying logic benefited also from a refinement of the guard condition for fixpoints by Pierre Boutillier, the point being that it is safe to propagate the information about structurally smaller arguments through β -redexes that are blocked by the “match” construction (blocked commutative cuts).

Relying on the added permissiveness of the guard condition, Hugo Herbelin could extend the pattern-matching compilation algorithm so that matching over a sequence of terms involving dependencies of a term or of the indices of the type of a term in the type of other terms is systematically supported.

Regarding the high-level specification language, Pierre Boutillier introduced the ability to give implicit arguments to anonymous functions, Hugo Herbelin introduced the ability to define notations with several binders (e.g. `exists x y z, P`), Matthieu Sozeau made the type classes inference mechanism more robust and predictable, Enrico Tassi introduced a command `Arguments` that generalizes `Implicit Arguments` and `Arguments Scope` for assigning various properties to arguments of constants. Various improvements in the type inference algorithm were provided by Matthieu Sozeau and Hugo Herbelin with contributions from Enrico Tassi.

Regarding tactics, Hugo Herbelin introduced support for referring to expressions occurring in the goal by pattern in tactics such as `set` or `destruct`. Hugo Herbelin also relied on ideas from Chung-Kil Hur's `Heq` plugin to introduce automatic computation of occurrences to generalize when using `destruct` and `induction` on types with indices. Stéphane Glondou introduced new tactics `constr_eq`, `is_evar` and `has_evar` to be used when writing complex tactics. Enrico Tassi added support to fine-tuning the behavior of `simpl`. Enrico Tassi added the ability to specify over which variables of a section a lemma has to be exactly generalized. Pierre Letouzey added a tactic `timeout` and the interruptibility of `vm_compute`. Bug fixes and miscellaneous improvements of the tactic language came from Hugo Herbelin, Pierre Letouzey and Matthieu Sozeau.

Regarding decision tactics, Loïc Pottier maintained `Nsatz`, moving in particular to a type-class based reification of goals while Frédéric Besson maintained `Micromega`, adding in particular support for division.

Regarding vernacular commands, Stéphane Glondou provided new commands to analyze the structure of type universes.

Regarding libraries, a new library about lists of a given length (called vectors) has been provided by Pierre Boutillier. A new instance of finite sets based on Red-Black trees and provided by Andrew Appel has been adapted for the standard library by Pierre Letouzey. In the library of real analysis, Yves Bertot changed the definition of π and provided a proof of the long-standing fact yet remaining unproved in this library, namely that $\sin \frac{\pi}{2} = 1$.

Pierre Corbineau maintained the Mathematical Proof Language (C-zar).

Bruno Barras and Benjamin Grégoire maintained the call-by-value reduction machines.

The extraction mechanism benefited from several improvements provided by Pierre Letouzey.

Pierre Letouzey maintained the module system, with contributions from Élie Soubiran.

Julien Forest maintained the `Function` command.

Matthieu Sozeau maintained the setoid rewriting mechanism.

COQ related tools have been upgraded too. In particular, `coq_makefile` has been largely revised by Pierre Boutillier. Also, patches from Adam Chlipala for `coqdoc` have been integrated by Pierre Boutillier.

Bruno Barras and Pierre Letouzey maintained the `coqchk` checker.

Pierre Courtieu and Arnaud Spiwack contributed new features for using COQ through Proof General.

The `Dp` plugin has been removed. Use the plugin provided with `Why3` instead (<http://why3.lri.fr>).

Under the hood, the COQ architecture benefited from improvements in terms of efficiency and robustness, especially regarding universes management and existential variables management, thanks to Pierre Letouzey and Yann Régis-Gianas with contributions from Stéphane Glondou and Matthias Puech. The build system is maintained by Pierre Letouzey with contributions from Stéphane Glondou and Pierre Boutillier.

A new backtracking mechanism simplifying the task of external interfaces has been designed by Pierre Letouzey.

The general maintenance was done by Pierre Letouzey, Hugo Herbelin, Pierre Boutillier, Matthieu Sozeau and Stéphane Glondou with also specific contributions from Guillaume Melquiond, Julien Narboux and Pierre-Marie Pédrot.

Packaging tools were provided by Pierre Letouzey (Windows), Pierre Boutillier (MacOS), Stéphane Glondou (Debian). Releasing, testing and benchmarking support was provided by Jean-Marc Notin.

Many suggestions for improvements were motivated by feedback from users, on either the bug tracker or the coq-club mailing list. Special thanks are going to the users who contributed patches, starting with Tom Prince. Other patch contributors include Cédric Auger, David Baelde, Dan Grayson, Paolo Herms, Robbert Krebbers, Marc Lasson, Hendrik Tews and Eelis van der Weegen.

Paris, December 2011
Hugo Herbelin

Credits: version 8.5

COQ version 8.5 contains the result of five specific long-term projects:

- A new asynchronous evaluation and compilation mode by Enrico Tassi with help from Bruno Barras and Carst Tankink.
- Full integration of the new proof engine by Arnaud Spiwack helped by Pierre-Marie Pédrot,
- Addition of conversion and reduction based on native compilation by Maxime Dénès and Benjamin Grégoire.
- Full universe polymorphism for definitions and inductive types by Matthieu Sozeau.
- An implementation of primitive projections with η -conversion bringing significant performance improvements when using records by Matthieu Sozeau.

The full integration of the proof engine, by Arnaud Spiwack and Pierre-Marie Pédrot, brings to primitive tactics and the user level Ltac language dependent subgoals, deep backtracking and multiple goal handling, along with miscellaneous features and an improved potential for future modifications. Dependent subgoals allow statements in a goal to mention the proof of another. Proofs of unsolved subgoals appear as existential variables. Primitive backtracking makes it possible to write a tactic with several possible outcomes which are tried successively when subsequent tactics fail. Primitives are also available to control the backtracking behavior of tactics. Multiple goal handling paves the way for smarter automation tactics. It is currently used for simple goal manipulation such as goal reordering.

The way COQ processes a document in batch and interactive mode has been redesigned by Enrico Tassi with help from Bruno Barras. Opaque proofs, the text between Proof and Qed, can be processed asynchronously, decoupling the checking of definitions and statements from the checking of proofs. It improves the responsiveness of interactive development, since proofs can be processed in the background. Similarly, compilation of a file can be split into two phases: the first one checking only definitions and statements and the second one checking proofs. A file resulting from the first phase – with the .vio extension – can be already Required. All .vio files can be turned into complete .vo files in parallel. The same infrastructure also allows terminating tactics to be run in parallel on a set of goals via the `par` goal selector.

COQIDE was modified to cope with asynchronous checking of the document. Its source code was also made separate from that of COQ, so that COQIDE no longer has a special status among user interfaces, paving the way for decoupling its release cycle from that of COQ in the future.

Carst Tankink developed a COQ back-end for user interfaces built on Makarius Wenzel's Prover IDE framework (PIDE), like PIDE/jEdit (with help from Makarius Wenzel) or PIDE/Coqoon (with help from Alexander Faithfull and Jesper Bengtson). The development of such features was funded by the Paral-ITP French ANR project.

The full universe polymorphism extension was designed by Matthieu Sozeau. It conservatively extends the universes system and core calculus with definitions and inductive declarations parameterized by universes and constraints. It is based on a modification of the kernel architecture to handle constraint checking only, leaving the generation of constraints to the refinement/type inference engine. Accordingly, tactics are now fully universe aware, resulting in more localized error messages in case of inconsistencies and allowing higher-level algorithms like unification to be entirely type safe. The internal representation of universes has been modified but this is invisible to the user.

The underlying logic has been extended with η -conversion for records defined with primitive projections by Matthieu Sozeau. This additional form of η -conversion is justified using the same principle than the previously added η -conversion for function types, based on formulations of the Calculus of Inductive Constructions with typed equality. Primitive projections, which do not carry the parameters of the record and are rigid names (not defined as a pattern-matching construct), make working with nested records more manageable in terms of time and space consumption. This extension and universe polymorphism were carried out partly while Matthieu Sozeau was working at the IAS in Princeton.

The guard condition has been made compliant with extensional equality principles such as propositional extensionality and univalence, thanks to Maxime Dénès and Bruno Barras. To ensure compatibility with the univalence axiom, a new flag “-indices-matter” has been implemented, taking into account the universe levels of indices when computing the levels of inductive types. This supports using COQ as a tool to explore the relations between homotopy theory and type theory.

Maxime Dénès and Benjamin Grégoire developed an implementation of conversion test and normal form computation using the OCaml native compiler. It complements the virtual machine conversion offering much faster computation for expensive functions.

COQ 8.5 also comes with a bunch of many various smaller-scale changes and improvements regarding the different components of the system. We shall only list a few of them.

Pierre Boutillier developed an improved tactic for simplification of expressions called `cbn`.

Maxime Dénès maintained the bytecode-based reduction machine. Pierre Letouzey maintained the extraction mechanism.

Pierre-Marie Pédrot has extended the syntax of terms to, experimentally, allow holes in terms to be solved by a locally specified tactic.

Existential variables are referred to by identifiers rather than mere numbers, thanks to Hugo Herbelin who also improved the tactic language here and there.

Error messages for universe inconsistencies have been improved by Matthieu Sozeau. Error messages for unification and type inference failures have been improved by Hugo Herbelin, Pierre-Marie Pédrot and Arnaud Spiwack.

Pierre Courtieu contributed new features for using COQ through Proof General and for better interactive experience (bullets, Search, etc).

The efficiency of the whole system has been significantly improved thanks to contributions from Pierre-Marie Pédrot.

A distribution channel for COQ packages using the OPAM tool has been initiated by Thomas Braibant and developed by Guillaume Claret, with contributions by Enrico Tassi and feedback from

Hugo Herbelin.

Packaging tools were provided by Pierre Letouzey and Enrico Tassi (Windows), Pierre Boutillier, Matthieu Sozeau and Maxime Dénès (MacOS X). Maxime Dénès improved significantly the testing and benchmarking support.

Many power users helped to improve the design of the new features via the bug tracker, the coq development mailing list or the coq-club mailing list. Special thanks are going to the users who contributed patches and intensive brain-storming, starting with Jason Gross, Jonathan Leivent, Greg Malecha, Clément Pit-Claudel, Marc Lasson, Lionel Rieg. It would however be impossible to mention with precision all names of people who to some extent influenced the development.

Version 8.5 is one of the most important release of COQ. Its development spanned over about 3 years and a half with about one year of beta-testing. General maintenance during part or whole of this period has been done by Pierre Boutillier, Pierre Courtieu, Maxime Dénès, Hugo Herbelin, Pierre Letouzey, Guillaume Melquiond, Pierre-Marie Pédro, Matthieu Sozeau, Arnaud Spiwack, Enrico Tassi as well as Bruno Barras, Yves Bertot, Frédéric Besson, Xavier Clerc, Pierre Corbineau, Jean-Christophe Filliâtre, Julien Forest, Sébastien Hinderer, Assia Mahboubi, Jean-Marc Notin, Yann Régis-Gianas, François Ripault, Carst Tankink. Maxime Dénès coordinated the release process.

Paris, January 2015, revised December 2015,
Hugo Herbelin, Matthieu Sozeau and the COQ development team

Credits: version 8.6

COQ version 8.6 contains the result of refinements, stabilization of 8.5's features and cleanups of the internals of the system. Over the year of (now time-based) development, about 450 bugs were resolved and over 100 contributions integrated. The main user visible changes are:

- A new, faster state-of-the-art universe constraint checker, by Jacques-Henri Jourdan.
- In CoqIDE and other asynchronous interfaces, more fine-grained asynchronous processing and error reporting by Enrico Tassi, making COQ capable of recovering from errors and continue processing the document.
- More access to the proof engine features from Ltac: goal management primitives, range selectors and a `typeclasses eauto` engine handling multiple goals and multiple successes, by Cyprien Mangin, Matthieu Sozeau and Arnaud Spiwack.
- Tactic behavior uniformization and specification, generalization of intro-patterns by Hugo Herbelin and others.
- A brand new warning system allowing to control warnings, turn them into errors or ignore them selectively by Maxime Dénès, Guillaume Melquiond, Pierre-Marie Pédro and others.
- Irrefutable patterns in abstractions, by Daniel de Rauglaudre.
- The `ssreflect` subterm selection algorithm by Georges Gonthier and Enrico Tassi is now accessible to tactic writers through the `ssmatching` plugin.
- Integration of `LtacProf`, a profiler for Ltac by Jason Gross, Paul Steckler, Enrico Tassi and Tobias Tebbi.

COQ 8.6 also comes with a bunch of smaller-scale changes and improvements regarding the different components of the system. We shall only list a few of them.

The `iota` reduction flag is now a shorthand for `match`, `fix` and `cofix` flags controlling the corresponding reduction rules (by Hugo Herbelin and Maxime Dénès).

Maxime Dénès maintained the native compilation machinery.

Pierre-Marie Pédrot separated the Ltac code from general purpose tactics, and generalized and rationalized the handling of generic arguments, allowing to create new versions of Ltac more easily in the future.

In patterns and terms, `@`, abbreviations and notations are now interpreted the same way, by Hugo Herbelin.

Name handling for universes has been improved by Pierre-Marie Pédrot and Matthieu Sozeau. The minimization algorithm has been improved by Matthieu Sozeau.

The unifier has been improved by Hugo Herbelin and Matthieu Sozeau, fixing some incompatibilities introduced in Coq 8.5. Unification constraints can now be left floating around and be seen by the user thanks to a new option. The `Keyed Unification` mode has been improved by Matthieu Sozeau.

The typeclass resolution engine and associated proof-search tactic have been reimplemented on top of the proof-engine monad, providing better integration in tactics, and new options have been introduced to control it, by Matthieu Sozeau with help from Théo Zimmermann.

The efficiency of the whole system has been significantly improved thanks to contributions from Pierre-Marie Pédrot, Maxime Dénès and Matthieu Sozeau and performance issue tracking by Jason Gross and Paul Steckler.

Standard library improvements by Jason Gross, Sébastien Hinderer, Pierre Letouzey and others.

Emilio Jesús Gallego Arias contributed many cleanups and refactorings of the pretty-printing and user interface communication components.

Frédéric Besson maintained the micromega tactic.

The OPAM repository for COQ packages has been maintained by Guillaume Claret, Guillaume Melquiond, Matthieu Sozeau, Enrico Tassi and others. A list of packages is now available at <https://coq.inria.fr/opam/www/>.

Packaging tools and software development kits were prepared by Michael Soegtrop with the help of Maxime Dénès and Enrico Tassi for Windows, and Maxime Dénès and Matthieu Sozeau for MacOS X. Packages are now regularly built on the continuous integration server. COQ now comes with a `META` file usable with `ocamlfind`, contributed by Emilio Jesús Gallego Arias, Gregory Malecha, and Matthieu Sozeau.

Matej Košík maintained and greatly improved the continuous integration setup and the testing of COQ contributions. He also contributed many API improvement and code cleanups throughout the system.

The contributors for this version are Bruno Barras, C.J. Bell, Yves Bertot, Frédéric Besson, Pierre Boutillier, Tej Chajed, Guillaume Claret, Xavier Clerc, Pierre Corbineau, Pierre Courtieu, Maxime Dénès, Ricky Elrod, Emilio Jesús Gallego Arias, Jason Gross, Hugo Herbelin, Sébastien Hinderer, Jacques-Henri Jourdan, Matej Kosik, Xavier Leroy, Pierre Letouzey, Gregory Malecha, Cyprien Mangin, Erik Martin-Dorel, Guillaume Melquiond, Clément Pit-Claudel, Pierre-Marie Pédrot, Daniel de Rauglaudre, Lionel Rieg, Gabriel Scherer, Thomas Sibut-Pinote, Matthieu Sozeau, Arnaud Spiwack, Paul Steckler, Enrico Tassi, Laurent Théry, Nickolai Zeldovich and Théo Zimmermann. The development process was coordinated by Hugo Herbelin and Matthieu Sozeau with the help of Maxime Dénès, who was also in charge of the release process.

Many power users helped to improve the design of the new features via the bug tracker, the pull request system, the COQ development mailing list or the coq-club mailing list. Special thanks to the users

who contributed patches and intensive brain-storming and code reviews, starting with Cyril Cohen, Jason Gross, Robbert Krebbers, Jonathan Leivent, Xavier Leroy, Gregory Malecha, Clément Pit-Claudel, Gabriel Scherer and Beta Ziliani. It would however be impossible to mention exhaustively the names of everybody who to some extent influenced the development.

Version 8.6 is the first release of COQ developed on a time-based development cycle. Its development spanned 10 months from the release of COQ 8.5 and was based on a public roadmap. To date, it contains more external contributions than any previous COQ system. Code reviews were systematically done before integration of new features, with an important focus given to compatibility and performance issues, resulting in a hopefully more robust release than COQ 8.5.

Coq Enhancement Proposals (CEPs for short) were introduced by Enrico Tassi to provide more visibility and a discussion period on new features, they are publicly available <https://github.com/coq/ceps>.

Started during this period, an effort is led by Yves Bertot and Maxime Dénès to put together a COQ consortium.

Paris, November 2016,
Matthieu Sozeau and the COQ development team

Credits: version 8.7

COQ version 8.7 contains the result of refinements, stabilization of features and cleanups of the internals of the system along with a few new features. The main user visible changes are:

- New tactics: variants of tactics supporting existential variables `eassert`, `eenough`, etc... by Hugo Herbelin. Tactics `extensionality in H` and `inversion_sigma` by Jason Gross, `specialize with ...` accepting partial bindings by Pierre Courtieu.
- Cumulative Polymorphic Inductive Types, allowing cumulativity of universes to go through applied inductive types, by Amin Timany and Matthieu Sozeau.
- Integration of the `SSReflect` plugin and its documentation in the reference manual, by Enrico Tassi, Assia Mahboubi and Maxime Dénès.
- The `coq_makefile` tool was completely redesigned to improve its maintainability and the extensibility of generated Makefiles, and to make `_CoqProject` files more palatable to IDEs by Enrico Tassi.

COQ 8.7 involved a large amount of work on cleaning and speeding up the code base, notably the work of Pierre-Marie Pédrot on making the tactic-level system insensitive to existential variable expansion, providing a safer API to plugin writers and making the code more robust. The `dev/doc/changes.txt` file documents the numerous changes to the implementation and improvements of interfaces. An effort to provide an official, streamlined API to plugin writers is in progress, thanks to the work of Matej Košík.

Version 8.7 also comes with a bunch of smaller-scale changes and improvements regarding the different components of the system. We shall only list a few of them.

The efficiency of the whole system has been significantly improved thanks to contributions from Pierre-Marie Pédrot, Maxime Dénès and Matthieu Sozeau and performance issue tracking by Jason Gross and Paul Steckler.

Thomas Sibut-Pinote and Hugo Herbelin added support for side effects hooks in `cbv`, `cbn` and `simpl`. The side effects are provided via a plugin available at <https://github.com/herbelin/reduction-effects/>.

The `BigN`, `BigZ`, `BigQ` libraries are no longer part of the COQ standard library, they are now provided by a separate repository <https://github.com/coq/bignums>, maintained by Pierre Letouzey.

In the `Reals` library, `IZR` has been changed to produce a compact representation of integers and real constants are now represented using `IZR` (work by Guillaume Melquiond).

Standard library additions and improvements by Jason Gross, Pierre Letouzey and others, documented in the `CHANGES` file.

The mathematical proof language/declarative mode plugin was removed from the archive.

The OPAM repository for COQ packages has been maintained by Guillaume Melquiond, Matthieu Sozeau, Enrico Tassi with contributions from many users. A list of packages is available at <https://coq.inria.fr/opam/www/>.

Packaging tools and software development kits were prepared by Michael Soegtrop with the help of Maxime Dénès and Enrico Tassi for Windows, and Maxime Dénès for MacOS X. Packages are regularly built on the Travis continuous integration server.

The contributors for this version are Abhishek Anand, C.J. Bell, Yves Bertot, Frédéric Besson, Tej Chajed, Pierre Courtieu, Maxime Dénès, Julien Forest, Gaëtan Gilbert, Jason Gross, Hugo Herbelin, Emilio Jesús Gallego Arias, Ralf Jung, Matej Košík, Xavier Leroy, Pierre Letouzey, Assia Mahboubi, Cyprien Mangin, Erik Martin-Dorel, Olivier Marty, Guillaume Melquiond, Sam Pablo Kuper, Benjamin Pierce, Pierre-Marie Pédro, Lars Rasmusson, Lionel Rieg, Valentin Robert, Yann Régis-Gianas, Thomas Sibut-Pinote, Michael Soegtrop, Matthieu Sozeau, Arnaud Spiwack, Paul Steckler, George Stelle, Pierre-Yves Strub, Enrico Tassi, Hendrik Tews, Amin Timany, Laurent Théry, Vadim Zaliva and Théo Zimmermann.

The development process was coordinated by Matthieu Sozeau with the help of Maxime Dénès, who was also in charge of the release process. Théo Zimmermann is the maintainer of this release.

Many power users helped to improve the design of the new features via the bug tracker, the pull request system, the COQ development mailing list or the coq-club mailing list. Special thanks to the users who contributed patches and intensive brain-storming and code reviews, starting with Jason Gross, Ralf Jung, Robbert Krebbers, Xavier Leroy, Clément Pit-Claudel and Gabriel Scherer. It would however be impossible to mention exhaustively the names of everybody who to some extent influenced the development.

Version 8.7 is the second release of COQ developed on a time-based development cycle. Its development spanned 9 months from the release of COQ 8.6 and was based on a public road-map. It attracted many external contributions. Code reviews and continuous integration testing were systematically used before integration of new features, with an important focus given to compatibility and performance issues, resulting in a hopefully more robust release than COQ 8.6 while maintaining compatibility.

Coq Enhancement Proposals (CEPs for short) and open pull-requests discussions were used to discuss publicly the new features.

The COQ consortium, an organization directed towards users and supporters of the system, is now upcoming and will rely on Inria's newly created Foundation.

Paris, August 2017,
Matthieu Sozeau and the COQ development team

Table of contents

I	The language	39
1	The GALLINA specification language	41
1.1	Lexical conventions	41
1.2	Terms	43
1.2.1	Syntax of terms	43
1.2.2	Types	43
1.2.3	Qualified identifiers and simple identifiers	43
1.2.4	Numerals	43
1.2.5	Sorts	43
1.2.6	Binders	43
1.2.7	Abstractions	45
1.2.8	Products	46
1.2.9	Applications	46
1.2.10	Type cast	46
1.2.11	Inferable subterms	46
1.2.12	Let-in definitions	46
1.2.13	Definition by case analysis	46
1.2.14	Recursive functions	48
1.3	The Vernacular	48
1.3.1	Assumptions	49
1.3.2	Definitions	51
1.3.3	Inductive definitions	52
1.3.4	Definition of recursive functions	57
1.3.5	Assertions and proofs	61
2	Extensions of GALLINA	65
2.1	Record types	65
2.1.1	Primitive Projections	69
2.2	Variants and extensions of <code>match</code>	70
2.2.1	Multiple and nested pattern-matching	70
2.2.2	Pattern-matching on boolean values: the <code>if</code> expression	70
2.2.3	Irrefutable patterns: the destructuring <code>let</code> variants	71
2.2.4	Controlling pretty-printing of <code>match</code> expressions	72
2.2.5	Printing <code>match</code> templates	75
2.3	Advanced recursive functions	75
2.4	Section mechanism	77

2.4.1	Section <i>ident</i>	78
2.4.2	End <i>ident</i>	78
2.5	Module system	78
2.5.1	Module <i>ident</i>	79
2.5.2	End <i>ident</i>	80
2.5.3	Module <i>ident</i> := <i>module_expression</i>	80
2.5.4	Module Type <i>ident</i>	80
2.5.5	End <i>ident</i>	81
2.5.6	Module Type <i>ident</i> := <i>module_type</i>	81
2.5.7	Declare Module <i>ident</i> : <i>module_type</i>	81
2.5.8	Import <i>qualid</i>	84
2.5.9	Print Module <i>ident</i>	86
2.5.10	Print Module Type <i>ident</i>	86
2.5.11	Locate Module <i>qualid</i>	86
2.6	Libraries and qualified names	86
2.6.1	Names of libraries	86
2.6.2	Qualified names	86
2.6.3	Libraries and filesystem	87
2.7	Implicit arguments	88
2.7.1	The different kinds of implicit arguments	88
2.7.2	Maximal or non maximal insertion of implicit arguments	89
2.7.3	Casual use of implicit arguments	90
2.7.4	Declaration of implicit arguments	90
2.7.5	Automatic declaration of implicit arguments	92
2.7.6	Mode for automatic declaration of implicit arguments	94
2.7.7	Controlling strict implicit arguments	94
2.7.8	Controlling contextual implicit arguments	94
2.7.9	Controlling reversible-pattern implicit arguments	94
2.7.10	Controlling the insertion of implicit arguments not followed by explicit arguments	95
2.7.11	Explicit applications	95
2.7.12	Renaming implicit arguments	95
2.7.13	Displaying what the implicit arguments are	96
2.7.14	Explicit displaying of implicit arguments for pretty-printing	96
2.7.15	Interaction with subtyping	96
2.7.16	Deactivation of implicit arguments for parsing	97
2.7.17	Canonical structures	97
2.7.18	Implicit types of variables	98
2.7.19	Implicit generalization	99
2.8	Coercions	100
2.9	Printing constructions in full	100
2.10	Printing universes	101
2.11	Existential variables	101
2.11.1	Explicit displaying of existential instances for pretty-printing	102
2.11.2	Solving existential variables using tactics	103

3	The COQ library	105
3.1	The basic library	105
3.1.1	Notations	105
3.1.2	Logic	105
3.1.3	Datatypes	109
3.1.4	Specification	110
3.1.5	Basic Arithmetics	111
3.1.6	Well-founded recursion	113
3.1.7	Accessing the <code>Type</code> level	114
3.1.8	Tactics	114
3.2	The standard library	114
3.2.1	Survey	114
3.2.2	Notations for integer arithmetics	115
3.2.3	Peano's arithmetic (<code>nat</code>)	115
3.2.4	Real numbers library	116
3.2.5	List library	118
3.3	Users' contributions	119
4	Calculus of Inductive Constructions	121
4.1	The terms	121
4.1.1	Sorts	121
4.1.2	Terms	122
4.2	Typing rules	123
4.3	Conversion rules	125
4.4	Subtyping rules	127
4.5	Inductive definitions	128
4.5.1	Types of inductive objects	129
4.5.2	Well-formed inductive definitions	129
4.5.3	Destructors	134
4.5.4	Fixpoint definitions	138
4.6	Admissible rules for global environments	141
4.7	Co-inductive types	142
4.8	The Calculus of Inductive Construction with impredicative <code>Set</code>	142
5	The Module System	145
5.1	Modules and module types	145
5.2	Typing Modules	146
II	The proof engine	151
6	Vernacular commands	153
6.1	Displaying	153
6.1.1	<code>Print <i>qualid</i>.</code>	153
6.1.2	<code>Print All.</code>	153
6.2	Flags, Options and Tables	154
6.2.1	<code>Set <i>flag</i>.</code>	154
6.2.2	<code>Unset <i>flag</i>.</code>	154

6.2.3	Test <i>flag</i> .	154
6.2.4	Set <i>option value</i> .	154
6.2.5	Unset <i>option</i> .	155
6.2.6	Test <i>option</i> .	155
6.2.7	Tables	155
6.2.8	Print Options.	155
6.3	Requests to the environment	155
6.3.1	Check <i>term</i> .	155
6.3.2	Eval <i>convtactic</i> in <i>term</i> .	155
6.3.3	Compute <i>term</i> .	156
6.3.4	Extraction <i>term</i> .	156
6.3.5	Print Assumptions <i>qualid</i> .	156
6.3.6	Search <i>qualid</i> .	156
6.3.7	SearchHead <i>term</i> .	158
6.3.8	SearchPattern <i>term_pattern</i> .	159
6.3.9	SearchRewrite <i>term</i> .	160
6.3.10	Locate <i>qualid</i> .	161
6.4	Loading files	161
6.4.1	Load <i>ident</i> .	161
6.5	Compiled files	162
6.5.1	Require <i>qualid</i> .	162
6.5.2	Print Libraries.	163
6.5.3	Declare ML Module <i>string</i> ₁ .. <i>string</i> _{<i>n</i>} .	163
6.5.4	Print ML Modules.	164
6.6	Loadpath	164
6.6.1	Pwd.	164
6.6.2	Cd <i>string</i> .	164
6.6.3	Add LoadPath <i>string</i> as <i>dirpath</i> .	164
6.6.4	Add Rec LoadPath <i>string</i> as <i>dirpath</i> .	164
6.6.5	Remove LoadPath <i>string</i> .	165
6.6.6	Print LoadPath.	165
6.6.7	Add ML Path <i>string</i> .	165
6.6.8	Add Rec ML Path <i>string</i> .	165
6.6.9	Print ML Path <i>string</i> .	165
6.6.10	Locate File <i>string</i> .	165
6.6.11	Locate Library <i>dirpath</i> .	165
6.7	Backtracking	165
6.7.1	Reset <i>ident</i> .	165
6.7.2	Back.	166
6.7.3	BackTo <i>num</i> .	166
6.8	Quitting and debugging	167
6.8.1	Quit.	167
6.8.2	Drop.	167
6.8.3	Time <i>command</i> .	167
6.8.4	Redirect " <i>file</i> " <i>command</i> .	167
6.8.5	Timeout <i>int</i> <i>command</i> .	167
6.8.6	Set Default Timeout <i>int</i> .	167

6.8.7	Unset Default Timeout.....	167
6.8.8	Test Default Timeout.	168
6.9	Controlling display	168
6.9.1	Set Silent.	168
6.9.2	Unset Silent.....	168
6.9.3	Set Warnings " (w_1, \dots, w_n) ".	168
6.9.4	Set Search Output Name Only.....	168
6.9.5	Unset Search Output Name Only.....	168
6.9.6	Set Printing Width <i>integer</i>	168
6.9.7	Unset Printing Width.	168
6.9.8	Test Printing Width.	168
6.9.9	Set Printing Depth <i>integer</i>	168
6.9.10	Unset Printing Depth.	169
6.9.11	Test Printing Depth.	169
6.9.12	Unset Printing Compact Contexts.....	169
6.9.13	Set Printing Compact Contexts.	169
6.9.14	Test Printing Compact Contexts.	169
6.9.15	Unset Printing Unfocused.	169
6.9.16	Set Printing Unfocused.....	169
6.9.17	Test Printing Unfocused.	169
6.9.18	Set Printing Dependent Evars Line.	169
6.9.19	Unset Printing Dependent Evars Line.	169
6.10	Controlling the reduction strategies and the conversion algorithm	170
6.10.1	Opaque <i>qualid</i> ₁ ... <i>qualid</i> _n	170
6.10.2	Transparent <i>qualid</i> ₁ ... <i>qualid</i> _n	170
6.10.3	Strategy <i>level</i> [<i>qualid</i> ₁ ... <i>qualid</i> _n].	171
6.10.4	Print Strategy <i>qualid</i>	171
6.10.5	Declare Reduction <i>ident</i> := <i>convtactic</i>	171
6.11	Controlling the locality of commands	172
6.11.1	Local, Global	172
7	Proof handling	173
7.1	Switching on/off the proof editing mode	173
7.1.1	Goal <i>form</i>	173
7.1.2	Qed.	174
7.1.3	Admitted.	174
7.1.4	Proof <i>term</i>	174
7.1.5	Proof using <i>ident</i> ₁ ... <i>ident</i> _n	174
7.1.6	Abort.	176
7.1.7	Existential <i>num</i> := <i>term</i>	176
7.1.8	Grab Existential Variables.....	176
7.2	Navigation in the proof tree	176
7.2.1	Undo.	176
7.2.2	Restart.	176
7.2.3	Focus.	177
7.2.4	Unfocus.	177
7.2.5	Unfocused.	177

7.2.6	{ and }	177
7.2.7	Bullets	177
7.2.8	Set Bullet Behavior.	178
7.3	Requesting information	179
7.3.1	Show.	179
7.3.2	Guarded.	180
7.4	Controlling the effect of proof editing commands	180
7.4.1	Set Hyps Limit <i>num</i> .	180
7.4.2	Unset Hyps Limit.	180
7.4.3	Set Automatic Introduction.	181
7.5	Controlling memory usage	181
7.5.1	Optimize Proof.	181
7.5.2	Optimize Heap.	181
8	Tactics	183
8.1	Invocation of tactics	183
8.1.1	Set Default Goal Selector " <i>toplevel_selector</i> ".	183
8.1.2	Test Default Goal Selector.	184
8.1.3	Bindings list	184
8.1.4	Occurrences sets and occurrences clauses	184
8.2	Applying theorems	185
8.2.1	exact <i>term</i>	185
8.2.2	assumption	185
8.2.3	refine <i>term</i>	185
8.2.4	apply <i>term</i>	187
8.2.5	apply <i>term</i> in <i>ident</i>	190
8.2.6	constructor <i>num</i>	191
8.3	Managing the local context	193
8.3.1	intro	193
8.3.2	intros <i>intro_pattern_list</i>	194
8.3.3	clear <i>ident</i>	197
8.3.4	revert <i>ident</i> ₁ ... <i>ident</i> _{<i>n</i>}	197
8.3.5	move <i>ident</i> ₁ after <i>ident</i> ₂	198
8.3.6	rename <i>ident</i> ₁ into <i>ident</i> ₂	200
8.3.7	set (<i>ident</i> := <i>term</i>)	200
8.3.8	decompose [<i>qualid</i> ₁ ... <i>qualid</i> _{<i>n</i>}] <i>term</i>	202
8.4	Controlling the proof flow	202
8.4.1	assert (<i>ident</i> : <i>form</i>)	202
8.4.2	generalize <i>term</i>	204
8.4.3	eval (<i>ident</i> : <i>term</i>)	205
8.4.4	instantiate (<i>ident</i> := <i>term</i>)	205
8.4.5	admit	206
8.4.6	absurd <i>term</i>	206
8.4.7	contradiction	207
8.4.8	contradict <i>ident</i>	207
8.4.9	exfalso	207
8.5	Case analysis and induction	207

8.5.1	destruct <i>term</i>	207
8.5.2	induction <i>term</i>	209
8.5.3	double induction <i>ident</i> ₁ <i>ident</i> ₂	213
8.5.4	dependent induction <i>ident</i>	213
8.5.5	functional induction (<i>qualid term</i> ₁ . . . <i>term</i> _{<i>n</i>})	215
8.5.6	discriminate <i>term</i>	216
8.5.7	injection <i>term</i>	217
8.5.8	inversion <i>ident</i>	219
8.5.9	fix <i>ident num</i>	224
8.5.10	cofix <i>ident</i>	225
8.6	Rewriting expressions	225
8.6.1	rewrite <i>term</i>	225
8.6.2	replace <i>term</i> ₁ with <i>term</i> ₂	227
8.6.3	subst <i>ident</i>	228
8.6.4	stepl <i>term</i>	228
8.6.5	change <i>term</i>	229
8.7	Performing computations	229
8.7.1	cbv <i>flag</i> ₁ . . . <i>flag</i> _{<i>n</i>} , lazy <i>flag</i> ₁ . . . <i>flag</i> _{<i>n</i>} , and compute	230
8.7.2	red	231
8.7.3	hnf	232
8.7.4	cbn and simpl	232
8.7.5	unfold <i>qualid</i>	234
8.7.6	fold <i>term</i>	235
8.7.7	pattern <i>term</i>	235
8.7.8	Conversion tactics applied to hypotheses	236
8.8	Automation	236
8.8.1	auto	236
8.8.2	eauto	237
8.8.3	autounfold with <i>ident</i> ₁ . . . <i>ident</i> _{<i>n</i>}	238
8.8.4	autorewrite with <i>ident</i> ₁ . . . <i>ident</i> _{<i>n</i>}	238
8.8.5	easy	239
8.9	Controlling automation	239
8.9.1	The hints databases for auto and eauto	239
8.9.2	Hint databases defined in the COQ standard library	243
8.9.3	Remove Hints <i>term</i> ₁ . . . <i>term</i> _{<i>n</i>} : <i>ident</i> ₁ . . . <i>ident</i> _{<i>m</i>}	244
8.9.4	Print Hint	244
8.9.5	Hint Rewrite <i>term</i> ₁ . . . <i>term</i> _{<i>n</i>} : <i>ident</i> ₁ . . . <i>ident</i> _{<i>m</i>}	244
8.9.6	Hint locality	245
8.9.7	Setting implicit automation tactics	245
8.10	Decision procedures	246
8.10.1	tauto	246
8.10.2	intuition <i>tactic</i>	247
8.10.3	rtauto	248
8.10.4	firstorder	249
8.10.5	congruence	249
8.11	Checking properties of terms	251
8.11.1	constr_eq <i>term</i> ₁ <i>term</i> ₂	251

8.11.2	<code>unify term₁ term₂</code>	251
8.11.3	<code>is_evar term</code>	251
8.11.4	<code>has_evar term</code>	251
8.11.5	<code>is_var term</code>	251
8.12	Equality	251
8.12.1	<code>f_equal</code>	251
8.12.2	<code>reflexivity</code>	252
8.12.3	<code>symmetry</code>	252
8.12.4	<code>transitivity term</code>	252
8.13	Equality and inductive sets	252
8.13.1	<code>decide equality</code>	252
8.13.2	<code>compare term₁ term₂</code>	252
8.13.3	<code>simplify_eq term</code>	252
8.13.4	<code>dependent rewrite -> ident</code>	253
8.14	Inversion	253
8.14.1	<code>functional inversion ident</code>	253
8.14.2	<code>quote ident</code>	254
8.15	Classical tactics	254
8.15.1	<code>classical_left</code> and <code>classical_right</code>	254
8.16	Automatizing	254
8.16.1	<code>btauto</code>	254
8.16.2	<code>omega</code>	255
8.16.3	<code>ring</code> and <code>ring_simplify term₁ ... term_n</code>	255
8.16.4	<code>field</code> , <code>field_simplify term₁ ... term_n</code> , and <code>field_simplify_eq</code>	255
8.16.5	<code>fourier</code>	256
8.17	Non-logical tactics	256
8.17.1	<code>cycle num</code>	256
8.17.2	<code>swap num₁ num₂</code>	257
8.17.3	<code>revgoals</code>	258
8.17.4	<code>shelve</code>	258
8.17.5	<code>Unshelve</code>	259
8.17.6	<code>give_up</code>	259
8.18	Simple tactic macros	259
9	The tactic language	261
9.1	Syntax	261
9.2	Semantics	262
9.3	Tactic toplevel definitions	276
9.3.1	Defining \mathcal{L}_{tac} functions	276
9.3.2	Printing \mathcal{L}_{tac} tactics	277
9.4	Debugging \mathcal{L}_{tac} tactics	277
9.4.1	Info trace	277
9.4.2	Interactive debugger	278
9.4.3	Profiling \mathcal{L}_{tac} tactics	279

10 Detailed examples of tactics	281
10.1 dependent induction	281
10.1.1 A larger example	283
10.2 autorewrite	287
10.3 quote	288
10.3.1 Introducing variables map	289
10.3.2 Combining variables and constants	290
10.4 Using the tactical language	292
10.4.1 About the cardinality of the set of natural numbers	292
10.4.2 Permutation on closed lists	292
10.4.3 Deciding intuitionistic propositional logic	294
10.4.4 Deciding type isomorphisms	294
11 The SSReflect proof language	299
11.1 Introduction	299
11.2 Usage	301
11.2.1 Getting started	301
11.2.2 Compatibility issues	301
11.3 Gallina extensions	302
11.3.1 Pattern assignment	302
11.3.2 Pattern conditional	303
11.3.3 Parametric polymorphism	304
11.3.4 Anonymous arguments	305
11.3.5 Wildcards	305
11.4 Definitions	305
11.4.1 Definitions	305
11.4.2 Abbreviations	306
11.4.3 Localization	309
11.5 Basic tactics	310
11.5.1 Bookkeeping	310
11.5.2 The defective tactics	313
11.5.3 Discharge	314
11.5.4 Introduction	317
11.5.5 Generation of equations	320
11.5.6 Type families	320
11.6 Control flow	322
11.6.1 Indentation and bullets	322
11.6.2 Terminators	323
11.6.3 Selectors	324
11.6.4 Iteration	325
11.6.5 Localization	326
11.6.6 Structure	327
11.7 Rewriting	333
11.7.1 An extended rewrite tactic	333
11.7.2 Remarks and examples	336
11.7.3 Locking, unlocking	342
11.7.4 Congruence	344

11.8	Contextual patterns	345
11.8.1	Syntax	345
11.8.2	Matching contextual patterns	346
11.8.3	Examples	347
11.8.4	Patterns for recurrent contexts	348
11.9	Views and reflection	348
11.9.1	Interpreting eliminations	349
11.9.2	Interpreting assumptions	351
11.9.3	Interpreting goals	353
11.9.4	Boolean reflection	353
11.9.5	The <code>reflect</code> predicate	354
11.9.6	General mechanism for interpreting goals and assumptions	355
11.9.7	Interpreting equivalences	357
11.9.8	Declaring new Hint Views	358
11.9.9	Multiple views	358
11.10	SSREFLECT searching tool	359
11.11	Synopsis and Index	360
III	User extensions	363
12	Syntax extensions and interpretation scopes	365
12.1	Notations	365
12.1.1	Basic notations	365
12.1.2	Precedences and associativity	366
12.1.3	Complex notations	366
12.1.4	Simple factorization rules	367
12.1.5	Displaying symbolic notations	368
12.1.6	The <code>Infix</code> command	369
12.1.7	Reserving notations	369
12.1.8	Simultaneous definition of terms and notations	370
12.1.9	Displaying informations about notations	370
12.1.10	Locating notations	370
12.1.11	Notations and simple binders	371
12.1.12	Notations with recursive patterns	372
12.1.13	Notations with recursive patterns involving binders	373
12.1.14	Summary	373
12.2	Interpretation scopes	374
12.2.1	Global interpretation rules for notations	374
12.2.2	Local interpretation rules for notations	375
12.2.3	The <code>type_scope</code> interpretation scope	377
12.2.4	The <code>function_scope</code> interpretation scope	378
12.2.5	Interpretation scopes used in the standard library of COQ	378
12.2.6	Displaying informations about scopes	379
12.3	Abbreviations	380
12.4	Tactic Notations	381

13 Proof schemes	383
13.1 Generation of induction principles with Scheme	383
13.1.1 Automatic declaration of schemes	385
13.1.2 Combined Scheme	385
13.2 Generation of induction principles with Functional Scheme	386
13.3 Generation of inversion principles with Derive Inversion	389
 IV Practical tools	 391
14 The COQ commands	393
14.1 Interactive use (<code>coqtop</code>)	393
14.2 Batch compilation (<code>coqc</code>)	393
14.3 Customization	394
14.3.1 By resource file	394
14.3.2 By environment variables	394
14.3.3 By command line options	394
14.4 Compiled libraries checker (<code>coqchk</code>)	397
 15 Utilities	 399
15.1 Building a toplevel extended with user tactics	399
15.2 Building a COQ project with <code>coq_makefile</code>	400
15.3 Modules dependencies	405
15.4 Documenting COQ files with <code>coqdoc</code>	405
15.4.1 Principles	405
15.4.2 Usage	408
15.4.3 The <code>coqdoc</code> L ^A T _E X style file	412
15.5 Embedded COQ phrases inside L ^A T _E X documents	413
15.6 COQ and GNU EMACS	413
15.6.1 The COQ Emacs mode	413
15.6.2 PROOF GENERAL	413
15.7 Module specification	414
15.8 Man pages	414
 16 COQ Integrated Development Environment	 415
16.1 Managing files and buffers, basic edition	415
16.2 Interactive navigation into COQ scripts	416
16.3 Try tactics automatically	417
16.4 Proof folding	417
16.5 Vernacular commands, templates	417
16.6 Queries	417
16.7 Compilation	417
16.8 Customizations	418
16.9 Using Unicode symbols	419
16.9.1 Displaying Unicode symbols	419
16.9.2 Defining an input method for non ASCII symbols	419
16.9.3 Character encoding for saved files	420

V	Addendum to the Reference Manual	421
17	Extended pattern-matching	427
17.1	Patterns	427
17.2	About patterns of parametric types	430
17.3	Matching objects of dependent types	432
17.3.1	Understanding dependencies in patterns	432
17.3.2	When the elimination predicate must be provided	433
17.4	Using pattern matching to write proofs	434
17.5	Pattern-matching on inductive objects involving local definitions	435
17.6	Pattern-matching and coercions	436
17.7	When does the expansion strategy fail ?	436
18	Implicit Coercions	439
18.1	General Presentation	439
18.2	Classes	439
18.3	Coercions	440
18.4	Identity Coercions	440
18.5	Inheritance Graph	441
18.6	Declaration of Coercions	441
18.6.1	Coercion <i>qualid</i> : <i>class</i> ₁ \rightarrow <i>class</i> ₂ .	441
18.6.2	Identity Coercion <i>ident</i> : <i>class</i> ₁ \rightarrow <i>class</i> ₂ .	442
18.7	Displaying Available Coercions	443
18.7.1	Print Classes.	443
18.7.2	Print Coercions.	443
18.7.3	Print Graph.	443
18.7.4	Print Coercion Paths <i>class</i> ₁ <i>class</i> ₂ .	443
18.8	Activating the Printing of Coercions	443
18.8.1	Set Printing Coercions.	443
18.8.2	Add Printing Coercion <i>qualid</i> .	443
18.9	Classes as Records	444
18.10	Coercions and Sections	444
18.11	Coercions and Modules	444
18.12	Examples	444
19	Canonical Structures	449
19.1	Notation overloading	449
19.1.1	Derived Canonical Structures	451
19.2	Hierarchy of structures	451
19.2.1	Compact declaration of Canonical Structures	457
20	Type Classes	459
20.1	Class and Instance declarations	459
20.2	Binding classes	460
20.3	Parameterized Instances	461
20.4	Sections and contexts	461
20.5	Building hierarchies	462
20.5.1	Superclasses	462

20.5.2	Substructures	463
20.6	Summary of the commands	463
20.6.1	Class <i>ident binder₁ ... binder_n : sort := { field₁ ; ... ; field_k }</i>	463
20.6.2	Instance <i>ident binder₁ ... binder_n : Class t₁ ... t_n [priority] := { field₁ := b₁ ; ... ; field_i := b_i }</i>	464
20.6.3	Existing Instance <i>ident [priority]</i>	464
20.6.4	Context <i>binder₁ ... binder_n</i>	464
20.6.5	typeclasses eauto	465
20.6.6	autoapply <i>term</i> with <i>ident</i>	465
20.6.7	Typeclasses Transparent, Opaque <i>ident₁ ... ident_n</i>	465
20.6.8	Set Typeclasses Dependency Order	466
20.6.9	Set Typeclasses Filtered Unification	466
20.6.10	Set Typeclasses Legacy Resolution	466
20.6.11	Set Typeclasses Module Eta	466
20.6.12	Set Typeclasses Limit Intros	466
20.6.13	Set Typeclass Resolution After Apply	467
20.6.14	Set Typeclass Resolution For Conversion	467
20.6.15	Set Typeclasses Strict Resolution	467
20.6.16	Set Typeclasses Unique Solutions	467
20.6.17	Set Typeclasses Unique Instances	467
20.6.18	Typeclasses eauto := [debug] [(dfs) (bfs)] [depth]	467
20.6.19	Set Typeclasses Debug [Verbosity <i>num</i>]	467
20.6.20	Set Refine Instance Mode	468
21	Omega: a solver of quantifier-free problems in Presburger Arithmetic	469
21.1	Description of omega	469
21.1.1	Arithmetical goals recognized by omega	469
21.1.2	Messages from omega	470
21.2	Using omega	470
21.3	Options	471
21.4	Technical data	471
21.4.1	Overview of the tactic	471
21.4.2	Overview of the <i>OMEGA</i> decision procedure	471
21.5	Bugs	472
22	Micromega: tactics for solving arithmetic goals over ordered rings	473
22.1	Short description of the tactics	473
22.2	<i>Positivstellensatz</i> refutations	474
22.3	lra: a decision procedure for linear real and rational arithmetic	474
22.4	lia: a tactic for linear integer arithmetic	475
22.5	nra: a proof procedure for non-linear arithmetic	476
22.6	nia: a proof procedure for non-linear integer arithmetic	476
22.7	psatz: a proof procedure for non-linear arithmetic	476

23	Extraction of programs in Objective Caml and Haskell	477
23.1	Generating ML code	477
23.2	Extraction options	478
23.2.1	Setting the target language	478
23.2.2	Inlining and optimizations	479
23.2.3	Extra elimination of useless arguments	480
23.2.4	Realizing axioms	480
23.2.5	Avoiding conflicts with existing filenames	482
23.3	Differences between COQ and ML type systems	483
23.4	Some examples	484
23.4.1	A detailed example: Euclidean division	484
23.4.2	Extraction's horror museum	485
23.4.3	Users' Contributions	486
24	PROGRAM	487
24.1	Elaborating programs	487
24.1.1	Syntactic control over equalities	488
24.1.2	Program Definition <i>ident := term</i> .	489
24.1.3	Program Fixpoint <i>ident params {order} : type := term</i>	489
24.1.4	Program Lemma <i>ident : type</i> .	490
24.2	Solving obligations	490
24.3	Frequently Asked Questions	491
25	The <code>ring</code> and <code>field</code> tactic families	493
25.1	What does this tactic do?	493
25.2	The variables map	494
25.3	Is it automatic?	494
25.4	Concrete usage in COQ	494
25.5	Adding a ring structure	496
25.6	How does it work?	499
25.7	Dealing with fields	500
25.8	Adding a new field structure	501
25.9	History of <code>ring</code>	502
25.10	Discussion	503
26	Nsatz: tactics for proving equalities in integral domains	505
26.1	Using the basic tactic <code>nsatz</code>	505
26.2	More about <code>nsatz</code>	505
27	Generalized rewriting	507
27.1	Introduction to generalized rewriting	508
27.1.1	Relations and morphisms	508
27.1.2	Adding new relations and morphisms	509
27.1.3	Rewriting and non reflexive relations	511
27.1.4	Rewriting and non symmetric relations	512
27.1.5	Rewriting in ambiguous setoid contexts	512
27.2	Commands and tactics	513
27.2.1	First class setoids and morphisms	513

27.2.2	Tactics enabled on user provided relations	514
27.2.3	Printing relations and morphisms	514
27.2.4	Deprecated syntax and backward incompatibilities	514
27.3	Extensions	515
27.3.1	Rewriting under binders	515
27.3.2	Sub-relations	516
27.3.3	Constant unfolding	516
27.4	Strategies for rewriting	517
27.4.1	Definitions	517
27.4.2	Usage	518
28	Asynchronous and Parallel Proof Processing	519
28.1	Proof annotations	519
28.2	Proof blocks and error resilience	520
28.2.1	Caveats	520
28.3	Interactive mode	520
28.4	Batch mode	521
28.5	Limiting the number of parallel workers	522
29	Polymorphic Universes	523
29.1	General Presentation	523
29.2	Polymorphic, Monomorphic	525
29.3	Cumulative, NonCumulative	525
29.4	Global and local universes	528
29.5	Conversion and unification	528
29.6	Minimization	528
29.7	Explicit Universes	529
29.7.1	Universe <i>ident</i>	529
29.7.2	Constraint <i>ident ord ident</i>	529
29.7.3	Polymorphic definitions	529
29.7.4	Unset Strict Universe Declaration.	530
30	Miscellaneous extensions	531
30.1	Program derivation	531
30.1.1	Derive <i>ident</i> ₁ SuchThat <i>term</i> As <i>ident</i> ₂	531
	Bibliography	533
	Global Index	543
	Tactics Index	557
	Vernacular Commands Index	561
	Vernacular Options Index	565
	Index of Error Messages	567
	List of Figures	571

Part I

The language

Chapter 1

The GALLINA specification language

This chapter describes GALLINA, the specification language of COQ. It allows developing mathematical theories and proofs of specifications of programs. The theories are built from axioms, hypotheses, parameters, lemmas, theorems and definitions of constants, functions, predicates and sets. The syntax of logical objects involved in theories is described in Section 1.2. The language of commands, called *The Vernacular* is described in section 1.3.

In COQ, logical objects are typed to ensure their logical correctness. The rules implemented by the typing algorithm are described in Chapter 4.

About the grammars in the manual

Grammars are presented in Backus-Naur form (BNF). Terminal symbols are set in `typewriter font`. In addition, there are special notations for regular expressions.

An expression enclosed in square brackets `[...]` means at most one occurrence of this expression (this corresponds to an optional component).

The notation “`entry sep ... sep entry`” stands for a non empty sequence of expressions parsed by `entry` and separated by the literal “`sep`”¹.

Similarly, the notation “`entry ... entry`” stands for a non empty sequence of expressions parsed by the “`entry`” entry, without any separator between.

Finally, the notation “`[entry sep ... sep entry]`” stands for a possibly empty sequence of expressions parsed by the “`entry`” entry, separated by the literal “`sep`”.

1.1 Lexical conventions

Blanks Space, newline and horizontal tabulation are considered as blanks. Blanks are ignored but they separate tokens.

Comments Comments in COQ are enclosed between `(* and *)`, and can be nested. They can contain any character. However, string literals must be correctly closed. Comments are treated as blanks.

Identifiers and access identifiers Identifiers, written *ident*, are sequences of letters, digits, `_` and `'`, that do not start with a digit or `'`. That is, they are recognized by the following lexical class:

¹This is similar to the expression “`entry { sep entry }`” in standard BNF, or “`entry (sep entry)*`” in the syntax of regular expressions.

```

first_letter ::= a..z | A..Z | _ | unicode-letter
subsequent_letter ::= a..z | A..Z | 0..9 | _ | ' | unicode-letter | unicode-id-part
ident ::= first_letter [subsequent_letter...subsequent_letter]

```

All characters are meaningful. In particular, identifiers are case-sensitive. The entry `unicode-letter` non-exhaustively includes Latin, Greek, Gothic, Cyrillic, Arabic, Hebrew, Georgian, Hangul, Hiragana and Katakana characters, CJK ideographs, mathematical letter-like symbols, hyphens, non-breaking space, ... The entry `unicode-id-part` non-exhaustively includes symbols for prime letters and subscripts.

Access identifiers, written `access_ident`, are identifiers prefixed by `.` (dot) without blank. They are used in the syntax of qualified identifiers.

Natural numbers and integers Numerals are sequences of digits. Integers are numerals optionally preceded by a minus sign.

```

digit ::= 0..9
num ::= digit...digit
integer ::= [-]num

```

Strings Strings are delimited by `"` (double quote), and enclose a sequence of any characters different from `"` or the sequence `""` to denote the double quote character. In grammars, the entry for quoted strings is *string*.

Keywords The following identifiers are reserved keywords, and cannot be employed otherwise:

<code>_</code>	<code>as</code>	<code>at</code>	<code>cofix</code>	<code>else</code>	<code>end</code>
<code>exists</code>	<code>exists2</code>	<code>fix</code>	<code>for</code>	<code>forall</code>	<code>fun</code>
<code>if</code>	<code>IF</code>	<code>in</code>	<code>let</code>	<code>match</code>	<code>mod</code>
<code>Prop</code>	<code>return</code>	<code>Set</code>	<code>then</code>	<code>Type</code>	<code>using</code>
<code>where</code>	<code>with</code>				

Special tokens The following sequences of characters are special tokens:

```

!      %      &      &&     (      ( )     )
*      +      ++     ,      -      ->     .
.(     ..     /      /\     :      ::      :<
:=     :>     ;      <      <-     <->     <:
<=     <>     =      =>     =_D     >      >->
>=     ?      ?=     @      [      \ /     ]
^      {      |      |-     ||     }      ~

```

Lexical ambiguities are resolved according to the “longest match” rule: when a sequence of non alphanumerical characters can be decomposed into several different ways, then the first token is the longest possible one (among all tokens defined at this moment), and so on.

1.2 Terms

1.2.1 Syntax of terms

Figures 1.1 and 1.2 describe the basic syntax of the terms of the *Calculus of Inductive Constructions* (also called CIC). The formal presentation of CIC is given in Chapter 4. Extensions of this syntax are given in chapter 2. How to customize the syntax is described in Chapter 12.

1.2.2 Types

COQ terms are typed. COQ types are recognized by the same syntactic class as *term*. We denote by *type* the semantic subclass of types inside the syntactic class *term*.

1.2.3 Qualified identifiers and simple identifiers

Qualified identifiers (*qualid*) denote *global constants* (definitions, lemmas, theorems, remarks or facts), *global variables* (parameters or axioms), *inductive types* or *constructors of inductive types*. *Simple identifiers* (or shortly *ident*) are a syntactic subset of qualified identifiers. Identifiers may also denote *local variables*, what qualified identifiers do not.

1.2.4 Numerals

Numerals have no definite semantics in the calculus. They are mere notations that can be bound to objects through the notation mechanism (see Chapter 12 for details). Initially, numerals are bound to Peano's representation of natural numbers (see 3.1.3).

Note: negative integers are not at the same level as *num*, for this would make precedence unnatural.

1.2.5 Sorts

There are three sorts *Set*, *Prop* and *Type*.

- *Prop* is the universe of *logical propositions*. The logical propositions themselves are typing the proofs. We denote propositions by *form*. This constitutes a semantic subclass of the syntactic class *term*.
- *Set* is the universe of *program types* or *specifications*. The specifications themselves are typing the programs. We denote specifications by *specif*. This constitutes a semantic subclass of the syntactic class *term*.
- *Type* is the type of *Set* and *Prop*

More on sorts can be found in Section 4.1.1.

1.2.6 Binders

Various constructions such as *fun*, *forall*, *fix* and *cofix* *bind* variables. A binding is represented by an identifier. If the binding variable is not used in the expression, the identifier can be replaced by the symbol *_*. When the type of a bound variable cannot be synthesized by the system, it can be specified with the notation *(ident : type)*. There is also a notation for a sequence of binding variables sharing the same type: *(ident₁... ident_n : type)*. A binder can also be any pattern prefixed by a quote, e.g. *' (x, y)*.

<i>term</i>	<code>::= forall <i>binders</i> , <i>term</i></code>	(1.2.8)
	<code> fun <i>binders</i> => <i>term</i></code>	(1.2.7)
	<code> fix <i>fix_bodies</i></code>	(1.2.14)
	<code> cofix <i>cofix_bodies</i></code>	(1.2.14)
	<code> let <i>ident</i> [<i>binders</i>] [: <i>term</i>] := <i>term</i> in <i>term</i></code>	(1.2.12)
	<code> let fix <i>fix_body</i> in <i>term</i></code>	(1.2.14)
	<code> let cofix <i>cofix_body</i> in <i>term</i></code>	(1.2.14)
	<code> let ([<i>name</i> , ... , <i>name</i>]) [<i>dep_ret_type</i>] := <i>term</i> in <i>term</i></code>	(1.2.13, 2.2.1)
	<code> let ' <i>pattern</i> [in <i>term</i>] := <i>term</i> [<i>return_type</i>] in <i>term</i></code>	(1.2.13, 2.2.1)
	<code> if <i>term</i> [<i>dep_ret_type</i>] then <i>term</i> else <i>term</i></code>	(1.2.13, 2.2.1)
	<code> <i>term</i> : <i>term</i></code>	(1.2.10)
	<code> <i>term</i> <: <i>term</i></code>	(1.2.10)
	<code> <i>term</i> :></code>	(24.1.1)
	<code> <i>term</i> -> <i>term</i></code>	(1.2.8)
	<code> <i>term</i> arg ... arg</code>	(1.2.9)
	<code> @ <i>qualid</i> [<i>term</i> ... <i>term</i>]</code>	(2.7.11)
	<code> <i>term</i> % <i>ident</i></code>	(12.2.2)
	<code> match <i>match_item</i> , ... , <i>match_item</i> [<i>return_type</i>] with</code>	
	<code> [[] <i>equation</i> ... <i>equation</i>] end</code>	(1.2.13)
	<code> <i>qualid</i></code>	(1.2.3)
	<code> <i>sort</i></code>	(1.2.5)
	<code> <i>num</i></code>	(1.2.4)
	<code> _</code>	(1.2.11)
	<code> (<i>term</i>)</code>	
<i>arg</i>	<code>::= <i>term</i></code>	
	<code> (<i>ident</i> := <i>term</i>)</code>	(2.7.11)
<i>binders</i>	<code>::= <i>binder</i> ... <i>binder</i></code>	
<i>binder</i>	<code>::= <i>name</i></code>	(1.2.6)
	<code> (<i>name</i> ... <i>name</i> : <i>term</i>)</code>	
	<code> (<i>name</i> [: <i>term</i>] := <i>term</i>)</code>	
	<code> ' <i>pattern</i></code>	
<i>name</i>	<code>::= <i>ident</i></code>	
	<code> _</code>	
<i>qualid</i>	<code>::= <i>ident</i></code>	
	<code> <i>qualid</i> access_<i>ident</i></code>	
<i>sort</i>	<code>::= Prop Set Type</code>	

Figure 1.1: Syntax of terms

<i>fix_bodies</i>	<code>::=</code>	<i>fix_body</i>
		<i>fix_body</i> with <i>fix_body</i> with ... with <i>fix_body</i> for <i>ident</i>
<i>cofix_bodies</i>	<code>::=</code>	<i>cofix_body</i>
		<i>cofix_body</i> with <i>cofix_body</i> with ... with <i>cofix_body</i> for <i>ident</i>
<i>fix_body</i>	<code>::=</code>	<i>ident</i> binders [<i>annotation</i>] [<i>: term</i>] := <i>term</i>
<i>cofix_body</i>	<code>::=</code>	<i>ident</i> [<i>binders</i>] [<i>: term</i>] := <i>term</i>
<i>annotation</i>	<code>::=</code>	{ struct <i>ident</i> }
<i>match_item</i>	<code>::=</code>	<i>term</i> [<i>as name</i>] [<i>in qualid</i> [<i>pattern ... pattern</i>]]
<i>dep_ret_type</i>	<code>::=</code>	[<i>as name</i>] <i>return_type</i>
<i>return_type</i>	<code>::=</code>	return <i>term</i>
<i>equation</i>	<code>::=</code>	<i>mult_pattern</i> ... <i>mult_pattern</i> => <i>term</i>
<i>mult_pattern</i>	<code>::=</code>	<i>pattern</i> , ... , <i>pattern</i>
<i>pattern</i>	<code>::=</code>	<i>qualid pattern ... pattern</i>
		@ <i>qualid pattern ... pattern</i>
		<i>pattern</i> as <i>ident</i>
		<i>pattern</i> % <i>ident</i>
		<i>qualid</i>
		—
		<i>num</i>
		(<i>or_pattern</i> , ... , <i>or_pattern</i>)
<i>or_pattern</i>	<code>::=</code>	<i>pattern</i> ... <i>pattern</i>

Figure 1.2: Syntax of terms (continued)

Some constructions allow the binding of a variable to value. This is called a “let-binder”. The entry *binder* of the grammar accepts either an assumption binder as defined above or a let-binder. The notation in the latter case is (*ident* := *term*). In a let-binder, only one variable can be introduced at the same time. It is also possible to give the type of the variable as follows: (*ident* : *term* := *term*).

Lists of *binder* are allowed. In the case of *fun* and *forall*, it is intended that at least one binder of the list is an assumption otherwise *fun* and *forall* gets identical. Moreover, parentheses can be omitted in the case of a single sequence of bindings sharing the same type (e.g.: *fun* (*x y z* : *A*) => *t* can be shortened in *fun* *x y z* : *A* => *t*).

1.2.7 Abstractions

The expression “*fun ident* : *type* => *term*” defines the *abstraction* of the variable *ident*, of type *type*, over the term *term*. It denotes a function of the variable *ident* that evaluates to the expression *term* (e.g. *fun* *x* : *A* => *x* denotes the identity function on type *A*). The keyword *fun* can be followed by

several binders as given in Section 1.2.6. Functions over several variables are equivalent to an iteration of one-variable functions. For instance the expression “`fun ident1 ... identn : type => term`” denotes the same function as “`fun ident1 : type => ... fun identn : type => term`”. If a `let`-binder occurs in the list of binders, it is expanded to a `let-in` definition (see Section 1.2.12).

1.2.8 Products

The expression “`forall ident : type, term`” denotes the *product* of the variable *ident* of type *type*, over the term *term*. As for abstractions, `forall` is followed by a binder list, and products over several variables are equivalent to an iteration of one-variable products. Note that *term* is intended to be a type.

If the variable *ident* occurs in *term*, the product is called *dependent product*. The intention behind a dependent product `forall x : A, B` is twofold. It denotes either the universal quantification of the variable *x* of type *A* in the proposition *B* or the functional dependent product from *A* to *B* (a construction usually written $\Pi_{x:A}.B$ in set theory).

Non dependent product types have a special notation: “*A* -> *B*” stands for “`forall _ : A, B`”. The *non dependent product* is used both to denote the propositional implication and function types.

1.2.9 Applications

The expression `term0 term1` denotes the application of *term₀* to *term₁*.

The expression `term0 term1 ... termn` denotes the application of the term *term₀* to the arguments *term₁* ... then *term_n*. It is equivalent to `(... (term0 term1) ...) termn`: associativity is to the left.

The notation `(ident := term)` for arguments is used for making explicit the value of implicit arguments (see Section 2.7.11).

1.2.10 Type cast

The expression “*term* : *type*” is a type cast expression. It enforces the type of *term* to be *type*.

“*term* <: *type*” locally sets up the virtual machine for checking that *term* has type *type*.

1.2.11 Inferable subterms

Expressions often contain redundant pieces of information. Subterms that can be automatically inferred by COQ can be replaced by the symbol “`_`” and COQ will guess the missing piece of information.

1.2.12 Let-in definitions

`let ident := term1 in term2` denotes the local binding of *term₁* to the variable *ident* in *term₂*. There is a syntactic sugar for `let-in` definition of functions: `let ident binder1 ... bindern := term1 in term2` stands for `let ident := fun binder1 ... bindern => term1 in term2`.

1.2.13 Definition by case analysis

Objects of inductive types can be destructured by a case-analysis construction called *pattern-matching* expression. A pattern-matching expression is used to analyze the structure of an inductive objects and to apply specific treatments accordingly.

This paragraph describes the basic form of pattern-matching. See Section 2.2.1 and Chapter 17 for the description of the general form. The basic form of pattern-matching is characterized by a single

match_item expression, a *mult_pattern* restricted to a single *pattern* and *pattern* restricted to the form *qualid ident ... ident*.

The expression `match term0 return_type with pattern1 => term1 | ... | patternn => termn` end, denotes a *pattern-matching* over the term *term₀* (expected to be of an inductive type *I*). The terms *term₁...term_n* are the *branches* of the pattern-matching expression. Each of *pattern_i* has a form *qualid ident ... ident* where *qualid* must denote a constructor. There should be exactly one branch for every constructor of *I*.

The *return_type* expresses the type returned by the whole *match* expression. There are several cases. In the *non dependent* case, all branches have the same type, and the *return_type* is the common type of branches. In this case, *return_type* can usually be omitted as it can be inferred from the type of the branches².

In the *dependent* case, there are three subcases. In the first subcase, the type in each branch may depend on the exact value being matched in the branch. In this case, the whole pattern-matching itself depends on the term being matched. This dependency of the term being matched in the return type is expressed with an “as *ident*” clause where *ident* is dependent in the return type. For instance, in the following example:

```
Coq < Inductive bool : Type := true : bool | false : bool.
Coq < Inductive eq (A:Type) (x:A) : A -> Prop := eq_refl : eq A x x.
Coq < Inductive or (A:Prop) (B:Prop) : Prop :=
  | or_introl : A -> or A B
  | or_intror : B -> or A B.
Coq < Definition bool_case (b:bool) : or (eq bool b true) (eq bool b false)
:= match b as x return or (eq bool x true) (eq bool x false) with
  | true => or_introl (eq bool true true) (eq bool true false)
  | false => or_intror (eq bool false true) (eq bool false false)
end.
```

the branches have respective types `or (eq bool true true) (eq bool true false)` and `or (eq bool false true) (eq bool false false)` while the whole pattern-matching expression has type `or (eq bool b true) (eq bool b false)`, the identifier *x* being used to represent the dependency. Remark that when the term being matched is a variable, the *as* clause can be omitted and the term being matched can serve itself as binding name in the return type. For instance, the following alternative definition is accepted and has the same meaning as the previous one.

```
Coq < Definition bool_case (b:bool) : or (eq bool b true) (eq bool b false)
:= match b return or (eq bool b true) (eq bool b false) with
  | true => or_introl (eq bool true true) (eq bool true false)
  | false => or_intror (eq bool false true) (eq bool false false)
end.
```

The second subcase is only relevant for annotated inductive types such as the equality predicate (see Section 3.1.2), the order predicate on natural numbers or the type of lists of a given length (see Section 17.3). In this configuration, the type of each branch can depend on the type dependencies specific to the branch and the whole pattern-matching expression has a type determined by the specific dependencies in the type of the term being matched. This dependency of the return type in the annotations of the inductive type is expressed using a “in *I* _ ... _ *pattern₁ ... pattern_n*” clause, where

²Except if the inductive type is empty in which case there is no equation that can be used to infer the return type.

- I is the inductive type of the term being matched;
- the $_$'s are matching the parameters of the inductive type: the return type is not dependent on them.
- the $pattern_i$'s are matching the annotations of the inductive type: the return type is dependent on them
- in the basic case which we describe below, each $pattern_i$ is a name $ident_i$; see 17.3.2 for the general case

For instance, in the following example:

```
Coq < Definition eq_sym (A:Type) (x y:A) (H:eq A x y) : eq A y x :=
  match H in eq _ _ z return eq A z x with
  | eq_refl _ _ => eq_refl A x
end.
```

the type of the branch has type $eq\ A\ x\ x$ because the third argument of eq is x in the type of the pattern $refl_equal$. On the contrary, the type of the whole pattern-matching expression has type $eq\ A\ y\ x$ because the third argument of eq is y in the type of H . This dependency of the case analysis in the third argument of eq is expressed by the identifier z in the return type.

Finally, the third subcase is a combination of the first and second subcase. In particular, it only applies to pattern-matching on terms in a type with annotations. For this third subcase, both the clauses `as` and `in` are available.

There are specific notations for case analysis on types with one or two constructors: “`if ... then ... else ...`” and “`let (... , ... , ...) := ... in ...`” (see Sections 2.2.2 and 2.2.3).

1.2.14 Recursive functions

The expression “`fix $ident_1\ binder_1 : type_1 := term_1$ with ... with $ident_n\ binder_n : type_n := term_n$ for $ident_i$` ” denotes the i^{th} component of a block of functions defined by mutual well-founded recursion. It is the local counterpart of the `Fixpoint` command. See Section 1.3.4 for more details. When $n = 1$, the “`for $ident_i$` ” clause is omitted.

The expression “`cofix $ident_1\ binder_1 : type_1$ with ... with $ident_n\ binder_n : type_n$ for $ident_i$` ” denotes the i^{th} component of a block of terms defined by a mutual guarded co-recursion. It is the local counterpart of the `CoFixpoint` command. See Section 1.3.4 for more details. When $n = 1$, the “`for $ident_i$` ” clause is omitted.

The association of a single fixpoint and a local definition have a special syntax: “`let fix $f \dots := \dots$ in ...`” stands for “`let $f := \text{fix } f \dots := \dots$ in ...`”. The same applies for co-fixpoints.

1.3 The Vernacular

Figure 1.3 describes *The Vernacular* which is the language of commands of GALLINA. A sentence of the vernacular language, like in many natural languages, begins with a capital letter and ends with a dot.

The different kinds of command are described hereafter. They all suppose that the terms occurring in the sentences are well-typed.

<i>sentence</i>	::=	<i>assumption</i> <i>definition</i> <i>inductive</i> <i>fixpoint</i> <i>assertion proof</i>
<i>assumption</i>	::=	<i>assumption_keyword</i> <i>assums</i> .
<i>assumption_keyword</i>	::=	Axiom Conjecture Parameter Parameters Variable Variables Hypothesis Hypotheses
<i>assums</i>	::=	<i>ident</i> ... <i>ident</i> : <i>term</i> (<i>ident</i> ... <i>ident</i> : <i>term</i>) ... (<i>ident</i> ... <i>ident</i> : <i>term</i>)
<i>definition</i>	::=	[Local] Definition <i>ident</i> [<i>binders</i>] [: <i>term</i>] := <i>term</i> . Let <i>ident</i> [<i>binders</i>] [: <i>term</i>] := <i>term</i> .
<i>inductive</i>	::=	Inductive <i>ind_body</i> with... with <i>ind_body</i> . CoInductive <i>ind_body</i> with... with <i>ind_body</i> .
<i>ind_body</i>	::=	<i>ident</i> [<i>binders</i>] [: <i>term</i>] := [[] <i>ident</i> [<i>binders</i>] [: <i>term</i>] ... <i>ident</i> [<i>binders</i>] [: <i>term</i>]]
<i>fixpoint</i>	::=	Fixpoint <i>fix_body</i> with... with <i>fix_body</i> . CoFixpoint <i>cofix_body</i> with... with <i>cofix_body</i> .
<i>assertion</i>	::=	<i>assertion_keyword</i> <i>ident</i> [<i>binders</i>] : <i>term</i> .
<i>assertion_keyword</i>	::=	Theorem Lemma Remark Fact Corollary Proposition Definition Example
<i>proof</i>	::=	Proof ... Qed . Proof ... Defined . Proof ... Admitted .

Figure 1.3: Syntax of sentences

1.3.1 Assumptions

Assumptions extend the environment with axioms, parameters, hypotheses or variables. An assumption binds an *ident* to a *type*. It is accepted by COQ if and only if this *type* is a correct type in the environment preexisting the declaration and if *ident* was not previously defined in the same module. This *type* is

considered to be the type (or specification, or statement) assumed by *ident* and we say that *ident* has type *type*.

`Axiom ident : term .`

This command links *term* to the name *ident* as its specification in the global context. The fact asserted by *term* is thus assumed as a postulate.

Error messages:

1. *ident* already exists

Variants:

1. `Parameter ident : term .`
Is equivalent to `Axiom ident : term`
2. `Parameter ident1 ... identn : term .`
Adds *n* parameters with specification *term*
3. `Parameter (ident1,1 ... ident1,k1 : term1) ... (identn,1 ... identn,kn : termn) .`
Adds *n* blocks of parameters with different specifications.
4. `Local Axiom ident : term .`
Such axioms are never made accessible through their unqualified name by `Import` and its variants (see 2.5.8). You have to explicitly give their fully qualified name to refer to them.
5. `Conjecture ident : term .`
Is equivalent to `Axiom ident : term`.

Remark: It is possible to replace `Parameter` by `Parameters`.

`Variable ident : term .`

This command links *term* to the name *ident* in the context of the current section (see Section 2.4 for a description of the section mechanism). When the current section is closed, name *ident* will be unknown and every object using this variable will be explicitly parametrized (the variable is *discharged*). Using the `Variable` command out of any section is equivalent to using `Local Parameter`.

Error messages:

1. *ident* already exists

Variants:

1. `Variable ident1 ... identn : term .`
Links *term* to names *ident*₁ ... *ident*_{*n*}.
2. `Variable (ident1,1 ... ident1,k1 : term1) ... (identn,1 ... identn,kn : termn) .`
Adds *n* blocks of variables with different specifications.
3. `Hypothesis ident : term .`
`Hypothesis` is a synonymous of `Variable`

Remark: It is possible to replace `Variable` by `Variables` and `Hypothesis` by `Hypotheses`.

It is advised to use the keywords `Axiom` and `Hypothesis` for logical postulates (i.e. when the assertion *term* is of sort `Prop`), and to use the keywords `Parameter` and `Variable` in other cases (corresponding to the declaration of an abstract mathematical entity).

1.3.2 Definitions

Definitions extend the environment with associations of names to terms. A definition can be seen as a way to give a meaning to a name or as a way to abbreviate a term. In any case, the name can later be replaced at any time by its definition.

The operation of unfolding a name into its definition is called δ -conversion (see Section 4.3). A definition is accepted by the system if and only if the defined term is well-typed in the current context of the definition and if the name is not already used. The name defined by the definition is called a *constant* and the term it refers to is its *body*. A definition has a type which is the type of its body.

A formal presentation of constants and environments is given in Section 4.2.

Definition *ident* := *term*.

This command binds *term* to the name *ident* in the environment, provided that *term* is well-typed.

Error messages:

1. *ident* already exists

Variants:

1. Definition *ident* : *term*₁ := *term*₂.
It checks that the type of *term*₂ is definitionally equal to *term*₁, and registers *ident* as being of type *term*₁, and bound to value *term*₂.
2. Definition *ident* *binder*₁ ... *binder*_{*n*} : *term*₁ := *term*₂.
This is equivalent to
Definition *ident* : forall *binder*₁ ... *binder*_{*n*}, *term*₁ := fun *binder*₁ ... *binder*_{*n*} => *term*₂.
3. Local Definition *ident* := *term*.
Such definitions are never made accessible through their unqualified name by `Import` and its variants (see 2.5.8). You have to explicitly give their fully qualified name to refer to them.
4. Example *ident* := *term*.
Example *ident* : *term*₁ := *term*₂.
Example *ident* *binder*₁ ... *binder*_{*n*} : *term*₁ := *term*₂.
These are synonyms of the `Definition` forms.

Error messages:

1. The term *term* has type *type* while it is expected to have type *type*

See also: Sections 6.10.1, 6.10.2, 8.7.5.

`Let ident := term.`

This command binds the value *term* to the name *ident* in the environment of the current section. The name *ident* disappears when the current section is eventually closed, and, all persistent objects (such as theorems) defined within the section and depending on *ident* are prefixed by the let-in definition `let ident := term in`. Using the `Let` command out of any section is equivalent to using `Local Definition`.

Error messages:

1. *ident* already exists

Variants:

1. `Let ident : term1 := term2.`
2. `Let Fixpoint ident fix_body with ... with fix_body ..`
3. `Let CoFixpoint ident cofix_body with ... with cofix_body ..`

See also: Sections 2.4 (section mechanism), 6.10.1, 6.10.2 (opaque/transparent constants), 8.7.5 (tactic `unfold`).

1.3.3 Inductive definitions

We gradually explain simple inductive types, simple annotated inductive types, simple parametric inductive types, mutually inductive types. We explain also co-inductive types.

Simple inductive types

The definition of a simple inductive type has the following form:

```
Inductive ident : sort :=
  ident1 : type1
| ...
| identn : typen
```

The name *ident* is the name of the inductively defined type and *sort* is the universes where it lives. The names *ident*₁, ..., *ident*_{*n*} are the names of its constructors and *type*₁, ..., *type*_{*n*} their respective types. The types of the constructors have to satisfy a *positivity condition* (see Section 4.5.2) for *ident*. This condition ensures the soundness of the inductive definition. If this is the case, the names *ident*, *ident*₁, ..., *ident*_{*n*} are added to the environment with their respective types. Accordingly to the universe where the inductive type lives (e.g. its type *sort*), COQ provides a number of destructors for *ident*. Destructors are named *ident_ind*, *ident_rec* or *ident_rect* which respectively correspond to elimination principles on `Prop`, `Set` and `Type`. The type of the destructors expresses structural induction/recursion principles over objects of *ident*. We give below two examples of the use of the Inductive definitions.

The set of natural numbers is defined as:

```
Coq < Inductive nat : Set :=
  | O : nat
  | S : nat -> nat.
```

```

nat is defined
nat_rect is defined
nat_ind is defined
nat_rec is defined

```

The type `nat` is defined as the least `Set` containing `O` and closed by the `S` constructor. The names `nat`, `O` and `S` are added to the environment.

Now let us have a look at the elimination principles. They are three of them: `nat_ind`, `nat_rec` and `nat_rect`. The type of `nat_ind` is:

```

Coq < Check nat_ind.
nat_ind
  : forall P : nat -> Prop,
    P O -> (forall n : nat, P n -> P (S n)) -> forall n : nat, P n

```

This is the well known structural induction principle over natural numbers, i.e. the second-order form of Peano's induction principle. It allows proving some universal property of natural numbers (`forall n:nat, P n`) by induction on `n`.

The types of `nat_rec` and `nat_rect` are similar, except that they pertain to $(P:\text{nat} \rightarrow \text{Set})$ and $(P:\text{nat} \rightarrow \text{Type})$ respectively. They correspond to primitive induction principles (allowing dependent types) respectively over sorts `Set` and `Type`. The constant `ident_ind` is always provided, whereas `ident_rec` and `ident_rect` can be impossible to derive (for example, when `ident` is a proposition).

Variants:

1. `Coq < Inductive nat : Set := O | S (_:nat).`

In the case where inductive types have no annotations (next section gives an example of such annotations), a constructor can be defined by only giving the type of its arguments.

Simple annotated inductive types

In an annotated inductive types, the universe where the inductive type is defined is no longer a simple sort, but what is called an arity, which is a type whose conclusion is a sort.

As an example of annotated inductive types, let us define the *even* predicate:

```

Coq < Inductive even : nat -> Prop :=
  | even_0 : even O
  | even_SS : forall n:nat, even n -> even (S (S n)).
even is defined
even_ind is defined

```

The type `nat->Prop` means that `even` is a unary predicate (inductively defined) over natural numbers. The type of its two constructors are the defining clauses of the predicate `even`. The type of `even_ind` is:

```

Coq < Check even_ind.
even_ind
  : forall P : nat -> Prop,
    P O ->
    (forall n : nat, even n -> P n -> P (S (S n))) ->
    forall n : nat, even n -> P n

```

From a mathematical point of view it asserts that the natural numbers satisfying the predicate `even` are exactly in the smallest set of naturals satisfying the clauses `even_0` or `even_SS`. This is why, when we want to prove any predicate `P` over elements of `even`, it is enough to prove it for `0` and to prove that if any natural number `n` satisfies `P` its double successor $(S (S n))$ satisfies also `P`. This is indeed analogous to the structural induction principle we got for `nat`.

Error messages:

1. Non strictly positive occurrence of *ident* in *type*
2. The conclusion of *type* is not valid; it must be built from *ident*

Parametrized inductive types

In the previous example, each constructor introduces a different instance of the predicate `even`. In some cases, all the constructors introduces the same generic instance of the inductive definition, in which case, instead of an annotation, we use a context of parameters which are binders shared by all the constructors of the definition.

The general scheme is:

Inductive *ident* *binder*₁...*binder*_k : *term* := *ident*₁: *term*₁ | ... | *ident*_n: *term*_n .

Parameters differ from inductive type annotations in the fact that the conclusion of each type of constructor *term*_i invoke the inductive type with the same values of parameters as its specification.

A typical example is the definition of polymorphic lists:

```
Coq < Inductive list (A:Set) : Set :=
  | nil : list A
  | cons : A -> list A -> list A.
```

Note that in the type of `nil` and `cons`, we write `(list A)` and not just `list`.

The constructors `nil` and `cons` will have respectively types:

```
Coq < Check nil.
nil
  : forall A : Set, list A

Coq < Check cons.
cons
  : forall A : Set, A -> list A -> list A
```

Types of destructors are also quantified with `(A:Set)`.

Variants:

1. `Coq < Inductive list (A:Set) : Set := nil | cons (_:A) (_:list A)`.
This is an alternative definition of lists where we specify the arguments of the constructors rather than their full type.
2. `Coq < Variant sum (A B:Set) : Set := left : A -> sum A B | right : B -> sum A B`.
The `Variant` keyword is identical to the `Inductive` keyword, except that it disallows recursive definition of types (in particular lists cannot be defined with the `Variant` keyword). No induction scheme is generated for this variant, unless the option `Nonrecursive Elimination Schemes` is set (see 13.1.1).

Error messages:

1. The *numth* argument of *ident* must be *ident'* in *type*

New from COQ V8.1 The condition on parameters for inductive definitions has been relaxed since COQ V8.1. It is now possible in the type of a constructor, to invoke recursively the inductive definition on an argument which is not the parameter itself.

One can define :

```
Coq < Inductive list2 (A:Set) : Set :=
  | nil2 : list2 A
  | cons2 : A -> list2 (A*A) -> list2 A.
list2 is defined
list2_rect is defined
list2_ind is defined
list2_rec is defined
```

that can also be written by specifying only the type of the arguments:

```
Coq < Inductive list2 (A:Set) : Set := nil2 | cons2 (_:A) (_:list2 (A*A)).
```

But the following definition will give an error:

```
Coq < Fail Inductive listw (A:Set) : Set :=
  | nilw : listw (A*A)
  | consw : A -> listw (A*A) -> listw (A*A).
The command has indeed failed with message:
Last occurrence of "listw" must have "A" as 1st argument in
"listw (A * A)%type".
```

Because the conclusion of the type of constructors should be `listw A` in both cases.

A parametrized inductive definition can be defined using annotations instead of parameters but it will sometimes give a different (bigger) sort for the inductive definition and will produce a less convenient rule for case elimination.

See also: Sections [4.5](#) and [8.5.2](#).

Mutually defined inductive types

The definition of a block of mutually inductive types has the form:

```
Inductive ident1 : type1 :=
  ident11 : type11
  | ...
  | identn11 : typen11
with
  ...
with identm : typem :=
  ident1m : type1m
  | ...
  | identnmm : typenmm.
```

It has the same semantics as the above Inductive definition for each *ident*₁, ..., *ident*_m. All names *ident*₁, ..., *ident*_m and *ident*₁¹, ..., *ident*_{n_m}^m are simultaneously added to the environment. Then well-typing of constructors can be checked. Each one of the *ident*₁, ..., *ident*_m can be used on its own.

It is also possible to parametrize these inductive definitions. However, parameters correspond to a local context in which the whole set of inductive declarations is done. For this reason, the parameters must be strictly the same for each inductive types. The extended syntax is:

```
Inductive ident1 params : type1 :=
  ident11 : type11
  | ...
  | identn11 : typen11
with
...
with identm params : typem :=
  ident1m : type1m
  | ...
  | identnmm : typenmm.
```

Example: The typical example of a mutual inductive data type is the one for trees and forests. We assume given two types *A* and *B* as variables. It can be declared the following way.

```
Coq < Variables A B : Set.
Coq < Inductive tree : Set :=
  node : A -> forest -> tree
  with forest : Set :=
    | leaf : B -> forest
    | cons : tree -> forest -> forest.
```

This declaration generates automatically six induction principles. They are respectively called *tree_rec*, *tree_ind*, *tree_rect*, *forest_rec*, *forest_ind*, *forest_rect*. These ones are not the most general ones but are just the induction principles corresponding to each inductive part seen as a single inductive definition.

To illustrate this point on our example, we give the types of *tree_rec* and *forest_rec*.

```
Coq < Check tree_rec.
tree_rec
  : forall P : tree -> Set,
    (forall (a : A) (f : forest), P (node a f)) -> forall t : tree, P t
Coq < Check forest_rec.
forest_rec
  : forall P : forest -> Set,
    (forall b : B, P (leaf b)) ->
    (forall (t : tree) (f0 : forest), P f0 -> P (cons t f0)) ->
    forall f1 : forest, P f1
```

Assume we want to parametrize our mutual inductive definitions with the two type variables *A* and *B*, the declaration should be done the following way:

```
Coq < Inductive tree (A B:Set) : Set :=
  node : A -> forest A B -> tree A B
  with forest (A B:Set) : Set :=
    | leaf : B -> forest A B
    | cons : tree A B -> forest A B -> forest A B.
```


Assume we define an inductive definition inside a section. When the section is closed, the variables declared in the section and occurring free in the declaration are added as parameters to the inductive definition.

See also: Section 2.4.

Co-inductive types

The objects of an inductive type are well-founded with respect to the constructors of the type. In other words, such objects contain only a *finite* number of constructors. Co-inductive types arise from relaxing this condition, and admitting types whose objects contain an infinity of constructors. Infinite objects are introduced by a non-ending (but effective) process of construction, defined in terms of the constructors of the type.

An example of a co-inductive type is the type of infinite sequences of natural numbers, usually called streams. It can be introduced in COQ using the `CoInductive` command:

```
Coq < CoInductive Stream : Set :=
    Seq : nat -> Stream -> Stream.
Stream is defined
```

The syntax of this command is the same as the command `Inductive` (see Section 1.3.3). Notice that no principle of induction is derived from the definition of a co-inductive type, since such principles only make sense for inductive ones. For co-inductive ones, the only elimination principle is case analysis. For example, the usual destructors on streams `hd:Stream->nat` and `tl:Str->Str` can be defined as follows:

```
Coq < Definition hd (x:Stream) := let (a,s) := x in a.
hd is defined

Coq < Definition tl (x:Stream) := let (a,s) := x in s.
tl is defined
```

Definition of co-inductive predicates and blocks of mutually co-inductive definitions are also allowed. An example of a co-inductive predicate is the extensional equality on streams:

```
Coq < CoInductive EqSt : Stream -> Stream -> Prop :=
    eqst :
        forall s1 s2:Stream,
            hd s1 = hd s2 -> EqSt (tl s1) (tl s2) -> EqSt s1 s2.
EqSt is defined
```

In order to prove the extensionally equality of two streams s_1 and s_2 we have to construct an infinite proof of equality, that is, an infinite object of type $(EqSt\ s_1\ s_2)$. We will see how to introduce infinite objects in Section 1.3.4.

1.3.4 Definition of recursive functions

Definition of functions by recursion over inductive objects

This section describes the primitive form of definition by recursion over inductive objects. See Section 2.3 for more advanced constructions. The command:

$$\text{Fixpoint } \textit{ident} \textit{ params } \{\text{struct } \textit{ident}_0 \} : \text{type}_0 := \textit{term}_0$$

allows defining functions by pattern-matching over inductive objects using a fixed point construction. The meaning of this declaration is to define *ident* a recursive function with arguments specified by the binders in *params* such that *ident* applied to arguments corresponding to these binders has type *type₀*, and is equivalent to the expression *term₀*. The type of the *ident* is consequently *forall params , type₀* and the value is equivalent to *fun params => term₀*.

To be accepted, a `Fixpoint` definition has to satisfy some syntactical constraints on a special argument called the decreasing argument. They are needed to ensure that the `Fixpoint` definition always terminates. The point of the `{struct ident}` annotation is to let the user tell the system which argument decreases along the recursive calls. For instance, one can define the addition function as :

```
Coq < Fixpoint add (n m:nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (add p m)
  end.
add is defined
add is recursively defined (decreasing on 1st argument)
```

The `{struct ident}` annotation may be left implicit, in this case the system try successively arguments from left to right until it finds one that satisfies the decreasing condition. Note that some fixpoints may have several arguments that fit as decreasing arguments, and this choice influences the reduction of the fixpoint. Hence an explicit annotation must be used if the leftmost decreasing argument is not the desired one. Writing explicit annotations can also speed up type-checking of large mutual fixpoints.

The `match` operator matches a value (here *n*) with the various constructors of its (inductive) type. The remaining arguments give the respective values to be returned, as functions of the parameters of the corresponding constructor. Thus here when *n* equals 0 we return *m*, and when *n* equals (*S p*) we return (*S (add p m)*).

The `match` operator is formally described in detail in Section 4.5.3. The system recognizes that in the inductive call (*add p m*) the first argument actually decreases because it is a *pattern variable* coming from `match n with`.

Example: The following definition is not correct and generates an error message:

```
Coq < Fail Fixpoint wrongplus (n m:nat) {struct n} : nat :=
  match m with
  | 0 => n
  | S p => S (wrongplus n p)
  end.
```

The command has indeed failed with message:

Recursive definition of wrongplus is ill-formed.

In environment

wrongplus : nat -> nat -> nat

n : nat

m : nat

p : nat

Recursive call to wrongplus has principal argument equal to "n" instead of a subterm of "n".

Recursive definition is:

```
"fun n m : nat => match m with
  | 0 => n
```

```

| S p => S (wrongplus n p)
end".

```

because the declared decreasing argument n actually does not decrease in the recursive call. The function computing the addition over the second argument should rather be written:

```

Coq < Fixpoint plus (n m:nat) {struct m} : nat :=
  match m with
  | 0 => n
  | S p => S (plus n p)
  end.

```

The ordinary match operation on natural numbers can be mimicked in the following way.

```

Coq < Fixpoint nat_match
  (C:Set) (f0:C) (fS:nat -> C -> C) (n:nat) {struct n} : C :=
  match n with
  | 0 => f0
  | S p => fS p (nat_match C f0 fS p)
  end.

```

The recursive call may not only be on direct subterms of the recursive variable n but also on a deeper subterm and we can directly write the function `mod2` which gives the remainder modulo 2 of a natural number.

```

Coq < Fixpoint mod2 (n:nat) : nat :=
  match n with
  | 0 => 0
  | S p => match p with
    | 0 => S 0
    | S q => mod2 q
    end
  end.

```

In order to keep the strong normalization property, the fixed point reduction will only be performed when the argument in position of the decreasing argument (which type should be in an inductive definition) starts with a constructor.

The `Fixpoint` construction enjoys also the `with` extension to define functions over mutually defined inductive types or more generally any mutually recursive definitions.

Variants:

1. `Fixpoint` *ident*₁ *params*₁ : *type*₁ := *term*₁
 with ...
 with *ident*_{*m*} *params*_{*m*} : *type*_{*m*} := *term*_{*m*}
 Allows to define simultaneously *ident*₁, ..., *ident*_{*m*}.

Example: The size of trees and forests can be defined the following way:

```

Coq < Fixpoint tree_size (t:tree) : nat :=
  match t with
  | node a f => S (forest_size f)
  end
  with forest_size (f:forest) : nat :=

```

```

match f with
| leaf b => 1
| cons t f' => (tree_size t + forest_size f')
end.

```

A generic command `Scheme` is useful to build automatically various mutual induction principles. It is described in Section 13.1.

Definitions of recursive objects in co-inductive types

The command:

$$\text{CoFixpoint } \textit{ident} : \textit{type}_0 := \textit{term}_0$$

introduces a method for constructing an infinite object of a coinductive type. For example, the stream containing all natural numbers can be introduced applying the following method to the number 0 (see Section 1.3.3 for the definition of `Stream`, `hd` and `tl`):

```

Coq < CoFixpoint from (n:nat) : Stream := Seq n (from (S n)).
from is defined
from is corecursively defined

```

Oppositely to recursive ones, there is no decreasing argument in a co-recursive definition. To be admissible, a method of construction must provide at least one extra constructor of the infinite object for each iteration. A syntactical guard condition is imposed on co-recursive definitions in order to ensure this: each recursive call in the definition must be protected by at least one constructor, and only by constructors. That is the case in the former definition, where the single recursive call of `from` is guarded by an application of `Seq`. On the contrary, the following recursive function does not satisfy the guard condition:

```

Coq < Fail CoFixpoint filter (p:nat -> bool) (s:Stream) : Stream :=
  if p (hd s) then Seq (hd s) (filter p (tl s)) else filter p (tl s).
The command has indeed failed with message:
Recursive definition of filter is ill-formed.
In environment
filter : (nat -> bool) -> Stream -> Stream
p : nat -> bool
s : Stream
Unguarded recursive call in "filter p (tl s)".
Recursive definition is:
"fun (p : nat -> bool) (s : Stream) =>
  if p (hd s) then Seq (hd s) (filter p (tl s)) else filter p (tl s)".

```

The elimination of co-recursive definition is done lazily, i.e. the definition is expanded only when it occurs at the head of an application which is the argument of a case analysis expression. In any other context, it is considered as a canonical expression which is completely evaluated. We can test this using the command `Eval`, which computes the normal forms of a term:

```

Coq < Eval compute in (from 0).
= (cofix from (n : nat) : Stream := Seq n (from (S n))) 0
: Stream

Coq < Eval compute in (hd (from 0)).

```

```

      = 0
      : nat

Coq < Eval compute in (tl (from 0)).
      = (cofix from (n : nat) : Stream := Seq n (from (S n))) 1
      : Stream

```

Variants:

1. `CoFixpoint ident1 params : type1 := term1`
 As for most constructions, arguments of co-fixpoints expressions can be introduced before the `:=` sign.
2. `CoFixpoint ident1 : type1 := term1`
`with`
`...`
`with identm : typem := termm`
 As in the `Fixpoint` command (see Section 1.3.4), it is possible to introduce a block of mutually dependent methods.

1.3.5 Assertions and proofs

An assertion states a proposition (or a type) of which the proof (or an inhabitant of the type) is interactively built using tactics. The interactive proof mode is described in Chapter 7 and the tactics in Chapter 8. The basic assertion command is:

```
Theorem ident [binders] : type.
```

After the statement is asserted, COQ needs a proof. Once a proof of *type* under the assumptions represented by *binders* is given and validated, the proof is generalized into a proof of `forall [binders], type` and the theorem is bound to the name *ident* in the environment.

Error messages:

1. The term *form* has type ... which should be Set, Prop or Type
2. *ident* already exists
 The name you provided is already defined. You have then to choose another name.

Variants:

1. `Lemma ident [binders] : type.`
`Remark ident [binders] : type.`
`Fact ident [binders] : type.`
`Corollary ident [binders] : type.`
`Proposition ident [binders] : type.`

These commands are synonyms of `Theorem ident [binders] : type.`

2. Theorem *ident [binders]: type with ... with ident [binders]: type.*

This command is useful for theorems that are proved by simultaneous induction over a mutually inductive assumption, or that assert mutually dependent statements in some mutual co-inductive type. It is equivalent to `Fixpoint` or `CoFixpoint` (see Section 1.3.4) but using tactics to build the proof of the statements (or the body of the specification, depending on the point of view). The inductive or co-inductive types on which the induction or coinduction has to be done is assumed to be non ambiguous and is guessed by the system.

Like in a `Fixpoint` or `CoFixpoint` definition, the induction hypotheses have to be used on *structurally smaller* arguments (for a `Fixpoint`) or be *guarded by a constructor* (for a `CoFixpoint`). The verification that recursive proof arguments are correct is done only at the time of registering the lemma in the environment. To know if the use of induction hypotheses is correct at some time of the interactive development of a proof, use the command `Guarded` (see Section 7.3.2).

The command can be used also with `Lemma`, `Remark`, etc. instead of `Theorem`.

3. Definition *ident [binders] : type.*

This allows defining a term of type *type* using the proof editing mode. It behaves as `Theorem` but is intended to be used in conjunction with `Defined` (see 1) in order to define a constant of which the computational behavior is relevant.

The command can be used also with `Example` instead of `Definition`.

See also: Sections 6.10.1 and 6.10.2 (Opaque and Transparent) and 8.7.5 (tactic `unfold`).

4. Let *ident [binders] : type.*

Like `Definition ident [binders] : type.` except that the definition is turned into a let-in definition generalized over the declarations depending on it after closing the current section.

5. `Fixpoint ident binders [annotation] [: term] [:= term] with ... with ident binders [annotation] [: term] [:= term].`

This generalizes the syntax of `Fixpoint` so that one or more bodies can be defined interactively using the proof editing mode (when a body is omitted, its type is mandatory in the syntax). When the block of proofs is completed, it is intended to be ended by `Defined`.

6. `CoFixpoint ident [binders] [: term] [:= term] with ... with ident [binders] [: term] [:= term].`

This generalizes the syntax of `CoFixpoint` so that one or more bodies can be defined interactively using the proof editing mode.

`Proof. ... Qed.`

A proof starts by the keyword `Proof`. Then COQ enters the proof editing mode until the proof is completed. The proof editing mode essentially contains tactics that are described in chapter 8. Besides tactics, there are commands to manage the proof editing mode. They are described in Chapter 7. When the proof is completed it should be validated and put in the environment using the keyword `Qed`.

Error message:

1. *ident* already exists

Remarks:

1. Several statements can be simultaneously asserted.
2. Not only other assertions but any vernacular command can be given while in the process of proving a given assertion. In this case, the command is understood as if it would have been given before the statements still to be proved.
3. `Proof` is recommended but can currently be omitted. On the opposite side, `Qed` (or `Defined`, see below) is mandatory to validate a proof.
4. Proofs ended by `Qed` are declared opaque. Their content cannot be unfolded (see 8.7), thus realizing some form of *proof-irrelevance*. To be able to unfold a proof, the proof should be ended by `Defined` (see below).

Variants:

1. `Proof. ... Defined.`
Same as `Proof. ... Qed.` but the proof is then declared transparent, which means that its content can be explicitly used for type-checking and that it can be unfolded in conversion tactics (see 8.7, 6.10.1, 6.10.2).
2. `Proof. ... Admitted.`
Turns the current asserted statement into an axiom and exits the proof mode.

Chapter 2

Extensions of GALLINA

GALLINA is the kernel language of COQ. We describe here extensions of the Gallina’s syntax.

2.1 Record types

The `Record` construction is a macro allowing the definition of records as is done in many programming languages. Its syntax is described on Figure 2.1. In fact, the `Record` macro is more general than the usual record types, since it allows also for “manifest” expressions. In this sense, the `Record` construction allows defining “signatures”.

<i>sentence</i>	++=	<i>record</i>
<i>record</i>	::=	<i>record_keyword ident [binders] [: sort] := [ident] { [field ; ... ; field] } .</i>
<i>record_keyword</i>	::=	<code>Record Inductive CoInductive</code>
<i>field</i>	::=	<i>name [binders] : type [where notation] name [binders] [: type] := term</i>

Figure 2.1: Syntax for the definition of `Record`

In the expression

```
Record ident params : sort := ident0 {  
  ident1 binders1 : term1 ; ... ;  
  identn bindersn : termn } .
```

the identifier *ident* is the name of the defined record and *sort* is its type. The identifier *ident*₀ is the name of its constructor. If *ident*₀ is omitted, the default name `Build_ident` is used. If *sort* is omitted, the default sort is `Type`. The identifiers *ident*₁, ..., *ident*_{*n*} are the names of fields and forall *binders*₁, *term*₁, ..., forall *binders*_{*n*}, *term*_{*n*} their respective types. Remark that the type of *ident*_{*i*} may depend on the previous *ident*_{*j*} (for *j* < *i*). Thus the order of the fields is important. Finally, *params* are the parameters of the record.

$\begin{array}{l} \text{term} \quad \quad \quad ++= \quad \{ \mid [\text{field_def} ; \dots ; \text{field_def}] \mid \} \\ \\ \text{field_def} \quad ::= \quad \text{name} [\text{binders}] := \text{term} \end{array}$
--

Figure 2.2: Syntax for constructing elements of a `Record` using named fields

More generally, a record may have explicitly defined (a.k.a. manifest) fields. For instance, `Record ident [params] : sort := { ident1 : type1 ; ident2 := term2 ; ident3 : type3 }` in which case the correctness of `type3` may rely on the instance `term2` of `ident2` and `term2` in turn may depend on `ident1`.

Example: The set of rational numbers may be defined as:

```
Coq < Record Rat : Set := mkRat
  {sign : bool;
   top : nat;
   bottom : nat;
   Rat_bottom_cond : 0 <> bottom;
   Rat_irred_cond :
     forall x y z:nat, (x * y) = top /\ (x * z) = bottom -> x = 1}.
Rat is defined
sign is defined
top is defined
bottom is defined
Rat_bottom_cond is defined
Rat_irred_cond is defined
```

Remark here that the field `Rat_bottom_cond` depends on the field `bottom` and `Rat_irred_cond` depends on both `top` and `bottom`.

Let us now see the work done by the `Record` macro. First the macro generates a variant type definition with just one constructor:

```
Variant ident params : sort :=
  ident0 (ident1 : term1) ... (identn : termn).
```

To build an object of type `ident`, one should provide the constructor `ident0` with n terms filling the fields of the record.

As an example, let us define the rational 1/2:

```
Coq < Theorem one_two_irred :
  forall x y z:nat, x * y = 1 /\ x * z = 2 -> x = 1.

Coq < Admitted.

Coq < Definition half := mkRat true 1 2 (O_S 1) one_two_irred.
half is defined

Coq < Check half.
half
  : Rat
```

Alternatively, the following syntax allows creating objects by using named fields, as shown on Figure 2.2. The fields do not have to be in any particular order, nor do they have to be all present if the missing ones can be inferred or prompted for (see Section 24).

```
Coq < Definition half' :=
  {| sign := true;
    Rat_bottom_cond := O_S 1;
    Rat_irred_cond := one_two_irred |}.
half' is defined
```

This syntax can be disabled globally for printing by

```
Unset Printing Records.
```

For a given type, one can override this using either

```
Add Printing Record ident.
```

to get record syntax or

```
Add Printing Constructor ident.
```

to get constructor syntax.

This syntax can also be used for pattern matching.

```
Coq < Eval compute in (
  match half with
  | {| sign := true; top := n |} => n
  | _ => 0
  end).
= 1
: nat
```

The macro generates also, when it is possible, the projection functions for destructuring an object of type *ident*. These projection functions are given the names of the corresponding fields. If a field is named “_” then no projection is built for it. In our example:

```
Coq < Eval compute in top half.
= 1
: nat

Coq < Eval compute in bottom half.
= 2
: nat

Coq < Eval compute in Rat_bottom_cond half.
= O_S 1
: 0 <> bottom half
```

An alternative syntax for projections based on a dot notation is available:

```
Coq < Eval compute in half.(top).
= 1
: nat
```

It can be activated for printing with the command

```
Set Printing Projections.
```

$ \begin{array}{lcl} \text{term} & \text{++=} & \text{term} . (\text{qualid}) \\ & & \text{term} . (\text{qualid arg} \dots \text{arg}) \\ & & \text{term} . (@\text{qualid term} \dots \text{term}) \end{array} $
--

Figure 2.3: Syntax for Record projections

```
Coq < Set Printing Projections.
```

```
Coq < Check top half.
```

```
half.(top)
      : nat
```

The corresponding grammar rules are given in Figure 2.3. When *qualid* denotes a projection, the syntax *term* . (*qualid*) is equivalent to *qualid term*, the syntax *term* . (*qualid arg*₁ ... *arg*_{*n*}) to *qualid arg*₁ ... *arg*_{*n*} *term*, and the syntax *term* . (@*qualid term*₁ ... *term*_{*n*}) to @*qualid term*₁ ... *term*_{*n*} *term*. In each case, *term* is the object projected and the other arguments are the parameters of the inductive type.

Remarks:

1. Records defined with the `Record` keyword are not allowed to be recursive (references to the record's name in the type of its field raises an error). To define recursive records, one can use the `Inductive` and `CoInductive` keywords, resulting in an inductive or co-inductive record. A *caveat*, however, is that records cannot appear in mutually inductive (or co-inductive) definitions.
2. Induction schemes are automatically generated for inductive records. Automatic generation of induction schemes for non-recursive records defined with the `Record` keyword can be activated with the `Nonrecursive Elimination Schemes` option (see 13.1.1).
3. `Structure` is a synonym of the keyword `Record`.

Warnings:

1. *ident*_{*i*} cannot be defined.

It can happen that the definition of a projection is impossible. This message is followed by an explanation of this impossibility. There may be three reasons:

- (a) The name *ident*_{*i*} already exists in the environment (see Section 1.3.1).
- (b) The body of *ident*_{*i*} uses an incorrect elimination for *ident* (see Sections 1.3.4 and 4.5.3).
- (c) The type of the projections *ident*_{*i*} depends on previous projections which themselves could not be defined.

Error messages:

1. Records declared with the keyword `Record` or `Structure` cannot be recursive.

The record name *ident* appears in the type of its fields, but uses the keyword `Record`. Use the keyword `Inductive` or `CoInductive` instead.

2. Cannot handle mutually (co)inductive records.

Records cannot be defined as part of mutually inductive (or co-inductive) definitions, whether with records only or mixed with standard definitions.

3. During the definition of the one-constructor inductive definition, all the errors of inductive definitions, as described in Section 1.3.3, may also occur.

See also: Coercions and records in Section 18.9 of the chapter devoted to coercions.

2.1.1 Primitive Projections

The option `Set Primitive Projections` turns on the use of primitive projections when defining subsequent records (even through the `Inductive` and `CoInductive` commands). Primitive projections extended the Calculus of Inductive Constructions with a new binary term constructor `r . (p)` representing a primitive projection `p` applied to a record object `r` (i.e., primitive projections are always applied). Even if the record type has parameters, these do not appear at applications of the projection, considerably reducing the sizes of terms when manipulating parameterized records and typechecking time. On the user level, primitive projections can be used as a replacement for the usual defined ones, although there are a few notable differences.

The internally omitted parameters can be reconstructed at printing time even though they are absent in the actual AST manipulated by the kernel. This can be obtained by setting the `Printing Primitive Projection Parameters` flag. Another compatibility printing can be activated thanks to the `Printing Primitive Projection Compatibility` option which governs the printing of pattern-matching over primitive records.

Primitive Record Types

When the `Set Primitive Projections` option is on, definitions of record types change meaning. When a type is declared with primitive projections, its `match` construct is disabled (see 2.1.1 though). To eliminate the (co-)inductive type, one must use its defined primitive projections.

There are currently two ways to introduce primitive records types:

- Through the `Record` command, in which case the type has to be non-recursive. The defined type enjoys eta-conversion definitionally, that is the generalized form of surjective pairing for records: $r = \text{Build_R } (r.(p_1) \dots r.(p_n))$. Eta-conversion allows to define dependent elimination for these types as well.
- Through the `Inductive` and `CoInductive` commands, when the body of the definition is a record declaration of the form `Build_R { p1 : t1; .. ; pn : tn }`. In this case the types can be recursive and eta-conversion is disallowed. These kind of record types differ from their traditional versions in the sense that dependent elimination is not available for them and only non-dependent case analysis can be defined.

Reduction

The basic reduction rule of a primitive projection is $p_i (\text{Build_R } t_1 \dots t_n) \rightarrow_\iota t_i$. However, to take the δ flag into account, projections can be in two states: folded or unfolded. An unfolded primitive

projection application obeys the rule above, while the folded version delta-reduces to the unfolded version. This allows to precisely mimic the usual unfolding rules of constants. Projections obey the usual `simpl` flags of the `Arguments` command in particular.

There is currently no way to input unfolded primitive projections at the user-level, and one must use the `Printing Primitive Projection Compatibility` to display unfolded primitive projections as matches and distinguish them from folded ones.

Compatibility Projections and `match`

To ease compatibility with ordinary record types, each primitive projection is also defined as a ordinary constant taking parameters and an object of the record type as arguments, and whose body is an application of the unfolded primitive projection of the same name. These constants are used when elaborating partial applications of the projection. One can distinguish them from applications of the primitive projection if the `Printing Primitive Projection Parameters` option is off: for a primitive projection application, parameters are printed as underscores while for the compatibility projections they are printed as usual.

Additionally, user-written `match` constructs on primitive records are desugared into substitution of the projections, they cannot be printed back as `match` constructs.

2.2 Variants and extensions of `match`

2.2.1 Multiple and nested pattern-matching

The basic version of `match` allows pattern-matching on simple patterns. As an extension, multiple nested patterns or disjunction of patterns are allowed, as in ML-like languages.

The extension just acts as a macro that is expanded during parsing into a sequence of `match` on simple patterns. Especially, a construction defined using the extended `match` is generally printed under its expanded form (see `Set Printing Matching` in section 2.2.4).

See also: Chapter 17.

2.2.2 Pattern-matching on boolean values: the `if` expression

For inductive types with exactly two constructors and for pattern-matchings expressions which do not depend on the arguments of the constructors, it is possible to use a `if ... then ... else` notation. For instance, the definition

```
Coq < Definition not (b:bool) :=
  match b with
  | true => false
  | false => true
  end.
not is defined
```

can be alternatively written

```
Coq < Definition not (b:bool) := if b then false else true.
not is defined
```

More generally, for an inductive type with constructors C_1 and C_2 , we have the following equivalence

$$\text{if } \text{term } [\text{dep_ret_type}] \text{ then } \text{term}_1 \text{ else } \text{term}_2 \equiv \begin{array}{l} \text{match } \text{term } [\text{dep_ret_type}] \text{ with} \\ | C_1 \text{ } _ \dots \text{ } _ \Rightarrow \text{term}_1 \\ | C_2 \text{ } _ \dots \text{ } _ \Rightarrow \text{term}_2 \\ \text{end} \end{array}$$

Here is an example.

```
Coq < Check (fun x (H:{x=0}+{x<>0}) =>
  match H with
  | left _ => true
  | right _ => false
end).
fun (x : nat) (H : {x = 0} + {x <> 0}) => if H then true else false
: forall x : nat, {x = 0} + {x <> 0} -> bool
```

Notice that the printing uses the `if` syntax because `sumbool` is declared as such (see Section 2.2.4).

2.2.3 Irrefutable patterns: the destructuring `let` variants

Pattern-matching on terms inhabiting inductive type having only one constructor can be alternatively written using `let ... in ...` constructions. There are two variants of them.

First destructuring `let` syntax

The expression `let ($\text{ident}_1, \dots, \text{ident}_n$) := term_0 in term_1` performs case analysis on a term_0 which must be in an inductive type with one constructor having itself n arguments. Variables $\text{ident}_1 \dots \text{ident}_n$ are bound to the n arguments of the constructor in expression term_1 . For instance, the definition

```
Coq < Definition fst (A B:Set) (H:A * B) := match H with
  | pair x y => x
end.
fst is defined
```

can be alternatively written

```
Coq < Definition fst (A B:Set) (p:A * B) := let (x, _) := p in x.
fst is defined
```

Notice that reduction is different from regular `let ... in ...` construction since it happens only if term_0 is in constructor form. Otherwise, the reduction is blocked.

The pretty-printing of a definition by matching on a irrefutable pattern can either be done using `match` or the `let` construction (see Section 2.2.4).

If term inhabits an inductive type with one constructor C , we have an equivalence between

`let ($\text{ident}_1, \dots, \text{ident}_n$) [dep_ret_type] := term in term'`

and

`match term [dep_ret_type] with C $\text{ident}_1 \dots \text{ident}_n \Rightarrow \text{term}'$ end`

Second destructuring `let` syntax

Another destructuring `let` syntax is available for inductive types with one constructor by giving an arbitrary pattern instead of just a tuple for all the arguments. For example, the preceding example can be written:

```
Coq < Definition fst (A B:Set) (p:A*B) := let 'pair x _ := p in x.
fst is defined
```

This is useful to match deeper inside tuples and also to use notations for the pattern, as the syntax `let 'p := t in b` allows arbitrary patterns to do the deconstruction. For example:

```
Coq < Definition deep_tuple (A:Set) (x:(A*A)*(A*A)) : A*A*A*A :=
  let '((a,b), (c, d)) := x in (a,b,c,d).
deep_tuple is defined

Coq < Notation " x 'With' p " := (exist _ x p) (at level 20).
Identifier 'With' now a keyword

Coq < Definition proj1_sig' (A:Set) (P:A->Prop) (t:{ x:A | P x }) : A :=
  let 'x With p := t in x.
proj1_sig' is defined
```

When printing definitions which are written using this construct it takes precedence over `let` printing directives for the datatype under consideration (see Section 2.2.4).

2.2.4 Controlling pretty-printing of `match` expressions

The following commands give some control over the pretty-printing of `match` expressions.

Printing nested patterns

The Calculus of Inductive Constructions knows pattern-matching only over simple patterns. It is however convenient to re-factorize nested pattern-matching into a single pattern-matching over a nested pattern. COQ's printer try to do such limited re-factorization.

```
Set Printing Matching.
```

This tells COQ to try to use nested patterns. This is the default behavior.

```
Unset Printing Matching.
```

This tells COQ to print only simple pattern-matching problems in the same way as the COQ kernel handles them.

```
Test Printing Matching.
```

This tells if the printing matching mode is on or off. The default is on.

Printing of wildcard pattern

Some variables in a pattern may not occur in the right-hand side of the pattern-matching clause. There are options to control the display of these variables.

```
Set Printing Wildcard.
```

The variables having no occurrences in the right-hand side of the pattern-matching clause are just printed using the wildcard symbol “_”.

```
Unset Printing Wildcard.
```

The variables, even useless, are printed using their usual name. But some non dependent variables have no name. These ones are still printed using a “_”.

```
Test Printing Wildcard.
```

This tells if the wildcard printing mode is on or off. The default is to print wildcard for useless variables.

Printing of the elimination predicate

In most of the cases, the type of the result of a matched term is mechanically synthesizable. Especially, if the result type does not depend of the matched term.

```
Set Printing Synth.
```

The result type is not printed when COQ knows that it can re-synthesize it.

```
Unset Printing Synth.
```

This forces the result type to be always printed.

```
Test Printing Synth.
```

This tells if the non-printing of synthesizable types is on or off. The default is to not print synthesizable types.

Printing matching on irrefutable pattern

If an inductive type has just one constructor, pattern-matching can be written using the first destructuring let syntax.

```
Add Printing Let ident.
```

This adds *ident* to the list of inductive types for which pattern-matching is written using a `let` expression.

```
Remove Printing Let ident.
```

This removes *ident* from this list. Note that removing an inductive type from this list has an impact only for pattern-matching written using `match`. Pattern-matching explicitly written using a destructuring `let` are not impacted.

```
Test Printing Let for ident.
```

This tells if *ident* belongs to the list.

```
Print Table Printing Let.
```

This prints the list of inductive types for which pattern-matching is written using a `let` expression.

The list of inductive types for which pattern-matching is written using a `let` expression is managed synchronously. This means that it is sensible to the command `Reset`.

Printing matching on booleans

If an inductive type is isomorphic to the boolean type, pattern-matching can be written using `if ... then ... else ...`

```
Add Printing If ident.
```

This adds *ident* to the list of inductive types for which pattern-matching is written using an `if` expression.

```
Remove Printing If ident.
```

This removes *ident* from this list.

```
Test Printing If for ident.
```

This tells if *ident* belongs to the list.

```
Print Table Printing If.
```

This prints the list of inductive types for which pattern-matching is written using an `if` expression.

The list of inductive types for which pattern-matching is written using an `if` expression is managed synchronously. This means that it is sensible to the command `Reset`.

Example

This example emphasizes what the printing options offer.

```
Coq < Definition snd (A B:Set) (H:A * B) := match H with
                                     | pair x y => y
                                     end.
```

snd is defined

```
Coq < Test Printing Let for prod.
```

Cases on elements of prod are printed using a 'let' form

```
Coq < Print snd.
```

snd =

```
fun (A B : Set) (H : A * B) => let (_, y) := H in y
    : forall A B : Set, A * B -> B
```

Argument scopes are [type_scope type_scope _]

```
Coq < Remove Printing Let prod.
```

```

Coq < Unset Printing Synth.
Coq < Unset Printing Wildcard.
Coq < Print snd.
snd =
fun (A B : Set) (H : A * B) => match H return B with
    | (x, y) => y
    end
    : forall A B : Set, A * B -> B
Argument scopes are [type_scope type_scope _]

```

2.2.5 Printing match templates

The `Show Match` vernacular command prints a match template for a given type. See Section 7.3.1.

2.3 Advanced recursive functions

The following *experimental* command is available when the `FunInd` library has been loaded via `Require Import FunInd`:

```
Function ident binder1...bindern {decrease_annot} : type0 := term0
```

This command can be seen as a generalization of `Fixpoint`. It is actually a wrapper for several ways of defining a function *and other useful related objects*, namely: an induction principle that reflects the recursive structure of the function (see 8.5.5) and its fixpoint equality. The meaning of this declaration is to define a function *ident*, similarly to `Fixpoint`. Like in `Fixpoint`, the decreasing argument must be given (unless the function is not recursive), but it might not necessarily be *structurally* decreasing. The point of the `{}` annotation is to name the decreasing argument *and* to describe which kind of decreasing criteria must be used to ensure termination of recursive calls.

The `Function` construction also enjoys the `with` extension to define mutually recursive definitions. However, this feature does not work for non structurally recursive functions.

See the documentation of functional induction (see Section 8.5.5) and Functional Scheme (see Section 13.2 and 13.2) for how to use the induction principle to easily reason about the function.

Remark: To obtain the right principle, it is better to put rigid parameters of the function as first arguments. For example it is better to define `plus` like this:

```

Coq < Function plus (m n : nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (plus m p)
  end.

```

than like this:

```

Coq < Function plus (n m : nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (plus p m)
  end.

```

Limitations $term_0$ must be built as a *pure pattern-matching tree* (`match...with`) with applications only *at the end* of each branch.

Function does not support partial application of the function being defined. Thus, the following example cannot be accepted due to the presence of partial application of *identwrong* into the body of *identwrong*:

```
Coq < Fail Function wrong (C:nat) : nat :=
  List.hd 0 (List.map wrong (C::nil)).
```

For now dependent cases are not treated for non structurally terminating functions.

Error messages:

1. The recursive argument must be specified
2. No argument name *ident*
3. Cannot use mutual definition with well-founded recursion or measure
4. Cannot define graph for *ident...* (warning)

The generation of the graph relation (R_{ident}) used to compute the induction scheme of *ident* raised a typing error. Only the *ident* is defined; the induction scheme will not be generated.

This error happens generally when:

- the definition uses pattern matching on dependent types, which `Function` cannot deal with yet.
- the definition is not a *pattern-matching tree* as explained above.

5. Cannot define principle(s) for *ident...* (warning)

The generation of the graph relation (R_{ident}) succeeded but the induction principle could not be built. Only the *ident* is defined. Please report.

6. Cannot build functional inversion principle (warning)
functional inversion will not be available for the function.

See also: [13.2](#), [13.2](#), [8.5.5](#)

Depending on the $\{\dots\}$ annotation, different definition mechanisms are used by `Function`. More precise description given below.

Variants:

1. Function *ident* $binder_1 \dots binder_n : type_0 := term_0$

Defines the not recursive function *ident* as if declared with `Definition`. Moreover the following are defined:

- *ident_rect*, *ident_rec* and *ident_ind*, which reflect the pattern matching structure of $term_0$ (see the documentation of `Inductive` [1.3.3](#));
- The inductive R_{ident} corresponding to the graph of *ident* (silently);

- *ident_complete* and *ident_correct* which are inversion information linking the function and its graph.
2. Function *ident* *binder*₁...*binder*_{*n*} {struct *ident*₀} : type₀ := *term*₀
 Defines the structural recursive function *ident* as if declared with `Fixpoint`. Moreover the following are defined:
 - The same objects as above;
 - The fixpoint equation of *ident*: *ident_equation*.
 3. Function *ident* *binder*₁...*binder*_{*n*} {measure *term*₁ *ident*₀} : type₀ := *term*₀
 4. Function *ident* *binder*₁...*binder*_{*n*} {wf *term*₁ *ident*₀} : type₀ := *term*₀
 Defines a recursive function by well founded recursion. **The module `Recdef` of the standard library must be loaded for this feature.** The { } annotation is mandatory and must be one of the following:
 - {measure *term*₁ *ident*₀} with *ident*₀ being the decreasing argument and *term*₁ being a function from type of *ident*₀ to nat for which value on the decreasing argument decreases (for the `lt` order on nat) at each recursive call of *term*₀. Parameters of the function are bound in *term*₀;
 - {wf *term*₁ *ident*₀} with *ident*₀ being the decreasing argument and *term*₁ an ordering relation on the type of *ident*₀ (i.e. of type $T_{ident_0} \rightarrow T_{ident_0} \rightarrow Prop$) for which the decreasing argument decreases at each recursive call of *term*₀. The order must be well founded. Parameters of the function are bound in *term*₀.

Depending on the annotation, the user is left with some proof obligations that will be used to define the function. These proofs are: proofs that each recursive call is actually decreasing with respect to the given criteria, and (if the criteria is `wf`) a proof that the ordering relation is well founded.

Once proof obligations are discharged, the following objects are defined:

- The same objects as with the `struct`;
- The lemma *ident_tcc* which collects all proof obligations in one property;
- The lemmas *ident_terminate* and *ident_F* which is needed to be inlined during extraction of *ident*.

The way this recursive function is defined is the subject of several papers by Yves Bertot and Antonia Balaa on the one hand, and Gilles Barthe, Julien Forest, David Pichardie, and Vlad Rusu on the other hand.

Remark: Proof obligations are presented as several subgoals belonging to a Lemma *ident_tcc*.

2.4 Section mechanism

The sectioning mechanism can be used to to organize a proof in structured sections. Then local declarations become available (see Section 1.3.2).

2.4.1 Section *ident*

This command is used to open a section named *ident*.

2.4.2 End *ident*

This command closes the section named *ident*. After closing of the section, the local declarations (variables and local definitions) get *discharged*, meaning that they stop being visible and that all global objects defined in the section are generalized with respect to the variables and local definitions they each depended on in the section.

Here is an example :

```
Coq < Section s1.
Coq < Variables x y : nat.
x is declared
y is declared
Coq < Let y' := y.
y' is defined
Coq < Definition x' := S x.
x' is defined
Coq < Definition x'' := x' + y'.
x'' is defined
Coq < Print x'.
x' = S x
      : nat
Coq < End s1.
Coq < Print x'.
x' = fun x : nat => S x
      : nat -> nat
Argument scope is [nat_scope]
Coq < Print x''.
x'' = fun x y : nat => let y' := y in x' x + y'
      : nat -> nat -> nat
Argument scopes are [nat_scope nat_scope]
```

Notice the difference between the value of x' and x'' inside section `s1` and outside.

Error messages:

1. This is not the last opened section

Remarks:

1. Most commands, like `Hint`, `Notation`, option management, ... which appear inside a section are canceled when the section is closed.

2.5 Module system

The module system provides a way of packaging related elements together, as well as a means of massive abstraction.

In the syntax of module application, the `!` prefix indicates that any `Inline` directive in the type of the functor arguments will be ignored (see 2.5.4 below).

```

    module_type ::= qualid
                | module_type with Definition qualid := term
                | module_type with Module qualid := qualid
                | qualid qualid ... qualid
                | !qualid qualid ... qualid

    module_binding ::= ( [Import|Export] ident ... ident : module_type )

    module_bindings ::= module_binding ... module_binding

    module_expression ::= qualid ... qualid
                      | !qualid ... qualid

```

Figure 2.4: Syntax of modules

2.5.1 Module *ident*

This command is used to start an interactive module named *ident*.

Variants:

1. `Module ident module_bindings`
Starts an interactive functor with parameters given by *module_bindings*.
2. `Module ident : module_type`
Starts an interactive module specifying its module type.
3. `Module ident module_bindings : module_type`
Starts an interactive functor with parameters given by *module_bindings*, and output module type *module_type*.
4. `Module ident <: module_type1 <: ... <: module_typen`
Starts an interactive module satisfying each *module_type_i*.
5. `Module ident module_bindings <: module_type1 <: ... <: module_typen`
Starts an interactive functor with parameters given by *module_bindings*. The output module type is verified against each module type *module_type_i*.
6. `Module [Import|Export]`
Behaves like `Module`, but automatically imports or exports the module.

Reserved commands inside an interactive module:

1. `Include module`
Includes the content of *module* in the current interactive module. Here *module* can be a module expression or a module type expression. If *module* is a high-order module or module type expression then the system tries to instantiate *module* by the current interactive module.
2. `Include module1 <+ ... <+ modulen`
is a shortcut for `Include module1 ... Include modulen`

2.5.2 End *ident*

This command closes the interactive module *ident*. If the module type was given the content of the module is matched against it and an error is signaled if the matching fails. If the module is basic (is not a functor) its components (constants, inductive types, submodules etc) are now available through the dot notation.

Error messages:

1. No such label *ident*
2. Signature components for label *ident* do not match
3. This is not the last opened module

2.5.3 Module *ident* := *module_expression*

This command defines the module identifier *ident* to be equal to *module_expression*.

Variants:

1. Module *ident* *module_bindings* := *module_expression*
 Defines a functor with parameters given by *module_bindings* and body *module_expression*.
2. Module *ident* *module_bindings* : *module_type* := *module_expression*
 Defines a functor with parameters given by *module_bindings* (possibly none), and output module type *module_type*, with body *module_expression*.
3. Module *ident* *module_bindings* <: *module_type*₁ <: ... <: *module_type*_{*n*} := *module_expression*
 Defines a functor with parameters given by *module_bindings* (possibly none) with body *module_expression*. The body is checked against each *module_type*_{*i*}.
4. Module *ident* *module_bindings* := *module_expression*₁ <+ ... <+ *module_expression*_{*n*}
 is equivalent to an interactive module where each *module_expression*_{*i*} are included.

2.5.4 Module Type *ident*

This command is used to start an interactive module type *ident*.

Variants:

1. Module Type *ident* *module_bindings*
 Starts an interactive functor type with parameters given by *module_bindings*.

Reserved commands inside an interactive module type:

1. Include *module*
Same as Include inside a module.
2. Include *module*₁ <+ ... <+ *module*_{*n*}
is a shortcut for Include *module*₁ ... Include *module*_{*n*}
3. *assumption_keyword* Inline *assums*
The instance of this assumption will be automatically expanded at functor application, except when this functor application is prefixed by a ! annotation.

2.5.5 End *ident*

This command closes the interactive module type *ident*.

Error messages:

1. This is not the last opened module type

2.5.6 Module Type *ident* := *module_type*

Defines a module type *ident* equal to *module_type*.

Variants:

1. Module Type *ident module_bindings* := *module_type*
Defines a functor type *ident* specifying functors taking arguments *module_bindings* and returning *module_type*.
2. Module Type *ident module_bindings* := *module_type*₁ <+ ... <+ *module_type*_{*n*}
is equivalent to an interactive module type were each *module_type*_{*i*} are included.

2.5.7 Declare Module *ident* : *module_type*

Declares a module *ident* of type *module_type*.

Variants:

1. Declare Module *ident module_bindings* : *module_type*
Declares a functor with parameters *module_bindings* and output module type *module_type*.

Example

Let us define a simple module.

```
Coq < Module M.
Interactive Module M started
Coq <   Definition T := nat.
T is defined
Coq <   Definition x := 0.
```

```

x is defined

Coq < Definition y : bool.
1 subgoal

=====
bool

Coq < exact true.
No more subgoals.

Coq < Defined.
y is defined

Coq < End M.
Module M is defined

```

Inside a module one can define constants, prove theorems and do any other things that can be done in the toplevel. Components of a closed module can be accessed using the dot notation:

```

Coq < Print M.x.
M.x = 0
      : nat

```

A simple module type:

```

Coq < Module Type SIG.
Interactive Module Type SIG started

Coq < Parameter T : Set.
T is declared

Coq < Parameter x : T.
x is declared

Coq < End SIG.
Module Type SIG is defined

```

Now we can create a new module from M, giving it a less precise specification: the *y* component is dropped as well as the body of *x*.

```

Coq < Module N : SIG with Definition T := nat := M.
Module N is defined

Coq < Print N.T.
N.T = nat
      : Set

Coq < Print N.x.
*** [ N.x : N.T ]

Coq < Fail Print N.y.
The command has indeed failed with message:
N.y not a defined object.

```

The definition of N using the module type expression SIG with Definition T:=nat is equivalent to the following one:

```

Coq < Module Type SIG'.
Coq <   Definition T : Set := nat.
Coq <   Parameter x : T.
Coq < End SIG'.
Coq < Module N : SIG' := M.

```

If we just want to be sure that our implementation satisfies a given module type without restricting the interface, we can use a transparent constraint

```

Coq < Module P <: SIG := M.
Module P is defined
Coq < Print P.y.
P.y = true
      : bool

```

Now let us create a functor, i.e. a parametric module

```

Coq < Module Two (X Y: SIG).
Interactive Module Two started
Coq <   Definition T := (X.T * Y.T)%type.
Coq <   Definition x := (X.x, Y.x).
Coq < End Two.
Module Two is defined

```

and apply it to our modules and do some computations

```

Coq < Module Q := Two M N.
Module Q is defined
Coq < Eval compute in (fst Q.x + snd Q.x).
      = N.x
      : nat

```

In the end, let us define a module type with two sub-modules, sharing some of the fields and give one of its possible implementations:

```

Coq < Module Type SIG2.
Interactive Module Type SIG2 started
Coq <   Declare Module M1 : SIG.
Module M1 is declared
Coq <   Module M2 <: SIG.
Interactive Module M2 started
Coq <       Definition T := M1.T.
T is defined
Coq <       Parameter x : T.
x is declared
Coq <   End M2.
Module M2 is defined
Coq < End SIG2.

```

```

Module Type SIG2 is defined
Coq < Module Mod <: SIG2.
Coq <   Module M1.
Coq <       Definition T := nat.
Coq <       Definition x := 1.
Coq <   End M1.
Coq <   Module M2 := M.
Coq < End Mod.
Module Mod is defined

```

Notice that `M` is a correct body for the component `M2` since its `T` component is equal `nat` and hence `M1.T` as specified.

Remarks:

1. Modules and module types can be nested components of each other.
2. One can have sections inside a module or a module type, but not a module or a module type inside a section.
3. Commands like `Hint` or `Notation` can also appear inside modules and module types. Note that in case of a module definition like:

```
Module N : SIG := M.
```

or

```

Module N : SIG.
...
End N.

```

hints and the like valid for `N` are not those defined in `M` (or the module body) but the ones defined in `SIG`.

2.5.8 Import *qualid*

If *qualid* denotes a valid basic module (i.e. its module type is a signature), makes its components available by their short names.

Example:

```

Coq < Module Mod.
Interactive Module Mod started
Coq <   Definition T:=nat.
T is defined
Coq <   Check T.
T
      : Set
Coq < End Mod.
Module Mod is defined

```

```

Coq < Check Mod.T.
Mod.T
      : Set

Coq < Fail Check T. (* Incorrect! *)
The command has indeed failed with message:
The reference T was not found in the current environment.

Coq < Import Mod.

Coq < Check T. (* Now correct *)
T
      : Set

```

Some features defined in modules are activated only when a module is imported. This is for instance the case of notations (see Section 12.1).

Declarations made with the `Local` flag are never imported by the `Import` command. Such declarations are only accessible through their fully qualified name.

Example:

```

Coq < Module A.
Interactive Module A started

Coq < Module B.
Interactive Module B started

Coq < Local Definition T := nat.
T is defined

Coq < End B.
Module B is defined

Coq < End A.
Module A is defined

Coq < Import A.

Coq < Fail Check B.T.
The command has indeed failed with message:
The reference B.T was not found in the current environment.

```

Variants:

1. Export *qualid*

When the module containing the command `Export qualid` is imported, *qualid* is imported as well.

Error messages:

1. *qualid* is not a module

Warnings:

1. Trying to mask the absolute name *qualid* !

2.5.9 Print Module *ident*

Prints the module type and (optionally) the body of the module *ident*.

For this command and `Print Module Type`, the option `Short Module Printing` (off by default) disables the printing of the types of fields, leaving only their names.

2.5.10 Print Module Type *ident*

Prints the module type corresponding to *ident*.

2.5.11 Locate Module *qualid*

Prints the full name of the module *qualid*.

2.6 Libraries and qualified names

2.6.1 Names of libraries

The theories developed in COQ are stored in *library files* which are hierarchically classified into *libraries* and *sublibraries*. To express this hierarchy, library names are represented by qualified identifiers *qualid*, i.e. as list of identifiers separated by dots (see Section 1.2.3). For instance, the library file `Mult` of the standard COQ library `Arith` is named `Coq.Arith.Mult`. The identifier that starts the name of a library is called a *library root*. All library files of the standard library of COQ have the reserved root `Coq` but library file names based on other roots can be obtained by using COQ commands (`coqc`, `coqtop`, `coqdep`, ...) options `-Q` or `-R` (see Section 14.3.3). Also, when an interactive COQ session starts, a library of root `Top` is started, unless option `-top` is set (see Section 14.3.3).

2.6.2 Qualified names

Library files are modules which possibly contain submodules which eventually contain constructions (axioms, parameters, definitions, lemmas, theorems, remarks or facts). The *absolute name*, or *full name*, of a construction in some library file is a qualified identifier starting with the logical name of the library file, followed by the sequence of submodules names encapsulating the construction and ended by the proper name of the construction. Typically, the absolute name `Coq.Init.Logic.eq` denotes Leibniz' equality defined in the module `Logic` in the sublibrary `Init` of the standard library of COQ.

The proper name that ends the name of a construction is the *short name* (or sometimes *base name*) of the construction (for instance, the short name of `Coq.Init.Logic.eq` is `eq`). Any partial suffix of the absolute name is a *partially qualified name* (e.g. `Logic.eq` is a partially qualified name for `Coq.Init.Logic.eq`). Especially, the short name of a construction is its shortest partially qualified name.

COQ does not accept two constructions (definition, theorem, ...) with the same absolute name but different constructions can have the same short name (or even same partially qualified names as soon as the full names are different).

Notice that the notion of absolute, partially qualified and short names also applies to library file names.

Visibility Coq maintains a table called *name table* which maps partially qualified names of constructions to absolute names. This table is updated by the commands `Require` (see 6.5.1), `Import` and `Export` (see 2.5.8) and also each time a new declaration is added to the context. An absolute name is called *visible* from a given short or partially qualified name when this latter name is enough to denote it. This means that the short or partially qualified name is mapped to the absolute name in COQ name table. Definitions flagged as `Local` are only accessible with their fully qualified name (see 1.3.2).

It may happen that a visible name is hidden by the short name or a qualified name of another construction. In this case, the name that has been hidden must be referred to using one more level of qualification. To ensure that a construction always remains accessible, absolute names can never be hidden.

Examples:

```
Coq < Check 0.
0
      : nat

Coq < Definition nat := bool.
nat is defined

Coq < Check 0.
0
      : Datatypes.nat

Coq < Check Datatypes.nat.
Datatypes.nat
      : Set

Coq < Locate nat.
Constant Top.nat
Inductive Coq.Init.Datatypes.nat
  (shorter name to refer to it in current context is Datatypes.nat)
```

See also: Command `Locate` in Section 6.3.10 and `Locate Library` in Section 6.6.11.

2.6.3 Libraries and filesystem

Please note that the questions described here have been subject to redesign in Coq v8.5. Former versions of Coq use the same terminology to describe slightly different things.

Compiled files (`.vo` and `.vio`) store sub-libraries. In order to refer to them inside COQ, a translation from file-system names to COQ names is needed. In this translation, names in the file system are called *physical* paths while COQ names are contrastingly called *logical* names.

A logical prefix `Lib` can be associated to a physical path *path* using the command line option `-Q path Lib`. All subfolders of *path* are recursively associated to the logical path `Lib` extended with the corresponding suffix coming from the physical path. For instance, the folder `path/f00/Bar` maps to `Lib.f00.Bar`. Subdirectories corresponding to invalid COQ identifiers are skipped, and, by convention, subdirectories named `CVS` or `_darcs` are skipped too.

Thanks to this mechanism, `.vo` files are made available through the logical name of the folder they are in, extended with their own basename. For example, the name associated to the file `path/f00/Bar/File.vo` is `Lib.f00.Bar.File`. The same caveat applies for invalid identifiers. When compiling a source file, the `.vo` file stores its logical name, so that an error is issued if it is loaded with the wrong loadpath afterwards.

Some folders have a special status and are automatically put in the path. COQ commands associate automatically a logical path to files in the repository trees rooted at the directory from where the command is launched, *coqlib*/user-contrib/, the directories listed in the `$COQPATH`, `${XDG_DATA_HOME}/coq/` and `${XDG_DATA_DIRS}/coq/` environment variables (see <http://standards.freedesktop.org/basedir-spec/basedir-spec-latest.html>) with the same physical-to-logical translation and with an empty logical prefix.

The command line option `-R` is a variant of `-Q` which has the strictly same behavior regarding loadpaths, but which also makes the corresponding `.vo` files available through their short names in a way not unlike the `Import` command (see 2.5.8). For instance, `-R path Lib` associates to the file `path/foo/Bar/File.vo` the logical name `Lib.foo.Bar.File`, but allows this file to be accessed through the short names `foo.Bar.File`, `Bar.File` and `File`. If several files with identical base name are present in different subdirectories of a recursive loadpath, which of these files is found first may be system-dependent and explicit qualification is recommended. The `From` argument of the `Require` command can be used to bypass the implicit shortening by providing an absolute root to the required file (see 6.5.1).

There also exists another independent loadpath mechanism attached to OCAML object files (`.cmo` or `.cmxs`) rather than COQ object files as described above. The OCAML loadpath is managed using the option `-I path` (in the OCAML world, there is neither a notion of logical name prefix nor a way to access files in subdirectories of `path`). See the command `Declare ML Module` in Section 6.5 to understand the need of the OCAML loadpath.

See Section 14.3.3 for a more general view over the COQ command line options.

2.7 Implicit arguments

An implicit argument of a function is an argument which can be inferred from contextual knowledge. There are different kinds of implicit arguments that can be considered implicit in different ways. There are also various commands to control the setting or the inference of implicit arguments.

2.7.1 The different kinds of implicit arguments

Implicit arguments inferable from the knowledge of other arguments of a function

The first kind of implicit arguments covers the arguments that are inferable from the knowledge of the type of other arguments of the function, or of the type of the surrounding context of the application. Especially, such implicit arguments correspond to parameters dependent in the type of the function. Typical implicit arguments are the type arguments in polymorphic functions. There are several kinds of such implicit arguments.

Strict Implicit Arguments. An implicit argument can be either strict or non strict. An implicit argument is said *strict* if, whatever the other arguments of the function are, it is still inferable from the type of some other argument. Technically, an implicit argument is strict if it corresponds to a parameter which is not applied to a variable which itself is another parameter of the function (since this parameter may erase its arguments), not in the body of a `match`, and not itself applied or matched against patterns (since the original form of the argument can be lost by reduction).

For instance, the first argument of

```
cons: forall A:Set, A -> list A -> list A
```


in module `List.v` is strict because `list` is an inductive type and `A` will always be inferable from the type `list A` of the third argument of `cons`. On the contrary, the second argument of a term of type

```
forall P:nat->Prop, forall n:nat, P n -> ex nat P
```

is implicit but not strict, since it can only be inferred from the type `P n` of the third argument and if `P` is, e.g., `fun _ => True`, it reduces to an expression where `n` does not occur any longer. The first argument `P` is implicit but not strict either because it can only be inferred from `P n` and `P` is not canonically inferable from an arbitrary `n` and the normal form of `P n` (consider e.g. that `n` is 0 and the third argument has type `True`, then any `P` of the form `fun n => match n with 0 => True | _ => anything` end would be a solution of the inference problem).

Contextual Implicit Arguments. An implicit argument can be *contextual* or not. An implicit argument is said *contextual* if it can be inferred only from the knowledge of the type of the context of the current expression. For instance, the only argument of

```
nil : forall A:Set, list A
```

is contextual. Similarly, both arguments of a term of type

```
forall P:nat->Prop, forall n:nat, P n \/ n = 0
```

are contextual (moreover, `n` is strict and `P` is not).

Reversible-Pattern Implicit Arguments. There is another class of implicit arguments that can be reinferred unambiguously if all the types of the remaining arguments are known. This is the class of implicit arguments occurring in the type of another argument in position of reversible pattern, which means it is at the head of an application but applied only to uninstantiated distinct variables. Such an implicit argument is called *reversible-pattern implicit argument*. A typical example is the argument `P` of `nat_rec` in

```
nat_rec : forall P : nat -> Set, P 0 -> (forall n : nat, P
n -> P (S n)) -> forall x : nat, P x.
```

(`P` is reinferable by abstracting over `n` in the type `P n`).

See Section 2.7.9 for the automatic declaration of reversible-pattern implicit arguments.

Implicit arguments inferable by resolution

This corresponds to a class of non dependent implicit arguments that are solved based on the structure of their type only.

2.7.2 Maximal or non maximal insertion of implicit arguments

In case a function is partially applied, and the next argument to be applied is an implicit argument, two disciplines are applicable. In the first case, the function is considered to have no arguments furtherly: one says that the implicit argument is not maximally inserted. In the second case, the function is considered to be implicitly applied to the implicit arguments it is waiting for: one says that the implicit argument is maximally inserted.

Each implicit argument can be declared to have to be inserted maximally or non maximally. This can be governed argument per argument by the command `Implicit Arguments` (see 2.7.4) or globally by the command `Set Maximal Implicit Insertion` (see 2.7.10). See also Section 2.7.13.

2.7.3 Casual use of implicit arguments

In a given expression, if it is clear that some argument of a function can be inferred from the type of the other arguments, the user can force the given argument to be guessed by replacing it by “_”. If possible, the correct argument will be automatically generated.

Error messages:

1. Cannot infer a term for this placeholder
Coq was not able to deduce an instantiation of a “_”.

2.7.4 Declaration of implicit arguments

In case one wants that some arguments of a given object (constant, inductive types, constructors, assumptions, local or not) are always inferred by Coq, one may declare once and for all which are the expected implicit arguments of this object. There are two ways to do this, a priori and a posteriori.

Implicit Argument Binders

In the first setting, one wants to explicitly give the implicit arguments of a declared object as part of its definition. To do this, one has to surround the bindings of implicit arguments by curly braces:

```
Coq < Definition id {A : Type} (x : A) : A := x.
id is defined
```

This automatically declares the argument A of id as a maximally inserted implicit argument. One can then do as-if the argument was absent in every situation but still be able to specify it if needed:

```
Coq < Definition compose {A B C} (g : B -> C) (f : A -> B) :=
  fun x => g (f x).
compose is defined
```

```
Coq < Goal forall A, compose id id = id (A:=A).
1 subgoal
```

```
=====
forall A : Type, compose id id = id
```

The syntax is supported in all top-level definitions: Definition, Fixpoint, Lemma and so on. For (co-)inductive datatype declarations, the semantics are the following: an inductive parameter declared as an implicit argument need not be repeated in the inductive definition but will become implicit for the constructors of the inductive only, not the inductive type itself. For example:

```
Coq < Inductive list {A : Type} : Type :=
  | nil : list
  | cons : A -> list -> list.
list is defined
list_rect is defined
list_ind is defined
list_rec is defined

Coq < Print list.
Inductive list (A : Type) : Type := nil : list | cons : A -> list -> list
For list: Argument A is implicit and maximally inserted
```

For nil: Argument A is implicit and maximally inserted
For cons: Argument A is implicit and maximally inserted
For list: Argument scope is [type_scope]
For nil: Argument scope is [type_scope]
For cons: Argument scopes are [type_scope _ _]

One can always specify the parameter if it is not uniform using the usual implicit arguments disambiguation syntax.

Declaring Implicit Arguments

To set implicit arguments a posteriori, one can use the command:

```
Arguments qualid possibly_bracketed_ident ... possibly_bracketed_ident
```

where the list of *possibly_bracketed_ident* is a prefix of the list of arguments of *qualid* where the ones to be declared implicit are surrounded by square brackets and the ones to be declared as maximally inserted implicits are surrounded by curly braces.

After the above declaration is issued, implicit arguments can just (and have to) be skipped in any expression involving an application of *qualid*.

Implicit arguments can be cleared with the following syntax:

```
Arguments qualid : clear implicits
```

Variants:

1. Global Arguments *qualid* *possibly_bracketed_ident* ... *possibly_bracketed_ident*

Tell to recompute the implicit arguments of *qualid* after ending of the current section if any, enforcing the implicit arguments known from inside the section to be the ones declared by the command.

2. Local Arguments *qualid* *possibly_bracketed_ident* ... *possibly_bracketed_ident*

When in a module, tell not to activate the implicit arguments of *qualid* declared by this command to contexts that require the module.

3. [Global / Local] Arguments *qualid* [*possibly_bracketed_ident* ... *possibly_bracketed_ident* , ... , *possibly_bracketed_ident* ... *possibly_bracketed_ident*]

For names of constants, inductive types, constructors, lemmas which can only be applied to a fixed number of arguments (this excludes for instance constants whose type is polymorphic), multiple implicit arguments declarations can be given. Depending on the number of arguments *qualid* is applied to in practice, the longest applicable list of implicit arguments is used to select which implicit arguments are inserted.

For printing, the omitted arguments are the ones of the longest list of implicit arguments of the sequence.

Example:

```

Coq < Inductive list (A:Type) : Type :=
  | nil : list A
  | cons : A -> list A -> list A.

Coq < Check (cons nat 3 (nil nat)).
cons nat 3 (nil nat)
  : list nat

Coq < Arguments cons [A] _ _.
Coq < Arguments nil [A].
Coq < Check (cons 3 nil).
cons 3 nil
  : list nat

Coq < Fixpoint map (A B:Type) (f:A->B) (l:list A) : list B :=
  match l with nil => nil | cons a t => cons (f a) (map A B f t) end.
map is defined
map is recursively defined (decreasing on 4th argument)

Coq < Fixpoint length (A:Type) (l:list A) : nat :=
  match l with nil => 0 | cons _ m => S (length A m) end.
length is defined
length is recursively defined (decreasing on 2nd argument)

Coq < Arguments map [A B] f l.
Coq < Arguments length {A} l. (* A has to be maximally inserted *)
Coq < Check (fun l:list (list nat) => map length l).
fun l : list (list nat) => map length l
  : list (list nat) -> list nat

Coq < Arguments map [A B] f l, [A] B f l, A B f l.
Coq < Check (fun l => map length l = map (list nat) nat length l).
fun l : list (list nat) => map length l = map length l
  : list (list nat) -> Prop

```

Remark: To know which are the implicit arguments of an object, use the command `Print Implicit` (see [2.7.13](#)).

2.7.5 Automatic declaration of implicit arguments

COQ can also automatically detect what are the implicit arguments of a defined object. The command is just

```
Arguments qualid : default implicits
```

The auto-detection is governed by options telling if strict, contextual, or reversible-pattern implicit arguments must be considered or not (see Sections [2.7.7](#), [2.7.8](#), [2.7.9](#) and also [2.7.10](#)).

Variants:

1. Global Arguments *qualid* : default implicits

Tell to recompute the implicit arguments of *qualid* after ending of the current section if any.

2. Local Arguments *qualid* : default implicits

When in a module, tell not to activate the implicit arguments of *qualid* computed by this declaration to contexts that requires the module.

Example:

```
Coq < Inductive list (A:Set) : Set :=
  | nil : list A
  | cons : A -> list A -> list A.

Coq < Arguments cons : default implicits.

Coq < Print Implicit cons.
cons : forall A : Set, A -> list A -> list A
Argument A is implicit

Coq < Arguments nil : default implicits.

Coq < Print Implicit nil.
nil : forall A : Set, list A

Coq < Set Contextual Implicit.

Coq < Arguments nil : default implicits.

Coq < Print Implicit nil.
nil : forall A : Set, list A
Argument A is implicit and maximally inserted
```

The computation of implicit arguments takes account of the unfolding of constants. For instance, the variable *p* below has type (Transitivity *R*) which is reducible to $\text{forall } x, y : U, R \ x \ y \rightarrow \text{forall } z : U, R \ y \ z \rightarrow R \ x \ z$. As the variables *x*, *y* and *z* appear strictly in body of the type, they are implicit.

```
Coq < Variable X : Type.

Coq < Definition Relation := X -> X -> Prop.

Coq < Definition Transitivity (R:Relation) :=
  forall x y:X, R x y -> forall z:X, R y z -> R x z.

Coq < Variables (R : Relation) (p : Transitivity R).

Coq < Arguments p : default implicits.

Coq < Print p.
*** [ p : Transitivity R ]
Expanded type for implicit arguments
p : forall x y : X, R x y -> forall z : X, R y z -> R x z
Arguments x, y, z are implicit

Coq < Print Implicit p.
p : forall x y : X, R x y -> forall z : X, R y z -> R x z
Arguments x, y, z are implicit

Coq < Variables (a b c : X) (r1 : R a b) (r2 : R b c).

Coq < Check (p r1 r2).
p r1 r2
  : R a c
```

2.7.6 Mode for automatic declaration of implicit arguments

In case one wants to systematically declare implicit the arguments detectable as such, one may switch to the automatic declaration of implicit arguments mode by using the command

```
Set Implicit Arguments.
```

Conversely, one may unset the mode by using `Unset Implicit Arguments`. The mode is off by default. Auto-detection of implicit arguments is governed by options controlling whether strict and contextual implicit arguments have to be considered or not.

2.7.7 Controlling strict implicit arguments

When the mode for automatic declaration of implicit arguments is on, the default is to automatically set implicit only the strict implicit arguments plus, for historical reasons, a small subset of the non strict implicit arguments. To relax this constraint and to set implicit all non strict implicit arguments by default, use the command

```
Unset Strict Implicit.
```

Conversely, use the command `Set Strict Implicit` to restore the original mode that declares implicit only the strict implicit arguments plus a small subset of the non strict implicit arguments.

In the other way round, to capture exactly the strict implicit arguments and no more than the strict implicit arguments, use the command:

```
Set Strongly Strict Implicit.
```

Conversely, use the command `Unset Strongly Strict Implicit` to let the option “`Strict Implicit`” decide what to do.

Remark: In versions of COQ prior to version 8.0, the default was to declare the strict implicit arguments as implicit.

2.7.8 Controlling contextual implicit arguments

By default, COQ does not automatically set implicit the contextual implicit arguments. To tell COQ to infer also contextual implicit argument, use command

```
Set Contextual Implicit.
```

Conversely, use command `Unset Contextual Implicit` to unset the contextual implicit mode.

2.7.9 Controlling reversible-pattern implicit arguments

By default, COQ does not automatically set implicit the reversible-pattern implicit arguments. To tell COQ to infer also reversible-pattern implicit argument, use command

```
Set Reversible Pattern Implicit.
```

Conversely, use command `Unset Reversible Pattern Implicit` to unset the reversible-pattern implicit mode.

<i>term</i>	<code>++=</code>	<code>@ qualid term ... term</code>
		<code>@ qualid</code>
		<code>qualid argument ... argument</code>
<i>argument</i>	<code>::=</code>	<code>term</code>
		<code>(ident := term)</code>

Figure 2.5: Syntax for explicitly giving implicit arguments

2.7.10 Controlling the insertion of implicit arguments not followed by explicit arguments

Implicit arguments can be declared to be automatically inserted when a function is partially applied and the next argument of the function is an implicit one. In case the implicit arguments are automatically declared (with the command `Set Implicit Arguments`), the command

```
Set Maximal Implicit Insertion.
```

is used to tell to declare the implicit arguments with a maximal insertion status. By default, automatically declared implicit arguments are not declared to be insertable maximally. To restore the default mode for maximal insertion, use command `Unset Maximal Implicit Insertion`.

2.7.11 Explicit applications

In presence of non strict or contextual argument, or in presence of partial applications, the synthesis of implicit arguments may fail, so one may have to give explicitly certain implicit arguments of an application. The syntax for this is `(ident := term)` where *ident* is the name of the implicit argument and *term* is its corresponding explicit term. Alternatively, one can locally deactivate the hiding of implicit arguments of a function by using the notation `@qualid term1 . . termn`. This syntax extension is given Figure 2.5.

Example (continued):

```
Coq < Check (p r1 (z:=c)).
p r1 (z:=c)
  : R b c -> R a c

Coq < Check (p (x:=a) (y:=b) r1 (z:=c) r2).
p r1 r2
  : R a c
```

2.7.12 Renaming implicit arguments

Implicit arguments names can be redefined using the following syntax:

```
Arguments qualid name ... name : rename
```

With the `assert` flag, `Arguments` can be used to assert that a given object has the expected number of arguments and that these arguments are named as expected.

Example (continued):

```

Coq < Arguments p [s t] _ [u] _ : rename.
Coq < Check (p r1 (u:=c)).
p r1 (u:=c)
      : R b c -> R a c
Coq < Check (p (s:=a) (t:=b) r1 (u:=c) r2).
p r1 r2
      : R a c
Coq < Fail Arguments p [s t] _ [w] _ : assert.
The command has indeed failed with message:
To rename arguments the "rename" flag must be specified.
Argument u renamed to w.

```

2.7.13 Displaying what the implicit arguments are

To display the implicit arguments associated to an object, and to know if each of them is to be used maximally or not, use the command

```
Print Implicit qualid.
```

2.7.14 Explicit displaying of implicit arguments for pretty-printing

By default the basic pretty-printing rules hide the inferable implicit arguments of an application. To force printing all implicit arguments, use command

```
Set Printing Implicit.
```

Conversely, to restore the hiding of implicit arguments, use command

```
Unset Printing Implicit.
```

By default the basic pretty-printing rules display the implicit arguments that are not detected as strict implicit arguments. This “defensive” mode can quickly make the display cumbersome so this can be deactivated by using the command

```
Unset Printing Implicit Defensive.
```

Conversely, to force the display of non strict arguments, use command

```
Set Printing Implicit Defensive.
```

See also: Set Printing All in Section 2.9.

2.7.15 Interaction with subtyping

When an implicit argument can be inferred from the type of more than one of the other arguments, then only the type of the first of these arguments is taken into account, and not an upper type of all of them. As a consequence, the inference of the implicit argument of “=” fails in

```
Coq < Fail Check nat = Prop.
```

but succeeds in

```
Coq < Check Prop = nat.
```


2.7.16 Deactivation of implicit arguments for parsing

Use of implicit arguments can be deactivated by issuing the command:

```
Set Parsing Explicit.
```

In this case, all arguments of constants, inductive types, constructors, etc, including the arguments declared as implicit, have to be given as if none arguments were implicit. By symmetry, this also affects printing. To restore parsing and normal printing of implicit arguments, use:

```
Unset Parsing Explicit.
```

2.7.17 Canonical structures

A canonical structure is an instance of a record/structure type that can be used to solve unification problems involving a projection applied to an unknown structure instance (an implicit argument) and a value. The complete documentation of canonical structures can be found in Chapter 19, here only a simple example is given.

Assume that *qualid* denotes an object (*Build_struct* $c_1 \dots c_n$) in the structure *struct* of which the fields are x_1, \dots, x_n . Assume that *qualid* is declared as a canonical structure using the command

```
Canonical Structure qualid.
```

Then, each time an equation of the form $(x_i _) =_{\beta\delta\iota\zeta} c_i$ has to be solved during the type-checking process, *qualid* is used as a solution. Otherwise said, *qualid* is canonically used to extend the field c_i into a complete structure built on c_i .

Canonical structures are particularly useful when mixed with coercions and strict implicit arguments. Here is an example.

```
Coq < Require Import Relations.
Coq < Require Import EqNat.
Coq < Set Implicit Arguments.
Coq < Unset Strict Implicit.
Coq < Structure Setoid : Type :=
  {Carrier :> Set;
   Equal : relation Carrier;
   Prf_equiv : equivalence Carrier Equal}.
Coq < Definition is_law (A B:Setoid) (f:A -> B) :=
  forall x y:A, Equal x y -> Equal (f x) (f y).
Coq < Axiom eq_nat_equiv : equivalence nat eq_nat.
Coq < Definition nat_setoid : Setoid := Build_Setoid eq_nat_equiv.
Coq < Canonical Structure nat_setoid.
```

Thanks to *nat_setoid* declared as canonical, the implicit arguments A and B can be synthesized in the next statement.

```
Coq < Lemma is_law_S : is_law S.
1 subgoal

=====
is_law (A:=nat_setoid) (B:=nat_setoid) S
```

Remark: If a same field occurs in several canonical structure, then only the structure declared first as canonical is considered.

Variants:

1. Canonical Structure *ident* := *term* : *type* .
 Canonical Structure *ident* := *term* .
 Canonical Structure *ident* : *type* := *term* .

These are equivalent to a regular definition of *ident* followed by the declaration

Canonical Structure *ident*.

See also: more examples in user contribution category (Rocq/ALGEBRA).

Print Canonical Projections.

This displays the list of global names that are components of some canonical structure. For each of them, the canonical structure of which it is a projection is indicated. For instance, the above example gives the following output:

```
Coq < Print Canonical Projections.
nat <- Carrier ( nat_setoid )
eq_nat <- Equal ( nat_setoid )
eq_nat_equiv <- Prf_equiv ( nat_setoid )
```

2.7.18 Implicit types of variables

It is possible to bind variable names to a given type (e.g. in a development using arithmetic, it may be convenient to bind the names *n* or *m* to the type *nat* of natural numbers). The command for that is

```
Implicit Types ident ... ident : type
```

The effect of the command is to automatically set the type of bound variables starting with *ident* (either *ident* itself or *ident* followed by one or more single quotes, underscore or digits) to be *type* (unless the bound variable is already declared with an explicit type in which case, this latter type is considered).

Example:

```
Coq < Require Import List.
Coq < Implicit Types m n : nat.
Coq < Lemma cons_inj_nat : forall m n l, n :: l = m :: l -> n = m.
1 subgoal

=====
forall (m n : nat) (l : Datatypes.list nat), n :: l = m :: l -> n = m
Coq < intros m n.
1 subgoal

m, n : nat
=====
forall l : Datatypes.list nat, n :: l = m :: l -> n = m
```

```
Coq < Lemma cons_inj_bool : forall (m n:bool) l, n :: l = m :: l -> n = m.
1 subgoal
```

```
=====
forall (m n : bool) (l : Datatypes.list bool), n :: l = m :: l -> n = m
```

Variants:

1. `Implicit Type ident : type`
This is useful for declaring the implicit type of a single variable.
2. `Implicit Types (ident1,1...ident1,k1 : term1) ... (identn,1...identn,kn : termn) .`
Adds *n* blocks of implicit types with different specifications.

2.7.19 Implicit generalization

Implicit generalization is an automatic elaboration of a statement with free variables into a closed statement where these variables are quantified explicitly. Implicit generalization is done inside binders starting with a ``` and terms delimited by ``{ }` and ``()`, always introducing maximally inserted implicit arguments for the generalized variables. Inside implicit generalization delimiters, free variables in the current context are automatically quantified using a product or a lambda abstraction to generate a closed term. In the following statement for example, the variables *n* and *m* are automatically generalized and become explicit arguments of the lemma as we are using ``()`:

```
Coq < Generalizable All Variables.
Coq < Lemma nat_comm : `(n = n + 0).
1 subgoal
```

```
=====
forall n : nat, n = n + 0
```

One can control the set of generalizable identifiers with the `Generalizable` vernacular command to avoid unexpected generalizations when mistyping identifiers. There are three variants of the command:

```
Generalizable (All|No) Variable(s)? (ident1 identn)? .
```

Variants:

1. `Generalizable All Variables`. All variables are candidate for generalization if they appear free in the context under a generalization delimiter. This may result in confusing errors in case of typos. In such cases, the context will probably contain some unexpected generalized variable.
2. `Generalizable No Variables`. Disable implicit generalization entirely. This is the default behavior.
3. `Generalizable Variable(s)? ident1 identn`. Allow generalization of the given identifiers only. Calling this command multiple times adds to the allowed identifiers.
4. `Global Generalizable` Allows to export the choice of generalizable variables.

One can also use implicit generalization for binders, in which case the generalized variables are added as binders and set maximally implicit.

```
Coq < Definition id `(x : A) : A := x.

Coq < Print id.
id = fun (A : Type) (x : A) => x
      : forall A : Type, A -> A
Argument A is implicit and maximally inserted
Argument scopes are [type_scope _]
```

The generalizing binders ``{ }` and ``()` work similarly to their explicit counterparts, only binding the generalized variables implicitly, as maximally-inserted arguments. In these binders, the binding name for the bound object is optional, whereas the type is mandatory, dually to regular binders.

2.8 Coercions

Coercions can be used to implicitly inject terms from one *class* in which they reside into another one. A *class* is either a sort (denoted by the keyword `Sortclass`), a product type (denoted by the keyword `Funcclass`), or a type constructor (denoted by its name), e.g. an inductive type or any constant with a type of the form `forall (x1 : A1)..(xn : An), s` where *s* is a sort.

Then the user is able to apply an object that is not a function, but can be coerced to a function, and more generally to consider that a term of type A is of type B provided that there is a declared coercion between A and B. The main command is

```
Coercion qualid : class1 >-> class2.
```

which declares the construction denoted by *qualid* as a coercion between *class₁* and *class₂*.

More details and examples, and a description of the commands related to coercions are provided in Chapter 18.

2.9 Printing constructions in full

Coercions, implicit arguments, the type of pattern-matching, but also notations (see Chapter 12) can obfuscate the behavior of some tactics (typically the tactics applying to occurrences of subterms are sensitive to the implicit arguments). The command

```
Set Printing All.
```

deactivates all high-level printing features such as coercions, implicit arguments, returned type of pattern-matching, notations and various syntactic sugar for pattern-matching or record projections. Otherwise said, `Set Printing All` includes the effects of the commands `Set Printing Implicit`, `Set Printing Coercions`, `Set Printing Synth`, `Unset Printing Projections` and `Unset Printing Notations`. To reactivate the high-level printing features, use the command

```
Unset Printing All.
```

2.10 Printing universes

The following command:

```
Set Printing Universes
```

activates the display of the actual level of each occurrence of `Type`. See Section 4.1.1 for details. This wizard option, in combination with `Set Printing All` (see section 2.9) can help to diagnose failures to unify terms apparently identical but internally different in the Calculus of Inductive Constructions. To reactivate the display of the actual level of the occurrences of `Type`, use

```
Unset Printing Universes.
```

The constraints on the internal level of the occurrences of `Type` (see Section 4.1.1) can be printed using the command

```
Print [Sorted] Universes.
```

If the optional `Sorted` option is given, each universe will be made equivalent to a numbered label reflecting its level (with a linear ordering) in the universe hierarchy.

This command also accepts an optional output filename:

```
Print [Sorted] Universes string.
```

If *string* ends in `.dot` or `.gv`, the constraints are printed in the DOT language, and can be processed by Graphviz tools. The format is unspecified if *string* doesn't end in `.dot` or `.gv`.

2.11 Existential variables

Coq terms can include existential variables which represents unknown subterms to eventually be replaced by actual subterms.

Existential variables are generated in place of unsolvable implicit arguments or “`_`” placeholders when using commands such as `Check` (see Section 6.3.1) or when using tactics such as `refine` (see Section 8.2.3), as well as in place of unsolvable instances when using tactics such that `eapply` (see Section 8.2.4). An existential variable is defined in a context, which is the context of variables of the placeholder which generated the existential variable, and a type, which is the expected type of the placeholder.

As a consequence of typing constraints, existential variables can be duplicated in such a way that they possibly appear in different contexts than their defining context. Thus, any occurrence of a given existential variable comes with an instance of its original context. In the simple case, when an existential variable denotes the placeholder which generated it, or is used in the same context as the one in which it was generated, the context is not displayed and the existential variable is represented by “`?`” followed by an identifier.

```
Coq < Parameter identity : forall (X:Set), X -> X.
identity is declared

Coq < Check identity _ _.
identity ?y ?x
      : ?X@{x:=?x}
where
```

```

?y : [ |- forall x : ?P, ?X]
?P : [ |- Set]
?X : [x : ?P |- Set]
?x : [ |- ?P]

Coq < Check identity _ (fun x => _).
identity ?y (fun x : ?P => ?y0)
      : ?X@{x:=fun x : ?P => ?y0}
where
?y : [ |- forall x : forall x : ?P, ?P0, ?X]
?X : [x : forall x : ?P, ?P0 |- Set]
?P : [ |- Set]
?P0 : [x : ?P |- Set]
?y0 : [x : ?P |- ?P0]

```

In the general case, when an existential variable *?ident* appears outside of its context of definition, its instance, written under the form `@{id1:=term1; ...; idn:=termn}`, is appending to its name, indicating how the variables of its defining context are instantiated. The variables of the context of the existential variables which are instantiated by themselves are not written, unless the flag `Printing Existential Instances` is on (see Section 2.11.1), and this is why an existential variable used in the same context as its context of definition is written with no instance.

```

Coq < Check (fun x y => _) 0 1.
(fun x y : nat => ?y) 0 1
      : ?T@{x:=0; y:=1}
where
?T : [x : nat y : nat |- Type]
?y : [x : nat y : nat |- ?T]

Coq < Set Printing Existential Instances.

Coq < Check (fun x y => _) 0 1.
(fun x y : nat => ?y@{x:=x; y:=y}) 0 1
      : ?T@{x:=0; y:=1}
where
?T : [x : nat y : nat |- Type]
?y : [x : nat y : nat |- ?T@{x:=x; y:=y}]

```

Existential variables can be named by the user upon creation using the syntax `?[ident]`. This is useful when the existential variable needs to be explicitly handled later in the script (e.g. with a named-goal selector, see 9.2).

2.11.1 Explicit displaying of existential instances for pretty-printing

The command:

```
Set Printing Existential Instances
```

activates the full display of how the context of an existential variable is instantiated at each of the occurrences of the existential variable.

To deactivate the full display of the instances of existential variables, use

```
Unset Printing Existential Instances.
```

2.11.2 Solving existential variables using tactics

Instead of letting the unification engine try to solve an existential variable by itself, one can also provide an explicit hole together with a tactic to solve it. Using the syntax `ltac: (tacexpr)`, the user can put a tactic anywhere a term is expected. The order of resolution is not specified and is implementation-dependent. The inner tactic may use any variable defined in its scope, including repeated alternations between variables introduced by term binding as well as those introduced by tactic binding. The expression *tacexpr* can be any tactic expression as described at section 9.

```
Coq < Definition foo (x : nat) : nat := ltac:(exact x).
```

This construction is useful when one wants to define complicated terms using highly automated tactics without resorting to writing the proof-term by means of the interactive proof engine.

This mechanism is comparable to the `Declare Implicit Tactic` command defined at 8.9.7, except that the used tactic is local to each hole instead of being declared globally.

Chapter 3

The CoQ library

The CoQ library is structured into two parts:

The initial library: it contains elementary logical notions and data-types. It constitutes the basic state of the system directly available when running CoQ;

The standard library: general-purpose libraries containing various developments of CoQ axiomatizations about sets, lists, sorting, arithmetic, etc. This library comes with the system and its modules are directly accessible through the `Require` command (see Section 6.5.1);

In addition, user-provided libraries or developments are provided by CoQ users' community. These libraries and developments are available for download at <http://coq.inria.fr> (see Section 3.3).

The chapter briefly reviews the CoQ libraries whose contents can also be browsed at <http://coq.inria.fr/stdlib>.

3.1 The basic library

This section lists the basic notions and results which are directly available in the standard CoQ system¹.

3.1.1 Notations

This module defines the parsing and pretty-printing of many symbols (infixes, prefixes, etc.). However, it does not assign a meaning to these notations. The purpose of this is to define and fix once for all the precedence and associativity of very common notations. The main notations fixed in the initial state are listed on Figure 3.1.

3.1.2 Logic

The basic library of CoQ comes with the definitions of standard (intuitionistic) logical connectives (they are defined as inductive constructions). They are equipped with an appealing syntax enriching the (subclass *form*) of the syntactic class *term*. The syntax extension is shown on Figure 3.2.

Remark: Implication is not defined but primitive (it is a non-dependent product of a proposition over another proposition). There is also a primitive universal quantification (it is a dependent product over a

¹Most of these constructions are defined in the `Prelude` module in directory `theories/Init` at the CoQ root directory; this includes the modules `Notations`, `Logic`, `Datatypes`, `Specif`, `Peano`, `Wf` and `Tactics`. Module `Logic_Type` also makes it in the initial state

Notation	Precedence	Associativity
$_ \leftrightarrow _$	95	no
$_ \setminus / _$	85	right
$_ /\setminus _$	80	right
$_ \sim _$	75	right
$_ = _$	70	no
$_ = _ = _$	70	no
$_ = _ :> _$	70	no
$_ <> _$	70	no
$_ <> _ :> _$	70	no
$_ < _$	70	no
$_ > _$	70	no
$_ <= _$	70	no
$_ >= _$	70	no
$_ < _ < _$	70	no
$_ < _ <= _$	70	no
$_ <= _ < _$	70	no
$_ <= _ <= _$	70	no
$_ + _$	50	left
$_ _$	50	left
$_ - _$	50	left
$_ * _$	40	left
$_ \&\& _$	40	left
$_ / _$	40	left
$_ - _$	35	right
$_ / _$	35	right
$_ ^ _$	30	right

Figure 3.1: Notations in the initial state

<i>form</i>	<code>::=</code>	<code>True</code>	<code>(True)</code>
		<code>False</code>	<code>(False)</code>
		<code>~ form</code>	<code>(not)</code>
		<code>form /\ form</code>	<code>(and)</code>
		<code>form \ / form</code>	<code>(or)</code>
		<code>form -> form</code>	<i>(primitive implication)</i>
		<code>form <-> form</code>	<code>(iff)</code>
		<code>forall ident : type , form</code>	<i>(primitive for all)</i>
		<code>exists ident [: specif] , form</code>	<code>(ex)</code>
		<code>exists2 ident [: specif] , form & form</code>	<code>(ex2)</code>
		<code>term = term</code>	<code>(eq)</code>
		<code>term = term :> specif</code>	<code>(eq)</code>

Figure 3.2: Syntax of formulas

proposition). The primitive universal quantification allows both first-order and higher-order quantification.

Propositional Connectives

First, we find propositional calculus connectives:

```
Coq < Inductive True : Prop := I.
Coq < Inductive False : Prop := .
Coq < Definition not (A: Prop) := A -> False.
Coq < Inductive and (A B:Prop) : Prop := conj (_:A) (_:B).
Coq < Section Projections.
Coq < Variables A B : Prop.
Coq < Theorem proj1 : A /\ B -> A.
Coq < Theorem proj2 : A /\ B -> B.
Coq < End Projections.

Coq < Inductive or (A B:Prop) : Prop :=
  | or_introl (_:A)
  | or_intror (_:B).
Coq < Definition iff (P Q:Prop) := (P -> Q) /\ (Q -> P).
Coq < Definition IF_then_else (P Q R:Prop) := P /\ Q \/ ~ P /\ R.
```

Quantifiers

Then we find first-order quantifiers:

```
Coq < Definition all (A:Set) (P:A -> Prop) := forall x:A, P x.
Coq < Inductive ex (A: Set) (P:A -> Prop) : Prop :=
  ex_intro (x:A) (_:P x).
Coq < Inductive ex2 (A:Set) (P Q:A -> Prop) : Prop :=
  ex_intro2 (x:A) (_:P x) (_:Q x).
```

The following abbreviations are allowed:

exists x:A, P	ex A (fun x:A => P)
exists x, P	ex _ (fun x => P)
exists2 x:A, P & Q	ex2 A (fun x:A => P) (fun x:A => Q)
exists2 x, P & Q	ex2 _ (fun x => P) (fun x => Q)

The type annotation “:A” can be omitted when A can be synthesized by the system.

Equality

Then, we find equality, defined as an inductive relation. That is, given a type A and an x of type A , the predicate $(eq\ A\ x)$ is the smallest one which contains x . This definition, due to Christine Paulin-Mohring, is equivalent to define eq as the smallest reflexive relation, and it is also equivalent to Leibniz' equality.

```
Coq < Inductive eq (A:Type) (x:A) : A -> Prop :=
    eq_refl : eq A x x.
```

Lemmas

Finally, a few easy lemmas are provided.

```
Coq < Theorem absurd : forall A C:Prop, A -> ~ A -> C.
```

```
Coq < Section equality.
```

```
Coq < Variables A B : Type.
```

```
Coq < Variable f : A -> B.
```

```
Coq < Variables x y z : A.
```

```
Coq < Theorem eq_sym : x = y -> y = x.
```

```
Coq < Theorem eq_trans : x = y -> y = z -> x = z.
```

```
Coq < Theorem f_equal : x = y -> f x = f y.
```

```
Coq < Theorem not_eq_sym : x <> y -> y <> x.
```

```
Coq < End equality.
```

```
Coq < Definition eq_ind_r :
    forall (A:Type) (x:A) (P:A->Prop), P x -> forall y:A, y = x -> P y.
```

```
Coq < Definition eq_rec_r :
    forall (A:Type) (x:A) (P:A->Set), P x -> forall y:A, y = x -> P y.
```

```
Coq < Definition eq_rect_r :
    forall (A:Type) (x:A) (P:A->Type), P x -> forall y:A, y = x -> P y.
```

```
Coq < Hint Immediate eq_sym not_eq_sym : core.
```

The theorem `f_equal` is extended to functions with two to five arguments. The theorem are names `f_equal2`, `f_equal3`, `f_equal4` and `f_equal5`. For instance `f_equal3` is defined the following way.

```
Coq < Theorem f_equal3 :
    forall (A1 A2 A3 B:Type) (f:A1 -> A2 -> A3 -> B)
        (x1 y1:A1) (x2 y2:A2) (x3 y3:A3),
        x1 = y1 -> x2 = y2 -> x3 = y3 -> f x1 x2 x3 = f y1 y2 y3.
```

<i>specif</i>	::=	<i>specif</i> * <i>specif</i>	(prod)
		<i>specif</i> + <i>specif</i>	(sum)
		<i>specif</i> + { <i>specif</i> }	(sumor)
		{ <i>specif</i> } + { <i>specif</i> }	(sumbool)
		{ <i>ident</i> : <i>specif</i> <i>form</i> }	(sig)
		{ <i>ident</i> : <i>specif</i> <i>form</i> & <i>form</i> }	(sig2)
		{ <i>ident</i> : <i>specif</i> & <i>specif</i> }	(sigT)
		{ <i>ident</i> : <i>specif</i> & <i>specif</i> & <i>specif</i> }	(sigT2)
<i>term</i>	::=	(<i>term</i> , <i>term</i>)	(pair)

Figure 3.3: Syntax of data-types and specifications

3.1.3 Datatypes

In the basic library, we find the definition² of the basic data-types of programming, again defined as inductive constructions over the sort `Set`. Some of them come with a special syntax shown on Figure 3.3.

Programming

```
Coq < Inductive unit : Set := tt.
Coq < Inductive bool : Set := true | false.
Coq < Inductive nat : Set := O | S (n:nat).
Coq < Inductive option (A:Set) : Set := Some ( _:A) | None.
Coq < Inductive identity (A:Type) (a:A) : A -> Type :=
  refl_identity : identity A a a.
```

Note that zero is the letter `O`, and *not* the numeral `0`.

The predicate `identity` is logically equivalent to equality but it lives in sort `Type`. It is mainly maintained for compatibility.

We then define the disjoint sum of `A+B` of two sets `A` and `B`, and their product `A*B`.

```
Coq < Inductive sum (A B:Set) : Set := inl ( _:A) | inr ( _:B).
Coq < Inductive prod (A B:Set) : Set := pair ( _:A) ( _:B).
Coq < Section projections.
Coq < Variables A B : Set.
Coq < Definition fst (H: prod A B) := match H with
  | pair _ _ x y => x
end.
Coq < Definition snd (H: prod A B) := match H with
  | pair _ _ x y => y
end.
Coq < End projections.
```

Some operations on `bool` are also provided: `andb` (with infix notation `&&`), `orb` (with infix notation `||`), `xorb`, `implb` and `negb`.

²They are in `Datatypes.v`

3.1.4 Specification

The following notions³ allow to build new data-types and specifications. They are available with the syntax shown on Figure 3.3.

For instance, given $A : \text{Type}$ and $P : A \rightarrow \text{Prop}$, the construct $\{x : A \mid P \ x\}$ (in abstract syntax $(\text{sig } A \ P)$) is a Type . We may build elements of this set as $(\text{exist } x \ p)$ whenever we have a witness $x : A$ with its justification $p : P \ x$.

From such a $(\text{exist } x \ p)$ we may in turn extract its witness $x : A$ (using an elimination construct such as `match`) but *not* its justification, which stays hidden, like in an abstract data-type. In technical terms, one says that `sig` is a “weak (dependent) sum”. A variant `sig2` with two predicates is also provided.

```
Coq < Inductive sig (A:Set) (P:A -> Prop) : Set := exist (x:A) (_:P x).
Coq < Inductive sig2 (A:Set) (P Q:A -> Prop) : Set :=
    exist2 (x:A) (_:P x) (_:Q x).
```

A “strong (dependent) sum” $\{x : A \ \& \ P \ x\}$ may be also defined, when the predicate P is now defined as a constructor of types in Type .

```
Coq < Inductive sigT (A:Type) (P:A -> Type) : Type := existT (x:A) (_:P x).
Coq < Section Projections2.
Coq < Variable A : Type.
Coq < Variable P : A -> Type.
Coq < Definition projT1 (H:sigT A P) := let (x, h) := H in x.
Coq < Definition projT2 (H:sigT A P) :=
    match H return P (projT1 H) with
    existT _ _ x h => h
    end.
Coq < End Projections2.
Coq < Inductive sigT2 (A: Type) (P Q:A -> Type) : Type :=
    existT2 (x:A) (_:P x) (_:Q x).
```

A related non-dependent construct is the constructive sum $\{A\} + \{B\}$ of two propositions A and B .

```
Coq < Inductive sumbool (A B:Prop) : Set := left (_:A) | right (_:B).
```

This `sumbool` construct may be used as a kind of indexed boolean data-type. An intermediate between `sumbool` and `sum` is the mixed `sumor` which combines $A : \text{Set}$ and $B : \text{Prop}$ in the $\text{Set } A + \{B\}$.

```
Coq < Inductive sumor (A:Set) (B:Prop) : Set :=
    | inleft (_:A)
    | inright (_:B).
```

We may define variants of the axiom of choice, like in Martin-Löf’s Intuitionistic Type Theory.

³They are defined in module `Specif.v`

```

Coq < Lemma Choice :
  forall (S S':Set) (R:S -> S' -> Prop),
    (forall x:S, {y : S' | R x y}) ->
      {f : S -> S' | forall z:S, R z (f z)}.

Coq < Lemma Choice2 :
  forall (S S':Set) (R:S -> S' -> Set),
    (forall x:S, {y : S' & R x y}) ->
      {f : S -> S' & forall z:S, R z (f z)}.

Coq < Lemma bool_choice :
  forall (S:Set) (R1 R2:S -> Prop),
    (forall x:S, {R1 x} + {R2 x}) ->
      {f : S -> bool |
        forall x:S, f x = true /\ R1 x /\ f x = false /\ R2 x}.

```

The next construct builds a sum between a data-type $A:Type$ and an exceptional value encoding errors:

```

Coq < Definition Exc := option.
Coq < Definition value := Some.
Coq < Definition error := None.

```

This module ends with theorems, relating the sorts `Set` or `Type` and `Prop` in a way which is consistent with the realizability interpretation.

```

Coq < Definition except := False_rec.
Coq < Theorem absurd_set : forall (A:Prop) (C:Set), A -> ~ A -> C.
Coq < Theorem and_rect2 :
  forall (A B:Prop) (P:Type), (A -> B -> P) -> A /\ B -> P.

```

3.1.5 Basic Arithmetics

The basic library includes a few elementary properties of natural numbers, together with the definitions of predecessor, addition and multiplication⁴. It also provides a scope `nat_scope` gathering standard notations for common operations (+, *) and a decimal notation for numbers. That is he can write 3 for $(S (S (S O)))$. This also works on the left hand side of a `match` expression (see for example section 8.2.3). This scope is opened by default.

The following example is not part of the standard library, but it shows the usage of the notations:

```

Coq < Fixpoint even (n:nat) : bool :=
  match n with
  | 0 => true
  | 1 => false
  | S (S n) => even n
  end.

Coq < Theorem eq_S : forall x y:nat, x = y -> S x = S y.

```

⁴This is in module `Peano.v`

```

Coq < Definition pred (n:nat) : nat :=
  match n with
  | 0 => 0
  | S u => u
  end.

Coq < Theorem pred_Sn : forall m:nat, m = pred (S m).
Coq < Theorem eq_add_S : forall n m:nat, S n = S m -> n = m.
Coq < Hint Immediate eq_add_S : core.
Coq < Theorem not_eq_S : forall n m:nat, n <> m -> S n <> S m.

Coq < Definition IsSucc (n:nat) : Prop :=
  match n with
  | 0 => False
  | S p => True
  end.

Coq < Theorem O_S : forall n:nat, 0 <> S n.
Coq < Theorem n_Sn : forall n:nat, n <> S n.

Coq < Fixpoint plus (n m:nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (p + m)
  end
  where "n + m" := (plus n m) : nat_scope.

Coq < Lemma plus_n_0 : forall n:nat, n = n + 0.
Coq < Lemma plus_n_Sm : forall n m:nat, S (n + m) = n + S m.

Coq < Fixpoint mult (n m:nat) {struct n} : nat :=
  match n with
  | 0 => 0
  | S p => m + p * m
  end
  where "n * m" := (mult n m) : nat_scope.

Coq < Lemma mult_n_0 : forall n:nat, 0 = n * 0.
Coq < Lemma mult_n_Sm : forall n m:nat, n * m + n = n * (S m).

```

Finally, it gives the definition of the usual orderings `le`, `lt`, `ge`, and `gt`.

```

Coq < Inductive le (n:nat) : nat -> Prop :=
  | le_n : le n n
  | le_S : forall m:nat, n <= m -> n <= (S m)
  where "n <= m" := (le n m) : nat_scope.

Coq < Definition lt (n m:nat) := S n <= m.
Coq < Definition ge (n m:nat) := m <= n.
Coq < Definition gt (n m:nat) := m < n.

```


Properties of these relations are not initially known, but may be required by the user from modules `Le` and `Lt`. Finally, `Peano` gives some lemmas allowing pattern-matching, and a double induction principle.

```
Coq < Theorem nat_case :
  forall (n:nat) (P:nat -> Prop),
    P 0 -> (forall m:nat, P (S m)) -> P n.

Coq < Theorem nat_double_ind :
  forall R:nat -> nat -> Prop,
    (forall n:nat, R 0 n) ->
    (forall n:nat, R (S n) 0) ->
    (forall n m:nat, R n m -> R (S n) (S m)) -> forall n m:nat, R n m.
```

3.1.6 Well-founded recursion

The basic library contains the basics of well-founded recursion and well-founded induction⁵.

```
Coq < Section Well_founded.

Coq < Variable A : Type.

Coq < Variable R : A -> A -> Prop.

Coq < Inductive Acc (x:A) : Prop :=
  Acc_intro : (forall y:A, R y x -> Acc y) -> Acc x.

Coq < Lemma Acc_inv x : Acc x -> forall y:A, R y x -> Acc y.

Coq < Definition well_founded := forall a:A, Acc a.

Coq < Hypothesis Rwf : well_founded.

Coq < Theorem well_founded_induction :
  forall P:A -> Set,
    (forall x:A, (forall y:A, R y x -> P y) -> P x) -> forall a:A, P a.

Coq < Theorem well_founded_ind :
  forall P:A -> Prop,
    (forall x:A, (forall y:A, R y x -> P y) -> P x) -> forall a:A, P a.
```

The automatically generated scheme `Acc_rect` can be used to define functions by fixpoints using well-founded relations to justify termination. Assuming extensionality of the functional used for the recursive call, the fixpoint equation can be proved.

```
Coq < Section FixPoint.

Coq < Variable P : A -> Type.

Coq < Variable F : forall x:A, (forall y:A, R y x -> P y) -> P x.

Coq < Fixpoint Fix_F (x:A) (r:Acc x) {struct r} : P x :=
  F x (fun (y:A) (p:R y x) => Fix_F y (Acc_inv x r y p)).

Coq < Definition Fix (x:A) := Fix_F x (Rwf x).

Coq < Hypothesis F_ext :
  forall (x:A) (f g:forall y:A, R y x -> P y),
```

⁵This is defined in module `Wf.v`

```

      (forall (y:A) (p:R y x), f y p = g y p) -> F x f = F x g.

Coq < Lemma Fix_F_eq :
  forall (x:A) (r:Acc x),
    F x (fun (y:A) (p:R y x) => Fix_F y (Acc_inv x r y p)) = Fix_F x r.

Coq < Lemma Fix_F_inv : forall (x:A) (r s:Acc x), Fix_F x r = Fix_F x s.

Coq < Lemma fix_eq : forall x:A, Fix x = F x (fun (y:A) (p:R y x) => Fix y).

Coq < End FixPoint.

Coq < End Well_founded.

```

3.1.7 Accessing the Type level

The basic library includes the definitions⁶ of the counterparts of some data-types and logical quantifiers at the Type level: negation, pair, and properties of identity.

```

Coq < Definition notT (A:Type) := A -> False.

Coq < Inductive prodT (A B:Type) : Type := pairT (_:A) (_:B).

```

At the end, it defines data-types at the Type level.

3.1.8 Tactics

A few tactics defined at the user level are provided in the initial state⁷. They are listed at <http://coq.inria.fr/stdlib> (paragraph Init, link Tactics).

3.2 The standard library

3.2.1 Survey

The rest of the standard library is structured into the following subdirectories:

⁶This is in module `Logic_Type.v`

⁷This is in module `Tactics.v`

Logic	Classical logic and dependent equality
Arith	Basic Peano arithmetic
PArith	Basic positive integer arithmetic
NArith	Basic binary natural number arithmetic
ZArith	Basic relative integer arithmetic
Numbers	Various approaches to natural, integer and cyclic numbers (currently axiomatically and on top of 2^{31} binary words)
Bool	Booleans (basic functions and results)
Lists	Monomorphic and polymorphic lists (basic functions and results), Streams (infinite sequences defined with co-inductive types)
Sets	Sets (classical, constructive, finite, infinite, power set, etc.)
FSets	Specification and implementations of finite sets and finite maps (by lists and by AVL trees)
Reals	Axiomatization of real numbers (classical, basic functions, integer part, fractional part, limit, derivative, Cauchy series, power series and results,...)
Relations	Relations (definitions and basic results)
Sorting	Sorted list (basic definitions and heapsort correctness)
Strings	8-bits characters and strings
Wellfounded	Well-founded relations (basic results)

These directories belong to the initial load path of the system, and the modules they provide are compiled at installation time. So they are directly accessible with the command `Require` (see Chapter 6).

The different modules of the COQ standard library are described in the additional document `Library.dvi`. They are also accessible on the WWW through the COQ homepage ⁸.

3.2.2 Notations for integer arithmetics

On Figure 3.4 is described the syntax of expressions for integer arithmetics. It is provided by requiring and opening the module `ZArith` and opening scope `Z_scope`.

Figure 3.4 shows the notations provided by `Z_scope`. It specifies how notations are interpreted and, when not already reserved, the precedence and associativity.

```
Coq < Require Import ZArith.
Coq < Check   (2 + 3)%Z.
(2 + 3)%Z
  : Z

Coq < Open Scope Z_scope.
Coq < Check 2 + 3.
2 + 3
  : Z
```

3.2.3 Peano's arithmetic (`nat`)

While in the initial state, many operations and predicates of Peano's arithmetic are defined, further operations and results belong to other modules. For instance, the decidability of the basic predicates are defined here. This is provided by requiring the module `Arith`.

Figure 3.5 describes notation available in scope `nat_scope`.

⁸<http://coq.inria.fr>

Notation	Interpretation	Precedence	Associativity
$_ < _$	<code>Z.lt</code>	70	no
$x \leq y$	<code>Z.le</code>		
$_ > _$	<code>Z.gt</code>		
$x \geq y$	<code>Z.ge</code>		
$x < y < z$	<code>x < y /\ y < z</code>		
$x < y \leq z$	<code>x < y /\ y \leq z</code>		
$x \leq y < z$	<code>x \leq y /\ y < z</code>		
$x \leq y \leq z$	<code>x \leq y /\ y \leq z</code>		
$_ ?= _$	<code>Z.compare</code>		
$_ + _$	<code>Z.add</code>	40	no
$_ - _$	<code>Z.sub</code>		
$_ * _$	<code>Z.mul</code>		
$_ / _$	<code>Z.div</code>		
$_ \bmod _$	<code>Z.modulo</code>		
$_ _$	<code>Z.opp</code>		
$_ ^ _$	<code>Z.pow</code>		

Figure 3.4: Definition of the scope for integer arithmetics (`Z_scope`)

Notation	Interpretation
$_ < _$	<code>lt</code>
$x \leq y$	<code>le</code>
$_ > _$	<code>gt</code>
$x \geq y$	<code>ge</code>
$x < y < z$	<code>x < y /\ y < z</code>
$x < y \leq z$	<code>x < y /\ y \leq z</code>
$x \leq y < z$	<code>x \leq y /\ y < z</code>
$x \leq y \leq z$	<code>x \leq y /\ y \leq z</code>
$_ + _$	<code>plus</code>
$_ - _$	<code>minus</code>
$_ * _$	<code>mult</code>

Figure 3.5: Definition of the scope for natural numbers (`nat_scope`)

3.2.4 Real numbers library

Notations for real numbers

This is provided by requiring and opening the module `Reals` and opening scope `R_scope`. This set of notations is very similar to the notation for integer arithmetics. The inverse function was added.

```
Coq < Require Import Reals.
```

```
Coq < Check (2 + 3)%R.
(2 + 3)%R
: R
```

```
Coq < Open Scope R_scope.
```

Notation	Interpretation
$_ < _$	Rlt
$x \leq y$	Rle
$_ > _$	Rgt
$x \geq y$	Rge
$x < y < z$	$x < y \wedge y < z$
$x < y \leq z$	$x < y \wedge y \leq z$
$x \leq y < z$	$x \leq y \wedge y < z$
$x \leq y \leq z$	$x \leq y \wedge y \leq z$
$_ + _$	Rplus
$_ - _$	Rminus
$_ * _$	Rmult
$_ / _$	Rdiv
$- _$	Ropp
$/ _$	Rinv
$_ ^ _$	pow

Figure 3.6: Definition of the scope for real arithmetics (R_scope)

```
Coq < Check 2 + 3.
2 + 3
  : R
```

Some tactics

In addition to the `ring`, `field` and `fourier` tactics (see Chapter 8) there are:

- `discrR`

Proves that a real integer constant c_1 is different from another real integer constant c_2 .

```
Coq < Require Import DiscrR.
Coq < Goal 5 <> 0.
```

```
Coq < discrR.
No more subgoals.
```

- `split_Rabs` allows unfolding the `Rabs` constant and splits corresponding conjunctions.

```
Coq < Require Import SplitAbsolu.
Coq < Goal forall x:R, x <= Rabs x.
```

```
Coq < intro; split_Rabs.
2 subgoals
```

```
x : R
Hlt : x < 0
=====
```

Notation	Interpretation	Precedence	Associativity
<code>_ ++ _</code>	app	60	right
<code>_ :: _</code>	cons	60	right

Figure 3.7: Definition of the scope for lists (`list_scope`)

```

x <= - x
subgoal 2 is:
x <= x

```

- `split_Rmult` splits a condition that a product is non null into subgoals corresponding to the condition on each operand of the product.

```

Coq < Require Import SplitRmult.
Coq < Goal forall x y z : R, x * y * z <> 0.

Coq < intros; split_Rmult.
3 subgoals

x, y, z : R
=====
x <> 0
subgoal 2 is:
y <> 0
subgoal 3 is:
z <> 0

```

These tactics has been written with the tactic language Ltac described in Chapter 9.

3.2.5 List library

Some elementary operations on polymorphic lists are defined here. They can be accessed by requiring module `List`.

It defines the following notions:

<code>length</code>	<code>length</code>
<code>head</code>	first element (with default)
<code>tail</code>	all but first element
<code>app</code>	concatenation
<code>rev</code>	reverse
<code>nth</code>	accessing n -th element (with default)
<code>map</code>	applying a function
<code>flat_map</code>	applying a function returning lists
<code>fold_left</code>	iterator (from head to tail)
<code>fold_right</code>	iterator (from tail to head)

Table show notations available when opening scope `list_scope`.

3.3 Users' contributions

Numerous users' contributions have been collected and are available at URL <http://coq.inria.fr/contribs/>. On this web page, you have a list of all contributions with informations (author, institution, quick description, etc.) and the possibility to download them one by one. You will also find informations on how to submit a new contribution.

Chapter 4

Calculus of Inductive Constructions

The underlying formal language of COQ is a *Calculus of Inductive Constructions* (CIC) whose inference rules are presented in this chapter. The history of this formalism as well as pointers to related work are provided in a separate chapter; see *Credits*.

4.1 The terms

The expressions of the CIC are *terms* and all terms have a *type*. There are types for functions (or programs), there are atomic types (especially datatypes)... but also types for proofs and types for the types themselves. Especially, any object handled in the formalism must belong to a type. For instance, universal quantification is relative to a type and takes the form “*for all x of type T, P*”. The expression “*x of type T*” is written “ $x:T$ ”. Informally, “ $x:T$ ” can be thought as “*x belongs to T*”.

The types of types are *sorts*. Types and sorts are themselves terms so that terms, types and sorts are all components of a common syntactic language of terms which is described in Section 4.1.2 but, first, we describe sorts.

4.1.1 Sorts

All sorts have a type and there is an infinite well-founded typing hierarchy of sorts whose base sorts are `Prop` and `Set`.

The sort `Prop` intends to be the type of logical propositions. If M is a logical proposition then it denotes the class of terms representing proofs of M . An object m belonging to M witnesses the fact that M is provable. An object of type `Prop` is called a proposition.

The sort `Set` intends to be the type of small sets. This includes data types such as booleans and naturals, but also products, subsets, and function types over these data types.

`Prop` and `Set` themselves can be manipulated as ordinary terms. Consequently they also have a type. Because assuming simply that `Set` has type `Set` leads to an inconsistent theory [25], the language of CIC has infinitely many sorts. There are, in addition to `Set` and `Prop` a hierarchy of universes `Type(i)` for any integer i .

Like `Set`, all of the sorts `Type(i)` contain small sets such as booleans, natural numbers, as well as products, subsets and function types over small sets. But, unlike `Set`, they also contain large sets, namely the sorts `Set` and `Type(j)` for $j < i$, and all products, subsets and function types over these sorts.

Formally, we call \mathcal{S} the set of sorts which is defined by:

$$\mathcal{S} \equiv \{\text{Prop}, \text{Set}, \text{Type}(i) \mid i \in \mathbb{N}\}$$

Their properties, such as: $\text{Prop}:\text{Type}(1)$, $\text{Set}:\text{Type}(1)$, and $\text{Type}(i):\text{Type}(i+1)$, are defined in Section 4.4.

The user does not have to mention explicitly the index i when referring to the universe $\text{Type}(i)$. One only writes Type . The system itself generates for each instance of Type a new index for the universe and checks that the constraints between these indexes can be solved. From the user point of view we consequently have $\text{Type}:\text{Type}$. We shall make precise in the typing rules the constraints between the indexes.

Implementation issues In practice, the Type hierarchy is implemented using *algebraic universes*. An algebraic universe u is either a variable (a qualified identifier with a number) or a successor of an algebraic universe (an expression $u+1$), or an upper bound of algebraic universes (an expression $\max(u_1, \dots, u_n)$), or the base universe (the expression 0) which corresponds, in the arity of template polymorphic inductive types (see Section 4.5.2), to the predicative sort Set . A graph of constraints between the universe variables is maintained globally. To ensure the existence of a mapping of the universes to the positive integers, the graph of constraints must remain acyclic. Typing expressions that violate the acyclicity of the graph of constraints results in a `Universe inconsistency` error (see also Section 2.10).

4.1.2 Terms

Terms are built from sorts, variables, constants, abstractions, applications, local definitions, and products. From a syntactic point of view, types cannot be distinguished from terms, except that they cannot start by an abstraction or a constructor. More precisely the language of the *Calculus of Inductive Constructions* is built from the following rules.

1. the sorts Set , Prop , $\text{Type}(i)$ are terms.
2. variables, hereafter ranged over by letters x, y , etc., are terms
3. constants, hereafter ranged over by letters c, d , etc., are terms.
4. if x is a variable and T, U are terms then $\forall x : T, U$ (`forall $x : T, U$ in COQ concrete syntax`) is a term. If x occurs in U , $\forall x : T, U$ reads as “for all x of type T , U ”. As U depends on x , one says that $\forall x : T, U$ is a *dependent product*. If x does not occur in U then $\forall x : T, U$ reads as “if T then U ”. A *non dependent product* can be written: $T \rightarrow U$.
5. if x is a variable and T, u are terms then $\lambda x : T. u$ (`fun $x : T \Rightarrow u$ in COQ concrete syntax`) is a term. This is a notation for the λ -abstraction of λ -calculus [8]. The term $\lambda x : T. u$ is a function which maps elements of T to the expression u .
6. if t and u are terms then $(t\ u)$ is a term (`$t\ u$ in COQ concrete syntax`). The term $(t\ u)$ reads as “ t applied to u ”.
7. if x is a variable, and t, T and u are terms then `let $x := t : T$ in u` is a term which denotes the term u where the variable x is locally bound to t of type T . This stands for the common “let-in” construction of functional programs such as ML or Scheme.

Free variables. The notion of free variables is defined as usual. In the expressions $\lambda x : T. U$ and $\forall x : T, U$ the occurrences of x in U are bound.

Substitution. The notion of substituting a term t to free occurrences of a variable x in a term u is defined as usual. The resulting term is written $u\{x/t\}$.

The logical vs programming readings. The constructions of the CIC can be used to express both logical and programming notions, accordingly to the Curry-Howard correspondence between proofs and programs, and between propositions and types [38, 81, 39].

For instance, let us assume that `nat` is the type of natural numbers with zero element written `0` and that `True` is the always true proposition. Then \rightarrow is used both to denote `nat \rightarrow nat` which is the type of functions from `nat` to `nat`, to denote `True \rightarrow True` which is an implicative proposition, to denote `nat \rightarrow Prop` which is the type of unary predicates over the natural numbers, etc.

Let us assume that `mult` is a function of type `nat \rightarrow nat \rightarrow nat` and `eqnat` a predicate of type `nat \rightarrow nat \rightarrow Prop`. The λ -abstraction can serve to build “ordinary” functions as in $\lambda x : \text{nat}.(\text{mult } x \ x)$ (i.e. `fun x : nat => mult x x` in COQ notation) but may build also predicates over the natural numbers. For instance $\lambda x : \text{nat}.(\text{eqnat } x \ 0)$ (i.e. `fun x : nat => eqnat x 0` in COQ notation) will represent the predicate of one variable x which asserts the equality of x with `0`. This predicate has type `nat \rightarrow Prop` and it can be applied to any expression of type `nat`, say t , to give an object $P \ t$ of type `Prop`, namely a proposition.

Furthermore `forall x : nat, P x` will represent the type of functions which associate to each natural number n an object of type $(P \ n)$ and consequently represent the type of proofs of the formula “ $\forall x. P(x)$ ”.

4.2 Typing rules

As objects of type theory, terms are subjected to *type discipline*. The well typing of a term depends on a global environment and a local context.

Local context. A *local context* is an ordered list of *local declarations* of names which we call *variables*. The declaration of some variable x is either a *local assumption*, written $x : T$ (T is a type) or a *local definition*, written $x := t : T$. We use brackets to write local contexts. A typical example is $[x : T; y := u : U; z : V]$. Notice that the variables declared in a local context must be distinct. If Γ declares some x , we write $x \in \Gamma$. By writing $(x : T) \in \Gamma$ we mean that either $x : T$ is an assumption in Γ or that there exists some t such that $x := t : T$ is a definition in Γ . If Γ defines some $x := t : T$, we also write $(x := t : T) \in \Gamma$. For the rest of the chapter, the $\Gamma :: (y : T)$ denotes the local context Γ enriched with the local assumption $y : T$. Similarly, $\Gamma :: (y := t : T)$ denotes the local context Γ enriched with the local definition $(y := t : T)$. The notation $[]$ denotes the empty local context. By $\Gamma_1; \Gamma_2$ we mean concatenation of the local context Γ_1 and the local context Γ_2 .

Global environment. A *global environment* is an ordered list of *global declarations*. Global declarations are either *global assumptions* or *global definitions*, but also declarations of inductive objects. Inductive objects themselves declare both inductive or coinductive types and constructors (see Section 4.5).

A *global assumption* will be represented in the global environment as $(c : T)$ which assumes the name c to be of some type T . A *global definition* will be represented in the global environment as $c := t : T$ which defines the name c to have value t and type T . We shall call such names *constants*. For the rest of the chapter, the $E; c : T$ denotes the global environment E enriched with the global

assumption $c : T$. Similarly, $E; c := t : T$ denotes the global environment E enriched with the global definition $(c := t : T)$.

The rules for inductive definitions (see Section 4.5) have to be considered as assumption rules to which the following definitions apply: if the name c is declared in E , we write $c \in E$ and if $c : T$ or $c := t : T$ is declared in E , we write $(c : T) \in E$.

Typing rules. In the following, we define simultaneously two judgments. The first one $E[\Gamma] \vdash t : T$ means the term t is well-typed and has type T in the global environment E and local context Γ . The second judgment $\mathcal{WF}(E)[\Gamma]$ means that the global environment E is well-formed and the local context Γ is a valid local context in this global environment.

A term t is well typed in a global environment E iff there exists a local context Γ and a term T such that the judgment $E[\Gamma] \vdash t : T$ can be derived from the following rules.

W-Empty

$$\mathcal{WF}(\square)\square$$

W-Local-Assum

$$\frac{E[\Gamma] \vdash T : s \quad s \in \mathcal{S} \quad x \notin \Gamma}{\mathcal{WF}(E)[\Gamma :: (x : T)]}$$

W-Local-Def

$$\frac{E[\Gamma] \vdash t : T \quad x \notin \Gamma}{\mathcal{WF}(E)[\Gamma :: (x := t : T)]}$$

W-Global-Assum

$$\frac{E\square \vdash T : s \quad s \in \mathcal{S} \quad c \notin E}{\mathcal{WF}(E; c : T)\square}$$

W-Global-Def

$$\frac{E\square \vdash t : T \quad c \notin E}{\mathcal{WF}(E; c := t : T)\square}$$

Ax-Prop

$$\frac{\mathcal{WF}(E)[\Gamma]}{E[\Gamma] \vdash \text{Prop} : \text{Type}(1)}$$

Ax-Set

$$\frac{\mathcal{WF}(E)[\Gamma]}{E[\Gamma] \vdash \text{Set} : \text{Type}(1)}$$

Ax-Type

$$\frac{\mathcal{WF}(E)[\Gamma]}{E[\Gamma] \vdash \text{Type}(i) : \text{Type}(i + 1)}$$

Var

$$\frac{\mathcal{WF}(E)[\Gamma] \quad (x : T) \in \Gamma \text{ or } (x := t : T) \in \Gamma \text{ for some } t}{E[\Gamma] \vdash x : T}$$

Const

$$\frac{\mathcal{WF}(E)[\Gamma] \quad (c : T) \in E \text{ or } (c := t : T) \in E \text{ for some } t}{E[\Gamma] \vdash c : T}$$

Prod-Prop

$$\frac{E[\Gamma] \vdash T : s \quad s \in \mathcal{S} \quad E[\Gamma :: (x : T)] \vdash U : \mathbf{Prop}}{E[\Gamma] \vdash \forall x : T, U : \mathbf{Prop}}$$

Prod-Set

$$\frac{E[\Gamma] \vdash T : s \quad s \in \{\mathbf{Prop}, \mathbf{Set}\} \quad E[\Gamma :: (x : T)] \vdash U : \mathbf{Set}}{E[\Gamma] \vdash \forall x : T, U : \mathbf{Set}}$$

Prod-Type

$$\frac{E[\Gamma] \vdash T : \mathbf{Type}(i) \quad E[\Gamma :: (x : T)] \vdash U : \mathbf{Type}(i)}{E[\Gamma] \vdash \forall x : T, U : \mathbf{Type}(i)}$$

Lam

$$\frac{E[\Gamma] \vdash \forall x : T, U : s \quad E[\Gamma :: (x : T)] \vdash t : U}{E[\Gamma] \vdash \lambda x : T. t : \forall x : T, U}$$

App

$$\frac{E[\Gamma] \vdash t : \forall x : U, T \quad E[\Gamma] \vdash u : U}{E[\Gamma] \vdash (t \ u) : T\{x/u\}}$$

Let

$$\frac{E[\Gamma] \vdash t : T \quad E[\Gamma :: (x := t : T)] \vdash u : U}{E[\Gamma] \vdash \mathbf{let} \ x := t : T \ \mathbf{in} \ u : U\{x/t\}}$$

Remark: Prod_1 and Prod_2 typing-rules make sense if we consider the semantic difference between \mathbf{Prop} and \mathbf{Set} :

- All values of a type that has a sort \mathbf{Set} are extractable.
- No values of a type that has a sort \mathbf{Prop} are extractable.

Remark: We may have $\mathbf{let} \ x := t : T \ \mathbf{in} \ u$ well-typed without having $((\lambda x : T. u) \ t)$ well-typed (where T is a type of t). This is because the value t associated to x may be used in a conversion rule (see Section 4.3).

4.3 Conversion rules

In CIC, there is an internal reduction mechanism. In particular, it can decide if two programs are *intentionally* equal (one says *convertible*). Convertibility is described in this section.

β -reduction. We want to be able to identify some terms as we can identify the application of a function to a given argument with its result. For instance the identity function over a given type T can be written $\lambda x : T. x$. In any global environment E and local context Γ , we want to identify any object a (of type T) with the application $((\lambda x : T. x) \ a)$. We define for this a *reduction* (or a *conversion*) rule we call β :

$$E[\Gamma] \vdash ((\lambda x : T. t) \ u) \triangleright_{\beta} t\{x/u\}$$

We say that $t\{x/u\}$ is the β -contraction of $((\lambda x : T. t) \ u)$ and, conversely, that $((\lambda x : T. t) \ u)$ is the β -expansion of $t\{x/u\}$.

According to β -reduction, terms of the *Calculus of Inductive Constructions* enjoy some fundamental properties such as confluence, strong normalization, subject reduction. These results are theoretically of great importance but we will not detail them here and refer the interested reader to [24].

ι -reduction. A specific conversion rule is associated to the inductive objects in the global environment. We shall give later on (see Section 4.5.3) the precise rules but it just says that a destructor applied to an object built from a constructor behaves as expected. This reduction is called ι -reduction and is more precisely studied in [126, 145].

δ -reduction. We may have variables defined in local contexts or constants defined in the global environment. It is legal to identify such a reference with its value, that is to expand (or unfold) it into its value. This reduction is called δ -reduction and shows as follows.

$$E[\Gamma] \vdash x \triangleright_{\delta} t \quad \text{if } (x := t : T) \in \Gamma \quad E[\Gamma] \vdash c \triangleright_{\delta} t \quad \text{if } (c := t : T) \in E$$

ζ -reduction. COQ allows also to remove local definitions occurring in terms by replacing the defined variable by its value. The declaration being destroyed, this reduction differs from δ -reduction. It is called ζ -reduction and shows as follows.

$$E[\Gamma] \vdash \text{let } x := u \text{ in } t \triangleright_{\zeta} t\{x/u\}$$

η -expansion. Another important concept is η -expansion. It is legal to identify any term t of functional type $\forall x : T, U$ with its so-called η -expansion $\lambda x : T. (t \ x)$ for x an arbitrary variable name fresh in t .

Remark: We deliberately do not define η -reduction:

$$\lambda x : T. (t \ x) \not\triangleright_{\eta} t$$

This is because, in general, the type of t need not to be convertible to the type of $\lambda x : T. (t \ x)$. E.g., if we take f such that:

$$f : \forall x : \text{Type}(2), \text{Type}(1)$$

then

$$\lambda x : \text{Type}(1), (f \ x) : \forall x : \text{Type}(1), \text{Type}(1)$$

We could not allow

$$\lambda x : \text{Type}(1), (f \ x) \not\triangleright_{\eta} f$$

because the type of the reduced term $\forall x : \text{Type}(2), \text{Type}(1)$ would not be convertible to the type of the original term $\forall x : \text{Type}(1), \text{Type}(1)$.

Convertibility. Let us write $E[\Gamma] \vdash t \triangleright u$ for the contextual closure of the relation t reduces to u in the global environment E and local context Γ with one of the previous reduction β, ι, δ or ζ .

We say that two terms t_1 and t_2 are $\beta\iota\delta\zeta\eta$ -convertible, or simply *convertible*, or *equivalent*, in the global environment E and local context Γ iff there exist terms u_1 and u_2 such that $E[\Gamma] \vdash t_1 \triangleright \dots \triangleright u_1$ and $E[\Gamma] \vdash t_2 \triangleright \dots \triangleright u_2$ and either u_1 and u_2 are identical, or they are convertible up to η -expansion, i.e. u_1 is $\lambda x : T. u'_1$ and $u_2 \ x$ is recursively convertible to u'_1 , or, symmetrically, u_2 is $\lambda x : T. u'_2$ and $u_1 \ x$ is recursively convertible to u'_2 . We then write $E[\Gamma] \vdash t_1 =_{\beta\iota\delta\zeta\eta} t_2$.

Apart from this we consider two instances of polymorphic and cumulative (see Chapter 29) inductive types (see below) convertible $E[\Gamma] \vdash t \ w_1 \dots w_m =_{\beta\iota\delta\zeta\eta} t \ w'_1 \dots w'_m$ if we have subtypings (see below) in both directions, i.e., $E[\Gamma] \vdash t \ w_1 \dots w_m \leq_{\beta\iota\delta\zeta\eta} t \ w'_1 \dots w'_m$ and $E[\Gamma] \vdash t \ w'_1 \dots w'_m \leq_{\beta\iota\delta\zeta\eta} t \ w_1 \dots w_m$. Furthermore, we consider $E[\Gamma] \vdash c \ v_1 \dots v_m =_{\beta\iota\delta\zeta\eta} c' \ v'_1 \dots v'_m$ convertible if $E[\Gamma] \vdash v_i =_{\beta\iota\delta\zeta\eta} v'_i$ and we have that c and

c' are the same constructors of different instances the same inductive types (differing only in universe levels) such that $E[\Gamma] \vdash c \ v_1 \dots v_m : t \ w_1 \dots w_m$ and $E[\Gamma] \vdash c' \ v'_1 \dots v'_m : t' \ w'_1 \dots w'_m$ and we have $E[\Gamma] \vdash t \ w_1 \dots w_m =_{\beta\delta\iota\zeta\eta} t' \ w'_1 \dots w'_m$.

The convertibility relation allows introducing a new typing rule which says that two convertible well-formed types have the same inhabitants.

4.4 Subtyping rules

At the moment, we did not take into account one rule between universes which says that any term in a universe of index i is also a term in the universe of index $i + 1$ (this is the *cumulativity* rule of CIC). This property extends the equivalence relation of convertibility into a *subtyping* relation inductively defined by:

1. if $E[\Gamma] \vdash t =_{\beta\delta\iota\zeta\eta} u$ then $E[\Gamma] \vdash t \leq_{\beta\delta\iota\zeta\eta} u$,
2. if $i \leq j$ then $E[\Gamma] \vdash \text{Type}(i) \leq_{\beta\delta\iota\zeta\eta} \text{Type}(j)$,
3. for any i , $E[\Gamma] \vdash \text{Set} \leq_{\beta\delta\iota\zeta\eta} \text{Type}(i)$,
4. $E[\Gamma] \vdash \text{Prop} \leq_{\beta\delta\iota\zeta\eta} \text{Set}$, hence, by transitivity, $E[\Gamma] \vdash \text{Prop} \leq_{\beta\delta\iota\zeta\eta} \text{Type}(i)$, for any i
5. if $E[\Gamma] \vdash T =_{\beta\delta\iota\zeta\eta} U$ and $E[\Gamma :: (x : T)] \vdash T' \leq_{\beta\delta\iota\zeta\eta} U'$ then $E[\Gamma] \vdash \forall x : T, T' \leq_{\beta\delta\iota\zeta\eta} \forall x : U, U'$.
6. if $\text{Ind}[p](\Gamma_I := \Gamma_C)$ is a universe polymorphic and cumulative (see Chapter 29) inductive type (see below) and $(t : \forall \Gamma_P, \forall \Gamma_{Arr(t)}, \mathcal{S}) \in \Gamma_I$ and $(t' : \forall \Gamma'_P, \forall \Gamma'_{Arr(t)}, \mathcal{S}') \in \Gamma_I$ are two different instances of *the same* inductive type (differing only in universe levels) with constructors

$$[c_1 : \forall \Gamma_P, \forall T_{1,1} \dots T_{1,n_1}, t \ v_{1,1} \dots v_{1,m}; \dots; c_k : \forall \Gamma_P, \forall T_{k,1} \dots T_{k,n_k}, t \ v_{n,1} \dots v_{n,m}]$$

and

$$[c_1 : \forall \Gamma'_P, \forall T'_{1,1} \dots T'_{1,n_1}, t' \ v'_{1,1} \dots v'_{1,m}; \dots; c_k : \forall \Gamma'_P, \forall T'_{k,1} \dots T'_{k,n_k}, t' \ v'_{n,1} \dots v'_{n,m}]$$

respectively then $E[\Gamma] \vdash t \ w_1 \dots w_m \leq_{\beta\delta\iota\zeta\eta} t' \ w'_1 \dots w'_m$ (notice that t and t' are both fully applied, i.e., they have a sort as a type) if $E[\Gamma] \vdash w_i =_{\beta\delta\iota\zeta\eta} w'_i$ for $1 \leq i \leq m$ and we have

$$E[\Gamma] \vdash T_{i,j} \leq_{\beta\delta\iota\zeta\eta} T'_{i,j} \text{ and } E[\Gamma] \vdash A_i \leq_{\beta\delta\iota\zeta\eta} A'_i$$

where $\Gamma_{Arr(t)} = [a_1 : A_1; a_1 : A_l]$ and $\Gamma_{Arr(t')} = [a_1 : A'_1; a_1 : A'_l]$.

The conversion rule up to subtyping is now exactly:

Conv

$$\frac{E[\Gamma] \vdash U : s \quad E[\Gamma] \vdash t : T \quad E[\Gamma] \vdash T \leq_{\beta\delta\iota\zeta\eta} U}{E[\Gamma] \vdash t : U}$$

Normal form. A term which cannot be any more reduced is said to be in *normal form*. There are several ways (or strategies) to apply the reduction rules. Among them, we have to mention the *head reduction* which will play an important role (see Chapter 8). Any term can be written as $\lambda x_1 : T_1. \dots \lambda x_k : T_k. (t_0 t_1 \dots t_n)$ where t_0 is not an application. We say then that t_0 is the *head* of t . If we assume that t_0 is $\lambda x : T. u_0$ then one step of β -head reduction of t is:

$$\lambda x_1 : T_1. \dots \lambda x_k : T_k. (\lambda x : T. u_0 t_1 \dots t_n) \triangleright \lambda(x_1 : T_1) \dots (x_k : T_k). (u_0\{x/t_1\} t_2 \dots t_n)$$

Iterating the process of head reduction until the head of the reduced term is no more an abstraction leads to the *β -head normal form* of t :

$$t \triangleright \dots \triangleright \lambda x_1 : T_1. \dots \lambda x_k : T_k. (v u_1 \dots u_m)$$

where v is not an abstraction (nor an application). Note that the head normal form must not be confused with the normal form since some u_i can be reducible. Similar notions of head-normal forms involving δ , ι and ζ reductions or any combination of those can also be defined.

4.5 Inductive Definitions

Formally, we can represent any *inductive definition* as $\text{Ind}[p](\Gamma_I := \Gamma_C)$ where:

- Γ_I determines the names and types of inductive types;
- Γ_C determines the names and types of constructors of these inductive types;
- p determines the number of parameters of these inductive types.

These inductive definitions, together with global assumptions and global definitions, then form the global environment. Additionally, for any p there always exists $\Gamma_P = [a_1 : A_1; \dots; a_p : A_p]$ such that each T in $(t : T) \in \Gamma_I \cup \Gamma_C$ can be written as: $\forall \Gamma_P, T'$ where Γ_P is called the *context of parameters*. Furthermore, we must have that each T in $(t : T) \in \Gamma_I$ can be written as: $\forall \Gamma_P, \forall \Gamma_{\text{Arr}(t)}, S$ where $\Gamma_{\text{Arr}(t)}$ is called the *Arity* of the inductive type t and S is called the sort of the inductive type t .

Examples The declaration for parameterized lists is:

$$\text{Ind } [1] \left([\text{list} : \text{Set} \rightarrow \text{Set}] := \left[\begin{array}{l} \text{nil} : \forall A : \text{Set}, \text{list } A \\ \text{cons} : \forall A : \text{Set}, A \rightarrow \text{list } A \rightarrow \text{list } A \end{array} \right] \right)$$

which corresponds to the result of the COQ declaration:

```
Coq < Inductive list (A:Set) : Set :=
  | nil : list A
  | cons : A -> list A -> list A.
```

The declaration for a mutual inductive definition of *tree* and *forest* is:

$$\text{Ind } [] \left(\left[\begin{array}{l} \text{tree} : \text{Set} \\ \text{forest} : \text{Set} \end{array} \right] := \left[\begin{array}{l} \text{node} : \text{forest} \rightarrow \text{tree} \\ \text{emptyf} : \text{forest} \\ \text{consf} : \text{tree} \rightarrow \text{forest} \rightarrow \text{forest} \end{array} \right] \right)$$

which corresponds to the result of the COQ declaration:


```

Coq < Inductive tree : Set :=
  node : forest -> tree
with forest : Set :=
  | emptyf : forest
  | consf : tree -> forest -> forest.

```

The declaration for a mutual inductive definition of **even** and **odd** is:

$$\text{Ind } [1] \left(\left[\begin{array}{l} \text{even} : \text{nat} \rightarrow \text{Prop} \\ \text{odd} : \text{nat} \rightarrow \text{Prop} \end{array} \right] := \left[\begin{array}{l} \text{even_O} : \text{even } 0 \\ \text{even_S} : \forall n : \text{nat}, \text{odd } n \rightarrow \text{even } (S \ n) \\ \text{odd_S} : \forall n : \text{nat}, \text{even } n \rightarrow \text{odd } (S \ n) \end{array} \right] \right)$$

which corresponds to the result of the COQ declaration:

```

Coq < Inductive even : nat -> Prop :=
  | even_O : even 0
  | even_S : forall n, odd n -> even (S n)
with odd : nat -> Prop :=
  | odd_S : forall n, even n -> odd (S n).

```

4.5.1 Types of inductive objects

We have to give the type of constants in a global environment E which contains an inductive declaration.

Ind

$$\frac{\mathcal{WF}(E)[\Gamma] \quad \text{Ind}[p](\Gamma_I := \Gamma_C) \in E \quad (a : A) \in \Gamma_I}{E[\Gamma] \vdash a : A}$$

Constr

$$\frac{\mathcal{WF}(E)[\Gamma] \quad \text{Ind}[p](\Gamma_I := \Gamma_C) \in E \quad (c : C) \in \Gamma_C}{E[\Gamma] \vdash c : C}$$

Example. Provided that our environment E contains inductive definitions we showed before, these two inference rules above enable us to conclude that:

```

E[Γ] ⊢ even : nat → Prop
E[Γ] ⊢ odd : nat → Prop
E[Γ] ⊢ even_O : even 0
E[Γ] ⊢ even_S : ∀ n : nat, odd n → even (S n)
E[Γ] ⊢ odd_S : ∀ n : nat, even n → odd (S n)

```

4.5.2 Well-formed inductive definitions

We cannot accept any inductive declaration because some of them lead to inconsistent systems. We restrict ourselves to definitions which satisfy a syntactic criterion of positivity. Before giving the formal rules, we need a few definitions:

Definition A type T is an *arity of sort* s if it converts to the sort s or to a product $\forall x : T, U$ with U an arity of sort s .

Examples $A \rightarrow \text{Set}$ is an arity of sort **Set**. $\forall A : \text{Prop}, A \rightarrow \text{Prop}$ is an arity of sort **Prop**.

Definition A type T is an *arity* if there is a $s \in \mathcal{S}$ such that T is an arity of sort s .

Examples $A \rightarrow \text{Set}$ and $\forall A : \text{Prop}, A \rightarrow \text{Prop}$ are arities.

Definition We say that T is a *type of constructor of I* in one of the following two cases:

- T is $(I \ t_1 \dots t_n)$
- T is $\forall x : U, T'$ where T' is also a type of constructor of I

Examples nat and $\text{nat} \rightarrow \text{nat}$ are types of constructors of nat .

$\forall A : \text{Type}, \text{list } A$ and $\forall A : \text{Type}, A \rightarrow \text{list } A \rightarrow \text{list } A$ are constructors of list .

Definition The type of constructor T will be said to *satisfy the positivity condition* for a constant X in the following cases:

- $T = (X \ t_1 \dots t_n)$ and X does not occur free in any t_i
- $T = \forall x : U, V$ and X occurs only strictly positively in U and the type V satisfies the positivity condition for X

The constant X *occurs strictly positively* in T in the following cases:

- X does not occur in T
- T converts to $(X \ t_1 \dots t_n)$ and X does not occur in any of t_i
- T converts to $\forall x : U, V$ and X does not occur in type U but occurs strictly positively in type V
- T converts to $(I \ a_1 \dots a_m \ t_1 \dots t_p)$ where I is the name of an inductive declaration of the form $\text{Ind}[m](I : A := c_1 : \forall p_1 : P_1, \dots \forall p_m : P_m, C_1; \dots; c_n : \forall p_1 : P_1, \dots \forall p_m : P_m, C_n)$ (in particular, it is not mutually defined and it has m parameters) and X does not occur in any of the t_i , and the (instantiated) types of constructor $C_i\{p_j/a_j\}_{j=1\dots m}$ of I satisfy the nested positivity condition for X

The type of constructor T of I *satisfies the nested positivity condition* for a constant X in the following cases:

- $T = (I \ b_1 \dots b_m \ u_1 \dots u_p)$, I is an inductive definition with m parameters and X does not occur in any u_i
- $T = \forall x : U, V$ and X occurs only strictly positively in U and the type V satisfies the nested positivity condition for X

For instance, if one considers the following variant of a tree type branching over the natural numbers

```
Inductive nattree (A:Type) : Type :=
| leaf : nattree A
| node : A -> (nat -> nattree A) -> nattree A
```

Then every instantiated constructor of `nattree` A satisfies the nested positivity condition for `nattree`

- concerning type `nattree` A of constructor `leaf`:
 Type `nattree` A of constructor `leaf` satisfies the positivity condition for `nattree` because `nattree` does not appear in any (real) arguments of the type of that constructor (primarily because `nattree` does not have any (real) arguments) ... (*bullet 1*)
- concerning type $\forall A \rightarrow (\mathbb{N} \rightarrow \text{nattree } A) \rightarrow \text{nattree } A$ of constructor `node`:
 Type $\forall A : \text{Type}, A \rightarrow (\mathbb{N} \rightarrow \text{nattree } A) \rightarrow \text{nattree } A$ of constructor `node` satisfies the positivity condition for `nattree` because:
 - `nattree` occurs only strictly positively in `Type` ... (*bullet 1*)
 - `nattree` occurs only strictly positively in A ... (*bullet 1*)
 - `nattree` occurs only strictly positively in $\mathbb{N} \rightarrow \text{nattree } A$... (*bullet 3+2*)
 - `nattree` satisfies the positivity condition for `nattree` A ... (*bullet 1*)

Correctness rules. We shall now describe the rules allowing the introduction of a new inductive definition.

W-Ind Let E be a global environment and $\Gamma_P, \Gamma_I, \Gamma_C$ are contexts such that Γ_I is $[I_1 : \forall \Gamma_P, A_1; \dots; I_k : \forall \Gamma_P, A_k]$ and Γ_C is $[c_1 : \forall \Gamma_P, C_1; \dots; c_n : \forall \Gamma_P, C_n]$.

$$\frac{(E[\Gamma_P] \vdash A_j : s'_j)_{j=1 \dots k} \quad (E[\Gamma_I; \Gamma_P] \vdash C_i : s_{q_i})_{i=1 \dots n}}{\mathcal{WF}(E; \text{Ind}[p](\Gamma_I := \Gamma_C))[\Gamma]}$$

provided that the following side conditions hold:

- $k > 0$ and all of I_j and c_i are distinct names for $j = 1 \dots k$ and $i = 1 \dots n$,
- p is the number of parameters of $\text{Ind}(\Gamma_I := \Gamma_C)$ and Γ_P is the context of parameters,
- for $j = 1 \dots k$ we have that A_j is an arity of sort s_j and $I_j \notin E$,
- for $i = 1 \dots n$ we have that C_i is a type of constructor of I_{q_i} which satisfies the positivity condition for $I_1 \dots I_k$ and $c_i \notin \Gamma \cup E$.

One can remark that there is a constraint between the sort of the arity of the inductive type and the sort of the type of its constructors which will always be satisfied for the impredicative sort `Prop` but may fail to define inductive definition on sort `Set` and generate constraints between universes for inductive definitions in the `Type` hierarchy.

Examples. It is well known that existential quantifier can be encoded as an inductive definition. The following declaration introduces the second-order existential quantifier $\exists X.P(X)$.

```
Coq < Inductive exProp (P:Prop->Prop) : Prop :=
  exP_intro : forall X:Prop, P X -> exProp P.
```

The same definition on `Set` is not allowed and fails:

```
Coq < Fail Inductive exSet (P:Set->Prop) : Set :=
  exS_intro : forall X:Set, P X -> exSet P.
The command has indeed failed with message:
Large non-propositional inductive types must be in Type.
```

It is possible to declare the same inductive definition in the universe **Type**. The `exType` inductive definition has type $(\text{Type}_i \rightarrow \text{Prop}) \rightarrow \text{Type}_j$ with the constraint that the parameter X of `exT_intro` has type Type_k with $k < j$ and $k \leq i$.

```
Coq < Inductive exType (P:Type->Prop) : Type :=
  exT_intro : forall X:Type, P X -> exType P.
```

Template polymorphism. Inductive types declared in **Type** are polymorphic over their arguments in **Type**. If A is an arity of some sort and s is a sort, we write A/s for the arity obtained from A by replacing its sort with s . Especially, if A is well-typed in some global environment and local context, then A/s is typable by typability of all products in the Calculus of Inductive Constructions. The following typing rule is added to the theory.

Ind-Family Let $\text{Ind}[p](\Gamma_I := \Gamma_C)$ be an inductive definition. Let $\Gamma_P = [p_1 : P_1; \dots; p_p : P_p]$ be its context of parameters, $\Gamma_I = [I_1 : \forall \Gamma_P, A_1; \dots; I_k : \forall \Gamma_P, A_k]$ its context of definitions and $\Gamma_C = [c_1 : \forall \Gamma_P, C_1; \dots; c_n : \forall \Gamma_P, C_n]$ its context of constructors, with c_i a constructor of I_{q_i} .

Let $m \leq p$ be the length of the longest prefix of parameters such that the m first arguments of all occurrences of all I_j in all C_k (even the occurrences in the hypotheses of C_k) are exactly applied to $p_1 \dots p_m$ (m is the number of *recursively uniform parameters* and the $p-m$ remaining parameters are the *recursively non-uniform parameters*). Let q_1, \dots, q_r , with $0 \leq r \leq m$, be a (possibly) partial instantiation of the recursively uniform parameters of Γ_P . We have:

$$\frac{\left\{ \begin{array}{l} \text{Ind}[p](\Gamma_I := \Gamma_C) \in E \\ (E[] \vdash q_l : P'_l)_{l=1\dots r} \\ (E[] \vdash P'_l \leq_{\beta\delta\iota\zeta\eta} P_l\{p_u/q_u\}_{u=1\dots l-1})_{l=1\dots r} \\ 1 \leq j \leq k \end{array} \right.}{E[] \vdash I_j q_1 \dots q_r : \forall [p_{r+1} : P_{r+1}; \dots; p_p : P_p], (A_j)_{/s_j}}$$

provided that the following side conditions hold:

- $\Gamma_{P'}$ is the context obtained from Γ_P by replacing each P_l that is an arity with P'_l for $1 \leq l \leq r$ (notice that P_l arity implies P'_l arity since $E[] \vdash P'_l \leq_{\beta\delta\iota\zeta\eta} P_l\{p_u/q_u\}_{u=1\dots l-1}$);
- there are sorts s_i , for $1 \leq i \leq k$ such that, for $\Gamma_{I'} = [I_1 : \forall \Gamma_{P'}, (A_1)_{/s_1}; \dots; I_k : \forall \Gamma_{P'}, (A_k)_{/s_k}]$ we have $(E[\Gamma_{I'}; \Gamma_{P'}] \vdash C_i : s_{q_i})_{i=1\dots n}$;
- the sorts s_i are such that all eliminations, to **Prop**, **Set** and **Type**(j), are allowed (see Section 4.5.3).

Notice that if $I_j q_1 \dots q_r$ is typable using the rules **Ind-Const** and **App**, then it is typable using the rule **Ind-Family**. Conversely, the extended theory is not stronger than the theory without **Ind-Family**. We get an equiconsistency result by mapping each $\text{Ind}[p](\Gamma_I := \Gamma_C)$ occurring into a given derivation into as many different inductive types and constructors as the number of different (partial) replacements of sorts, needed for this derivation, in the parameters that are arities (this is possible because $\text{Ind}[p](\Gamma_I := \Gamma_C)$ well-formed implies that $\text{Ind}[p](\Gamma_{I'} := \Gamma_{C'})$ is well-formed and has the same allowed eliminations,

where $\Gamma_{I'}$ is defined as above and $\Gamma_{C'} = [c_1 : \forall \Gamma_{P'}, C_1; \dots; c_n : \forall \Gamma_{P'}, C_n]$. That is, the changes in the types of each partial instance $q_1 \dots q_r$ can be characterized by the ordered sets of arity sorts among the types of parameters, and to each signature is associated a new inductive definition with fresh names. Conversion is preserved as any (partial) instance $I_j q_1 \dots q_r$ or $C_i q_1 \dots q_r$ is mapped to the names chosen in the specific instance of $\text{Ind}[p](\Gamma_I := \Gamma_C)$.

In practice, the rule **Ind-Family** is used by COQ only when all the inductive types of the inductive definition are declared with an arity whose sort is in the **Type** hierarchy. Then, the polymorphism is over the parameters whose type is an arity of sort in the **Type** hierarchy. The sort s_j are chosen canonically so that each s_j is minimal with respect to the hierarchy $\text{Prop} \subset \text{Set}_p \subset \text{Type}$ where Set_p is predicative **Set**. More precisely, an empty or small singleton inductive definition (i.e. an inductive definition of which all inductive types are singleton – see paragraph 4.5.3) is set in **Prop**, a small non-singleton inductive type is set in **Set** (even in case **Set** is impredicative – see Section 4.8), and otherwise in the **Type** hierarchy.

Note that the side-condition about allowed elimination sorts in the rule **Ind-Family** is just to avoid to recompute the allowed elimination sorts at each instance of a pattern-matching (see section 4.5.3). As an example, let us consider the following definition:

```
Coq < Inductive option (A:Type) : Type :=
  | None : option A
  | Some : A -> option A.
```

As the definition is set in the **Type** hierarchy, it is used polymorphically over its parameters whose types are arities of a sort in the **Type** hierarchy. Here, the parameter A has this property, hence, if `option` is applied to a type in **Set**, the result is in **Set**. Note that if `option` is applied to a type in **Prop**, then, the result is not set in **Prop** but in **Set** still. This is because `option` is not a singleton type (see section 4.5.3) and it would lose the elimination to **Set** and **Type** if set in **Prop**.

```
Coq < Check (fun A:Set => option A).
fun A : Set => option A
      : Set -> Set

Coq < Check (fun A:Prop => option A).
fun A : Prop => option A
      : Prop -> Set
```

Here is another example.

```
Coq < Inductive prod (A B:Type) : Type := pair : A -> B -> prod A B.
```

As `prod` is a singleton type, it will be in **Prop** if applied twice to propositions, in **Set** if applied twice to at least one type in **Set** and none in **Type**, and in **Type** otherwise. In all cases, the three kind of eliminations schemes are allowed.

```
Coq < Check (fun A:Set => prod A).
fun A : Set => prod A
      : Set -> Type -> Type

Coq < Check (fun A:Prop => prod A A).
fun A : Prop => prod A A
      : Prop -> Prop

Coq < Check (fun (A:Prop) (B:Set) => prod A B).
fun (A : Prop) (B : Set) => prod A B
```

```

      : Prop -> Set -> Set
Coq < Check (fun (A:Type) (B:Prop) => prod A B) .
fun (A : Type) (B : Prop) => prod A B
      : Type -> Prop -> Type

```

Remark: Template polymorphism used to be called “sort-polymorphism of inductive types” before universe polymorphism (see Chapter 29) was introduced.

4.5.3 Destructors

The specification of inductive definitions with arities and constructors is quite natural. But we still have to say how to use an object in an inductive type.

This problem is rather delicate. There are actually several different ways to do that. Some of them are logically equivalent but not always equivalent from the computational point of view or from the user point of view.

From the computational point of view, we want to be able to define a function whose domain is an inductively defined type by using a combination of case analysis over the possible constructors of the object and recursion.

Because we need to keep a consistent theory and also we prefer to keep a strongly normalizing reduction, we cannot accept any sort of recursion (even terminating). So the basic idea is to restrict ourselves to primitive recursive functions and functionals.

For instance, assuming a parameter $A : \mathbf{Set}$ exists in the local context, we want to build a function length of type $\text{list } A \rightarrow \mathbf{nat}$ which computes the length of the list, so such that $(\text{length } (\text{nil } A)) = \mathbf{O}$ and $(\text{length } (\text{cons } A a l)) = (\mathbf{S} (\text{length } l))$. We want these equalities to be recognized implicitly and taken into account in the conversion rule.

From the logical point of view, we have built a type family by giving a set of constructors. We want to capture the fact that we do not have any other way to build an object in this type. So when trying to prove a property about an object m in an inductive definition it is enough to enumerate all the cases where m starts with a different constructor.

In case the inductive definition is effectively a recursive one, we want to capture the extra property that we have built the smallest fixed point of this recursive equation. This says that we are only manipulating finite objects. This analysis provides induction principles. For instance, in order to prove $\forall l : \text{list } A, (\text{has_length } A l (\text{length } l))$ it is enough to prove:

- $(\text{has_length } A (\text{nil } A) (\text{length } (\text{nil } A)))$
- $\forall a : A, \forall l : \text{list } A, (\text{has_length } A l (\text{length } l)) \rightarrow$
 $\rightarrow (\text{has_length } A (\text{cons } A a l) (\text{length } (\text{cons } A a l)))$

which given the conversion equalities satisfied by length is the same as proving:

- $(\text{has_length } A (\text{nil } A) \mathbf{O})$
- $\forall a : A, \forall l : \text{list } A, (\text{has_length } A l (\text{length } l)) \rightarrow$
 $\rightarrow (\text{has_length } A (\text{cons } A a l) (\mathbf{S} (\text{length } l)))$

One conceptually simple way to do that, following the basic scheme proposed by Martin-Löf in his Intuitionistic Type Theory, is to introduce for each inductive definition an elimination operator. At the logical level it is a proof of the usual induction principle and at the computational level it implements a generic operator for doing primitive recursion over the structure.

But this operator is rather tedious to implement and use. We choose in this version of COQ to factorize the operator for primitive recursion into two more primitive operations as was first suggested by Th. Coquand in [28]. One is the definition by pattern-matching. The second one is a definition by guarded fixpoints.

The `match...with ...end` construction.

The basic idea of this operator is that we have an object m in an inductive type I and we want to prove a property which possibly depends on m . For this, it is enough to prove the property for $m = (c_i u_1 \dots u_{p_i})$ for each constructor of I . The COQ term for this proof will be written:

$$\text{match } m \text{ with } (c_1 x_{11} \dots x_{1p_1}) \Rightarrow f_1 \mid \dots \mid (c_n x_{n1} \dots x_{np_n}) \Rightarrow f_n \text{ end}$$

In this expression, if m eventually happens to evaluate to $(c_i u_1 \dots u_{p_i})$ then the expression will behave as specified in its i -th branch and it will reduce to f_i where the $x_{i1} \dots x_{ip_i}$ are replaced by the $u_1 \dots u_{p_i}$ according to the ι -reduction.

Actually, for type-checking a `match...with...end` expression we also need to know the predicate P to be proved by case analysis. In the general case where I is an inductively defined n -ary relation, P is a predicate over $n+1$ arguments: the n first ones correspond to the arguments of I (parameters excluded), and the last one corresponds to object m . COQ can sometimes infer this predicate but sometimes not. The concrete syntax for describing this predicate uses the `as...in...return` construction. For instance, let us assume that I is an unary predicate with one parameter and one argument. The predicate is made explicit using the syntax:

$$\text{match } m \text{ as } x \text{ in } I _ a \text{ return } P \text{ with } (c_1 x_{11} \dots x_{1p_1}) \Rightarrow f_1 \mid \dots \mid (c_n x_{n1} \dots x_{np_n}) \Rightarrow f_n \text{ end}$$

The `as` part can be omitted if either the result type does not depend on m (non-dependent elimination) or m is a variable (in this case, m can occur in P where it is considered a bound variable). The `in` part can be omitted if the result type does not depend on the arguments of I . Note that the arguments of I corresponding to parameters *must* be `_`, because the result type is not generalized to all possible values of the parameters. The other arguments of I (sometimes called indices in the literature) have to be variables (a above) and these variables can occur in P . The expression after `in` must be seen as an *inductive type pattern*. Notice that expansion of implicit arguments and notations apply to this pattern. For the purpose of presenting the inference rules, we use a more compact notation:

$$\text{case}(m, (\lambda ax. P), \lambda x_{11} \dots x_{1p_1}. f_1 \mid \dots \mid \lambda x_{n1} \dots x_{np_n}. f_n)$$

Allowed elimination sorts. An important question for building the typing rule for `match` is what can be the type of $\lambda ax. P$ with respect to the type of m . If $m : I$ and $I : A$ and $\lambda ax. P : B$ then by $[I : A \mid B]$ we mean that one can use $\lambda ax. P$ with m in the above `match`-construct.

Notations. The $[I : A \mid B]$ is defined as the smallest relation satisfying the following rules: We write $[I \mid B]$ for $[I : A \mid B]$ where A is the type of I .

The case of inductive definitions in sorts **Set** or **Type** is simple. There is no restriction on the sort of the predicate to be eliminated.

Prod

$$\frac{[(I x) : A' \mid B']}{[I : \forall x : A, A' \mid \forall x : A, B']}$$

Set & Type

$$\frac{s_1 \in \{\text{Set}, \text{Type}(j)\} \quad s_2 \in \mathcal{S}}{[I : s_1 | I \rightarrow s_2]}$$

The case of Inductive definitions of sort **Prop** is a bit more complicated, because of our interpretation of this sort. The only harmless allowed elimination, is the one when predicate P is also of sort **Prop**.

Prop

$$[I : \text{Prop} | I \rightarrow \text{Prop}]$$

Prop is the type of logical propositions, the proofs of properties P in **Prop** could not be used for computation and are consequently ignored by the extraction mechanism. Assume A and B are two propositions, and the logical disjunction $A \vee B$ is defined inductively by:

```
Coq < Inductive or (A B:Prop) : Prop :=
  or_introl : A -> or A B | or_intror : B -> or A B.
```

The following definition which computes a boolean value by case over the proof of `or A B` is not accepted:

```
Coq < Fail Definition choice (A B: Prop) (x:or A B) :=
  match x with or_introl _ _ a => true | or_intror _ _ b => false end.
The command has indeed failed with message:
Incorrect elimination of "x" in the inductive type "or":
the return type has sort "Set" while it should be "Prop".
Elimination of an inductive object of sort Prop
is not allowed on a predicate in sort Set
because proofs can be eliminated only to build proofs.
```

From the computational point of view, the structure of the proof of `(or A B)` in this term is needed for computing the boolean value.

In general, if I has type **Prop** then P cannot have type $I \rightarrow \text{Set}$, because it will mean to build an informative proof of type $(P\ m)$ doing a case analysis over a non-computational object that will disappear in the extracted program. But the other way is safe with respect to our interpretation we can have I a computational object and P a non-computational one, it just corresponds to proving a logical property of a computational object.

In the same spirit, elimination on P of type $I \rightarrow \text{Type}$ cannot be allowed because it trivially implies the elimination on P of type $I \rightarrow \text{Set}$ by cumulativity. It also implies that there are two proofs of the same property which are provably different, contradicting the proof-irrelevance property which is sometimes a useful axiom:

```
Coq < Axiom proof_irrelevance : forall (P : Prop) (x y : P), x=y.
proof_irrelevance is declared
```

The elimination of an inductive definition of type **Prop** on a predicate P of type $I \rightarrow \text{Type}$ leads to a paradox when applied to impredicative inductive definition like the second-order existential quantifier `exProp` defined above, because it give access to the two projections on this type.

Empty and singleton elimination There are special inductive definitions in **Prop** for which more eliminations are allowed.

Prop-extended

$$\frac{I \text{ is an empty or singleton definition} \quad s \in \mathcal{S}}{[I : \text{Prop} | I \rightarrow s]}$$

A *singleton definition* has only one constructor and all the arguments of this constructor have type `Prop`. In that case, there is a canonical way to interpret the informative extraction on an object in that type, such that the elimination on any sort s is legal. Typical examples are the conjunction of non-informative propositions and the equality. If there is an hypothesis $h : a = b$ in the local context, it can be used for rewriting not only in logical propositions but also in any type.

```
Coq < Print eq_rec.
eq_rec =
fun (A : Type) (x : A) (P : A -> Set) => eq_rect x P
    : forall (A : Type) (x : A) (P : A -> Set),
      P x -> forall y : A, x = y -> P y
Argument A is implicit
Argument scopes are [type_scope _ function_scope _ _ _]

Coq < Extraction eq_rec.
(** val eq_rec : 'a1 -> 'a2 -> 'a1 -> 'a2 **)
let eq_rec _ f _ =
  f
```

An empty definition has no constructors, in that case also, elimination on any sort is allowed.

Type of branches. Let c be a term of type C , we assume C is a type of constructor for an inductive type I . Let P be a term that represents the property to be proved. We assume r is the number of parameters and p is the number of arguments.

We define a new type $\{c : C\}^P$ which represents the type of the branch corresponding to the $c : C$ constructor.

$$\begin{aligned} \{c : (I \ p_1 \dots p_r \ t_1 \dots t_p)\}^P &\equiv (P \ t_1 \dots t_p \ c) \\ \{c : \forall x : T, C\}^P &\equiv \forall x : T, \{(c \ x) : C\}^P \end{aligned}$$

We write $\{c\}^P$ for $\{c : C\}^P$ with C the type of c .

Example. The following term in concrete syntax:

```
match t as l return P' with
| nil _ => t1
| cons _ hd t1 => t2
end
```

can be represented in abstract syntax as

$$\text{case}(t, P, f_1 \mid f_2)$$

where

$$\begin{aligned} P &= \lambda l . P' \\ f_1 &= t_1 \\ f_2 &= \lambda (hd : \text{nat}) . \lambda (tl : \text{list nat}) . t_2 \end{aligned}$$

According to the definition:

$$\{(\text{nil nat})\}^P \equiv \{(\text{nil nat}) : (\text{list nat})\}^P \equiv (P \ (\text{nil nat}))$$

$$\begin{aligned}
\{(\text{cons nat})\}^P &\equiv \{(\text{cons nat}) : (\text{nat} \rightarrow \text{list nat} \rightarrow \text{list nat})\}^P \equiv \\
&\equiv \forall n : \text{nat}, \{(\text{cons nat } n) : \text{list nat} \rightarrow \text{list nat}\}^P \equiv \\
&\equiv \forall n : \text{nat}, \forall l : \text{list nat}, \{(\text{cons nat } n \ l) : \text{list nat}\}^P \equiv \\
&\equiv \forall n : \text{nat}, \forall l : \text{list nat}, (P (\text{cons nat } n \ l)).
\end{aligned}$$

Given some P , then $\{(\text{nil nat})\}^P$ represents the expected type of f_1 , and $\{(\text{cons nat})\}^P$ represents the expected type of f_2 .

Typing rule. Our very general destructor for inductive definition enjoys the following typing rule

match

$$\frac{E[\Gamma] \vdash c : (I \ q_1 \dots q_r \ t_1 \dots t_s) \quad E[\Gamma] \vdash P : B \quad [(I \ q_1 \dots q_r) | B] \quad (E[\Gamma] \vdash f_i : \{(c_{p_i} \ q_1 \dots q_r)\}^P)_{i=1 \dots l}}{E[\Gamma] \vdash \text{case}(c, P, f_1 | \dots | f_l) : (P \ t_1 \dots t_s \ c)}$$

provided I is an inductive type in a definition $\text{Ind}[r](\Gamma_I := \Gamma_C)$ with $\Gamma_C = [c_1 : C_1; \dots; c_n : C_n]$ and $c_{p_1} \dots c_{p_l}$ are the only constructors of I .

Example. Below is a typing rule for the term shown in the previous example:

$$\frac{E[\Gamma] \vdash t : (\text{list nat}) \quad E[\Gamma] \vdash P : B \quad [(\text{list nat}) | B] \quad E[\Gamma] \vdash f_1 : \{(\text{nil nat})\}^P \quad E[\Gamma] \vdash f_2 : \{(\text{cons nat})\}^P}{E[\Gamma] \vdash \text{case}(t, P, f_1 | f_2) : (P \ t)}$$

Definition of ι -reduction. We still have to define the ι -reduction in the general case.

A ι -redex is a term of the following form:

$$\text{case}((c_{p_i} \ q_1 \dots q_r \ a_1 \dots a_m), P, f_1 | \dots | f_l)$$

with c_{p_i} the i -th constructor of the inductive type I with r parameters.

The ι -contraction of this term is $(f_i \ a_1 \dots a_m)$ leading to the general reduction rule:

$$\text{case}((c_{p_i} \ q_1 \dots q_r \ a_1 \dots a_m), P, f_1 | \dots | f_n) \triangleright_{\iota} (f_i \ a_1 \dots a_m)$$

4.5.4 Fixpoint definitions

The second operator for elimination is fixpoint definition. This fixpoint may involve several mutually recursive definitions. The basic concrete syntax for a recursive set of mutually recursive declarations is (with Γ_i contexts):

$$\text{fix } f_1(\Gamma_1) : A_1 := t_1 \text{ with } \dots \text{ with } f_n(\Gamma_n) : A_n := t_n$$

The terms are obtained by projections from this set of declarations and are written

$$\text{fix } f_1(\Gamma_1) : A_1 := t_1 \text{ with } \dots \text{ with } f_n(\Gamma_n) : A_n := t_n \text{ for } f_i$$

In the inference rules, we represent such a term by

$$\text{Fix } f_i \{f_1 : A'_1 := t'_1 \dots f_n : A'_n := t'_n\}$$

with t'_i (resp. A'_i) representing the term t_i abstracted (resp. generalized) with respect to the bindings in the context Γ_i , namely $t'_i = \lambda \Gamma_i. t_i$ and $A'_i = \forall \Gamma_i. A_i$.

Typing rule

The typing rule is the expected one for a fixpoint.

Fix

$$\frac{(E[\Gamma] \vdash A_i : s_i)_{i=1\dots n} \quad (E[\Gamma, f_1 : A_1, \dots, f_n : A_n] \vdash t_i : A_i)_{i=1\dots n}}{E[\Gamma] \vdash \text{Fix } f_i \{f_1 : A_1 := t_1 \dots f_n : A_n := t_n\} : A_i}$$

Any fixpoint definition cannot be accepted because non-normalizing terms allow proofs of absurdity. The basic scheme of recursion that should be allowed is the one needed for defining primitive recursive functionals. In that case the fixpoint enjoys a special syntactic restriction, namely one of the arguments belongs to an inductive type, the function starts with a case analysis and recursive calls are done on variables coming from patterns and representing subterms. For instance in the case of natural numbers, a proof of the induction principle of type

$$\forall P : \text{nat} \rightarrow \text{Prop}, (P \text{ O}) \rightarrow (\forall n : \text{nat}, (P \ n) \rightarrow (P \ (\text{S } n))) \rightarrow \forall n : \text{nat}, (P \ n)$$

can be represented by the term:

$$\lambda P : \text{nat} \rightarrow \text{Prop}. \lambda f : (P \text{ O}). \lambda g : (\forall n : \text{nat}, (P \ n) \rightarrow (P \ (\text{S } n))). \\ \text{Fix } h \{h : \forall n : \text{nat}, (P \ n) := \lambda n : \text{nat}. \text{case}(n, P, f \mid \lambda p : \text{nat}. (g \ p \ (h \ p)))\}$$

Before accepting a fixpoint definition as being correctly typed, we check that the definition is “guarded”. A precise analysis of this notion can be found in [67]. The first stage is to precise on which argument the fixpoint will be decreasing. The type of this argument should be an inductive definition. For doing this, the syntax of fixpoints is extended and becomes

$$\text{Fix } f_i \{f_1/k_1 : A_1 := t_1 \dots f_n/k_n : A_n := t_n\}$$

where k_i are positive integers. Each k_i represents the index of parameter of f_i , on which f_i is decreasing. Each A_i should be a type (reducible to a term) starting with at least k_i products $\forall y_1 : B_1, \dots \forall y_{k_i} : B_{k_i}, A'_i$ and B_{k_i} is an inductive type.

Now in the definition t_i , if f_j occurs then it should be applied to at least k_j arguments and the k_j -th argument should be syntactically recognized as structurally smaller than y_{k_i} .

The definition of being structurally smaller is a bit technical. One needs first to define the notion of *recursive arguments of a constructor*. For an inductive definition $\text{Ind}[r](\Gamma_I := \Gamma_C)$, if the type of a constructor c has the form $\forall p_1 : P_1, \dots \forall p_r : P_r, \forall x_1 : T_1, \dots \forall x_r : T_r, (I_j \ p_1 \dots p_r \ t_1 \dots t_s)$, then the recursive arguments will correspond to T_i in which one of the I_l occurs.

The main rules for being structurally smaller are the following:

Given a variable y of type an inductive definition in a declaration $\text{Ind}[r](\Gamma_I := \Gamma_C)$ where Γ_I is $[I_1 : A_1; \dots; I_k : A_k]$, and Γ_C is $[c_1 : C_1; \dots; c_n : C_n]$. The terms structurally smaller than y are:

- $(t \ u)$ and $\lambda x : u. t$ when t is structurally smaller than y .
- $\text{case}(c, P, f_1 \dots f_n)$ when each f_i is structurally smaller than y .
If c is y or is structurally smaller than y , its type is an inductive definition I_p part of the inductive declaration corresponding to y . Each f_i corresponds to a type of constructor $C_q \equiv \forall p_1 : P_1, \dots, \forall p_r : P_r, \forall y_1 : B_1, \dots \forall y_k : B_k, (I \ a_1 \dots a_k)$ and can consequently be written $\lambda y_1 : B'_1. \dots \lambda y_k : B'_k. g_i$. (B'_i is obtained from B_i by substituting parameters variables) the variables y_j occurring in g_i corresponding to recursive arguments B_i (the ones in which one of the I_l occurs) are structurally smaller than y .

The following definitions are correct, we enter them using the `Fixpoint` command as described in Section 1.3.4 and show the internal representation.

```

Coq < Fixpoint plus (n m:nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (plus p m)
  end.
plus is defined
plus is recursively defined (decreasing on 1st argument)

Coq < Print plus.
plus =
fix plus (n m : nat) {struct n} : nat :=
  match n with
  | 0 => m
  | S p => S (plus p m)
  end
  : nat -> nat -> nat
Argument scopes are [nat_scope nat_scope]

Coq < Fixpoint lgth (A:Set) (l:list A) {struct l} : nat :=
  match l with
  | nil _ => 0
  | cons _ a l' => S (lgth A l')
  end.
lgth is defined
lgth is recursively defined (decreasing on 2nd argument)

Coq < Print lgth.
lgth =
fix lgth (A : Set) (l : list A) {struct l} : nat :=
  match l with
  | nil _ => 0
  | cons _ _ l' => S (lgth A l')
  end
  : forall A : Set, list A -> nat
Argument scopes are [type_scope _]

Coq < Fixpoint sizet (t:tree) : nat := let (f) := t in S (sizef f)
  with sizef (f:forest) : nat :=
  match f with
  | emptyf => 0
  | consf t f => plus (sizet t) (sizef f)
  end.
sizet is defined
sizef is defined
sizet, sizef are recursively defined (decreasing respectively on 1st,
1st arguments)

Coq < Print sizet.
sizet =
fix sizet (t : tree) : nat := let (f) := t in S (sizef f)
with sizef (f : forest) : nat :=
  match f with
  | emptyf => 0

```

```

| consf t f0 => plus (sized t) (sized f0)
end
for sized
  : tree -> nat

```

Reduction rule

Let F be the set of declarations: $f_1/k_1 : A_1 := t_1 \dots f_n/k_n : A_n := t_n$. The reduction for fixpoints is:

$$(\text{Fix } f_i\{F\} a_1 \dots a_{k_i}) \triangleright_\iota t_i\{(f_k/\text{Fix } f_k\{F\})_{k=1\dots n}\} a_1 \dots a_{k_i}$$

when a_{k_i} starts with a constructor. This last restriction is needed in order to keep strong normalization and corresponds to the reduction for primitive recursive operators. The following reductions are now possible:

$$\begin{aligned} \text{plus } (S (S O)) (S O) &\triangleright_\iota S (\text{plus } (S O) (S O)) \\ &\triangleright_\iota S (S (\text{plus } O (S O))) \\ &\triangleright_\iota S (S (S O)) \end{aligned}$$

Mutual induction

The principles of mutual induction can be automatically generated using the `Scheme` command described in Section 13.1.

4.6 Admissible rules for global environments

From the original rules of the type system, one can show the admissibility of rules which change the local context of definition of objects in the global environment. We show here the admissible rules that are used in the discharge mechanism at the end of a section.

Abstraction. One can modify a global declaration by generalizing it over a previously assumed constant c . For doing that, we need to modify the reference to the global declaration in the subsequent global environment and local context by explicitly applying this constant to the constant c' .

Below, if Γ is a context of the form $[y_1 : A_1; \dots; y_n : A_n]$, we write $\forall x : U, \Gamma\{c/x\}$ to mean $[y_1 : \forall x : U, A_1\{c/x\}; \dots; y_n : \forall x : U, A_n\{c/x\}]$ and $E\{|\Gamma|/|\Gamma|c\}$ to mean the parallel substitution $E\{y_1/(y_1 c) \dots \{y_n/(y_n c)\}$.

First abstracting property:

$$\begin{aligned} &\frac{\mathcal{WF}(E; c : U; E'; c' := t : T; E'')[\Gamma]}{\mathcal{WF}(E; c : U; E'; c' := \lambda x : U. t\{c/x\} : \forall x : U, T\{c/x\}; E''\{c'/(c' c)\})[\Gamma\{c/(c' c)\}]} \\ &\frac{\mathcal{WF}(E; c : U; E'; c' : T; E'')[\Gamma]}{\mathcal{WF}(E; c : U; E'; c' : \forall x : U, T\{c/x\}; E''\{c'/(c' c)\})[\Gamma\{c/(c' c)\}]} \\ &\frac{\mathcal{WF}(E; c : U; E'; \text{Ind}[p](\Gamma_I := \Gamma_C); E'')[\Gamma]}{\mathcal{WF} \left(\frac{(E; c : U; E'; \text{Ind}[p+1](\forall x : U, \Gamma_I\{c/x\} := \forall x : U, \Gamma_C\{c/x\}); E''\{|\Gamma_I, \Gamma_C|/|\Gamma_I, \Gamma_C| c\})}{[\Gamma\{|\Gamma_I, \Gamma_C|/|\Gamma_I, \Gamma_C| c\}]} \right)} \end{aligned}$$

One can similarly modify a global declaration by generalizing it over a previously defined constant c' . Below, if Γ is a context of the form $[y_1 : A_1; \dots; y_n : A_n]$, we write $\Gamma\{c/u\}$ to mean $[y_1 : A_1\{c/u\}; \dots; y_n : A_n\{c/u\}]$.

Second abstracting property:

$$\frac{\mathcal{WF}(E; c := u : U; E'; c' := t : T; E'')[\Gamma]}{\mathcal{WF}(E; c := u : U; E'; c' := (\text{let } x := u : U \text{ in } t\{c/x\}) : T\{c/u\}; E'')[\Gamma]}$$

$$\frac{\mathcal{WF}(E; c := u : U; E'; c' : T; E'')[\Gamma]}{\mathcal{WF}(E; c := u : U; E'; c' : T\{c/u\}; E'')[\Gamma]}$$

$$\frac{\mathcal{WF}(E; c := u : U; E'; \text{Ind}[p](\Gamma_I := \Gamma_C); E'')[\Gamma]}{\mathcal{WF}(E; c := u : U; E'; \text{Ind}[p](\Gamma_I\{c/u\} := \Gamma_C\{c/u\}); E'')[\Gamma]}$$

Pruning the local context. If one abstracts or substitutes constants with the above rules then it may happen that some declared or defined constant does not occur any more in the subsequent global environment and in the local context. One can consequently derive the following property.

First pruning property:

$$\frac{\mathcal{WF}(E; c : U; E')[\Gamma] \quad c \text{ does not occur in } E' \text{ and } \Gamma}{\mathcal{WF}(E; E')[\Gamma]}$$

Second pruning property:

$$\frac{\mathcal{WF}(E; c := u : U; E')[\Gamma] \quad c \text{ does not occur in } E' \text{ and } \Gamma}{\mathcal{WF}(E; E')[\Gamma]}$$

4.7 Co-inductive types

The implementation contains also co-inductive definitions, which are types inhabited by infinite objects. More information on co-inductive definitions can be found in [68, 70, 71].

4.8 The Calculus of Inductive Construction with impredicative Set

COQ can be used as a type-checker for the Calculus of Inductive Constructions with an impredicative sort **Set** by using the compiler option `-impredicative-set`. For example, using the ordinary `coqtop` command, the following is rejected.

```
Coq < Fail Definition id: Set := forall X:Set, X->X.
The command has indeed failed with message:
The term "forall X : Set, X -> X" has type "Type"
while it is expected to have type "Set".
```

while it will type-check, if one uses instead the `coqtop -impredicative-set` command.

The major change in the theory concerns the rule for product formation in the sort **Set**, which is extended to a domain in any sort:

Prod

$$\frac{E[\Gamma] \vdash T : s \quad s \in \mathcal{S} \quad E[\Gamma :: (x : T)] \vdash U : \mathbf{Set}}{E[\Gamma] \vdash \forall x : T, U : \mathbf{Set}}$$

This extension has consequences on the inductive definitions which are allowed. In the impredicative system, one can build so-called *large inductive definitions* like the example of second-order existential quantifier (`exSet`).

There should be restrictions on the eliminations which can be performed on such definitions. The eliminations rules in the impredicative system for sort **Set** become:

Set

$$\frac{s \in \{\mathbf{Prop}, \mathbf{Set}\}}{[I : \mathbf{Set} | I \rightarrow s]} \quad \frac{I \text{ is a small inductive definition} \quad s \in \{\mathbf{Type}(i)\}}{[I : \mathbf{Set} | I \rightarrow s]}$$

Chapter 5

The Module System

The module system extends the Calculus of Inductive Constructions providing a convenient way to structure large developments as well as a means of massive abstraction.

5.1 Modules and module types

Access path. It is denoted by p , it can be either a module variable X or, if p' is an access path and id an identifier, then $p'.id$ is an access path.

Structure element. It is denoted by e and is either a definition of a constant, an assumption, a definition of an inductive, a definition of a module, an alias of module or a module type abbreviation.

Structure expression. It is denoted by S and can be:

- an access path p
- a plain structure `Struct $e; \dots; e$ End`
- a functor `Functor($X : S$) S'` , where X is a module variable, S and S' are structure expression
- an application $S p$, where S is a structure expression and p an access path
- a refined structure S with $p := p'$ or S with $p := t : T$ where S is a structure expression, p and p' are access paths, t is a term and T is the type of t .

Module definition, is written `Mod($X : S [:= S']$)` and consists of a module variable X , a module type S which can be any structure expression and optionally a module implementation S' which can be any structure expression except a refined structure.

Module alias, is written `ModA($X == p$)` and consists of a module variable X and a module path p .

Module type abbreviation, is written `ModType($Y := S$)`, where Y is an identifier and S is any structure expression .

5.2 Typing Modules

In order to introduce the typing system we first slightly extend the syntactic class of terms and environments given in section 4.1. The environments, apart from definitions of constants and inductive types now also hold any other structure elements. Terms, apart from variables, constants and complex terms, include also access paths.

We also need additional typing judgments:

- $E[] \vdash \mathcal{WF}(S)$, denoting that a structure S is well-formed,
- $E[] \vdash p : S$, denoting that the module pointed by p has type S in environment E .
- $E[] \vdash S \longrightarrow \bar{S}$, denoting that a structure S is evaluated to a structure \bar{S} in weak head normal form.
- $E[] \vdash S_1 <: S_2$, denoting that a structure S_1 is a subtype of a structure S_2 .
- $E[] \vdash e_1 <: e_2$, denoting that a structure element e_1 is more precise than a structure element e_2 .

The rules for forming structures are the following:

WF-STR

$$\frac{\mathcal{WF}(E; E')[]}{E[] \vdash \mathcal{WF}(\text{Struct } E' \text{ End})}$$

WF-FUN

$$\frac{E; \text{Mod}(X : S)[] \vdash \mathcal{WF}(\bar{S}')}{E[] \vdash \mathcal{WF}(\text{Functor}(X : S) S')}$$

Evaluation of structures to weak head normal form:

WEVAL-APP

$$\frac{\begin{array}{c} E[] \vdash S \longrightarrow \text{Functor}(X : S_1) S_2 \quad E[] \vdash S_1 \longrightarrow \bar{S}_1 \\ E[] \vdash p : S_3 \quad E[] \vdash S_3 <: \bar{S}_1 \end{array}}{E[] \vdash S p \longrightarrow S_2\{p/X, t_1/p_1.c_1, \dots, t_n/p_n.c_n\}}$$

In the last rule, $\{t_1/p_1.c_1, \dots, t_n/p_n.c_n\}$ is the resulting substitution from the inlining mechanism. We substitute in S the inlined fields $p_i.c_i$ from $\text{Mod}(X : S_1)$ by the corresponding delta-reduced term t_i in p .

WEVAL-WITH-MOD

$$\frac{\begin{array}{c} E[] \vdash S \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X : S_1); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash S_1 \longrightarrow \bar{S}_1 \\ E[] \vdash p : S_2 \quad E; e_1; \dots; e_i[] \vdash S_2 <: \bar{S}_1 \end{array}}{E[] \vdash S \text{ with } x := p \longrightarrow \text{Struct } e_1; \dots; e_i; \text{ModA}(X == p); e_{i+2}\{p/X\}; \dots; e_n\{p/X\} \text{ End}}$$

WEVAL-WITH-MOD-REC

$$\frac{\begin{array}{c} E[] \vdash S \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X_1 : S_1); e_{i+2}; \dots; e_n \text{ End} \\ E; e_1; \dots; e_i[] \vdash S_1 \text{ with } p := p_1 \longrightarrow \bar{S}_2 \end{array}}{E[] \vdash S \text{ with } X_1.p := p_1 \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X : \bar{S}_2); e_{i+2}\{p_1/X_1.p\}; \dots; e_n\{p_1/X_1.p\} \text{ End}}$$

WEVAL-WITH-DEF

$$\frac{E[] \vdash S \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Assum}()(c : T_1); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash \text{Def}()(c := t : T) <: \text{Assum}()(c : T_1)}{E[] \vdash S \text{ with } c := t : T \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Def}()(c := t : T); e_{i+2}; \dots; e_n \text{ End}}$$

WEVAL-WITH-DEF-REC

$$\frac{E[] \vdash S \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X_1 : S_1); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash S_1 \text{ with } p := p_1 \longrightarrow \bar{S}_2}{E[] \vdash S \text{ with } X_1.p := t : T \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X : \bar{S}_2); e_{i+2}; \dots; e_n \text{ End}}$$

WEVAL-PATH-MOD

$$\frac{E[] \vdash p \longrightarrow \text{Struct } e_1; \dots; e_i; \text{Mod}(X : S [:= S_1]); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash S \longrightarrow \bar{S}}{E[] \vdash p.X \longrightarrow \bar{S}}$$

$$\frac{\mathcal{WF}(E)[] \quad \text{Mod}(X : S [:= S_1]) \in E \quad E[] \vdash S \longrightarrow \bar{S}}{E[] \vdash X \longrightarrow \bar{S}}$$

WEVAL-PATH-ALIAS

$$\frac{E[] \vdash p \longrightarrow \text{Struct } e_1; \dots; e_i; \text{ModA}(X == p_1); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash p_1 \longrightarrow \bar{S}}{E[] \vdash p.X \longrightarrow \bar{S}}$$

$$\frac{\mathcal{WF}(E)[] \quad \text{ModA}(X == p_1) \in E \quad E[] \vdash p_1 \longrightarrow \bar{S}}{E[] \vdash X \longrightarrow \bar{S}}$$

WEVAL-PATH-TYPE

$$\frac{E[] \vdash p \longrightarrow \text{Struct } e_1; \dots; e_i; \text{ModType}(Y := S); e_{i+2}; \dots; e_n \text{ End} \quad E; e_1; \dots; e_i[] \vdash S \longrightarrow \bar{S}}{E[] \vdash p.Y \longrightarrow \bar{S}}$$

WEVAL-PATH-TYPE

$$\frac{\mathcal{WF}(E)[] \quad \text{ModType}(Y := S) \in E \quad E[] \vdash S \longrightarrow \bar{S}}{E[] \vdash Y \longrightarrow \bar{S}}$$

Rules for typing module:

MT-EVAL

$$\frac{E[] \vdash p \longrightarrow \bar{S}}{E[] \vdash p : \bar{S}}$$

MT-STR

$$\frac{E[] \vdash p : S}{E[] \vdash p : S/p}$$

The last rule, called strengthening is used to make all module fields manifestly equal to themselves. The notation S/p has the following meaning:

- if $S \longrightarrow \text{Struct } e_1; \dots; e_n \text{ End}$ then $S/p = \text{Struct } e_1/p; \dots; e_n/p \text{ End}$ where e/p is defined as follows:
 - $\text{Def}()(c := t : T)/p^1 = \text{Def}()(c := t : T)$
 - $\text{Assum}()(c : U)/p = \text{Def}()(c := p.c : U)$
 - $\text{Mod}(X : S)/p = \text{ModA}(X == p.X)$
 - $\text{ModA}(X == p')/p = \text{ModA}(X == p')$
 - $\text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I)/p = \text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I)$
 - $\text{Ind}_{p'}()[\Gamma_P](\Gamma_C := \Gamma_I)/p = \text{Ind}_{p'}()[\Gamma_P](\Gamma_C := \Gamma_I)$
- if $S \longrightarrow \text{Functor}(X : S') S''$ then $S/p = S$

The notation $\text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I)$ denotes an inductive definition that is definitionally equal to the inductive definition in the module denoted by the path p . All rules which have $\text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I)$ as premises are also valid for $\text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I)$. We give the formation rule for $\text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I)$ below as well as the equality rules on inductive types and constructors.

The module subtyping rules:

MSUB-STR

$$\frac{\begin{array}{c} E; e_1; \dots; e_n[] \vdash e_{\sigma(i)} <: e'_i \text{ for } i = 1..m \\ \sigma : \{1 \dots m\} \rightarrow \{1 \dots n\} \text{ injective} \end{array}}{E[] \vdash \text{Struct } e_1; \dots; e_n \text{ End} <: \text{Struct } e'_1; \dots; e'_m \text{ End}}$$

MSUB-FUN

$$\frac{E[] \vdash \overline{S'_1} <: \overline{S_1} \quad E; \text{Mod}(X : S'_1)[] \vdash \overline{S_2} <: \overline{S'_2}}{E[] \vdash \text{Functor}(X : S_1) S_2 <: \text{Functor}(X : S'_1) S'_2}$$

Structure element subtyping rules:

ASSUM-ASSUM

$$\frac{E[] \vdash T_1 \leq_{\beta\delta\iota\zeta\eta} T_2}{E[] \vdash \text{Assum}()(c : T_1) <: \text{Assum}()(c : T_2)}$$

DEF-ASSUM

$$\frac{E[] \vdash T_1 \leq_{\beta\delta\iota\zeta\eta} T_2}{E[] \vdash \text{Def}()(c := t : T_1) <: \text{Assum}()(c : T_2)}$$

ASSUM-DEF

$$\frac{E[] \vdash T_1 \leq_{\beta\delta\iota\zeta\eta} T_2 \quad E[] \vdash c =_{\beta\delta\iota\zeta\eta} t_2}{E[] \vdash \text{Assum}()(c : T_1) <: \text{Def}()(c := t_2 : T_2)}$$

¹Opaque definitions are processed as assumptions.

DEF-DEF

$$\frac{E[] \vdash T_1 \leq_{\beta\delta\iota\zeta\eta} T_2 \quad E[] \vdash t_1 =_{\beta\delta\iota\zeta\eta} t_2}{E[] \vdash \text{Def}()(c := t_1 : T_1) <: \text{Def}()(c := t_2 : T_2)}$$

IND-IND

$$\frac{E[] \vdash \Gamma_P =_{\beta\delta\iota\zeta\eta} \Gamma'_P \quad E[\Gamma_P] \vdash \Gamma_C =_{\beta\delta\iota\zeta\eta} \Gamma'_C \quad E[\Gamma_P; \Gamma_C] \vdash \Gamma_I =_{\beta\delta\iota\zeta\eta} \Gamma'_I}{E[] \vdash \text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I) <: \text{Ind}[\Gamma'_P](\Gamma'_C := \Gamma'_I)}$$

INDP-IND

$$\frac{E[] \vdash \Gamma_P =_{\beta\delta\iota\zeta\eta} \Gamma'_P \quad E[\Gamma_P] \vdash \Gamma_C =_{\beta\delta\iota\zeta\eta} \Gamma'_C \quad E[\Gamma_P; \Gamma_C] \vdash \Gamma_I =_{\beta\delta\iota\zeta\eta} \Gamma'_I}{E[] \vdash \text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I) <: \text{Ind}[\Gamma'_P](\Gamma'_C := \Gamma'_I)}$$

INDP-INDP

$$\frac{E[] \vdash \Gamma_P =_{\beta\delta\iota\zeta\eta} \Gamma'_P \quad E[\Gamma_P] \vdash \Gamma_C =_{\beta\delta\iota\zeta\eta} \Gamma'_C \quad E[\Gamma_P; \Gamma_C] \vdash \Gamma_I =_{\beta\delta\iota\zeta\eta} \Gamma'_I \quad E[] \vdash p =_{\beta\delta\iota\zeta\eta} p'}{E[] \vdash \text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I) <: \text{Ind}_{p'}()[\Gamma'_P](\Gamma'_C := \Gamma'_I)}$$

MOD-MOD

$$\frac{E[] \vdash S_1 <: S_2}{E[] \vdash \text{Mod}(X : S_1) <: \text{Mod}(X : S_2)}$$

ALIAS-MOD

$$\frac{E[] \vdash p : S_1 \quad E[] \vdash S_1 <: S_2}{E[] \vdash \text{ModA}(X == p) <: \text{Mod}(X : S_2)}$$

MOD-ALIAS

$$\frac{E[] \vdash p : S_2 \quad E[] \vdash S_1 <: S_2 \quad E[] \vdash X =_{\beta\delta\iota\zeta\eta} p}{E[] \vdash \text{Mod}(X : S_1) <: \text{ModA}(X == p)}$$

ALIAS-ALIAS

$$\frac{E[] \vdash p_1 =_{\beta\delta\iota\zeta\eta} p_2}{E[] \vdash \text{ModA}(X == p_1) <: \text{ModA}(X == p_2)}$$

MODTYPE-MODTYPE

$$\frac{E[] \vdash S_1 <: S_2 \quad E[] \vdash S_2 <: S_1}{E[] \vdash \text{ModType}(Y := S_1) <: \text{ModType}(Y := S_2)}$$

New environment formation rules

WF-MOD

$$\frac{\mathcal{WF}(E)[] \quad E[] \vdash \mathcal{WF}(S)}{\mathcal{WF}(E; \text{Mod}(X : S))[]}$$

WF-MOD

$$\frac{E[] \vdash S_2 <: S_1 \quad \mathcal{WF}(E)[] \quad E[] \vdash \mathcal{WF}(S_1) \quad E[] \vdash \mathcal{WF}(S_2)}{\mathcal{WF}(E; \text{Mod}(X : S_1 [:= S_2]))[]}$$

WF-ALIAS

$$\frac{\mathcal{WF}(E)[] \quad E[] \vdash p : S}{\mathcal{WF}(E, \text{ModA}(X == p))[]}$$

WF-MODTYPE

$$\frac{\mathcal{WF}(E)[] \quad E[] \vdash \mathcal{WF}(S)}{\mathcal{WF}(E, \text{ModType}(Y := S))[]}$$

WF-IND

$$\frac{\begin{array}{c} \mathcal{WF}(E; \text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I))[] \\ E[] \vdash p : \text{Struct } e_1; \dots; e_n; \text{Ind}[\Gamma'_P](\Gamma'_C := \Gamma'_I); \dots \text{ End} : \\ E[] \vdash \text{Ind}[\Gamma'_P](\Gamma'_C := \Gamma'_I) <: \text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I) \end{array}}{\mathcal{WF}(E; \text{Ind}_p()[\Gamma_P](\Gamma_C := \Gamma_I))[]}$$

Component access rules

ACC-TYPE

$$\frac{E[\Gamma] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Assum}()(c : T); \dots \text{ End}}{E[\Gamma] \vdash p.c : T}$$

$$\frac{E[\Gamma] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Def}()(c := t : T); \dots \text{ End}}{E[\Gamma] \vdash p.c : T}$$

ACC-DELTA Notice that the following rule extends the delta rule defined in section 4.3

$$\frac{E[\Gamma] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Def}()(c := t : U); \dots \text{ End}}{E[\Gamma] \vdash p.c \triangleright_\delta t}$$

In the rules below we assume Γ_P is $[p_1 : P_1; \dots; p_r : P_r]$, Γ_I is $[I_1 : A_1; \dots; I_k : A_k]$, and Γ_C is $[c_1 : C_1; \dots; c_n : C_n]$

ACC-IND

$$\frac{E[\Gamma] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I); \dots \text{ End}}{E[\Gamma] \vdash p.I_j : (p_1 : P_1) \dots (p_r : P_r) A_j}$$

$$\frac{E[\Gamma] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Ind}[\Gamma_P](\Gamma_C := \Gamma_I); \dots \text{ End}}{E[\Gamma] \vdash p.c_m : (p_1 : P_1) \dots (p_r : P_r) C_m I_j (I_j p_1 \dots p_r)_{j=1 \dots k}}$$

ACC-INDP

$$\frac{E[] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Ind}_{p'}()[\Gamma_P](\Gamma_C := \Gamma_I) ; \dots \text{ End}}{E[] \vdash p.I_i \triangleright_\delta p'.I_i}$$

$$\frac{E[] \vdash p : \text{Struct } e_1; \dots; e_i; \text{Ind}_{p'}()[\Gamma_P](\Gamma_C := \Gamma_I) ; \dots \text{ End}}{E[] \vdash p.c_i \triangleright_\delta p'.c_i}$$

Part II

The proof engine

Chapter 6

Vernacular commands

6.1 Displaying

6.1.1 `Print qualid`.

This command displays on the screen information about the declared or defined object referred by *qualid*.

Error messages:

1. *qualid* not a defined object

Variants:

1. `Print Term qualid`.
This is a synonym to `Print qualid` when *qualid* denotes a global constant.
2. `About qualid`.
This displays various information about the object denoted by *qualid*: its kind (module, constant, assumption, inductive, constructor, abbreviation, ...), long name, type, implicit arguments and argument scopes. It does not print the body of definitions or proofs.

6.1.2 `Print All`.

This command displays information about the current state of the environment, including sections and modules.

Variants:

1. `Inspect num`.
This command displays the *num* last objects of the current environment, including sections and modules.
2. `Print Section ident`.
should correspond to a currently open section, this command displays the objects defined since the beginning of this section.

6.2 Flags, Options and Tables

COQ configurability is based on flags (e.g. Set Printing All in Section 2.9), options (e.g. Set Printing Width *integer* in Section 6.9.6), or tables (e.g. Add Printing Record *ident*, in Section 2.2.4). The names of flags, options and tables are made of non-empty sequences of identifiers (conventionally with capital initial letter). The general commands handling flags, options and tables are given below.

6.2.1 Set *flag*.

This command switches *flag* on. The original state of *flag* is restored when the current module ends.

Variants:

1. Local Set *flag*.
This command switches *flag* on. The original state of *flag* is restored when the current *section* ends.
2. Global Set *flag*.
This command switches *flag* on. The original state of *flag* is *not* restored at the end of the module. Additionally, if set in a file, *flag* is switched on when the file is Require-d.

6.2.2 Unset *flag*.

This command switches *flag* off. The original state of *flag* is restored when the current module ends.

Variants:

1. Local Unset *flag*.
This command switches *flag* off. The original state of *flag* is restored when the current *section* ends.
2. Global Unset *flag*.
This command switches *flag* off. The original state of *flag* is *not* restored at the end of the module. Additionally, if set in a file, *flag* is switched off when the file is Require-d.

6.2.3 Test *flag*.

This command prints whether *flag* is on or off.

6.2.4 Set *option value*.

This command sets *option* to *value*. The original value of *option* is restored when the current module ends.

Variants:

1. Local Set *option value*. This command sets *option* to *value*. The original value of *option* is restored at the end of the module.
2. Global Set *option value*. This command sets *option* to *value*. The original value of *option* is *not* restored at the end of the module. Additionally, if set in a file, *option* is set to *value* when the file is Require-d.

6.2.5 Unset *option*.

This command resets *option* to its default value.

Variants:

1. Local Unset *option*.
This command resets *option* to its default value. The original state of *option* is restored when the current *section* ends.
2. Global Unset *option*.
This command resets *option* to its default value. The original state of *option* is *not* restored at the end of the module. Additionally, if unset in a file, *option* is reset to its default value when the file is Require-d.

6.2.6 Test *option*.

This command prints the current value of *option*.

6.2.7 Tables

The general commands for tables are Add **table** *value*, Remove **table** *value*, Test **table**, Test **table** for *value* and Print Table **table**.

6.2.8 Print Options.

This command lists all available flags, options and tables.

Variants:

1. Print Tables.
This is a synonymous of Print Options.

6.3 Requests to the environment

6.3.1 Check *term*.

This command displays the type of *term*. When called in proof mode, the term is checked in the local context of the current subgoal.

Variants:

1. selector: Check *term*.
specifies on which subgoal to perform typing (see Section 8.1).

6.3.2 Eval *convtactic* in *term*.

This command performs the specified reduction on *term*, and displays the resulting term with its type. The term to be reduced may depend on hypothesis introduced in the first subgoal (if a proof is in progress).

See also: Section 8.7.

6.3.3 Compute *term*.

This command performs a call-by-value evaluation of *term* by using the bytecode-based virtual machine. It is a shortcut for `Eval vm_compute in term`.

See also: Section 8.7.

6.3.4 Extraction *term*.

This command displays the extracted term from *term*. The extraction is processed according to the distinction between **Set** and **Prop**; that is to say, between logical and computational content (see Section 4.1.1). The extracted term is displayed in OCAML syntax, where global identifiers are still displayed as in COQ terms.

Variants:

1. Recursive Extraction *qualid*₁ ... *qualid*_{*n*}.
Recursively extracts all the material needed for the extraction of global *qualid*₁, ..., *qualid*_{*n*}.

See also: Chapter 23.

6.3.5 Print Assumptions *qualid*.

This commands display all the assumptions (axioms, parameters and variables) a theorem or definition depends on. Especially, it informs on the assumptions with respect to which the validity of a theorem relies.

Variants:

1. Print Opaque Dependencies *qualid*.
Displays the set of opaque constants *qualid* relies on in addition to the assumptions.
2. Print Transparent Dependencies *qualid*.
Displays the set of transparent constants *qualid* relies on in addition to the assumptions.
3. Print All Dependencies *qualid*.
Displays all assumptions and constants *qualid* relies on.

6.3.6 Search *qualid*.

This command displays the name and type of all objects (hypothesis of the current goal, theorems, axioms, etc) of the current context whose statement contains *qualid*. This command is useful to remind the user of the name of library lemmas.

Error messages:

1. The reference *qualid* was not found in the current environment
There is no constant in the environment named *qualid*.

Variants:

1. Search *string*.

If *string* is a valid identifier, this command displays the name and type of all objects (theorems, axioms, etc) of the current context whose name contains *string*. If *string* is a notation's string denoting some reference *qualid* (referred to by its main symbol as in "+" or by its notation's string as in "_ + _" or "_ ' U' _", see Section 12.1), the command works like Search *qualid*.

2. Search *string*%*key*.

The string *string* must be a notation or the main symbol of a notation which is then interpreted in the scope bound to the delimiting key *key* (see Section 12.2.2).

3. Search *term_pattern*.

This searches for all statements or types of definition that contains a subterm that matches the pattern *term_pattern* (holes of the pattern are either denoted by "_" or by "?ident" when non linear patterns are expected).

4. Search [-]*term_pattern-string* ... [-]*term_pattern-string*.

where *term_pattern-string* is a *term_pattern* or a *string*, or a *string* followed by a scope delimiting key %*key*.

This generalization of Search searches for all objects whose statement or type contains a subterm matching *term_pattern* (or *qualid* if *string* is the notation for a reference *qualid*) and whose name contains all *string* of the request that correspond to valid identifiers. If a *term_pattern* or a *string* is prefixed by "-", the search excludes the objects that mention that *term_pattern* or that *string*.

5. Search *term_pattern-string* ... *term_pattern-string* inside *module*₁ ... *module*_{*n*}.

This restricts the search to constructions defined in modules *module*₁ ... *module*_{*n*}.

6. Search *term_pattern-string* ... *term_pattern-string* outside *module*₁ ... *module*_{*n*}.

This restricts the search to constructions not defined in modules *module*₁ ... *module*_{*n*}.

7. selector: Search [-]*term_pattern-string* ... [-]*term_pattern-string*.

This specifies the goal on which to search hypothesis (see Section 8.1). By default the 1st goal is searched. This variant can be combined with other variants presented here.

Examples:

```
Coq < Require Import ZArith.

Coq < Search Z.mul Z.add "distr".
Z.mul_add_distr_l:
  forall n m p : Z, (n * (m + p))%Z = (n * m + n * p)%Z
Z.mul_add_distr_r:
  forall n m p : Z, ((n + m) * p)%Z = (n * p + m * p)%Z
fast_Zmult_plus_distr_l:
  forall (n m p : Z) (P : Z -> Prop),
  P (n * p + m * p)%Z -> P ((n + m) * p)%Z
```

```

Coq < Search "+"%Z "*"%Z "distr" -positive -Prop.
Z.mul_add_distr_l:
  forall n m p : Z, (n * (m + p))%Z = (n * m + n * p)%Z
Z.mul_add_distr_r:
  forall n m p : Z, ((n + m) * p)%Z = (n * p + m * p)%Z
Coq < Search (?x * _ + ?x * _)%Z outside OmegaLemmas.
Z.mul_add_distr_l:
  forall n m p : Z, (n * (m + p))%Z = (n * m + n * p)%Z

```

Warning: Up to COQ version 8.4, `Search` had the behavior of current `SearchHead` and the behavior of current `Search` was obtained with command `SearchAbout`. For compatibility, the deprecated name `SearchAbout` can still be used as a synonym of `Search`. For compatibility, the list of objects to search when using `SearchAbout` may also be enclosed by optional `[]` delimiters.

6.3.7 SearchHead *term* .

This command displays the name and type of all hypothesis of the current goal (if any) and theorems of the current context whose statement's conclusion has the form $(term \ t1 \ \dots \ tn)$. This command is useful to remind the user of the name of library lemmas.

```

Coq < SearchHead le.
le_n: forall n : nat, n <= n
le_0_n: forall n : nat, 0 <= n
le_S: forall n m : nat, n <= m -> n <= S m
le_pred: forall n m : nat, n <= m -> Nat.pred n <= Nat.pred m
le_n_S: forall n m : nat, n <= m -> S n <= S m
le_S_n: forall n m : nat, S n <= S m -> n <= m
Coq < SearchHead (@eq bool).
andb_true_intro:
  forall b1 b2 : bool, b1 = true /\ b2 = true -> (b1 && b2)%bool = true

```

Variants:

1. `SearchHead term` inside `module1 ... modulen` .

This restricts the search to constructions defined in modules `module1 ... modulen`.

2. `SearchHead term` outside `module1 ... modulen` .

This restricts the search to constructions not defined in modules `module1 ... modulen`.

Error messages:

- (a) Module/section `module` not found No module `module` has been required (see Section 6.5.1).

3. `selector: SearchHead term` .

This specifies the goal on which to search hypothesis (see Section 8.1). By default the 1st goal is searched. This variant can be combined with other variants presented here.

Warning: Up to COQ version 8.4, `SearchHead` was named `Search`.

6.3.8 SearchPattern *term*.

This command displays the name and type of all hypothesis of the current goal (if any) and theorems of the current context whose statement's conclusion or last hypothesis and conclusion matches the expression *term* where holes in the latter are denoted by “_”. It is a variant of `Search term_pattern` that does not look for subterms but searches for statements whose conclusion has exactly the expected form, or whose statement finishes by the given series of hypothesis/conclusion.

```
Coq < Require Import Arith.

Coq < SearchPattern (_ + _ = _ + _).
Nat.add_comm: forall n m : nat, n + m = m + n
plus_Snm_nSm: forall n m : nat, S n + m = n + S m
Nat.add_succ_comm: forall n m : nat, S n + m = n + S m
Nat.add_shuffle3: forall n m p : nat, n + (m + p) = m + (n + p)
plus_assoc_reverse: forall n m p : nat, n + m + p = n + (m + p)
Nat.add_assoc: forall n m p : nat, n + (m + p) = n + m + p
Nat.add_shuffle0: forall n m p : nat, n + m + p = n + p + m
f_equal2_plus:
  forall x1 y1 x2 y2 : nat, x1 = y1 -> x2 = y2 -> x1 + x2 = y1 + y2
Nat.add_shuffle2: forall n m p q : nat, n + m + (p + q) = n + q + (m + p)
Nat.add_shuffle1: forall n m p q : nat, n + m + (p + q) = n + p + (m + q)

Coq < SearchPattern (nat -> bool).
Nat.odd: nat -> bool
Init.Nat.odd: nat -> bool
Nat.even: nat -> bool
Init.Nat.even: nat -> bool
Init.Nat.testbit: nat -> nat -> bool
Nat.leb: nat -> nat -> bool
Nat.eqb: nat -> nat -> bool
Init.Nat.eqb: nat -> nat -> bool
Nat.ltb: nat -> nat -> bool
Nat.testbit: nat -> nat -> bool
Init.Nat.leb: nat -> nat -> bool
Init.Nat.ltb: nat -> nat -> bool
BinNat.N.testbit_nat: BinNums.N -> nat -> bool
BinPosDef.Pos.testbit_nat: BinNums.positive -> nat -> bool
BinPos.Pos.testbit_nat: BinNums.positive -> nat -> bool
BinNatDef.N.testbit_nat: BinNums.N -> nat -> bool

Coq < SearchPattern (forall l : list _, _ l l).
List.incl_refl: forall (A : Type) (l : list A), List.incl l l
List.lel_refl: forall (A : Type) (l : list A), List.lel l l
```

Patterns need not be linear: you can express that the same expression must occur in two places by using pattern variables “*?ident*”.

```
Coq < SearchPattern (?X1 + _ = _ + ?X1).
Nat.add_comm: forall n m : nat, n + m = m + n
```

Variants:

1. `SearchPattern term` inside *module*₁ ... *module*_n.

This restricts the search to constructions defined in modules *module*₁ ... *module*_n.

2. `SearchPattern term` outside `module1 ... modulen`.

This restricts the search to constructions not defined in modules `module1 ... modulen`.

3. `selector: SearchPattern term`.

This specifies the goal on which to search hypothesis (see Section 8.1). By default the 1st goal is searched. This variant can be combined with other variants presented here.

6.3.9 SearchRewrite term.

This command displays the name and type of all hypothesis of the current goal (if any) and theorems of the current context whose statement's conclusion is an equality of which one side matches the expression `term`. Holes in `term` are denoted by “_”.

```
Coq < Require Import Arith.
```

```
Coq < SearchRewrite (_ + _ + _).
Nat.add_shuffle0: forall n m p : nat, n + m + p = n + p + m
plus_assoc_reverse: forall n m p : nat, n + m + p = n + (m + p)
Nat.add_assoc: forall n m p : nat, n + (m + p) = n + m + p
Nat.add_shuffle1: forall n m p q : nat, n + m + (p + q) = n + p + (m + q)
Nat.add_shuffle2: forall n m p q : nat, n + m + (p + q) = n + q + (m + p)
Nat.add_carry_div2:
  forall (a b : nat) (c0 : bool),
    (a + b + Nat.b2n c0) / 2 =
      a / 2 + b / 2 +
      Nat.b2n
        (Nat.testbit a 0 && Nat.testbit b 0
         || c0 && (Nat.testbit a 0 || Nat.testbit b 0))
```

Variants:

1. `SearchRewrite term` inside `module1 ... modulen`.

This restricts the search to constructions defined in modules `module1 ... modulen`.

2. `SearchRewrite term` outside `module1 ... modulen`.

This restricts the search to constructions not defined in modules `module1 ... modulen`.

3. `selector: SearchRewrite term`.

This specifies the goal on which to search hypothesis (see Section 8.1). By default the 1st goal is searched. This variant can be combined with other variants presented here.

Nota Bene:

For the `Search`, `SearchHead`, `SearchPattern` and `SearchRewrite` queries, it is possible to globally filter the search results via the command `Add Search Blacklist "substring1"`. A lemma whose fully-qualified name contains any of the declared substrings will be removed from the search results. The default blacklisted substrings are `"_subproof"` `"Private_"`. The command `Remove Search Blacklist ...` allows expunging this blacklist.

6.3.10 `Locate qualid .`

This command displays the full name of objects whose name is a prefix of the qualified identifier *qualid*, and consequently the COQ module in which they are defined. It searches for objects from the different qualified name spaces of COQ: terms, modules, Ltac, etc.

```
Coq < Locate nat.
Inductive Coq.Init.Datatypes.nat

Coq < Locate Datatypes.O.
Constructor Coq.Init.Datatypes.O
  (shorter name to refer to it in current context is O)

Coq < Locate Init.Datatypes.O.
Constructor Coq.Init.Datatypes.O
  (shorter name to refer to it in current context is O)

Coq < Locate Coq.Init.Datatypes.O.
Constructor Coq.Init.Datatypes.O
  (shorter name to refer to it in current context is O)

Coq < Locate I.Dont.Exist.
No object of suffix I.Dont.Exist
```

Variants:

1. `Locate Term qualid .`
As `Locate` but restricted to terms.
2. `Locate Module qualid .` As `Locate` but restricted to modules.
3. `Locate Ltac qualid .`
As `Locate` but restricted to tactics.

See also: Section [12.1.10](#)

6.4 Loading files

COQ offers the possibility of loading different parts of a whole development stored in separate files. Their contents will be loaded as if they were entered from the keyboard. This means that the loaded files are ASCII files containing sequences of commands for COQ's toplevel. This kind of file is called a *script* for COQ. The standard (and default) extension of COQ's script files is `.v`.

6.4.1 `Load ident .`

This command loads the file named *ident* .v, searching successively in each of the directories specified in the *loadpath*. (see Section [2.6.3](#))

Variants:

1. `Load string .`
Loads the file denoted by the string *string*, where *string* is any complete filename. Then the `~` and `..` abbreviations are allowed as well as shell variables. If no extension is specified, COQ will use the default extension `.v`

2. Load Verbose *ident* ., Load Verbose *string*
 Display, while loading, the answers of COQ to each command (including tactics) contained in the loaded file **See also:** Section 6.9.1

Error messages:

1. Can't find file *ident* on loadpath

6.5 Compiled files

This section describes the commands used to load compiled files (see Chapter 14 for documentation on how to compile a file). A compiled file is a particular case of module called *library file*.

6.5.1 Require *qualid* .

This command looks in the loadpath for a file containing module *qualid* and adds the corresponding module to the environment of COQ. As library files have dependencies in other library files, the command `Require qualid` recursively requires all library files the module *qualid* depends on and adds the corresponding modules to the environment of COQ too. COQ assumes that the compiled files have been produced by a valid COQ compiler and their contents are then not replayed nor rechecked.

To locate the file in the file system, *qualid* is decomposed under the form *dirpath* . *ident* and the file *ident* .vo is searched in the physical directory of the file system that is mapped in COQ loadpath to the logical path *dirpath* (see Section 2.6.3). The mapping between physical directories and logical names at the time of requiring the file must be consistent with the mapping used to compile the file. If several files match, one of them is picked in an unspecified fashion.

Variants:

1. Require Import *qualid* .

This loads and declares the module *qualid* and its dependencies then imports the contents of *qualid* as described in Section 2.5.8.

It does not import the modules on which *qualid* depends unless these modules were itself required in module *qualid* using `Require Export`, as described below, or recursively required through a sequence of `Require Export`.

If the module required has already been loaded, `Require Import qualid` simply imports it, as `Import qualid` would.

2. Require Export *qualid* .

This command acts as `Require Import qualid`, but if a further module, say *A*, contains a command `Require Export B`, then the command `Require Import A` also imports the module *B*.

3. Require [*Import* | *Export*] *qualid*₁ ... *qualid*_{*n*} .

This loads the modules *qualid*₁, ..., *qualid*_{*n*} and their recursive dependencies. If *Import* or *Export* is given, it also imports *qualid*₁, ..., *qualid*_{*n*} and all the recursive dependencies that were marked or transitively marked as *Export*.

4. From *dirpath* Require *qualid*.

This command acts as `Require`, but picks any library whose absolute name is of the form *dirpath* . *dirpath* ' . *qualid* for some *dirpath* '. This is useful to ensure that the *qualid* library comes from a given package by making explicit its absolute root.

Error messages:

1. Cannot load *qualid*: no physical path bound to *dirpath*

2. Cannot find library *foo* in loadpath

The command did not find the file *foo.v*. Either *foo.v* exists but is not compiled or *foo.v* is in a directory which is not in your `LoadPath` (see Section 2.6.3).

3. Compiled library *ident.v* makes inconsistent assumptions over library *qualid*

The command tried to load library file *ident.v* that depends on some specific version of library *qualid* which is not the one already loaded in the current COQ session. Probably *ident.v* was not properly recompiled with the last version of the file containing module *qualid*.

4. Bad magic number

The file *ident.v* was found but either it is not a COQ compiled module, or it was compiled with an older and incompatible version of COQ.

5. The file *ident.v* contains library *dirpath* and not library *dirpath* '

The library file *dirpath* ' is indirectly required by the `Require` command but it is bound in the current loadpath to the file *ident.v* which was bound to a different library name *dirpath* at the time it was compiled.

6. Require is not allowed inside a module or a module type

This command is not allowed inside a module or a module type being defined. It is meant to describe a dependency between compilation units. Note however that the commands `Import` and `Export` alone can be used inside modules (see Section 2.5.8).

See also: Chapter 14

6.5.2 Print Libraries.

This command displays the list of library files loaded in the current COQ session. For each of these libraries, it also tells if it is imported.

6.5.3 Declare ML Module *string*₁ .. *string*_{*n*}.

This commands loads the OCAML compiled files *string*₁ ... *string*_{*n*} (dynamic link). It is mainly used to load tactics dynamically. The files are searched into the current OCAML loadpath (see the command `Add ML Path` in the Section 2.6.3). Loading of OCAML files is only possible under the bytecode version of `coqtop` (i.e. `coqtop.byte`, see chapter 14), or when COQ has been compiled with a version of OCAML that supports native `Dynlink` (≥ 3.11).

Variants:

1. Local Declare ML Module *string*₁ .. *string*_{*n*}.

This variant is not exported to the modules that import the module where they occur, even if outside a section.

Error messages:

1. File not found on loadpath : *string*
2. Loading of ML object file forbidden in a native Coq

6.5.4 Print ML Modules.

This print the name of all OCAML modules loaded with `Declare ML Module`. To know from where these module were loaded, the user should use the command `Locate File` (see Section 6.6.10)

6.6 Loadpath

Loadpaths are preferably managed using COQ command line options (see Section 2.6.3) but there remain vernacular commands to manage them for practical purposes. Such commands are only meant to be issued in the toplevel, and using them in source files is discouraged.

6.6.1 Pwd.

This command displays the current working directory.

6.6.2 Cd *string*.

This command changes the current directory according to *string* which can be any valid path.

Variants:

1. Cd.
- Is equivalent to `Pwd`.

6.6.3 Add LoadPath *string* as *dirpath*.

This command is equivalent to the command line option `-Q string dirpath`. It adds the physical directory *string* to the current COQ loadpath and maps it to the logical directory *dirpath*.

Variants:

1. Add LoadPath *string*.
- Performs as `Add LoadPath string as dirpath` but for the empty directory path.

6.6.4 Add Rec LoadPath *string* as *dirpath*.

This command is equivalent to the command line option `-R string dirpath`. It adds the physical directory *string* and all its subdirectories to the current COQ loadpath.

Variants:

1. Add Rec LoadPath *string*.
- Works as `Add Rec LoadPath string as dirpath` but for the empty logical directory path.

6.6.5 `Remove LoadPath string .`

This command removes the path *string* from the current COQ loadpath.

6.6.6 `Print LoadPath .`

This command displays the current COQ loadpath.

Variants:1. `Print LoadPath dirpath .`

Works as `Print LoadPath` but displays only the paths that extend the *dirpath* prefix.

6.6.7 `Add ML Path string .`

This command adds the path *string* to the current OCAML loadpath (see the command `Declare ML Module` in the Section 6.5).

6.6.8 `Add Rec ML Path string .`

This command adds the directory *string* and all its subdirectories to the current OCAML loadpath (see the command `Declare ML Module` in the Section 6.5).

6.6.9 `Print ML Path string .`

This command displays the current OCAML loadpath. This command makes sense only under the bytecode version of `coqtop`, i.e. `coqtop.byte` (see the command `Declare ML Module` in the section 6.5).

6.6.10 `Locate File string .`

This command displays the location of file *string* in the current loadpath. Typically, *string* is a `.cmo` or `.vo` or `.v` file.

6.6.11 `Locate Library dirpath .`

This command gives the status of the COQ module *dirpath*. It tells if the module is loaded and if not searches in the load path for a module of logical name *dirpath*.

6.7 Backtracking

The backtracking commands described in this section can only be used interactively, they cannot be part of a vernacular file loaded via `Load` or compiled by `coqc`.

6.7.1 `Reset ident .`

This command removes all the objects in the environment since *ident* was introduced, including *ident*. *ident* may be the name of a defined or declared object as well as the name of a section. One cannot reset over the name of a module or of an object inside a module.

Error messages:

1. `ident`: no such entry

Variants:

1. `Reset Initial`.
Goes back to the initial state, just after the start of the interactive session.

6.7.2 `Back`.

This command undoes all the effects of the last vernacular command. Commands read from a vernacular file via a `Load` are considered as a single command. Proof management commands are also handled by this command (see Chapter 7). For that, `Back` may have to undo more than one command in order to reach a state where the proof management information is available. For instance, when the last command is a `Qed`, the management information about the closed proof has been discarded. In this case, `Back` will then undo all the proof steps up to the statement of this proof.

Variants:

1. `Back n`
Undoes n vernacular commands. As for `Back`, some extra commands may be undone in order to reach an adequate state. For instance `Back n` will not re-enter a closed proof, but rather go just before that proof.

Error messages:

1. `Invalid backtrack`
The user wants to undo more commands than available in the history.

6.7.3 `BackTo num`.

This command brings back the system to the state labeled num , forgetting the effect of all commands executed after this state. The state label is an integer which grows after each successful command. It is displayed in the prompt when in `-emacs` mode. Just as `Back` (see above), the `BackTo` command now handles proof states. For that, it may have to undo some extra commands and end on a state $num' \leq num$ if necessary.

Variants:

1. `Backtrack num1 num2 num3`.
`Backtrack` is a *deprecated* form of `BackTo` which allows explicitly manipulating the proof environment. The three numbers num_1 , num_2 and num_3 represent the following:
 - num_3 : Number of `Abort` to perform, i.e. the number of currently opened nested proofs that must be canceled (see Chapter 7).
 - num_2 : *Proof state number* to unbury once aborts have been done. COQ will compute the number of `Undo` to perform (see Chapter 7).
 - num_1 : State label to reach, as for `BackTo`.

Error messages:

1. `Invalid backtrack`
The destination state label is unknown.

6.8 Quitting and debugging

6.8.1 `Quit.`

This command permits to quit COQ.

6.8.2 `Drop.`

This is used mostly as a debug facility by COQ's implementors and does not concern the casual user. This command permits to leave COQ temporarily and enter the OCAML toplevel. The OCAML command:

```
#use "include";;
```

add the right loadpaths and loads some toplevel printers for all abstract types of COQ- `section_path`, identifiers, terms, judgments, You can also use the file `base_include` instead, that loads only the pretty-printers for `section_paths` and identifiers. You can return back to COQ with the command:

```
go();;
```

Warnings:

1. It only works with the bytecode version of COQ (i.e. `coqtop` called with option `-byte`, see the contents of Section 14.1).
2. You must have compiled COQ from the source package and set the environment variable `COQTOP` to the root of your copy of the sources (see Section 14.3.2).

6.8.3 `Time command.`

This command executes the vernacular command *command* and display the time needed to execute it.

6.8.4 `Redirect "file" command.`

This command executes the vernacular command *command*, redirecting its output to "*file.out*".

6.8.5 `Timeout int command.`

This command executes the vernacular command *command*. If the command has not terminated after the time specified by the integer (time expressed in seconds), then it is interrupted and an error message is displayed.

6.8.6 `Set Default Timeout int.`

After using this command, all subsequent commands behave as if they were passed to a `Timeout` command. Commands already starting by a `Timeout` are unaffected.

6.8.7 `Unset Default Timeout.`

This command turns off the use of a default timeout.

6.8.8 `Test Default Timeout.`

This command displays whether some default timeout has been set or not.

6.9 Controlling display**6.9.1** `Set Silent.`

This command turns off the normal displaying.

6.9.2 `Unset Silent.`

This command turns the normal display on.

6.9.3 `Set Warnings "(w1, ..., wn)".`

This command configures the display of warnings. It is experimental, and expects, between quotes, a comma-separated list of warning names or categories. Adding `-` in front of a warning or category disables it, adding `+` makes it an error. It is possible to use the special categories `all` and `default`, the latter containing the warnings enabled by default. The flags are interpreted from left to right, so in case of an overlap, the flags on the right have higher priority, meaning that `A, -A` is equivalent to `-A`.

6.9.4 `Set Search Output Name Only.`

This command restricts the output of search commands to identifier names; turning it on causes invocations of `Search`, `SearchHead`, `SearchPattern`, `SearchRewrite` etc. to omit types from their output, printing only identifiers.

6.9.5 `Unset Search Output Name Only.`

This command turns type display in search results back on.

6.9.6 `Set Printing Width integer.`

This command sets which left-aligned part of the width of the screen is used for display.

6.9.7 `Unset Printing Width.`

This command resets the width of the screen used for display to its default value (which is 78 at the time of writing this documentation).

6.9.8 `Test Printing Width.`

This command displays the current screen width used for display.

6.9.9 `Set Printing Depth integer.`

This command sets the nesting depth of the formatter used for pretty-printing. Beyond this depth, display of subterms is replaced by dots.

6.9.10 `Unset Printing Depth.`

This command resets the nesting depth of the formatter used for pretty-printing to its default value (at the time of writing this documentation, the default value is 50).

6.9.11 `Test Printing Depth.`

This command displays the current nesting depth used for display.

6.9.12 `Unset Printing Compact Contexts.`

This command resets the displaying of goals contexts to non compact mode (default at the time of writing this documentation). Non compact means that consecutive variables of different types are printed on different lines.

6.9.13 `Set Printing Compact Contexts.`

This command sets the displaying of goals contexts to compact mode. The printer tries to reduce the vertical size of goals contexts by putting several variables (even if of different types) on the same line provided it does not exceed the printing width (See `Set Printing Width` above).

6.9.14 `Test Printing Compact Contexts.`

This command displays the current state of compaction of goal.

6.9.15 `Unset Printing Unfocused.`

This command resets the displaying of goals to focused goals only (default). Unfocused goals are created by focusing other goals with bullets(see [7.2.7](#)) or curly braces (see [7.2.6](#)).

6.9.16 `Set Printing Unfocused.`

This command enables the displaying of unfocused goals. The goals are displayed after the focused ones and are distinguished by a separator.

6.9.17 `Test Printing Unfocused.`

This command displays the current state of unfocused goals display.

6.9.18 `Set Printing Dependent Evars Line.`

This command enables the printing of the “(dependent evars: ...)” line when `-emacs` is passed.

6.9.19 `Unset Printing Dependent Evars Line.`

This command disables the printing of the “(dependent evars: ...)” line when `-emacs` is passed.

6.10 Controlling the reduction strategies and the conversion algorithm

COQ provides reduction strategies that the tactics can invoke and two different algorithms to check the convertibility of types. The first conversion algorithm lazily compares applicative terms while the other is a brute-force but efficient algorithm that first normalizes the terms before comparing them. The second algorithm is based on a bytecode representation of terms similar to the bytecode representation used in the ZINC virtual machine [98]. It is especially useful for intensive computation of algebraic values, such as numbers, and for reflection-based tactics. The commands to fine-tune the reduction strategies and the lazy conversion algorithm are described first.

6.10.1 Opaque *qualid*₁ ... *qualid*_n.

This command has an effect on unfoldable constants, i.e. on constants defined by `Definition` or `Let` (with an explicit body), or by a command assimilated to a definition such as `Fixpoint`, `Program Definition`, etc, or by a proof ended by `Defined`. The command tells not to unfold the constants *qualid*₁ ... *qualid*_n in tactics using δ -conversion (unfolding a constant is replacing it by its definition).

`Opaque` has also an effect on the conversion algorithm of COQ, telling it to delay the unfolding of a constant as much as possible when COQ has to check the conversion (see Section 4.3) of two distinct applied constants.

The scope of `Opaque` is limited to the current section, or current file, unless the variant `Global Opaque qualid1 ... qualidn` is used.

See also: sections 8.7, 8.16, 7.1

Error messages:

1. The reference *qualid* was not found in the current environment
There is no constant referred by *qualid* in the environment. Nevertheless, if you asked `Opaque foo bar` and if `bar` does not exist, `foo` is set opaque.

6.10.2 Transparent *qualid*₁ ... *qualid*_n.

This command is the converse of `Opaque` and it applies on unfoldable constants to restore their unfoldability after an `Opaque` command.

Note in particular that constants defined by a proof ended by `Qed` are not unfoldable and `Transparent` has no effect on them. This is to keep with the usual mathematical practice of *proof irrelevance*: what matters in a mathematical development is the sequence of lemma statements, not their actual proofs. This distinguishes lemmas from the usual defined constants, whose actual values are of course relevant in general.

The scope of `Transparent` is limited to the current section, or current file, unless the variant `Global Transparent qualid1 ... qualidn` is used.

Error messages:

1. The reference *qualid* was not found in the current environment
There is no constant referred by *qualid* in the environment.

See also: sections 8.7, 8.16, 7.1

6.10.3 Strategy level [*qualid*₁ ... *qualid*_{*n*}] .

This command generalizes the behavior of `Opaque` and `Transparent` commands. It is used to fine-tune the strategy for unfolding constants, both at the tactic level and at the kernel level. This command associates a level to *qualid*₁ ... *qualid*_{*n*}. Whenever two expressions with two distinct head constants are compared (for instance, this comparison can be triggered by a type cast), the one with lower level is expanded first. In case of a tie, the second one (appearing in the cast type) is expanded.

Levels can be one of the following (higher to lower):

opaque : level of opaque constants. They cannot be expanded by tactics (behaves like $+\infty$, see next item).

num : levels indexed by an integer. Level 0 corresponds to the default behavior, which corresponds to transparent constants. This level can also be referred to as **transparent**. Negative levels correspond to constants to be expanded before normal transparent constants, while positive levels correspond to constants to be expanded after normal transparent constants.

expand : level of constants that should be expanded first (behaves like $-\infty$)

These directives survive section and module closure, unless the command is prefixed by `Local`. In the latter case, the behavior regarding sections and modules is the same as for the `Transparent` and `Opaque` commands.

6.10.4 Print Strategy *qualid* .

This command prints the strategy currently associated to *qualid*. It fails if *qualid* is not an unfoldable reference, that is, neither a variable nor a constant.

Error messages:

1. The reference is not unfoldable.

Variants:

1. `Print Strategies`
Print all the currently non-transparent strategies.

6.10.5 Declare Reduction *ident* := *convtactic* .

This command allows giving a short name to a reduction expression, for instance `lazy beta delta [foo bar]`. This short name can then be used in `Eval ident in ...` or `eval` directives. This command accepts the `Local` modifier, for discarding this reduction name at the end of the file or module. For the moment the name cannot be qualified. In particular declaring the same name in several modules or in several functor applications will be refused if these declarations are not local. The name *ident* cannot be used directly as an Ltac tactic, but nothing prevent the user to also perform a `Ltac ident := convtactic`.

See also: sections [8.7](#)

6.11 Controlling the locality of commands

6.11.1 Local, Global

Some commands support a `Local` or `Global` prefix modifier to control the scope of their effect. There are four kinds of commands:

- Commands whose default is to extend their effect both outside the section and the module or library file they occur in.

For these commands, the `Local` modifier limits the effect of the command to the current section or module it occurs in.

As an example, the `Coercion` (see Section 2.8) and `Strategy` (see Section 6.10.3) commands belong to this category.

- Commands whose default behavior is to stop their effect at the end of the section they occur in but to extent their effect outside the module or library file they occur in.

For these commands, the `Local` modifier limits the effect of the command to the current module if the command does not occur in a section and the `Global` modifier extends the effect outside the current sections and current module if the command occurs in a section.

As an example, the `Implicit Arguments` (see Section 2.7), `Ltac` (see Chapter 9) or `Notation` (see Section 12.1) commands belong to this category.

Notice that a subclass of these commands do not support extension of their scope outside sections at all and the `Global` is not applicable to them.

- Commands whose default behavior is to stop their effect at the end of the section or module they occur in.

For these commands, the `Global` modifier extends their effect outside the sections and modules they occurs in.

The `Transparent` and `Opaque` (see Section 6.10) commands belong to this category.

- Commands whose default behavior is to extend their effect outside sections but not outside modules when they occur in a section and to extend their effect outside the module or library file they occur in when no section contains them.

For these commands, the `Local` modifier limits the effect to the current section or module while the `Global` modifier extends the effect outside the module even when the command occurs in a section.

The `Set` and `Unset` commands belong to this category.

Chapter 7

Proof handling

In COQ's proof editing mode all top-level commands documented in Chapter 6 remain available and the user has access to specialized commands dealing with proof development pragmas documented in this section. He can also use some other specialized commands called *tactics*. They are the very tools allowing the user to deal with logical reasoning. They are documented in Chapter 8.

When switching in editing proof mode, the prompt `Coq <` is changed into `ident <` where *ident* is the declared name of the theorem currently edited.

At each stage of a proof development, one has a list of goals to prove. Initially, the list consists only in the theorem itself. After having applied some tactics, the list of goals contains the subgoals generated by the tactics.

To each subgoal is associated a number of hypotheses called the *local_context* of the goal. Initially, the local context contains the local variables and hypotheses of the current section (see Section 1.3.1) and the local variables and hypotheses of the theorem statement. It is enriched by the use of certain tactics (see e.g. `intro` in Section 8.3.1).

When a proof is completed, the message `Proof completed` is displayed. One can then register this proof as a defined constant in the environment. Because there exists a correspondence between proofs and terms of λ -calculus, known as the *Curry-Howard isomorphism* [81, 6, 75, 85], COQ stores proofs as terms of CIC. Those terms are called *proof terms*.

Error message: When one attempts to use a proof editing command out of the proof editing mode, COQ raises the error message `:No focused proof`.

7.1 Switching on/off the proof editing mode

The proof editing mode is entered by asserting a statement, which typically is the assertion of a theorem:

```
Theorem ident [binders] : form .
```

The list of assertion commands is given in Section 1.3.5. The command `Goal` can also be used.

7.1.1 `Goal form .`

This is intended for quick assertion of statements, without knowing in advance which name to give to the assertion, typically for quick testing of the provability of a statement. If the proof of the statement is eventually completed and validated, the statement is then bound to the name `Unnamed_thm` (or a variant of this name not already used for another statement).

7.1.2 Qed.

This command is available in interactive editing proof mode when the proof is completed. Then `Qed` extracts a proof term from the proof script, switches back to COQ top-level and attaches the extracted proof term to the declared name of the original goal. This name is added to the environment as an `Opaque` constant.

Error messages:

1. Attempt to save an incomplete proof
2. Sometimes an error occurs when building the proof term, because tactics do not enforce completely the term construction constraints.

The user should also be aware of the fact that since the proof term is completely rechecked at this point, one may have to wait a while when the proof is large. In some exceptional cases one may even incur a memory overflow.

Variants:

1. `Defined.`
Defines the proved term as a transparent constant.
2. `Save ident.`
Forces the name of the original goal to be *ident*. This command (and the following ones) can only be used if the original goal has been opened using the `Goal` command.

7.1.3 Admitted.

This command is available in interactive editing proof mode to give up the current proof and declare the initial goal as an axiom.

7.1.4 Proof term.

This command applies in proof editing mode. It is equivalent to `exact term. Qed.` That is, you have to give the full proof in one gulp, as a proof term (see Section 8.2.1).

Variant: `Proof.`

Is a noop which is useful to delimit the sequence of tactic commands which start a proof, after a `Theorem` command. It is a good practice to use `Proof.` as an opening parenthesis, closed in the script with a closing `Qed.`

See also: `Proof with tactic.` in Section 8.9.7.

7.1.5 Proof using *ident*₁ ... *ident*_n.

This command applies in proof editing mode. It declares the set of section variables (see 1.3.1) used by the proof. At `Qed` time, the system will assert that the set of section variables actually used in the proof is a subset of the declared one.

The set of declared variables is closed under type dependency. For example if `T` is variable and `a` is a variable of type `T`, the commands `Proof using a` and `Proof using T a` are actually equivalent.

Variant: Proof using $ident_1 \dots ident_n$ with *tactic* . in Section 8.9.7.

Variant: Proof using All.

Use all section variables.

Variant: Proof using Type. **Variant:** Proof using.

Use only section variables occurring in the statement.

Variant: Proof using Type*.

The * operator computes the forward transitive closure. E.g. if the variable H has type $p < 5$ then H is in p^* since p occurs in the type of H. Type* is the forward transitive closure of the entire set of section variables occurring in the statement.

Variant: Proof using $-(ident_1 \dots ident_n)$.

Use all section variables except $ident_1 \dots ident_n$.

Variant: Proof using $collection_1 + collection_2$.

Variant: Proof using $collection_1 - collection_2$.

Variant: Proof using $collection - (ident_1 \dots ident_n)$.

Variant: Proof using $collection *$.

Use section variables being, respectively, in the set union, set difference, set complement, set forward transitive closure. See Section 7.1.5 to know how to form a named collection. The * operator binds stronger than + and -.

Proof using options

The following options modify the behavior of Proof using.

Variant: Set Default Proof Using "expression".

Use expression as the default Proof using value. E.g. Set Default Proof Using "a b" . will complete all Proof commands not followed by a using part with using a b.

Variant: Set Suggest Proof Using.

When Qed is performed, suggest a using annotation if the user did not provide one.

Name a set of section hypotheses for Proof using

The command Collection can be used to name a set of section hypotheses, with the purpose of making Proof using annotations more compact.

Variant: Collection Some := x y z.

Define the collection named "Some" containing x y and z

Variant: Collection Fewer := Some - x.

Define the collection named "Fewer" containing only x y

Variant: Collection Many := Fewer + Some. **Variant:** Collection Many := Fewer - Some.

Define the collection named "Many" containing the set union or set difference of "Fewer" and "Some".

Variant: Collection Many := Fewer - (x y).

Define the collection named "Many" containing the set difference of "Fewer" and the unnamed collection x y.

7.1.6 `Abort`.

This command cancels the current proof development, switching back to the previous proof development, or to the COQ toplevel if no other proof was edited.

Error messages:

1. No focused proof (No proof-editing in progress)

Variants:

1. `Abort ident`.
Aborts the editing of the proof named *ident*.
2. `Abort All`.
Aborts all current goals, switching back to the COQ toplevel.

7.1.7 `Existential num := term`.

This command instantiates an existential variable. *num* is an index in the list of uninstantiated existential variables displayed by `Show Existentials` (described in Section 7.3.1).

This command is intended to be used to instantiate existential variables when the proof is completed but some uninstantiated existential variables remain. To instantiate existential variables during proof edition, you should use the tactic `instantiate`.

See also: `instantiate (num := term)`. in Section 8.4.4. **See also:** `Grab Existential Variables`. below.

7.1.8 `Grab Existential Variables`.

This command can be run when a proof has no more goal to be solved but has remaining uninstantiated existential variables. It takes every uninstantiated existential variable and turns it into a goal.

7.2 Navigation in the proof tree

7.2.1 `Undo`.

This command cancels the effect of the last command. Thus, it backtracks one step.

Variants:

1. `Undo num`.
Repeats `Undo num` times.

7.2.2 `Restart`.

This command restores the proof editing process to the original goal.

Error messages:

1. No focused proof to restart

7.2.3 `Focus`.

This focuses the attention on the first subgoal to prove and the printing of the other subgoals is suspended until the focused subgoal is solved or unfocused. This is useful when there are many current subgoals which clutter your screen.

Variant:

1. `Focus num`.
This focuses the attention on the num^{th} subgoal to prove.

7.2.4 `Unfocus`.

This command restores to focus the goal that were suspended by the last `Focus` command.

7.2.5 `Unfocused`.

Succeeds in the proof is fully unfocused, fails if there are some goals out of focus.

7.2.6 `{` and `}`

The command `{` (without a terminating period) focuses on the first goal, much like `Focus` does, however, the subproof can only be unfocused when it has been fully solved (*i.e.* when there is no focused goal left). Unfocusing is then handled by `}` (again, without a terminating period). See also example in next section.

Note that when a focused goal is proved a message is displayed together with a suggestion about the right bullet or `}` to unfocus it or focus the next one.

Error messages:

1. This proof is focused, but cannot be unfocused this way You are trying to use `}` but the current subproof has not been fully solved.
2. see also error message about bullets below.

7.2.7 Bullets

Alternatively to `{` and `}`, proofs can be structured with bullets. The use of a bullet b for the first time focuses on the first goal g , the same bullet cannot be used again until the proof of g is completed, then it is mandatory to focus the next goal with b . The consequence is that g and all goals present when g was focused are focused with the same bullet b . See the example below.

Different bullets can be used to nest levels. The scope of bullet does not go beyond enclosing `{` and `}`, so bullets can be reused as further nesting levels provided they are delimited by these. Available bullets are `-`, `+`, `*`, `-`, `++`, `**`, `--`, `+++`, `***`, ... (without a terminating period).

Note again that when a focused goal is proved a message is displayed together with a suggestion about the right bullet or `}` to unfocus it or focus the next one.

Remark: In `PROOF GENERAL` (Emacs interface to `COQ`), you must use bullets with the priority ordering shown above to have a correct indentation. For example `-` must be the outer bullet and `**` the inner one in the example below.

The following example script illustrates all these features:

```

Coq < Goal (((True/\True)/\True)/\True)/\True.
Coq < Proof.
Coq <   split.
Coq <   - split.
Coq <     + split.
Coq <       ** { split.
Coq <         - trivial.
Coq <         - trivial.
Coq <       }
Coq <     ** trivial.
Coq <   + trivial.
Coq < - assert True.
Coq <   { trivial. }
Coq <   assumption.

```

Error messages:

1. Wrong bullet *bullet1* : Current bullet *bullet2* is not finished.
Before using bullet *bullet1* again, you should first finish proving the current focused goal. Note that *bullet1* and *bullet2* may be the same.
2. Wrong bullet *bullet1* : Bullet *bullet2* is mandatory here. You must put *bullet2* to focus next goal. No other bullet is allowed here.
3. No such goal. Focus next goal with bullet *bullet*.
You tried to applied a tactic but no goal where under focus. Using *bullet* is mandatory here.
4. No such goal. Try unfocusing with `"}`". You just finished a goal focused by `{`, you must unfocus it with `"}`".

7.2.8 Set Bullet Behavior.

The bullet behavior can be controlled by the following commands.

Set Bullet Behavior "None".

This makes bullets inactive.

Set Bullet Behavior "Strict Subproofs".

This makes bullets active (this is the default behavior).

7.3 Requesting information

7.3.1 Show.

This command displays the current goals.

Variants:

1. Show *num*.

Displays only the *num*-th subgoal.

Error messages:

- (a) No such goal
- (b) No focused proof

2. Show *ident*.

Displays the named goal *ident*. This is useful in particular to display a shelved goal but only works if the corresponding existential variable has been named by the user (see 2.11) as in the following example.

```
Coq < Goal exists n, n = 0.
Coq <   eexists ?[n].
Coq <   Show n.
subgoal n is:
```

```
=====
nat
```

3. Show Script.

Displays the whole list of tactics applied from the beginning of the current proof. This tactics script may contain some holes (subgoals not yet proved). They are printed under the form <Your Tactic Text here>.

4. Show Proof.

It displays the proof term generated by the tactics that have been applied. If the proof is not completed, this term contain holes, which correspond to the sub-terms which are still to be constructed. These holes appear as a question mark indexed by an integer, and applied to the list of variables in the context, since it may depend on them. The types obtained by abstracting away the context from the type of each hole-placer are also printed.

5. Show Conjectures.

It prints the list of the names of all the theorems that are currently being proved. As it is possible to start proving a previous lemma during the proof of a theorem, this list may contain several names.

6. Show Intro.

If the current goal begins by at least one product, this command prints the name of the first product, as it would be generated by an anonymous *intro*. The aim of this command is to ease the writing of more robust scripts. For example, with an appropriate *PROOF GENERAL* macro, it is possible to transform any anonymous *intro* into a qualified one such as *intro y13*. In the case of a non-product goal, it prints nothing.

7. Show Intros.

This command is similar to the previous one, it simulates the naming process of an `intros`.

8. Show Existentials.

It displays the set of all uninstantiated existential variables in the current proof tree, along with the type and the context of each variable.

9. Show Match *ident*.

This variant displays a template of the Gallina `match` construct with a branch for each constructor of the type *ident*.

Example:

```
Coq < Show Match nat.
match # with
| 0 =>
| S x =>
end
```

Error messages:

(a) Unknown inductive type

10. Show Universes.

It displays the set of all universe constraints and its normalized form at the current stage of the proof, useful for debugging universe inconsistencies.

7.3.2 Guarded.

Some tactics (e.g. `refine` 8.2.3) allow to build proofs using fixpoint or co-fixpoint constructions. Due to the incremental nature of interactive proof construction, the check of the termination (or guardedness) of the recursive calls in the fixpoint or cofixpoint constructions is postponed to the time of the completion of the proof.

The command `Guarded` allows checking if the guard condition for fixpoint and cofixpoint is violated at some time of the construction of the proof without having to wait the completion of the proof."

7.4 Controlling the effect of proof editing commands

7.4.1 Set Hyps Limit *num*.

This command sets the maximum number of hypotheses displayed in goals after the application of a tactic. All the hypotheses remains usable in the proof development.

7.4.2 Unset Hyps Limit.

This command goes back to the default mode which is to print all available hypotheses.

7.4.3 Set Automatic Introduction.

The option `Automatic Introduction` controls the way binders are handled in assertion commands such as `Theorem ident [binders] : form`. When the option is set, which is the default, *binders* are automatically put in the local context of the goal to prove.

The option can be unset by issuing `Unset Automatic Introduction`. When the option is unset, *binders* are discharged on the statement to be proved and a tactic such as `intro` (see Section 8.3.1) has to be used to move the assumptions to the local context.

7.5 Controlling memory usage

When experiencing high memory usage the following commands can be used to force Coq to optimize some of its internal data structures.

7.5.1 Optimize Proof.

This command forces Coq to shrink the data structure used to represent the ongoing proof.

7.5.2 Optimize Heap.

This command forces the OCaml runtime to perform a heap compaction. This is in general an expensive operation. See: <http://caml.inria.fr/pub/docs/manual-ocaml/libref/Gc.html#VALcompact>

Chapter 8

Tactics

A deduction rule is a link between some (unique) formula, that we call the *conclusion* and (several) formulas that we call the *premises*. A deduction rule can be read in two ways. The first one says: “*if I know this and this then I can deduce this*”. For instance, if I have a proof of A and a proof of B then I have a proof of $A \wedge B$. This is forward reasoning from premises to conclusion. The other way says: “*to prove this I have to prove this and this*”. For instance, to prove $A \wedge B$, I have to prove A and I have to prove B . This is backward reasoning from conclusion to premises. We say that the conclusion is the *goal* to prove and premises are the *subgoals*. The tactics implement *backward reasoning*. When applied to a goal, a tactic replaces this goal with the subgoals it generates. We say that a tactic reduces a goal to its subgoal(s).

Each (sub)goal is denoted with a number. The current goal is numbered 1. By default, a tactic is applied to the current goal, but one can address a particular goal in the list by writing $n:tactic$ which means “*apply tactic tactic to goal number n*”. We can show the list of subgoals by typing `Show` (see Section 7.3.1).

Since not every rule applies to a given statement, every tactic cannot be used to reduce any goal. In other words, before applying a tactic to a given goal, the system checks that some *preconditions* are satisfied. If it is not the case, the tactic raises an error message.

Tactics are built from atomic tactics and tactic expressions (which extends the folklore notion of tactical) to combine those atomic tactics. This chapter is devoted to atomic tactics. The tactic language will be described in Chapter 9.

8.1 Invocation of tactics

A tactic is applied as an ordinary command. It may be preceded by a goal selector (see Section 9.2). If no selector is specified, the default selector (see Section 8.1.1) is used.

$$\begin{array}{lcl} tactic_invocation & ::= & toplevel_selector : tactic . \\ & & | tactic . \end{array}$$

8.1.1 Set Default Goal Selector “*toplevel_selector*”.

After using this command, the default selector – used when no selector is specified when applying a tactic – is set to the chosen value. The initial value is 1, hence the tactics are, by default, applied to the first goal. Using `Set Default Goal Selector “all”` will make it so that tactics are, by default, applied to every goal simultaneously. Then, to apply a tactic `tac` to the first goal only, you can

write `1:tac`. Although more selectors are available, only “all” or a single natural number are valid default goal selectors.

8.1.2 Test Default Goal Selector.

This command displays the current default selector.

8.1.3 Bindings list

Tactics that take a term as argument may also support a bindings list, so as to instantiate some parameters of the term by name or position. The general form of a term equipped with a bindings list is *term* with *bindings_list* where *bindings_list* may be of two different forms:

- In a bindings list of the form $(ref_1 := term_1) \dots (ref_n := term_n)$, *ref* is either an *ident* or a *num*. The references are determined according to the type of *term*. If *ref_i* is an identifier, this identifier has to be bound in the type of *term* and the binding provides the tactic with an instance for the parameter of this name. If *ref_i* is some number *n*, this number denotes the *n*-th non dependent premise of the *term*, as determined by the type of *term*.

Error message: No such binder

- A bindings list can also be a simple list of terms *term₁ ... term_n*. In that case the references to which these terms correspond are determined by the tactic. In case of `induction`, `destruct`, `elim` and `case` (see Section 9) the terms have to provide instances for all the dependent products in the type of *term* while in the case of `apply`, or of `constructor` and its variants, only instances for the dependent products that are not bound in the conclusion of the type are required.

Error message: Not the right number of missing arguments

8.1.4 Occurrences sets and occurrences clauses

An occurrences clause is a modifier to some tactics that obeys the following syntax:

```

occurrence_clause ::= in goal_occurrences
goal_occurrences ::= [ident1 [at_occurrences] ,
                      ... ,
                      identm [at_occurrences]]
                      [| - [* [at_occurrences]]]
                      | * |- [* [at_occurrences]]
                      | *
at_occurrences   ::= at occurrences
occurrences      ::= [-] num1 ... numn
```

The role of an occurrence clause is to select a set of occurrences of a *term* in a goal. In the first case, the *ident_i [at num₁ⁱ ... num_{n_i}ⁱ]* parts indicate that occurrences have to be selected in the hypotheses named *ident_i*. If no numbers are given for hypothesis *ident_i*, then all the occurrences of *term* in the hypothesis are selected. If numbers are given, they refer to occurrences of *term* when the term is printed using option `Set Printing All` (see Section 2.9), counting from left to right. In particular, occurrences of *term* in implicit arguments (see Section 2.7) or coercions (see Section 2.8) are counted.

If a minus sign is given between `at` and the list of occurrences, it negates the condition so that the clause denotes all the occurrences except the ones explicitly mentioned after the minus sign.

As an exception to the left-to-right order, the occurrences in the `return` subexpression of a `match` are considered *before* the occurrences in the matched term.

In the second case, the `*` on the left of `| -` means that all occurrences of *term* are selected in every hypothesis.

In the first and second case, if `*` is mentioned on the right of `| -`, the occurrences of the conclusion of the goal have to be selected. If some numbers are given, then only the occurrences denoted by these numbers are selected. In no numbers are given, all occurrences of *term* in the goal are selected.

Finally, the last notation is an abbreviation for `* | - *`. Note also that `| -` is optional in the first case when no `*` is given.

Here are some tactics that understand occurrences clauses: `set`, `remember`, `induction`, `destruct`.

See also: Sections [8.3.7](#), [8.5.2](#), [2.9](#).

8.2 Applying theorems

8.2.1 `exact term`

This tactic applies to any goal. It gives directly the exact proof term of the goal. Let `T` be our goal, let `p` be a term of type `U` then `exact p` succeeds iff `T` and `U` are convertible (see Section [4.3](#)).

Error messages:

1. `Not an exact proof`

Variants:

1. `eexact term`

This tactic behaves like `exact` but is able to handle terms and goals with meta-variables.

8.2.2 `assumption`

This tactic looks in the local context for an hypothesis which type is equal to the goal. If it is the case, the subgoal is proved. Otherwise, it fails.

Error messages:

1. `No such assumption`

Variants:

1. `eassumption`

This tactic behaves like `assumption` but is able to handle goals with meta-variables.

8.2.3 `refine term`

This tactic applies to any goal. It behaves like `exact` with a big difference: the user can leave some holes (denoted by `_` or `(_: type)`) in the term. `refine` will generate as many subgoals as there are holes in the term. The type of holes must be either synthesized by the system or declared by an explicit cast like `(_: nat -> Prop)`. Any subgoal that occurs in other subgoals is automatically shelved, as if

calling `shelve_unifiable` (see Section 8.17.4). This low-level tactic can be useful to advanced users.

Example:

```
Coq < Inductive Option : Set :=
  | Fail : Option
  | Ok : bool -> Option.

Coq < Definition get : forall x:Option, x <> Fail -> bool.
1 subgoal

=====
forall x : Option, x <> Fail -> bool
Coq < refine
  (fun x:Option =>
    match x return x <> Fail -> bool with
    | Fail => _
    | Ok b => fun _ => b
  end).
1 subgoal

x : Option
=====
Fail <> Fail -> bool
Coq < intros; absurd (Fail = Fail); trivial.
No more subgoals.

Coq < Defined.
```

Error messages:

1. `invalid argument: the tactic refine does not know what to do with the term you gave.`
2. `Refine passed ill-formed term: the term you gave is not a valid proof (not easy to debug in general).` This message may also occur in higher-level tactics that call `refine` internally.
3. `Cannot infer a term for this placeholder: there is a hole in the term you gave which type cannot be inferred. Put a cast around it.`

Variants:

1. `simple refine term`
This tactic behaves like `refine`, but it does not shelve any subgoal. It does not perform any beta-reduction either.
2. `notypeclasses refine term`
This tactic behaves like `refine` except it performs typechecking without resolution of typeclasses.
3. `simple notypeclasses refine term`
This tactic behaves like `simple refine` except it performs typechecking without resolution of typeclasses.

8.2.4 `apply term`

This tactic applies to any goal. The argument *term* is a term well-formed in the local context. The tactic `apply` tries to match the current goal against the conclusion of the type of *term*. If it succeeds, then the tactic returns as many subgoals as the number of non-dependent premises of the type of *term*. If the conclusion of the type of *term* does not match the goal *and* the conclusion is an inductive type isomorphic to a tuple type, then each component of the tuple is recursively matched to the goal in the left-to-right order.

The tactic `apply` relies on first-order unification with dependent types unless the conclusion of the type of *term* is of the form $(P \ t_1 \ \dots \ t_n)$ with P to be instantiated. In the latter case, the behavior depends on the form of the goal. If the goal is of the form $(\text{fun } x \Rightarrow Q) \ u_1 \ \dots \ u_n$ and the t_i and u_i unifies, then P is taken to be $(\text{fun } x \Rightarrow Q)$. Otherwise, `apply` tries to define P by abstracting over $t_1 \dots t_n$ in the goal. See `pattern` in Section 8.7.7 to transform the goal so that it gets the form $(\text{fun } x \Rightarrow Q) \ u_1 \ \dots \ u_n$.

Error messages:

1. Unable to unify ... with ...

The `apply` tactic failed to match the conclusion of *term* and the current goal. You can help the `apply` tactic by transforming your goal with the `change` or `pattern` tactics (see sections 8.7.7, 8.6.5).

2. Unable to find an instance for the variables *ident* ... *ident*

This occurs when some instantiations of the premises of *term* are not deducible from the unification. This is the case, for instance, when you want to apply a transitivity property. In this case, you have to use one of the variants below:

Variants:

1. `apply term with term1 ... termn`

Provides `apply` with explicit instantiations for all dependent premises of the type of *term* that do not occur in the conclusion and consequently cannot be found by unification. Notice that *term₁ ... term_n* must be given according to the order of these dependent premises of the type of *term*.

Error message: Not the right number of missing arguments

2. `apply term with (ref1 := term1) ... (refn := termn)`

This also provides `apply` with values for instantiating premises. Here, variables are referred by names and non-dependent products by increasing numbers (see syntax in Section 8.1.3).

3. `apply term1 , ... , termn`

This is a shortcut for `apply term1 ; [.. | ... ; [.. | apply termn] ...]`, i.e. for the successive applications of *term_{i+1}* on the last subgoal generated by `apply termi`, starting from the application of *term₁*.

4. `eapply term`

The tactic `eapply` behaves like `apply` but it does not fail when no instantiations are deducible for some variables in the premises. Rather, it turns these variables into existential variables which are variables still to instantiate (see Section 2.11). The instantiation is intended to be found later in the proof.

5. `simple apply term`

This behaves like `apply` but it reasons modulo conversion only on subterms that contain no variables to instantiate. For instance, the following example does not succeed because it would require the conversion of `id ?foo` and `0`.

```
Coq < Definition id (x : nat) := x.
Coq < Hypothesis H : forall y, id y = y.
Coq < Goal 0 = 0.
Coq < Fail simple apply H.
The command has indeed failed with message:
Unable to unify "id ?M158 = ?M158" with "0 = 0".
1 subgoal

=====
0 = 0
```

Because it reasons modulo a limited amount of conversion, `simple apply` fails quicker than `apply` and it is then well-suited for uses in used-defined tactics that backtrack often. Moreover, it does not traverse tuples as `apply` does.

6. `[simple] apply term1 [with bindings_list1] , ... , termn [with bindings_listn]`
`[simple] eapply term1 [with bindings_list1] , ... , termn [with bindings_listn]`

This summarizes the different syntaxes for `apply` and `eapply`.

7. `lapply term`

This tactic applies to any goal, say `G`. The argument `term` has to be well-formed in the current context, its type being reducible to a non-dependent product `A -> B` with `B` possibly containing products. Then it generates two subgoals `B->G` and `A`. Applying `lapply H` (where `H` has type `A->B` and `B` does not start with a product) does the same as giving the sequence `cut B. 2:apply H.` where `cut` is described below.

Warning: When `term` contains more than one non dependent product the tactic `lapply` only takes into account the first product.

Example: Assume we have a transitive relation `R` on `nat`:

```
Coq < Variable R : nat -> nat -> Prop.
Coq < Hypothesis Rtrans : forall x y z:nat, R x y -> R y z -> R x z.
Coq < Variables n m p : nat.
Coq < Hypothesis Rnm : R n m.
Coq < Hypothesis Rmp : R m p.
```

Consider the goal `(R n p)` provable using the transitivity of `R`:

```
Coq < Goal R n p.
```

The direct application of `Rtrans` with `apply` fails because no value for `y` in `Rtrans` is found by `apply`:

```
Coq < Fail apply Rtrans.
The command has indeed failed with message:
Unable to find an instance for the variable y.
1 subgoal
```

```
=====
R n p
```

A solution is to apply `(Rtrans n m p)` or `(Rtrans n m)`.

```
Coq < apply (Rtrans n m p).
2 subgoals
```

```
=====
R n m
subgoal 2 is:
R m p
```

Note that `n` can be inferred from the goal, so the following would work too.

```
Coq < apply (Rtrans _ m).
```

More elegantly, apply `Rtrans` with `(y:=m)` allows only mentioning the unknown `m`:

```
Coq < apply Rtrans with (y := m).
```

Another solution is to mention the proof of `(R x y)` in `Rtrans`...

```
Coq < apply Rtrans with (1 := Rnm).
1 subgoal
```

```
=====
R m p
```

...or the proof of `(R y z)`.

```
Coq < apply Rtrans with (2 := Rmp).
1 subgoal
```

```
=====
R n m
```

On the opposite, one can use `eapply` which postpones the problem of finding `m`. Then one can apply the hypotheses `Rnm` and `Rmp`. This instantiates the existential variable and completes the proof.

```
Coq < eapply Rtrans.
2 focused subgoals
(shelved: 1)
```

```
=====
```

```

    R n ?y
subgoal 2 is:
    R ?y p

Coq < apply Rnm.
1 subgoal

=====
    R m p

Coq < apply Rmp.
No more subgoals.

```

Remark: When the conclusion of the type of the term to apply is an inductive type isomorphic to a tuple type and *apply* looks recursively whether a component of the tuple matches the goal, it excludes components whose statement would result in applying an universal lemma of the form *forall* A, ... -> A. Excluding this kind of lemma can be avoided by setting the following option:

```
Set Universal Lemma Under Conjunction
```

This option, which preserves compatibility with versions of COQ prior to 8.4 is also available for *apply term in ident* (see Section 8.2.5).

8.2.5 *apply term in ident*

This tactic applies to any goal. The argument *term* is a term well-formed in the local context and the argument *ident* is an hypothesis of the context. The tactic *apply term in ident* tries to match the conclusion of the type of *ident* against a non-dependent premise of the type of *term*, trying them from right to left. If it succeeds, the statement of hypothesis *ident* is replaced by the conclusion of the type of *term*. The tactic also returns as many subgoals as the number of other non-dependent premises in the type of *term* and of the non-dependent premises of the type of *ident*. If the conclusion of the type of *term* does not match the goal *and* the conclusion is an inductive type isomorphic to a tuple type, then the tuple is (recursively) decomposed and the first component of the tuple of which a non-dependent premise matches the conclusion of the type of *ident*. Tuples are decomposed in a width-first left-to-right order (for instance if the type of H1 is a $A \leftrightarrow B$ statement, and the type of H2 is A then *apply H1 in H2* transforms the type of H2 into B). The tactic *apply* relies on first-order pattern-matching with dependent types.

Error messages:

1. Statement without assumptions

This happens if the type of *term* has no non dependent premise.

2. Unable to apply

This happens if the conclusion of *ident* does not match any of the non dependent premises of the type of *term*.

Variants:

1. *apply term , ... , term in ident*

This applies each of *term* in sequence in *ident*.

2. `apply term with bindings_list , ... , term with bindings_list in ident`

This does the same but uses the bindings in each *bindings_list* to instantiate the parameters of the corresponding type of *term* (see syntax of bindings in Section 8.1.3).

3. `eapply term with bindings_list , ... , term with bindings_list in ident`

This works as `apply term with bindings_list , ... , term with bindings_list in ident` but turns unresolved bindings into existential variables, if any, instead of failing.

4. `apply term with bindings_list , ... , term with bindings_list in ident as intro_pattern`

This works as `apply term with bindings_list , ... , term with bindings_list in ident` then applies the *intro_pattern* to the hypothesis *ident*.

5. `eapply term with bindings_list , ... , term with bindings_list in ident as intro_pattern`

This works as `apply term with bindings_list , ... , term with bindings_list in ident as intro_pattern` but using `eapply`.

6. `simple apply term in ident`

This behaves like `apply term in ident` but it reasons modulo conversion only on subterms that contain no variables to instantiate. For instance, if `id := fun x:nat => x` and `H : forall y, id y = y -> True` and `H0 : 0 = 0` then `simple apply H in H0` does not succeed because it would require the conversion of `id ?1234` and `0` where `?1234` is a variable to instantiate. Tactic `simple apply term in ident` does not either traverse tuples as `apply term in ident` does.

7. `[simple] apply term [with bindings_list] , ... , term [with bindings_list] in ident [as intro_pattern]`
`[simple] eapply term [with bindings_list] , ... , term [with bindings_list] in ident [as intro_pattern]`

This summarizes the different syntactic variants of `apply term in ident` and `eapply term in ident`.

8.2.6 constructor *num*

This tactic applies to a goal such that its conclusion is an inductive type (say *I*). The argument *num* must be less or equal to the numbers of constructor(s) of *I*. Let *ci* be the *i*-th constructor of *I*, then `constructor i` is equivalent to `intros; apply ci`.

Error messages:

1. Not an inductive product
2. Not enough constructors

Variants:

1. `constructor`

This tries constructor 1 then constructor 2, ..., then constructor *n* where *n* is the number of constructors of the head of the goal.

2. `constructor num` with *bindings_list*

Let c_i be the i -th constructor of I , then `constructor i` with *bindings_list* is equivalent to `intros`; apply c_i with *bindings_list*.

Warning: the terms in the *bindings_list* are checked in the context where `constructor` is executed and not in the context where `apply` is executed (the introductions are not taken into account).

3. `split`

This applies only if I has a single constructor. It is then equivalent to `constructor 1`. It is typically used in the case of a conjunction $A \wedge B$.

Error message: Not an inductive goal with 1 constructor

4. `exists bindings_list`

This applies only if I has a single constructor. It is then equivalent to `intros`; `constructor 1` with *bindings_list*. It is typically used in the case of an existential quantification $\exists x, P(x)$.

Error message: Not an inductive goal with 1 constructor

5. `exists bindings_list , ... , bindings_list`

This iteratively applies `exists bindings_list`.

6. `left`
`right`

These tactics apply only if I has two constructors, for instance in the case of a disjunction $A \vee B$. Then, they are respectively equivalent to `constructor 1` and `constructor 2`.

Error message: Not an inductive goal with 2 constructors

7. `left with bindings_list`
`right with bindings_list`
`split with bindings_list`

As soon as the inductive type has the right number of constructors, these expressions are equivalent to calling `constructor i` with *bindings_list* for the appropriate i .

8. `econstructor`
`eexists`
`esplit`
`elleft`
`eright`

These tactics and their variants behave like `constructor`, `exists`, `split`, `left`, `right` and their variants but they introduce existential variables instead of failing when the instantiation of a variable cannot be found (cf `eapply` and Section 8.2.4).

8.3 Managing the local context

8.3.1 `intro`

This tactic applies to a goal that is either a product or starts with a let binder. If the goal is a product, the tactic implements the “Lam” rule given in Section 4.2¹. If the goal starts with a let binder, then the tactic implements a mix of the “Let” and “Conv”.

If the current goal is a dependent product $\forall x : T, U$ (resp `let $x := t$ in U`) then `intro` puts $x : T$ (resp $x := t$) in the local context. The new subgoal is U .

If the goal is a non-dependent product $T \rightarrow U$, then it puts in the local context either $Hn : T$ (if T is of type `Set` or `Prop`) or $Xn : T$ (if the type of T is `Type`). The optional index n is such that Hn or Xn is a fresh identifier. In both cases, the new subgoal is U .

If the goal is neither a product nor starting with a let definition, the tactic `intro` applies the tactic `hnf` until the tactic `intro` can be applied or the goal is not head-reducible.

Error messages:

1. No product even after head-reduction
2. `ident` is already used

Variants:

1. `intros`

This repeats `intro` until it meets the head-constant. It never reduces head-constants and it never fails.

2. `intro ident`

This applies `intro` but forces `ident` to be the name of the introduced hypothesis.

Error message: name `ident` is already used

Remark: If a name used by `intro` hides the base name of a global constant then the latter can still be referred to by a qualified name (see 2.6.2).

3. `intros ident1 ... identn`

This is equivalent to the composed tactic `intro ident1; ... ; intro identn`.

More generally, the `intros` tactic takes a pattern as argument in order to introduce names for components of an inductive definition or to clear introduced hypotheses. This is explained in 8.3.2.

4. `intros until ident`

This repeats `intro` until it meets a premise of the goal having form `(ident : term)` and discharges the variable named `ident` of the current goal.

Error message: No such hypothesis in current goal

¹Actually, only the second subgoal will be generated since the other one can be automatically checked.

5. `intros until num`

This repeats `intro` until the *num*-th non-dependent product. For instance, on the subgoal `forall x y:nat, x=y -> y=x` the tactic `intros until 1` is equivalent to `intros x y H, as x=y -> y=x` is the first non-dependent product. And on the subgoal `forall x y z:nat, x=y -> y=x` the tactic `intros until 1` is equivalent to `intros x y z` as the product on *z* can be rewritten as a non-dependent product: `forall x y:nat, nat -> x=y -> y=x`

Error message: No such hypothesis in current goal

This happens when *num* is 0 or is greater than the number of non-dependent products of the goal.

6. `intro after ident` `intro before ident` `intro at top` `intro at bottom`

These tactics apply `intro` and move the freshly introduced hypothesis respectively after the hypothesis *ident*, before the hypothesis *ident*, at the top of the local context, or at the bottom of the local context. All hypotheses on which the new hypothesis depends are moved too so as to respect the order of dependencies between hypotheses. Note that `intro at bottom` is a synonym for `intro` with no argument.

Error message: No such hypothesis: *ident*

7. `intro ident1 after ident2` `intro ident1 before ident2` `intro ident1 at top` `intro ident1 at bottom`

These tactics behave as previously but naming the introduced hypothesis *ident₁*. It is equivalent to `intro ident1` followed by the appropriate call to `move` (see Section 8.3.5).

8.3.2 `intros intro_pattern_list`

This extension of the tactic `intros` allows to apply tactics on the fly on the variables or hypotheses which have been introduced. An *introduction pattern list* `intro_pattern_list` is a list of introduction patterns possibly containing the filling introduction patterns `*` and `**`. An *introduction pattern* is either:

- a *naming introduction pattern*, i.e. either one of:
 - the pattern ?
 - the pattern ?*ident*
 - an identifier
- an *action introduction pattern* which itself classifies into:
 - a *disjunctive/conjunctive introduction pattern*, i.e. either one of:
 - * a disjunction of lists of patterns: `[intro_pattern_list1 | ... | intro_pattern_listn]`
 - * a conjunction of patterns: `(p1 , ... , pn)`

- * a list of patterns $(p_1 \ \& \ \dots \ \& \ p_n)$ for sequence of right-associative binary constructs
- an *equality introduction pattern*, i.e. either one of:
 - * a pattern for decomposing an equality: $[= \ p_1 \ \dots \ p_n]$
 - * the rewriting orientations: \rightarrow or \leftarrow
- the on-the-fly application of lemmas: $p\%term_1 \dots \%term_n$ where p itself is not a pattern for on-the-fly application of lemmas (note: syntax is in experimental stage)
- the wildcard: $_$

Assuming a goal of type $Q \rightarrow P$ (non-dependent product), or of type $\forall x : T, P$ (dependent product), the behavior of `intros p` is defined inductively over the structure of the introduction pattern p :

- introduction on $?$ performs the introduction, and lets COQ choose a fresh name for the variable;
- introduction on $?ident$ performs the introduction, and lets COQ choose a fresh name for the variable based on *ident*;
- introduction on *ident* behaves as described in Section 8.3.1;
- introduction over a disjunction of list of patterns $[intro_pattern_list_1 \mid \dots \mid intro_pattern_list_n]$ expects the product to be over an inductive type whose number of constructors is n (or more generally over a type of conclusion an inductive type built from n constructors, e.g. $C \rightarrow A \setminus B$ with $n = 2$ since $A \setminus B$ has 2 constructors): it destructs the introduced hypothesis as `destruct` (see Section 8.5.1) would and applies on each generated subgoal the corresponding tactic; `intros intro_pattern_listi`. The introduction patterns in $intro_pattern_list_i$ are expected to consume no more than the number of arguments of the i^{th} constructor. If it consumes less, then COQ completes the pattern so that all the arguments of the constructors of the inductive type are introduced (for instance, the list of patterns $[\mid]$ `H` applied on goal `forall x:nat, x=0 -> 0=x` behaves the same as the list of patterns $[\mid ?]$ `H`);
- introduction over a conjunction of patterns (p_1, \dots, p_n) expects the goal to be a product over an inductive type I with a single constructor that itself has at least n arguments: it performs a case analysis over the hypothesis, as `destruct` would, and applies the patterns $p_1 \dots p_n$ to the arguments of the constructor of I (observe that (p_1, \dots, p_n) is an alternative notation for $[p_1 \dots p_n]$);
- introduction via $(p_1 \ \& \ \dots \ \& \ p_n)$ is a shortcut for introduction via $(p_1, (\dots, (\dots, p_n) \dots))$; it expects the hypothesis to be a sequence of right-associative binary inductive constructors such as `conj` or `ex_intro`; for instance, an hypothesis with type $A \setminus (\text{exists } x, B \setminus C \setminus D)$ can be introduced via pattern $(a \ \& \ x \ \& \ b \ \& \ c \ \& \ d)$;
- if the product is over an equality type, then a pattern of the form $[= \ p_1 \ \dots \ p_n]$ applies either `injection` (see Section 8.5.7) or `discriminate` (see Section 8.5.6) instead of `destruct`; if `injection` is applicable, the patterns p_1, \dots, p_n are used on the hypotheses generated by `injection`; if the number of patterns is smaller than the number of hypotheses generated, the pattern $?$ is used to complete the list;

- introduction over \rightarrow (respectively \leftarrow) expects the hypothesis to be an equality and the right-hand-side (respectively the left-hand-side) is replaced by the left-hand-side (respectively the right-hand-side) in the conclusion of the goal; the hypothesis itself is erased; if the term to substitute is a variable, it is substituted also in the context of goal and the variable is removed too;
- introduction over a pattern $p\%term_1 \dots \%term_n$ first applies $term_1, \dots, term_n$ on the hypothesis to be introduced (as in `apply term_1, ..., term_n in`) prior to the application of the introduction pattern p ;
- introduction on the wildcard depends on whether the product is dependent or not: in the non-dependent case, it erases the corresponding hypothesis (i.e. it behaves as an `intro` followed by a `clear`, cf Section 8.3.3) while in the dependent case, it succeeds and erases the variable only if the wildcard is part of a more complex list of introduction patterns that also erases the hypotheses depending on this variable;
- introduction over $*$ introduces all forthcoming quantified variables appearing in a row; introduction over $**$ introduces all forthcoming quantified variables or hypotheses until the goal is not any more a quantification or an implication.

Example:

```
Coq < Goal forall A B C:Prop, A /\ B /\ C -> (A -> C) -> C.
1 subgoal

=====
forall A B C : Prop, A /\ B /\ C -> (A -> C) -> C
Coq < intros * [a | (_,c)] f.
2 subgoals

A, B, C : Prop
a : A
f : A -> C
=====
C
subgoal 2 is:
C
```

Remark: `intros $p_1 \dots p_n$` is not equivalent to `intros $p_1; \dots; p_n$` for the following reason: If one of the p_i is a wildcard pattern, he might succeed in the first case because the further hypotheses it depends in are eventually erased too while it might fail in the second case because of dependencies in hypotheses which are not yet introduced (and a fortiori not yet erased).

Remark: In `intros intro_pattern_list`, if the last introduction pattern is a disjunctive or conjunctive pattern `[intro_pattern_list1 | ... | intro_pattern_listn]`, the completion of `intro_pattern_listi` so that all the arguments of the i^{th} constructors of the corresponding inductive type are introduced can be controlled with the following option:

Set Bracketing Last Introduction Pattern

Force completion, if needed, when the last introduction pattern is a disjunctive or conjunctive pattern (this is the default).

Unset Bracketing Last Introduction Pattern

Deactivate completion when the last introduction pattern is a disjunctive or conjunctive pattern.

8.3.3 `clear ident`

This tactic erases the hypothesis named *ident* in the local context of the current goal. As a consequence, *ident* is no more displayed and no more usable in the proof development.

Error messages:

1. No such hypothesis
2. *ident* is used in the conclusion
3. *ident* is used in the hypothesis *ident'*

Variants:

1. `clear ident1 ... identn`

This is equivalent to `clear ident1. ... clear identn.`

2. `clearbody ident`

This tactic expects *ident* to be a local definition then clears its body. Otherwise said, this tactic turns a definition into an assumption.

Error message: *ident* is not a local definition

3. `clear - ident1 ... identn`

This tactic clears all the hypotheses except the ones depending in the hypotheses named *ident₁ ... ident_n* and in the goal.

4. `clear`

This tactic clears all the hypotheses except the ones the goal depends on.

5. `clear dependent ident`

This clears the hypothesis *ident* and all the hypotheses that depend on it.

8.3.4 `revert ident1 ... identn`

This applies to any goal with variables *ident₁ ... ident_n*. It moves the hypotheses (possibly defined) to the goal, if this respects dependencies. This tactic is the inverse of `intro`.

Error messages:

1. No such hypothesis
2. *ident* is used in the hypothesis *ident'*

Variants:

1. `revert dependent ident`

This moves to the goal the hypothesis *ident* and all the hypotheses that depend on it.

8.3.5 `move ident1 after ident2`

This moves the hypothesis named *ident₁* in the local context after the hypothesis named *ident₂*, where “after” is in reference to the direction of the move. The proof term is not changed.

If *ident₁* comes before *ident₂* in the order of dependencies, then all the hypotheses between *ident₁* and *ident₂* that (possibly indirectly) depend on *ident₁* are moved too, and all of them are thus moved after *ident₂* in the order of dependencies.

If *ident₁* comes after *ident₂* in the order of dependencies, then all the hypotheses between *ident₁* and *ident₂* that (possibly indirectly) occur in the type of *ident₁* are moved too, and all of them are thus moved before *ident₂* in the order of dependencies.

Variants:

1. `move ident1 before ident2`

This moves *ident₁* towards and just before the hypothesis named *ident₂*. As for `move ident1 after ident2`, dependencies over *ident₁* (when *ident₁* comes before *ident₂* in the order of dependencies) or in the type of *ident₁* (when *ident₁* comes after *ident₂* in the order of dependencies) are moved too.

2. `move ident at top`

This moves *ident* at the top of the local context (at the beginning of the context).

3. `move ident at bottom`

This moves *ident* at the bottom of the local context (at the end of the context).

Error messages:

1. No such hypothesis
2. Cannot move *ident₁* after *ident₂*: it occurs in the type of *ident₂*
3. Cannot move *ident₁* after *ident₂*: it depends on *ident₂*

Example:

```
Coq < Goal forall x : nat, x = 0 -> forall z y : nat, y = y -> 0 = x.
1 subgoal

=====
forall x : nat, x = 0 -> nat -> forall y : nat, y = y -> 0 = x

Coq < intros x H z y H0.
1 subgoal

x : nat
H : x = 0
z, y : nat
H0 : y = y
=====
0 = x

Coq < move x after H0.
```

```

1 subgoal

  z, y : nat
  H0 : y = y
  x : nat
  H : x = 0
  =====
  0 = x

Coq < Undo.
1 subgoal

  x : nat
  H : x = 0
  z, y : nat
  H0 : y = y
  =====
  0 = x

Coq < move x before H0.
1 subgoal

  z, y, x : nat
  H : x = 0
  H0 : y = y
  =====
  0 = x

Coq < Undo.
1 subgoal

  x : nat
  H : x = 0
  z, y : nat
  H0 : y = y
  =====
  0 = x

Coq < move H0 after H.
1 subgoal

  x, y : nat
  H0 : y = y
  H : x = 0
  z : nat
  =====
  0 = x

Coq < Undo.
1 subgoal

  x : nat
  H : x = 0
  z, y : nat
  H0 : y = y
  =====

```

```

    0 = x
Coq < move H0 before H.
1 subgoal

x : nat
H : x = 0
y : nat
H0 : y = y
z : nat
=====
0 = x

```

8.3.6 rename $ident_1$ into $ident_2$

This renames hypothesis $ident_1$ into $ident_2$ in the current context. The name of the hypothesis in the proof-term, however, is left unchanged.

Variants:

1. rename $ident_1$ into $ident_2$, ..., $ident_{2k-1}$ into $ident_{2k}$

This renames the variables $ident_1 \dots ident_{2k-1}$ into respectively $ident_2 \dots ident_{2k}$ in parallel. In particular, the target identifiers may contain identifiers that exist in the source context, as long as the latter are also renamed by the same tactic.

Error messages:

1. No such hypothesis
2. $ident_2$ is already used

8.3.7 set ($ident := term$)

This replaces $term$ by $ident$ in the conclusion of the current goal and adds the new definition $ident := term$ to the local context.

If $term$ has holes (i.e. subexpressions of the form “_”), the tactic first checks that all subterms matching the pattern are compatible before doing the replacement using the leftmost subterm matching the pattern.

Error messages:

1. The variable $ident$ is already defined

Variants:

1. set ($ident := term$) in $goal_occurrences$

This notation allows specifying which occurrences of $term$ have to be substituted in the context. The in $goal_occurrences$ clause is an occurrence clause whose syntax and behavior are described in Section 8.1.4.

2. set ($ident binder \dots binder := term$)

This is equivalent to set ($ident := fun binder \dots binder => term$).

3. `set term`

This behaves as `set (ident := term)` but `ident` is generated by COQ. This variant also supports an occurrence clause.

4. `set (ident0 binder ... binder := term) in goal_occurrences`
`set term in goal_occurrences`

These are the general forms that combine the previous possibilities.

5. `eset (ident0 binder ... binder := term) in goal_occurrences`
`eset term in goal_occurrences`

While the different variants of `set` expect that no existential variables are generated by the tactic, `eset` removes this constraint. In practice, this is relevant only when `eset` is used as a synonym of `epose`, i.e. when the term does not occur in the goal.

6. `remember term as ident`

This behaves as `set (ident := term) in *` and using a logical (Leibniz's) equality instead of a local definition.

7. `remember term as ident eqn:ident`

This behaves as `remember term as ident`, except that the name of the generated equality is also given.

8. `remember term as ident in goal_occurrences`

This is a more general form of `remember` that remembers the occurrences of `term` specified by an occurrences set.

9. `eremember term as ident`
`eremember term as ident in goal_occurrences`
`eremember term as ident eqn:ident`

While the different variants of `remember` expect that no existential variables are generated by the tactic, `eremember` removes this constraint.

10. `pose (ident := term)`

This adds the local definition `ident := term` to the current context without performing any replacement in the goal or in the hypotheses. It is equivalent to `set (ident := term) in |-`.

11. `pose (ident binder ... binder := term)`

This is equivalent to `pose (ident := fun binder ... binder => term)`.

12. `pose term`

This behaves as `pose (ident := term)` but `ident` is generated by COQ.

13. `epose (ident := term)`
`epose (ident binder ... binder := term)`
`epose term`

While the different variants of `pose` expect that no existential variables are generated by the tactic, `epose` removes this constraint.

8.3.8 `decompose [qualid1 ... qualidn] term`

This tactic recursively decomposes a complex proposition in order to obtain atomic ones.

Example:

```
Coq < Goal forall A B C:Prop, A /\ B /\ C \/ B /\ C \/ C /\ A -> C.
1 subgoal
```

```
=====
```

```
forall A B C : Prop, A /\ B /\ C \/ B /\ C \/ C /\ A -> C
```

```
Coq < intros A B C H; decompose [and or] H; assumption.
No more subgoals.
```

```
Coq < Qed.
```

`decompose` does not work on right-hand sides of implications or products.

Variants:

1. `decompose sum term`

This decomposes sum types (like `or`).

2. `decompose record term`

This decomposes record types (inductive types with one constructor, like `and` and `exists` and those defined with the `Record` macro, see Section 2.1).

8.4 Controlling the proof flow

8.4.1 `assert (ident : form)`

This tactic applies to any goal. `assert (H : U)` adds a new hypothesis of name `H` asserting `U` to the current goal and opens a new subgoal ². The subgoal `U` comes first in the list of subgoals remaining to prove.

Error messages:

1. Not a proposition or a type

Arises when the argument *form* is neither of type `Prop`, `Set` nor `Type`.

Variants:

1. `assert form`

This behaves as `assert (ident : form)` but *ident* is generated by `COQ`.

2. `assert form by tactic`

This tactic behaves like `assert` but applies *tactic* to solve the subgoals generated by `assert`.

Error message: Proof is not complete

²This corresponds to the cut rule of sequent calculus.

3. `assert form as intro_pattern`

If `intro_pattern` is a naming introduction pattern (see Section 8.3.2), the hypothesis is named after this introduction pattern (in particular, if `intro_pattern` is `ident`, the tactic behaves like `assert (ident : form)`).

If `intro_pattern` is an action introduction pattern, the tactic behaves like `assert form` followed by the action done by this introduction pattern.

4. `assert form as intro_pattern by tactic`

This combines the two previous variants of `assert`.

5. `assert (ident := term)`

This behaves as `assert (ident : type) by exact term` where `type` is the type of `term`. This is deprecated in favor of `pose proof`.

If the head of `term` is `ident`, the tactic behaves as `specialize term`.

Error message: Variable `ident` is already declared

6. `eassert form as intro_pattern by tactic`

`assert (ident := term)`

While the different variants of `assert` expect that no existential variables are generated by the tactic, `eassert` removes this constraint. This allows not to specify the asserted statement completely before starting to prove it.

7. `pose proof term [as intro_pattern]`

This tactic behaves like `assert T [as intro_pattern] by exact term` where `T` is the type of `term`.

In particular, `pose proof term as ident` behaves as `assert (ident := term)` and `pose proof term as intro_pattern` is the same as applying the `intro_pattern` to `term`.

8. `epose proof term [as intro_pattern]`

While `pose proof` expects that no existential variables are generated by the tactic, `epose proof` removes this constraint.

9. `enough (ident : form)`

This adds a new hypothesis of name `ident` asserting `form` to the goal the tactic `enough` is applied to. A new subgoal stating `form` is inserted after the initial goal rather than before it as `assert` would do.

10. `enough form`

This behaves like `enough (ident : form)` with the name `ident` of the hypothesis generated by Coq.

11. `enough form as intro_pattern`

This behaves like `enough form` using `intro_pattern` to name or destruct the new hypothesis.

12. `enough (ident : form) by tactic`
`enough form by tactic`
`enough form as intro_pattern by tactic`

This behaves as above but with *tactic* expected to solve the initial goal after the extra assumption *form* is added and possibly destructed. If the `as intro_pattern` clause generates more than one subgoal, *tactic* is applied to all of them.

13. `eenough (ident : form) by tactic`
`eenough form by tactic`
`eenough form as intro_pattern by tactic`

While the different variants of `enough` expect that no existential variables are generated by the *tactic*, `eenough` removes this constraint.

14. `cut form`

This tactic applies to any goal. It implements the non-dependent case of the “App” rule given in Section 4.2. (This is Modus Ponens inference rule.) `cut U` transforms the current goal *T* into the two following subgoals: $U \rightarrow T$ and *U*. The subgoal $U \rightarrow T$ comes first in the list of remaining subgoal to prove.

15. `specialize (ident term1 ... termn) [as intro_pattern]`
`specialize ident with bindings_list [as intro_pattern]`

The tactic `specialize` works on local hypothesis *ident*. The premises of this hypothesis (either universal quantifications or non-dependent implications) are instantiated by concrete terms coming either from arguments *term₁ ... term_n* or from a bindings list (see Section 8.1.3 for more about bindings lists). In the first form the application to *term₁ ... term_n* can be partial. The first form is equivalent to `assert (ident := ident term1 ... termn)`.

In the second form, instantiation elements can also be partial. In this case the uninstantiated arguments are inferred by unification if possible or left quantified in the hypothesis otherwise.

With the `as` clause, the local hypothesis *ident* is left unchanged and instead, the modified hypothesis is introduced as specified by the *intro_pattern*.

The name *ident* can also refer to a global lemma or hypothesis. In this case, for compatibility reasons, the behavior of `specialize` is close to that of `generalize`: the instantiated statement becomes an additional premise of the goal. The `as` clause is especially useful in this case to immediately introduce the instantiated statement as a local hypothesis.

Error messages:

- (a) *ident* is used in hypothesis *ident'*
- (b) *ident* is used in conclusion

8.4.2 generalize term

This tactic applies to any goal. It generalizes the conclusion with respect to some term.

Example:

```

Coq < Show.
1 subgoal

  x, y : nat
  =====
  0 <= x + y + y

Coq < generalize (x + y + y).
1 subgoal

  x, y : nat
  =====
  forall n : nat, 0 <= n

```

If the goal is G and t is a subterm of type T in the goal, then `generalize t` replaces the goal by `forall (x:T), G'` where G' is obtained from G by replacing all occurrences of t by x . The name of the variable (here n) is chosen based on T .

Variants:

1. `generalize $term_1$, ... , $term_n$`
 This is equivalent to `generalize $term_n$; ... ; generalize $term_1$` . Note that the sequence of $term_i$'s are processed from n to 1.
2. `generalize $term$ at num_1 ... num_i`
 This is equivalent to `generalize $term$` but it generalizes only over the specified occurrences of $term$ (counting from left to right on the expression printed using option `Set Printing All`).
3. `generalize $term$ as $ident$`
 This is equivalent to `generalize $term$` but it uses $ident$ to name the generalized hypothesis.
4. `generalize $term_1$ at num_{11} ... num_{1i_1} as $ident_1$, ... , $term_n$ at num_{n1} ... num_{ni_n} as $ident_2$`
 This is the most general form of `generalize` that combines the previous behaviors.
5. `generalize dependent $term$`
 This generalizes $term$ but also *all* hypotheses that depend on $term$. It clears the generalized hypotheses.

8.4.3 `evar ($ident$: $term$)`

The `evar` tactic creates a new local definition named $ident$ with type $term$ in the context. The body of this binding is a fresh existential variable.

8.4.4 `instantiate ($ident$:= $term$)`

The `instantiate` tactic refines (see Section 8.2.3) an existential variable $ident$ with the term $term$. It is equivalent to only `[$ident$]: refine $term$` (preferred alternative).

Remarks:

1. To be able to refer to an existential variable by name, the user must have given the name explicitly (see 2.11).
2. When you are referring to hypotheses which you did not name explicitly, be aware that Coq may make a different decision on how to name the variable in the current goal and in the context of the existential variable. This can lead to surprising behaviors.

Variants:

1. `instantiate (num := term)` This variant allows to refer to an existential variable which was not named by the user. The `num` argument is the position of the existential variable from right to left in the goal. Because this variant is not robust to slight changes in the goal, its use is strongly discouraged.
2. `instantiate (num := term) in ident`
3. `instantiate (num := term) in (Value of ident)`
4. `instantiate (num := term) in (Type of ident)`

These allow to refer respectively to existential variables occurring in a hypothesis or in the body or the type of a local definition.

5. `instantiate`

Without argument, the `instantiate` tactic tries to solve as many existential variables as possible, using information gathered from other tactics in the same tactical. This is automatically done after each complete tactic (i.e. after a dot in proof mode), but not, for example, between each tactic when they are sequenced by semicolons.

8.4.5 `admit`

The `admit` tactic allows temporarily skipping a subgoal so as to progress further in the rest of the proof. A proof containing admitted goals cannot be closed with `Qed` but only with `Admitted`.

Variants:

1. `give_up`
Synonym of `admit`.

8.4.6 `absurd term`

This tactic applies to any goal. The argument `term` is any proposition `P` of type `Prop`. This tactic applies `False` elimination, that is it deduces the current goal from `False`, and generates as subgoals $\sim P$ and `P`. It is very useful in proofs by cases, where some cases are impossible. In most cases, `P` or $\sim P$ is one of the hypotheses of the local context.

8.4.7 contradiction

This tactic applies to any goal. The `contradiction` tactic attempts to find in the current context (after all `intros`) an hypothesis that is equivalent to an empty inductive type (e.g. `False`), to the negation of a singleton inductive type (e.g. `True` or `x=x`), or two contradictory hypotheses.

Error messages:

1. No such assumption

Variants:

1. `contradiction ident`

The proof of `False` is searched in the hypothesis named *ident*.

8.4.8 contradict ident

This tactic allows manipulating negated hypothesis and goals. The name *ident* should correspond to a hypothesis. With `contradict H`, the current goal and context is transformed in the following way:

- $H : \neg A \vdash B$ becomes $\vdash A$
- $H : \neg A \vdash \neg B$ becomes $H : B \vdash A$
- $H : A \vdash B$ becomes $\vdash \neg A$
- $H : A \vdash \neg B$ becomes $H : B \vdash \neg A$

8.4.9 exfalso

This tactic implements the “ex falso quodlibet” logical principle: an elimination of `False` is performed on the current goal, and the user is then required to prove that `False` is indeed provable in the current context. This tactic is a macro for `elimtype False`.

8.5 Case analysis and induction

The tactics presented in this section implement induction or case analysis on inductive or co-inductive objects (see Section 4.5).

8.5.1 destruct term

This tactic applies to any goal. The argument *term* must be of inductive or co-inductive type and the tactic generates subgoals, one for each possible form of *term*, i.e. one for each constructor of the inductive or co-inductive type. Unlike `induction`, no induction hypothesis is generated by `destruct`.

There are special cases:

- If *term* is an identifier *ident* denoting a quantified variable of the conclusion of the goal, then `destruct ident` behaves as `intros until ident; destruct ident`. If *ident* is not anymore dependent in the goal after application of `destruct`, it is erased (to avoid erasure, use parentheses, as in `destruct (ident)`).

- If *term* is a *num*, then `destruct num` behaves as `intros` until *num* followed by `destruct` applied to the last introduced hypothesis. Remark: For destruction of a numeral, use syntax `destruct (num)` (not very interesting anyway).
- In case *term* is an hypothesis *ident* of the context, and *ident* is not anymore dependent in the goal after application of `destruct`, it is erased (to avoid erasure, use parentheses, as in `destruct (ident)`).
- The argument *term* can also be a pattern of which holes are denoted by “_”. In this case, the tactic checks that all subterms matching the pattern in the conclusion and the hypotheses are compatible and performs case analysis using this subterm.

Variants:

1. `destruct term1, ..., termn`
This is a shortcut for `destruct term1; ...; destruct termn`.
2. `destruct term as disj_conj_intro_pattern`
This behaves as `destruct term` but uses the names in *intro_pattern* to name the variables introduced in the context. The *intro_pattern* must have the form `[p11 ... p1n1 | ... | pm1 ... pmnm]` with *m* being the number of constructors of the type of *term*. Each variable introduced by `destruct` in the context of the *i*th goal gets its name from the list *p_{i1} ... p_{in_i}* in order. If there are not enough names, `destruct` invents names for the remaining variables to introduce. More generally, the *p_{ij}* can be any introduction pattern (see Section 8.3.2). This provides a concise notation for chaining destruction of an hypothesis.
3. `destruct term eqn:naming_intro_pattern`
This behaves as `destruct term` but adds an equation between *term* and the value that *term* takes in each of the possible cases. The name of the equation is specified by *naming_intro_pattern* (see Section 8.3.2), in particular ? can be used to let Coq generate a fresh name.
4. `destruct term with bindings_list`
This behaves like `destruct term` providing explicit instances for the dependent premises of the type of *term* (see syntax of bindings in Section 8.1.3).
5. `edestruct term`
This tactic behaves like `destruct term` except that it does not fail if the instance of a dependent premises of the type of *term* is not inferable. Instead, the unresolved instances are left as existential variables to be inferred later, in the same way as `eapply` does (see Section 8.2.4).
6. `destruct term1 using term2`
`destruct term1 using term2 with bindings_list`
These are synonyms of `induction term1 using term2` and `induction term1 using term2 with bindings_list`.
7. `destruct term in goal_occurrences`
This syntax is used for selecting which occurrences of *term* the case analysis has to be done on. The `in goal_occurrences` clause is an occurrence clause whose syntax and behavior is described in Section 8.1.4.

8. `destruct term1 with bindings_list1 as disj_conj_intro_pattern`
`eqn:naming_intro_pattern using term2 with bindings_list2 in goal_occurrences`
`edestruct term1 with bindings_list1 as disj_conj_intro_pattern`
`eqn:naming_intro_pattern using term2 with bindings_list2 in goal_occurrences`

These are the general forms of `destruct` and `edestruct`. They combine the effects of the `with`, `as`, `eqn:`, `using`, and `in` clauses.

9. `case term`

The tactic `case` is a more basic tactic to perform case analysis without recursion. It behaves as `elim term` but using a case-analysis elimination principle and not a recursive one.

10. `case term with bindings_list`

Analogous to `elim term with bindings_list` above.

11. `ecase term`
`ecase term with bindings_list`

In case the type of `term` has dependent premises, or dependent premises whose values are not inferable from the `with bindings_list` clause, `ecase` turns them into existential variables to be resolved later on.

12. `simple destruct ident`

This tactic behaves as `intros until ident; case ident` when `ident` is a quantified variable of the goal.

13. `simple destruct num`

This tactic behaves as `intros until num; case ident` where `ident` is the name given by `intros until num` to the `num`-th non-dependent premise of the goal.

14. `case_eq term`

The tactic `case_eq` is a variant of the `case` tactic that allow to perform case analysis on a term without completely forgetting its original form. This is done by generating equalities between the original form of the term and the outcomes of the case analysis.

8.5.2 induction term

This tactic applies to any goal. The argument `term` must be of inductive type and the tactic `induction` generates subgoals, one for each possible form of `term`, i.e. one for each constructor of the inductive type.

If the argument is dependent in either the conclusion or some hypotheses of the goal, the argument is replaced by the appropriate constructor form in each of the resulting subgoals and induction hypotheses are added to the local context using names whose prefix is `IH`.

There are particular cases:

- If `term` is an identifier `ident` denoting a quantified variable of the conclusion of the goal, then `induction ident` behaves as `intros until ident; induction ident`. If `ident` is not anymore dependent in the goal after application of `induction`, it is erased (to avoid erasure, use parentheses, as in `induction (ident)`).

- If *term* is a *num*, then `induction num` behaves as `intros` until *num* followed by `induction` applied to the last introduced hypothesis. Remark: For simple induction on a numeral, use syntax `induction (num)` (not very interesting anyway).
- In case *term* is an hypothesis *ident* of the context, and *ident* is not anymore dependent in the goal after application of `induction`, it is erased (to avoid erasure, use parentheses, as in `induction (ident)`).
- The argument *term* can also be a pattern of which holes are denoted by “_”. In this case, the tactic checks that all subterms matching the pattern in the conclusion and the hypotheses are compatible and performs induction using this subterm.

Example:

```
Coq < Lemma induction_test : forall n:nat, n = n -> n <= n.
1 subgoal

=====
forall n : nat, n = n -> n <= n
Coq < intros n H.
1 subgoal

n : nat
H : n = n
=====
n <= n
Coq < induction n.
2 subgoals

H : 0 = 0
=====
0 <= 0
subgoal 2 is:
S n <= S n
```

Error messages:

1. Not an inductive product
2. Unable to find an instance for the variables *ident* ... *ident*
Use in this case the variant `elim ... with ... below`.

Variants:

1. `induction term` as *disj_conj_intro_pattern*

This behaves as `induction term` but uses the names in *disj_conj_intro_pattern* to name the variables introduced in the context. The *disj_conj_intro_pattern* must typically be of the form `[p11 ... p1n1 | ... | pm1 ... pmnm]` with *m* being the number of constructors of the type of *term*. Each variable introduced by `induction` in the context of the *i*th goal gets its name

from the list $p_{i1} \dots p_{in_i}$ in order. If there are not enough names, `induction` invents names for the remaining variables to introduce. More generally, the p_{ij} can be any disjunctive/conjunctive introduction pattern (see Section 8.3.2). For instance, for an inductive type with one constructor, the pattern notation (p_1, \dots, p_n) can be used instead of $[p_1 \dots p_n]$.

2. `induction term with bindings_list`

This behaves like `induction term` providing explicit instances for the premises of the type of `term` (see the syntax of bindings in Section 8.1.3).

3. `einduction term`

This tactic behaves like `induction term` excepts that it does not fail if some dependent premise of the type of `term` is not inferable. Instead, the unresolved premises are posed as existential variables to be inferred later, in the same way as `eapply` does (see Section 8.2.4).

4. `induction term1 using term2`

This behaves as `induction term1` but using `term2` as induction scheme. It does not expect the conclusion of the type of `term1` to be inductive.

5. `induction term1 using term2 with bindings_list`

This behaves as `induction term1 using term2` but also providing instances for the premises of the type of `term2`.

6. `induction term1, ..., termn using qualid`

This syntax is used for the case `qualid` denotes an induction principle with complex predicates as the induction principles generated by `Function` or `Functional Scheme` may be.

7. `induction term in goal_occurrences`

This syntax is used for selecting which occurrences of `term` the induction has to be carried on. The `in goal_occurrences` clause is an occurrence clause whose syntax and behavior is described in Section 8.1.4. If variables or hypotheses not mentioning `term` in their type are listed in `goal_occurrences`, those are generalized as well in the statement to prove.

Example:

```
Coq < Lemma comm x y : x + y = y + x.
```

```
1 subgoal
```

```
  x, y : nat
```

```
  =====
```

```
  x + y = y + x
```

```
Coq < induction y in x |- *.
```

```
2 subgoals
```

```
  x : nat
```

```
  =====
```

```
  x + 0 = 0 + x
```

```
subgoal 2 is:
```

```
  x + S y = S y + x
```

```

Coq < Show 2.
subgoal 2 is:

  x, y : nat
  IHy : forall x : nat, x + y = y + x
  =====
  x + S y = S y + x

```

8. `induction term1 with bindings_list1 as disj_conj_intro_pattern using term2`
`with bindings_list2 in goal_occurrences`
`einduction term1 with bindings_list1 as disj_conj_intro_pattern using term2`
`with bindings_list2 in goal_occurrences`

These are the most general forms of induction and einduction. It combines the effects of the `with`, `as`, `using`, and `in` clauses.

9. `elim term`

This is a more basic induction tactic. Again, the type of the argument *term* must be an inductive type. Then, according to the type of the goal, the tactic `elim` chooses the appropriate destructor and applies it as the tactic `apply` would do. For instance, if the proof context contains `n:nat` and the current goal is `T` of type `Prop`, then `elim n` is equivalent to `apply nat_ind with (n:=n)`. The tactic `elim` does not modify the context of the goal, neither introduces the induction loading into the context of hypotheses.

More generally, `elim term` also works when the type of *term* is a statement with premises and whose conclusion is inductive. In that case the tactic performs induction on the conclusion of the type of *term* and leaves the non-dependent premises of the type as subgoals. In the case of dependent products, the tactic tries to find an instance for which the elimination lemma applies and fails otherwise.

10. `elim term with bindings_list`

Allows to give explicit instances to the premises of the type of *term* (see Section 8.1.3).

11. `eelim term`

In case the type of *term* has dependent premises, this turns them into existential variables to be resolved later on.

12. `elim term1 using term2`
`elim term1 using term2 with bindings_list`

Allows the user to give explicitly an elimination predicate *term₂* that is not the standard one for the underlying inductive type of *term₁*. The *bindings_list* clause allows instantiating premises of the type of *term₂*.

13. `elim term1 with bindings_list1 using term2 with bindings_list2`
`eelim term1 with bindings_list1 using term2 with bindings_list2`

These are the most general forms of `elim` and `eelim`. It combines the effects of the `using` clause and of the two uses of the `with` clause.

14. `elimtype form`

The argument *form* must be inductively defined. `elimtype I` is equivalent to `cut I. intro Hn; elim Hn; clear Hn`. Therefore the hypothesis *Hn* will not appear in the context(s) of the subgoal(s). Conversely, if *t* is a term of (inductive) type *I* that does not occur in the goal, then `elim t` is equivalent to `elimtype I; 2: exact t`.

15. simple induction *ident*

This tactic behaves as `intros until ident; elim ident` when *ident* is a quantified variable of the goal.

16. simple induction *num*

This tactic behaves as `intros until num; elim ident` where *ident* is the name given by `intros until num` to the *num*-th non-dependent premise of the goal.

8.5.3 double induction *ident*₁ *ident*₂

This tactic is deprecated and should be replaced by `induction ident1; induction ident2` (or `induction ident1; destruct ident2` depending on the exact needs).

Variant:1. double induction *num*₁ *num*₂

This tactic is deprecated and should be replaced by `induction num1; induction num3` where *num*₃ is the result of *num*₂-*num*₁.

8.5.4 dependent induction *ident*

The *experimental* tactic `dependent induction` performs induction-inversion on an instantiated inductive predicate. One needs to first require the `Coq.Program.Equality` module to use this tactic. The tactic is based on the `BasicElim` tactic by Conor McBride [107] and the work of Cristina Cornes around inversion [36]. From an instantiated inductive predicate and a goal, it generates an equivalent goal where the hypothesis has been generalized over its indexes which are then constrained by equalities to be the right instances. This permits to state lemmas without resorting to manually adding these equalities and still get enough information in the proofs.

Example:

```
Coq < Lemma le_minus : forall n:nat, n < 1 -> n = 0.
1 subgoal
```

```
=====
forall n : nat, n < 1 -> n = 0
```

```
Coq < intros n H ; induction H.
2 subgoals
```

```
  n : nat
  =====
  n = 0
subgoal 2 is:
  n = 0
```

Here we did not get any information on the indexes to help fulfill this proof. The problem is that, when we use the `induction` tactic, we lose information on the hypothesis instance, notably that the second argument is 1 here. Dependent induction solves this problem by adding the corresponding equality to the context.

```
Coq < Require Import Coq.Program.Equality.
Coq < Lemma le_minus : forall n:nat, n < 1 -> n = 0.
1 subgoal

=====
forall n : nat, n < 1 -> n = 0
Coq < intros n H ; dependent induction H.
2 subgoals

=====
0 = 0
subgoal 2 is:
n = 0
```

The subgoal is cleaned up as the tactic tries to automatically simplify the subgoals with respect to the generated equalities. In this enriched context, it becomes possible to solve this subgoal.

```
Coq < reflexivity.
1 subgoal
```

```
n : nat
H : S n <= 0
IHle : 0 = 1 -> n = 0
=====
n = 0
```

Now we are in a contradictory context and the proof can be solved.

```
Coq < inversion H.
No more subgoals.
```

This technique works with any inductive predicate. In fact, the `dependent induction` tactic is just a wrapper around the `induction` tactic. One can make its own variant by just writing a new tactic based on the definition found in `Coq.Program.Equality`.

Variants:

1. `dependent induction ident` generalizing `ident1 ... identn`

This performs dependent induction on the hypothesis `ident` but first generalizes the goal by the given variables so that they are universally quantified in the goal. This is generally what one wants to do with the variables that are inside some constructors in the induction hypothesis. The other ones need not be further generalized.

2. `dependent destruction ident`

This performs the generalization of the instance `ident` but uses `destruct` instead of `induction` on the generalized hypothesis. This gives results equivalent to `inversion` or `dependent inversion` if the hypothesis is dependent.

See also: 10.1 for a larger example of dependent induction and an explanation of the underlying technique.

8.5.5 functional induction (*qualid* $term_1 \dots term_n$)

The tactic `functional induction` performs case analysis and induction following the definition of a function. It makes use of a principle generated by `Function` (see Section 2.3) or `Functional Scheme` (see Section 13.2). Note that this tactic is only available after a `Require Import FunInd`.

```
Coq < Require Import FunInd.
[Loading ML file extraction_plugin.cmxs ... done]
[Loading ML file recdef_plugin.cmxs ... done]

Coq < Functional Scheme minus_ind := Induction for minus Sort Prop.
sub_equation is defined
minus_ind is defined

Coq < Check minus_ind.
minus_ind
  : forall P : nat -> nat -> nat -> Prop,
    (forall n m : nat, n = 0 -> P 0 m n) ->
    (forall n m k : nat, n = S k -> m = 0 -> P (S k) 0 n) ->
    (forall n m k : nat,
      n = S k ->
      forall l : nat, m = S l -> P k l (k - l) -> P (S k) (S l) (k - l)) ->
    forall n m : nat, P n m (n - m)

Coq < Lemma le_minus (n m:nat) : n - m <= n.
1 subgoal

  n, m : nat
  =====
  n - m <= n

Coq < functional induction (minus n m) using minus_ind; simpl; auto.
No more subgoals.

Coq < Qed.
```

Remark: (*qualid* $term_1 \dots term_n$) must be a correct full application of *qualid*. In particular, the rules for implicit arguments are the same as usual. For example use `@qualid` if you want to write implicit arguments explicitly.

Remark: Parentheses over *qualid*... $term_n$ are mandatory.

Remark: `functional induction (f x1 x2 x3)` is actually a wrapper for `induction x1, x2, x3, (f x1 x2 x3)` using *qualid* followed by a cleaning phase, where *qualid* is the induction principle registered for *f* (by the `Function` (see Section 2.3) or `Functional Scheme` (see Section 13.2) command) corresponding to the sort of the goal. Therefore `functional induction` may fail if the induction scheme *qualid* is not defined. See also Section 2.3 for the function terms accepted by `Function`.

Remark: There is a difference between obtaining an induction scheme for a function by using `Function` (see Section 2.3) and by using `Functional Scheme` after a normal definition using `Fixpoint` or `Definition`. See 2.3 for details.

See also: 2.3,13.2,13.2, 8.14.1

Error messages:

1. Cannot find induction information on *qualid*
2. Not the right number of induction arguments

Variants:

1. functional induction (*qualid term₁ ... term_n*) as *disj_conj_intro_pattern* using *term_{m+1}* with *bindings_list*

Similarly to `Induction` and `elim` (see Section 8.5.2), this allows giving explicitly the name of the introduced variables, the induction principle, and the values of dependent premises of the elimination scheme, including *predicates* for mutual induction when *qualid* is part of a mutually recursive definition.

8.5.6 discriminate *term*

This tactic proves any goal from an assumption stating that two structurally different terms of an inductive set are equal. For example, from $(S (S O)) = (S O)$ we can derive by absurdity any proposition.

The argument *term* is assumed to be a proof of a statement of conclusion $term_1 = term_2$ with *term₁* and *term₂* being elements of an inductive set. To build the proof, the tactic traverses the normal forms³ of *term₁* and *term₂* looking for a couple of subterms *u* and *w* (*u* subterm of the normal form of *term₁* and *w* subterm of the normal form of *term₂*), placed at the same positions and whose head symbols are two different constructors. If such a couple of subterms exists, then the proof of the current goal is completed, otherwise the tactic fails.

Remark: The syntax `discriminate ident` can be used to refer to a hypothesis quantified in the goal. In this case, the quantified hypothesis whose name is *ident* is first introduced in the local context using `intros until ident`.

Error messages:

1. No primitive equality found
2. Not a discriminable equality

Variants:

1. `discriminate num`

This does the same thing as `intros until num` followed by `discriminate ident` where *ident* is the identifier for the last introduced hypothesis.

2. `discriminate term with bindings_list`

This does the same thing as `discriminate term` but using the given bindings to instantiate parameters or hypotheses of *term*.

3. `ediscriminate num`
`ediscriminate term [with bindings_list]`

This works the same as `discriminate` but if the type of *term*, or the type of the hypothesis referred to by *num*, has uninstantiated parameters, these parameters are left as existential variables.

³Reminder: opaque constants will not be expanded by δ reductions.

4. `discriminate`

This behaves like `discriminate ident` if `ident` is the name of an hypothesis to which `discriminate` is applicable; if the current goal is of the form $term_1 <> term_2$, this behaves as `intro ident; discriminate ident`.

Error message: No discriminable equalities

8.5.7 `injection term`

The `injection` tactic exploits the property that constructors of inductive types are injective, i.e. that if c is a constructor of an inductive type and $c \vec{t}_1$ and $c \vec{t}_2$ are equal then \vec{t}_1 and \vec{t}_2 are equal too.

If `term` is a proof of a statement of conclusion $term_1 = term_2$, then `injection` applies the injectivity of constructors as deep as possible to derive the equality of all the subterms of $term_1$ and $term_2$ at positions where $term_1$ and $term_2$ start to differ. For example, from $(S\ p, S\ n) = (q, S\ (S\ m))$ we may derive $S\ p = q$ and $n = S\ m$. For this tactic to work, $term_1$ and $term_2$ should be typed with an inductive type and they should be neither convertible, nor having a different head constructor. If these conditions are satisfied, the tactic derives the equality of all the subterms of $term_1$ and $term_2$ at positions where they differ and adds them as antecedents to the conclusion of the current goal.

Example: Consider the following goal:

```
Coq < Inductive list : Set :=
      | nil : list
      | cons : nat -> list -> list.
```

```
Coq < Variable P : list -> Prop.
```

```
Coq < Show.
```

```
1 subgoal
```

```
l : list
n : nat
H : P nil
H0 : cons n l = cons 0 nil
=====
P l
```

```
Coq < injection H0.
```

```
1 subgoal
```

```
l : list
n : nat
H : P nil
H0 : cons n l = cons 0 nil
=====
l = nil -> n = 0 -> P l
```

Beware that `injection` yields an equality in a sigma type whenever the injected object has a dependent type P with its two instances in different types $(P\ t_1 \dots t_n)$ and $(P\ u_1 \dots u_n)$. If t_1 and u_1 are the same and have for type an inductive type for which a decidable equality has been declared using the command `Scheme Equality` (see 13.1), the use of a sigma type is avoided.

Remark: If some quantified hypothesis of the goal is named *ident*, then `injection ident` first introduces the hypothesis in the local context using `intros until ident`.

Error messages:

1. Not a projectable equality but a discriminable one
2. Nothing to do, it is an equality between convertible terms
3. Not a primitive equality
4. Nothing to inject

Variants:

1. `injection num`
This does the same thing as `intros until num` followed by `injection ident` where *ident* is the identifier for the last introduced hypothesis.
2. `injection term with bindings_list`
This does the same as `injection term` but using the given bindings to instantiate parameters or hypotheses of *term*.
3. `einjection num`
`einjection term [with bindings_list]`
This works the same as `injection` but if the type of *term*, or the type of the hypothesis referred to by *num*, has uninstantiated parameters, these parameters are left as existential variables.
4. `injection`
If the current goal is of the form $term_1 <> term_2$, this behaves as `intro ident; injection ident`.

Error message: goal does not satisfy the expected preconditions

5. `injection term [with bindings_list] as intro_pattern ... intro_pattern`
`injection num as intro_pattern ... intro_pattern`
`injection as intro_pattern ... intro_pattern`
`einjection term [with bindings_list] as intro_pattern ... intro_pattern`
`einjection num as intro_pattern ... intro_pattern`
`einjection as intro_pattern ... intro_pattern`

These variants apply `intros intro_pattern ... intro_pattern` after the call to `injection` or `einjection` so that all equalities generated are moved in the context of hypotheses. The number of *intro_pattern* must not exceed the number of equalities newly generated. If it is smaller, fresh names are automatically generated to adjust the list of *intro_pattern* to the number of new equalities. The original equality is erased if it corresponds to an hypothesis.

It is possible to ensure that `injection term` erases the original hypothesis and leaves the generated equalities in the context rather than putting them as antecedents of the current goal, as if giving `injection term` as (with an empty list of names). To obtain this behavior, the option `Set Structural Injection` must be activated. This option is off by default.

By default, `injection` only creates new equalities between terms whose type is in sort `Type` or `Set`, thus implementing a special behavior for objects that are proofs of a statement in `Prop`. This behavior can be turned off by setting the option `Set Keep Proof Equalities`.

8.5.8 inversion *ident*

Let the type of *ident* in the local context be $(I \vec{t})$, where I is a (co)inductive predicate. Then, *inversion* applied to *ident* derives for each possible constructor c_i of $(I \vec{t})$, **all** the necessary conditions that should hold for the instance $(I \vec{t})$ to be proved by c_i .

Remark: If *ident* does not denote a hypothesis in the local context but refers to a hypothesis quantified in the goal, then the latter is first introduced in the local context using `intros until ident`.

Remark: As inversion proofs may be large in size, we recommend the user to stock the lemmas whenever the same instance needs to be inverted several times. See Section 13.3.

Remark: Part of the behavior of the *inversion* tactic is to generate equalities between expressions that appeared in the hypothesis that is being processed. By default, no equalities are generated if they relate two proofs (i.e. equalities between terms whose type is in sort `Prop`). This behavior can be turned off by using the option `Set Keep Proof Equalities`.

Variants:

1. *inversion num*

This does the same thing as `intros until num then inversion ident` where *ident* is the identifier for the last introduced hypothesis.

2. *inversion_clear ident*

This behaves as *inversion* and then erases *ident* from the context.

3. *inversion ident as intro_pattern*

This generally behaves as *inversion* but using names in *intro_pattern* for naming hypotheses. The *intro_pattern* must have the form $[p_{11} \dots p_{1n_1} \mid \dots \mid p_{m1} \dots p_{mn_m}]$ with m being the number of constructors of the type of *ident*. Be careful that the list must be of length m even if *inversion* discards some cases (which is precisely one of its roles): for the discarded cases, just use an empty list (i.e. $n_i = 0$).

The arguments of the i^{th} constructor and the equalities that *inversion* introduces in the context of the goal corresponding to the i^{th} constructor, if it exists, get their names from the list $p_{i1} \dots p_{in_i}$ in order. If there are not enough names, *inversion* invents names for the remaining variables to introduce. In case an equation splits into several equations (because *inversion* applies *injection* on the equalities it generates), the corresponding name p_{ij} in the list must be replaced by a sublist of the form $[p_{ij1} \dots p_{ijq}]$ (or, equivalently, $(p_{ij1}, \dots, p_{ijq})$) where q is the number of subequalities obtained from splitting the original equation. Here is an example.

The *inversion ... as* variant of *inversion* generally behaves in a slightly more expectable way than *inversion* (no artificial duplication of some hypotheses referring to other hypotheses) To take benefit of these improvements, it is enough to use `inversion ... as []`, letting the names being finally chosen by COQ.

```
Coq < Inductive contains0 : list nat -> Prop :=
  | in_hd : forall l, contains0 (0 :: l)
  | in_tl : forall l b, contains0 l -> contains0 (b :: l).
contains0 is defined
contains0_ind is defined
```

```

Coq < Goal forall l:list nat, contains0 (1 :: l) -> contains0 l.
1 subgoal

=====
forall l : Datatypes.list nat, contains0 (1 :: l) -> contains0 l
Coq < intros l H; inversion H as [ | l' p Hl' [Heqp Heql'] ].
1 subgoal

l : Datatypes.list nat
H : contains0 (1 :: l)
l' : Datatypes.list nat
p : nat
Hl' : contains0 l
Heqp : p = 1
Heql' : l' = l
=====
contains0 l

```

4. `inversion num as intro_pattern`

This allows naming the hypotheses introduced by `inversion num` in the context.

5. `inversion_clear ident as intro_pattern`

This allows naming the hypotheses introduced by `inversion_clear` in the context. Notice that hypothesis names can be provided as if `inversion` were called, even though the `inversion_clear` will eventually erase the hypotheses.

6. `inversion ident in ident1 ... identn`

Let `ident1 ... identn`, be identifiers in the local context. This tactic behaves as generalizing `ident1 ... identn`, and then performing `inversion`.

7. `inversion ident as intro_pattern in ident1 ... identn`

This allows naming the hypotheses introduced in the context by `inversion ident` in `ident1 ... identn`.

8. `inversion_clear ident in ident1 ... identn`

Let `ident1 ... identn`, be identifiers in the local context. This tactic behaves as generalizing `ident1 ... identn`, and then performing `inversion_clear`.

9. `inversion_clear ident as intro_pattern in ident1 ... identn`

This allows naming the hypotheses introduced in the context by `inversion_clear ident` in `ident1 ... identn`.

10. `dependent inversion ident`

That must be used when `ident` appears in the current goal. It acts like `inversion` and then substitutes `ident` for the corresponding term in the goal.

11. `dependent inversion ident as intro_pattern`

This allows naming the hypotheses introduced in the context by dependent inversion `ident`.

12. `dependent inversion_clear ident`

Like `dependent inversion`, except that *ident* is cleared from the local context.

13. `dependent inversion_clear ident as intro_pattern`

This allows naming the hypotheses introduced in the context by `dependent inversion_clear ident`.

14. `dependent inversion ident with term`

This variant allows you to specify the generalization of the goal. It is useful when the system fails to generalize the goal automatically. If *ident* has type $(I \vec{t})$ and *I* has type $\forall(\vec{x} : \vec{T}), s$, then *term* must be of type $I : \forall(\vec{x} : \vec{T}), I \vec{x} \rightarrow s'$ where s' is the type of the goal.

15. `dependent inversion ident as intro_pattern with term`

This allows naming the hypotheses introduced in the context by `dependent inversion ident with term`.

16. `dependent inversion_clear ident with term`

Like `dependent inversion ... with` but clears *ident* from the local context.

17. `dependent inversion_clear ident as intro_pattern with term`

This allows naming the hypotheses introduced in the context by `dependent inversion_clear ident with term`.

18. `simple inversion ident`

It is a very primitive inversion tactic that derives all the necessary equalities but it does not simplify the constraints as `inversion` does.

19. `simple inversion ident as intro_pattern`

This allows naming the hypotheses introduced in the context by `simple inversion`.

20. `inversion ident using ident'`

Let *ident* have type $(I \vec{t})$ (*I* an inductive predicate) in the local context, and *ident'* be a (dependent) inversion lemma. Then, this tactic refines the current goal with the specified lemma.

21. `inversion ident using ident' in ident1... identn`

This tactic behaves as generalizing *ident₁... ident_n*, then doing `inversion ident using ident'`.

22. `inversion_sigma`

This tactic turns equalities of dependent pairs (e.g., `existT P x p = existT P y q`, frequently left over by `inversion` on a dependent type family) into pairs of equalities (e.g., a hypothesis `H : x = y` and a hypothesis of type `rew H in p = q`); these hypotheses can subsequently be simplified using `subst`, without ever invoking any kind of axiom asserting uniqueness of identity proofs. If you want to explicitly specify the hypothesis to be inverted, or name the generated hypotheses, you can invoke `induction H as [H1 H2] using eq_sigT_rect`. This tactic also works for `sig`, `sigT2`, and `sig2`, and there are similar `eq_sig*_rect` induction lemmas.

Example 1: Non-dependent inversion

Let us consider the relation `Le` over natural numbers and the following variables:

```
Coq < Inductive Le : nat -> nat -> Set :=
  | LeO : forall n:nat, Le 0 n
  | LeS : forall n m:nat, Le n m -> Le (S n) (S m).

Coq < Variable P : nat -> nat -> Prop.
Coq < Variable Q : forall n m:nat, Le n m -> Prop.
```

Let us consider the following goal:

```
Coq < Show.
1 subgoal

  n, m : nat
  H : Le (S n) m
  =====
  P n m
```

To prove the goal, we may need to reason by cases on `H` and to derive that `m` is necessarily of the form $(S\ m_0)$ for certain m_0 and that $(Le\ n\ m_0)$. Deriving these conditions corresponds to prove that the only possible constructor of $(Le\ (S\ n)\ m)$ is `LeS` and that we can invert the \rightarrow in the type of `LeS`. This inversion is possible because `Le` is the smallest set closed by the constructors `LeO` and `LeS`.

```
Coq < inversion_clear H.
1 subgoal

  n, m, m0 : nat
  H0 : Le n m0
  =====
  P n (S m0)
```

Note that `m` has been substituted in the goal for $(S\ m_0)$ and that the hypothesis $(Le\ n\ m_0)$ has been added to the context.

Sometimes it is interesting to have the equality $m = (S\ m_0)$ in the context to use it after. In that case we can use `inversion` that does not clear the equalities:

```
Coq < inversion H.
1 subgoal

  n, m : nat
  H : Le (S n) m
  n0, m0 : nat
  H1 : Le n m0
  H0 : n0 = n
  H2 : S m0 = m
  =====
  P n (S m0)
```

Example 2: Dependent inversion

Let us consider the following goal:

```
Coq < Show.
```

```
1 subgoal
```

```

n, m : nat
H : Le (S n) m
=====
Q (S n) m H

```

As H occurs in the goal, we may want to reason by cases on its structure and so, we would like inversion tactics to substitute H by the corresponding term in constructor form. Neither `Inversion` nor `Inversion_clear` make such a substitution. To have such a behavior we use the dependent inversion tactics:

```
Coq < dependent inversion_clear H.
```

```
1 subgoal
```

```

n, m, m0 : nat
l : Le n m0
=====
Q (S n) (S m0) (LeS n m0 l)

```

Note that H has been substituted by $(\text{LeS } n \ m0 \ l)$ and m by $(S \ m0)$.

Example 3: Using *inversion_sigma*

Let us consider the following inductive type of length-indexed lists, and a lemma about inverting equality of cons:

```
Coq < Require Coq.Logic.Eqdep_dec.
```

```
Coq < Inductive vec A : nat -> Type :=
```

```

| nil : vec A 0
| cons {n} (x : A) (xs : vec A n) : vec A (S n).

```

```

Coq < Lemma invert_cons : forall A n x xs y ys,
  @cons A n x xs = @cons A n y ys
  -> xs = ys.

```

```
Coq < Proof.
```

```
Coq < intros A n x xs y ys H.
```

```
1 subgoal
```

```

A : Type    n : nat    x : A    xs : vec A n    y : A    ys : vec A n
H : cons A x xs = cons A y ys
=====
xs = ys

```

After performing inversion, we are left with an equality of `existTs`:

```
Coq < inversion H.
```

```
1 subgoal
```

```

A : Type    n : nat    x : A    xs : vec A n    y : A    ys : vec A n
H : cons A x xs = cons A y ys
H1 : x = y

```

```

H2 : existT (fun n : nat => vec A n) n xs =
      existT (fun n : nat => vec A n) n ys
=====
xs = ys

```

We can turn this equality into a usable form with `inversion_sigma`:

```

Coq < inversion_sigma.
1 subgoal

A : Type    n : nat    x : A    xs : vec A n    y : A    ys : vec A n
H : cons A x xs = cons A y ys
H1 : x = y
H0 : n = n
H3 : eq_rect n (fun a : nat => vec A a) xs n H0 = ys
=====
xs = ys

```

To finish cleaning up the proof, we will need to use the fact that that all proofs of $n = n$ for n a nat are `eq_refl`:

```

Coq < let H := match goal with H : n = n |- _ => H end in
      pose proof (Eqdep_dec.UIP_refl_nat _ H); subst H.
1 subgoal

A : Type    n : nat    x : A    xs : vec A n    y : A    ys : vec A n
H : cons A x xs = cons A y ys
H1 : x = y
H3 : eq_rect n (fun a : nat => vec A a) xs n eq_refl = ys
=====
xs = ys

```

```

Coq < simpl in *.
1 subgoal

A : Type    n : nat    x : A    xs : vec A n    y : A    ys : vec A n
H : cons A x xs = cons A y ys
H1 : x = y
H3 : xs = ys
=====
xs = ys

```

Finally, we can finish the proof:

```

Coq < assumption.
No more subgoals.

Coq < Qed.
invert_cons is defined

```

8.5.9 fix ident num

This tactic is a primitive tactic to start a proof by induction. In general, it is easier to rely on higher-level induction tactics such as the ones described in Section 8.5.2.

In the syntax of the tactic, the identifier *ident* is the name given to the induction hypothesis. The natural number *num* tells on which premise of the current goal the induction acts, starting from 1, counting both dependent and non dependent products, but skipping local definitions. Especially, the current lemma must be composed of at least *num* products.

Like in a `fix` expression, the induction hypotheses have to be used on structurally smaller arguments. The verification that inductive proof arguments are correct is done only at the time of registering the lemma in the environment. To know if the use of induction hypotheses is correct at some time of the interactive development of a proof, use the command `Guarded` (see Section 7.3.2).

Variants:

1. `fix ident1 num with (ident2 binder2 ... binder2 [{ struct ident'2 }] : type2) ... (identn bindern ... bindern [{ struct ident'n }] : typen)`

This starts a proof by mutual induction. The statements to be simultaneously proved are respectively `forall binder2 ... binder2, type2, ..., forall bindern ... bindern, typen`. The identifiers *ident₁ ... ident_n* are the names of the induction hypotheses. The identifiers *ident'₂ ... ident'_n* are the respective names of the premises on which the induction is performed in the statements to be simultaneously proved (if not given, the system tries to guess itself what they are).

8.5.10 `cofix ident`

This tactic starts a proof by coinduction. The identifier *ident* is the name given to the coinduction hypothesis. Like in a `cofix` expression, the use of induction hypotheses have to be guarded by a constructor. The verification that the use of co-inductive hypotheses is correct is done only at the time of registering the lemma in the environment. To know if the use of coinduction hypotheses is correct at some time of the interactive development of a proof, use the command `Guarded` (see Section 7.3.2).

Variants:

1. `cofix ident1 with (ident2 binder2 ... binder2 : type2) ... (identn bindern ... bindern : typen)`

This starts a proof by mutual coinduction. The statements to be simultaneously proved are respectively `forall binder2 ... binder2, type2, ..., forall bindern ... bindern, typen`. The identifiers *ident₁ ... ident_n* are the names of the coinduction hypotheses.

8.6 Rewriting expressions

These tactics use the equality `eq: forall A: Type, A -> A -> Prop` defined in file `Logic.v` (see Section 3.1.2). The notation for `eq T t u` is simply `t=u` dropping the implicit type of *t* and *u*.

8.6.1 `rewrite term`

This tactic applies to any goal. The type of *term* must have the form

`forall (x1:A1) ... (xn:An) eq term1 term2.`

where `eq` is the Leibniz equality or a registered setoid equality.

Then `rewrite term` finds the first subterm matching *term₁* in the goal, resulting in instances *term'₁* and *term'₂* and then replaces every occurrence of *term'₁* by *term'₂*. Hence, some of the variables *x_i* are solved by unification, and some of the types *A₁, ..., A_n* become new subgoals.

Error messages:

1. The term provided does not end with an equation
2. Tactic generated a subgoal identical to the original goal
This happens if $term_1$ does not occur in the goal.

Variants:

1. `rewrite -> term`
Is equivalent to `rewrite term`
2. `rewrite <- term`
Uses the equality $term_1 = term_2$ from right to left
3. `rewrite term in clause`
Analogous to `rewrite term` but rewriting is done following *clause* (similarly to 8.7). For instance:
 - `rewrite H in H1` will rewrite *H* in the hypothesis *H1* instead of the current goal.
 - `rewrite H in H1 at 1, H2 at - 2 |- *` means `rewrite H; rewrite H in H1 at 1; rewrite H in H2 at - 2`. In particular a failure will happen if any of these three simpler tactics fails.
 - `rewrite H in * |-` will do `rewrite H in Hi` for all hypothesis $H_i <> H$. A success will happen as soon as at least one of these simpler tactics succeeds.
 - `rewrite H in *` is a combination of `rewrite H` and `rewrite H in * |-` that succeeds if at least one of these two tactics succeeds.

Orientation `->` or `<-` can be inserted before the term to rewrite.
4. `rewrite term at occurrences`
Rewrite only the given occurrences of $term'_1$. Occurrences are specified from left to right as for pattern (§8.7.7). The rewrite is always performed using setoid rewriting, even for Leibniz's equality, so one has to `Import Setoid` to use this variant.
5. `rewrite term by tactic`
Use *tactic* to completely solve the side-conditions arising from the rewrite.
6. `rewrite term1 , ... , termn`
Is equivalent to the *n* successive tactics `rewrite term1` up to `rewrite termn`, each one working on the first subgoal generated by the previous one. Orientation `->` or `<-` can be inserted before each term to rewrite. One unique *clause* can be added at the end after the keyword `in`; it will then affect all rewrite operations.
7. In all forms of `rewrite` described above, a term to rewrite can be immediately prefixed by one of the following modifiers:
 - `? : the tactic rewrite ?term` performs the rewrite of *term* as many times as possible (perhaps zero time). This form never fails.

- `n?` : works similarly, except that it will do at most n rewrites.
- `!` : works as `?`, except that at least one rewrite should succeed, otherwise the tactic fails.
- `n!` (or simply `n`) : precisely n rewrites of *term* will be done, leading to failure if these n rewrites are not possible.

8. `erewrite term`

This tactic works as `rewrite term` but turning unresolved bindings into existential variables, if any, instead of failing. It has the same variants as `rewrite` has.

8.6.2 `replace term1 with term2`

This tactic applies to any goal. It replaces all free occurrences of *term₁* in the current goal with *term₂* and generates the equality *term₂=term₁* as a subgoal. This equality is automatically solved if it occurs among the assumption, or if its symmetric form occurs. It is equivalent to `cut term2=term1; [intro Hn; rewrite <- Hn; clear Hn| assumption || symmetry; try assumption]`.

Error messages:

1. terms do not have convertible types

Variants:

1. `replace term1 with term2 by tactic`

This acts as `replace term1 with term2` but applies *tactic* to solve the generated subgoal *term₂=term₁*.

2. `replace term`

Replaces *term* with *term'* using the first assumption whose type has the form *term=term'* or *term'=term*.

3. `replace -> term`

Replaces *term* with *term'* using the first assumption whose type has the form *term=term'*

4. `replace <- term`

Replaces *term* with *term'* using the first assumption whose type has the form *term'=term*

5. `replace term1 with term2 in clause`
`replace term1 with term2 in clause by tactic`
`replace term in clause`
`replace -> term in clause`
`replace <- term in clause`

Acts as before but the replacements take place in *clause* (see Section 8.7) and not only in the conclusion of the goal. The *clause* argument must not contain any `type` or `nor value of`.

6. `cutrewrite <- (term1 = term2)`

This tactic is deprecated. It acts like `replace term2 with term1`, or, equivalently as enough *(term₁ = term₂)* as `<-`.

7. `cutrewrite -> (term1 = term2)`

This tactic is deprecated. It can be replaced by enough *(term₁ = term₂)* as `->`.

8.6.3 `subst ident`

This tactic applies to a goal that has *ident* in its context and (at least) one hypothesis, say *H*, of type $ident = t$ or $t = ident$ with *ident* not occurring in *t*. Then it replaces *ident* by *t* everywhere in the goal (in the hypotheses and in the conclusion) and clears *ident* and *H* from the context.

If *ident* is a local definition of the form $ident := t$, it is also unfolded and cleared.

Remark: When several hypotheses have the form $ident = t$ or $t = ident$, the first one is used.

Remark: If *H* is itself dependent in the goal, it is replaced by the proof of reflexivity of equality.

Variants:

1. `subst ident1 ... identn`

This is equivalent to `subst ident1; ...; subst identn`.

2. `subst`

This applies `subst` repeatedly from top to bottom to all identifiers of the context for which an equality of the form $ident = t$ or $t = ident$ or $ident := t$ exists, with *ident* not occurring in *t*.

Remark: The behavior of `subst` can be controlled using option `Set Regular Subst Tactic`. When this option is activated, `subst` also deals with the following corner cases:

- A context with ordered hypotheses $ident_1 = ident_2$ and $ident_1 = t$, or $t' = ident_1$ with *t'* not a variable, and no other hypotheses of the form $ident_2 = u$ or $u = ident_2$; without the option, a second call to `subst` would be necessary to replace $ident_2$ by *t* or *t'* respectively.
- The presence of a recursive equation which without the option would be a cause of failure of `subst`.
- A context with cyclic dependencies as with hypotheses $ident_1 = f ident_2$ and $ident_2 = g ident_1$ which without the option would be a cause of failure of `subst`.

Additionally, it prevents a local definition such as $ident := t$ to be unfolded which otherwise it would exceptionally unfold in configurations containing hypotheses of the form $ident = u$, or $u' = ident$ with *u'* not a variable.

Finally, it preserves the initial order of hypotheses, which without the option it may break.

The option is on by default.

8.6.4 `step1 term`

This tactic is for chaining rewriting steps. It assumes a goal of the form “*R term₁ term₂*” where *R* is a binary relation and relies on a database of lemmas of the form `forall x y z, R x y -> eq x z -> R z y` where *eq* is typically a setoid equality. The application of `step1 term` then replaces the goal by “*R term term₂*” and adds a new goal stating “*eq term term₁*”.

Lemmas are added to the database using the command

```
Declare Left Step term.
```

The tactic is especially useful for parametric setoids which are not accepted as regular setoids for `rewrite` and `setoid_replace` (see Chapter 27).

Variants:

1. `stepl term by tactic`

This applies `stepl term` then applies `tactic` to the second goal.

2. `stepr term`
`stepr term by tactic`

This behaves as `stepl` but on the right-hand-side of the binary relation. Lemmas are expected to be of the form “forall $x\ y\ z, R\ x\ y \rightarrow eq\ y\ z \rightarrow R\ x\ z$ ” and are registered using the command

```
Declare Right Step term.
```

8.6.5 change term

This tactic applies to any goal. It implements the rule “Conv” given in Section 4.4. `change U` replaces the current goal `T` with `U` providing that `U` is well-formed and that `T` and `U` are convertible.

Error messages:

1. Not convertible

Variants:

1. `change term1 with term2`

This replaces the occurrences of `term1` by `term2` in the current goal. The terms `term1` and `term2` must be convertible.

2. `change term1 at num1 ... numi with term2`

This replaces the occurrences numbered `num1 ... numi` of `term1` by `term2` in the current goal. The terms `term1` and `term2` must be convertible.

Error message: Too few occurrences

3. `change term in ident`

4. `change term1 with term2 in ident`

5. `change term1 at num1 ... numi with term2 in ident`

This applies the `change` tactic not to the goal but to the hypothesis `ident`.

See also: 8.7

8.7 Performing computations

This set of tactics implements different specialized usages of the tactic `change`.

All conversion tactics (including `change`) can be parameterized by the parts of the goal where the conversion can occur. This is done using *goal clauses* which consists in a list of hypotheses and, optionally, of a reference to the conclusion of the goal. For defined hypothesis it is possible to specify if the conversion should occur on the type part, the body part or both (default).

Goal clauses are written after a conversion tactic (tactics `set` 8.3.7, `rewrite` 8.6.1, `replace` 8.6.2 and `autorewrite` 8.8.4 also use goal clauses) and are introduced by the keyword `in`. If no goal clause is provided, the default is to perform the conversion only in the conclusion.

The syntax and description of the various goal clauses is the following:

```
in ident1 ... identn |- only in hypotheses ident1 ... identn
in ident1 ... identn |- * in hypotheses ident1 ... identn and in the conclusion
in * |- in every hypothesis
in * (equivalent to in * |- *) everywhere
in (type of ident1) (value of ident2) ... |- in type part of ident1, in the value part of
    ident2, etc.
```

For backward compatibility, the notation `in ident1...identn` performs the conversion in hypotheses `ident1...identn`.

8.7.1 `cbv flag1 ... flagn`, `lazy flag1 ... flagn`, and `compute`

These parameterized reduction tactics apply to any goal and perform the normalization of the goal according to the specified flags. In correspondence with the kinds of reduction considered in COQ namely β (reduction of functional application), δ (unfolding of transparent constants, see 6.10.2), ι (reduction of pattern-matching over a constructed term, and unfolding of `fix` and `cofix` expressions) and ζ (contraction of local definitions), the flags are either `beta`, `delta`, `match`, `fix`, `cofix`, `iota` or `zeta`. The `iota` flag is a shorthand for `match`, `fix` and `cofix`. The `delta` flag itself can be refined into `delta [qualid1...qualidk]` or `delta -[qualid1...qualidk]`, restricting in the first case the constants to unfold to the constants listed, and restricting in the second case the constant to unfold to all but the ones explicitly mentioned. Notice that the `delta` flag does not apply to variables bound by a `let-in` construction inside the term itself (use here the `zeta` flag). In any cases, opaque constants are not unfolded (see Section 6.10.1).

Normalization according to the flags is done by first evaluating the head of the expression into a *weak-head* normal form, i.e. until the evaluation is blocked by a variable (or an opaque constant, or an axiom), as e.g. in `x u1 ... un`, or `match x with ... end`, or `(fix f x {struct x} := ...) x`, or is a constructed form (a λ -expression, a constructor, a `cofixpoint`, an inductive type, a product type, a sort), or is a redex that the flags prevent to reduce. Once a weak-head normal form is obtained, subterms are recursively reduced using the same strategy.

Reduction to weak-head normal form can be done using two strategies: *lazy* (`lazy` tactic), or *call-by-value* (`cbv` tactic). The *lazy* strategy is a call-by-need strategy, with sharing of reductions: the arguments of a function call are weakly evaluated only when necessary, and if an argument is used several times then it is weakly computed only once. This reduction is efficient for reducing expressions with dead code. For instance, the proofs of a proposition `exists x. P(x)` reduce to a pair of a witness `t`, and a proof that `t` satisfies the predicate `P`. Most of the time, `t` may be computed without computing the proof of `P(t)`, thanks to the *lazy* strategy.

The *call-by-value* strategy is the one used in ML languages: the arguments of a function call are systematically weakly evaluated first. Despite the *lazy* strategy always performs fewer reductions than the *call-by-value* strategy, the latter is generally more efficient for evaluating purely computational expressions (i.e. with few dead code).

Variants:

1. `compute`
`cbv`

These are synonyms for `cbv beta delta iota zeta`.

2. `lazy`

This is a synonym for `lazy beta delta iota zeta`.

3. `compute [qualid1...qualidk]`
`cbv [qualid1...qualidk]`

These are synonyms of `cbv beta delta [qualid1...qualidk] iota zeta`.

4. `compute -[qualid1...qualidk]`
`cbv -[qualid1...qualidk]`

These are synonyms of `cbv beta delta -[qualid1...qualidk] iota zeta`.

5. `lazy [qualid1...qualidk]`
`lazy -[qualid1...qualidk]`

These are respectively synonyms of `lazy beta delta [qualid1...qualidk] iota zeta` and `lazy beta delta -[qualid1...qualidk] iota zeta`.

6. `vm_compute`

This tactic evaluates the goal using the optimized call-by-value evaluation bytecode-based virtual machine described in [77]. This algorithm is dramatically more efficient than the algorithm used for the `cbv` tactic, but it cannot be fine-tuned. It is specially interesting for full evaluation of algebraic objects. This includes the case of reflection-based tactics.

7. `native_compute`

This tactic evaluates the goal by compilation to OCAML as described in [16]. If COQ is running in native code, it can be typically two to five times faster than `vm_compute`. Note however that the compilation cost is higher, so it is worth using only for intensive computations.

Remark: The following option makes `cbv` (and its derivative `compute`) print information about the constants it encounters and the unfolding decisions it makes.

```
Set Debug Cbv
```

8.7.2 `red`

This tactic applies to a goal that has the form `forall (x:T1)...(xk:Tk), t` with t $\beta\iota\zeta$ -reducing to $c\ t_1 \dots t_n$ and c a constant. If c is transparent then it replaces c with its definition (say t) and then reduces $(t\ t_1 \dots t_n)$ according to $\beta\iota\zeta$ -reduction rules.

Error messages:

1. `Not reducible`

8.7.3 hnf

This tactic applies to any goal. It replaces the current goal with its head normal form according to the $\beta\delta\iota\zeta$ -reduction rules, i.e. it reduces the head of the goal until it becomes a product or an irreducible term. All inner $\beta\iota$ -redexes are also reduced.

Example: The term `forall n:nat, (plus (S n) (S n))` is not reduced by `hnf`.

Remark: The δ rule only applies to transparent constants (see Section 6.10.1 on transparency and opacity).

8.7.4 cbn and simpl

These tactics apply to any goal. They try to reduce a term to something still readable instead of fully normalizing it. They perform a sort of strong normalization with two key differences:

- They unfold a constant if and only if it leads to a ι -reduction, i.e. reducing a match or unfolding a fixpoint.
- While reducing a constant unfolding to (co)fixpoints, the tactics use the name of the constant the (co)fixpoint comes from instead of the (co)fixpoint definition in recursive calls.

The `cbn` tactic is claimed to be a more principled, faster and more predictable replacement for `simpl`.

The `cbn` tactic accepts the same flags as `cbv` and `lazy`. The behavior of both `simpl` and `cbn` can be tuned using the `Arguments` vernacular command as follows:

- A constant can be marked to be never unfolded by `cbn` or `simpl`:

```
Coq < Arguments minus n m : simpl never.
```

After that command an expression like `(minus (S x) y)` is left untouched by the tactics `cbn` and `simpl`.

- A constant can be marked to be unfolded only if applied to enough arguments. The number of arguments required can be specified using the `/` symbol in the arguments list of the `Arguments` vernacular command.

```
Coq < Definition fcomp A B C f (g : A -> B) (x : A) : C := f (g x).
Coq < Notation "f \o g" := (fcomp f g) (at level 50).
Coq < Arguments fcomp {A B C} f g x /.
```

After that command the expression `(f \o g)` is left untouched by `simpl` while `((f \o g) t)` is reduced to `(f (g t))`. The same mechanism can be used to make a constant volatile, i.e. always unfolded.

```
Coq < Definition volatile := fun x : nat => x.
Coq < Arguments volatile / x.
```

- A constant can be marked to be unfolded only if an entire set of arguments evaluates to a constructor. The `!` symbol can be used to mark such arguments.


```
Coq < Arguments minus !n !m.
```

After that command, the expression $(\text{minus } (S\ x)\ y)$ is left untouched by `simpl`, while $(\text{minus } (S\ x)\ (S\ y))$ is reduced to $(\text{minus } x\ y)$.

- A special heuristic to determine if a constant has to be unfolded can be activated with the following command:

```
Coq < Arguments minus n m : simpl nomatch.
```

The heuristic avoids to perform a simplification step that would expose a `match` construct in head position. For example the expression $(\text{minus } (S\ (S\ x))\ (S\ y))$ is simplified to $(\text{minus } (S\ x)\ y)$ even if an extra simplification is possible.

In detail, the tactic `simpl` first applies $\beta\iota$ -reduction. Then, it expands transparent constants and tries to reduce further using $\beta\iota$ -reduction. But, when no ι rule is applied after unfolding then δ -reductions are not applied. For instance trying to use `simpl` on $(\text{plus } n\ 0) = n$ changes nothing.

Notice that only transparent constants whose name can be reused in the recursive calls are possibly unfolded by `simpl`. For instance a constant defined by `plus' := plus` is possibly unfolded and reused in the recursive calls, but a constant such as `succ := plus (S 0)` is never unfolded. This is the main difference between `simpl` and `cbn`. The tactic `cbn` reduces whenever it will be able to reuse it or not: `succ t` is reduced to `S t`.

Variants:

1. `cbn [qualid1...qualidk]`
`cbn -[qualid1...qualidk]`

These are respectively synonyms of `cbn beta delta [qualid1...qualidk iota zeta]` and `cbn beta delta -[qualid1...qualidk iota zeta]` (see 8.7.1).

2. `simpl pattern`

This applies `simpl` only to the subterms matching *pattern* in the current goal.

3. `simpl pattern at num1 ... numi`

This applies `simpl` only to the num_1, \dots, num_i occurrences of the subterms matching *pattern* in the current goal.

Error message: Too few occurrences

4. `simpl qualid`
`simpl string`

This applies `simpl` only to the applicative subterms whose head occurrence is the unfoldable constant *qualid* (the constant can be referred to by its notation using *string* if such a notation exists).

5. `simpl qualid at num1 ... numi`
`simpl string at num1 ... numi`

This applies `simpl` only to the num_1, \dots, num_i applicative subterms whose head occurrence is *qualid* (or *string*).

Refolding Reduction

Deprecated since 8.7

This option (off by default) controls the use of the refolding strategy of `cbn` while doing reductions in unification, type inference and tactic applications. It can result in expensive unifications, as refolding currently uses a potentially exponential heuristic.

Set Debug RAKAM

This option makes `cbn` print various debugging information. `RAKAM` is the Refolding Algebraic Krivine Abstract Machine.

8.7.5 `unfold qualid`

This tactic applies to any goal. The argument *qualid* must denote a defined transparent constant or local definition (see Sections 1.3.2 and 6.10.2). The tactic `unfold` applies the δ rule to each occurrence of the constant to which *qualid* refers in the current goal and then replaces it with its $\beta\iota$ -normal form.

Error messages:

1. *qualid* does not denote an evaluable constant

Variants:

1. `unfold qualid in ident`

Replaces *qualid* in hypothesis *ident* with its definition and replaces the hypothesis with its $\beta\iota$ normal form.

2. `unfold qualid1, ..., qualidn`

Replaces *simultaneously* *qualid*₁, ..., *qualid*_n with their definitions and replaces the current goal with its $\beta\iota$ normal form.

3. `unfold qualid1 at num11, ..., numi1, ..., qualidn at num1n ... numjn`

The lists *num*₁¹, ..., *num*_i¹ and *num*₁ⁿ, ..., *num*_jⁿ specify the occurrences of *qualid*₁, ..., *qualid*_n to be unfolded. Occurrences are located from left to right.

Error message: bad occurrence number of *qualid*_i

Error message: *qualid*_i does not occur

4. `unfold string`

If *string* denotes the discriminating symbol of a notation (e.g. "+") or an expression defining a notation (e.g. "_ + _"), and this notation refers to an unfoldable constant, then the tactic unfolds it.

5. `unfold string%key`

This is variant of `unfold string` where *string* gets its interpretation from the scope bound to the delimiting key *key* instead of its default interpretation (see Section 12.2.2).

6. `unfold qualid_or_string1 at num11, ..., numi1, ..., qualid_or_stringn at num1n ... numjn`

This is the most general form, where *qualid_or_string* is either a *qualid* or a *string* referring to a notation.

8.7.6 `fold term`

This tactic applies to any goal. The term *term* is reduced using the `red` tactic. Every occurrence of the resulting term in the goal is then replaced by *term*.

Variants:

1. `fold term1 ... termn`

Equivalent to `fold term1; ...; fold termn`.

8.7.7 `pattern term`

This command applies to any goal. The argument *term* must be a free subterm of the current goal. The command `pattern` performs β -expansion (the inverse of β -reduction) of the current goal (say \mathbb{T}) by

1. replacing all occurrences of *term* in \mathbb{T} with a fresh variable
2. abstracting this variable
3. applying the abstracted goal to *term*

For instance, if the current goal T is expressible as $\phi(t)$ where the notation captures all the instances of t in $\phi(t)$, then `pattern t` transforms it into $(\text{fun } x:A \Rightarrow \phi(x)) \ t$. This command can be used, for instance, when the tactic `apply` fails on matching.

Variants:

1. `pattern term at num1 ... numn`

Only the occurrences $\text{num}_1 \dots \text{num}_n$ of *term* are considered for β -expansion. Occurrences are located from left to right.

2. `pattern term at - num1 ... numn`

All occurrences except the occurrences of indexes $\text{num}_1 \dots \text{num}_n$ of *term* are considered for β -expansion. Occurrences are located from left to right.

3. `pattern term1, ..., termm`

Starting from a goal $\phi(t_1 \dots t_m)$, the tactic `pattern t1, ..., tm` generates the equivalent goal $(\text{fun } (x_1:A_1) \dots (x_m:A_m) \Rightarrow \phi(x_1 \dots x_m)) \ t_1 \dots t_m$. If t_i occurs in one of the generated types A_j these occurrences will also be considered and possibly abstracted.

4. `pattern term1 at num11 ... numn11, ..., termm at num1m ... numnmm`

This behaves as above but processing only the occurrences $\text{num}_1^1, \dots, \text{num}_i^1$ of $\text{term}_1, \dots, \text{num}_1^m, \dots, \text{num}_j^m$ of term_m starting from term_m .

5. `pattern term1 [at [-] num11 ... numn11], ..., termm [at [-] num1m ... numnmm]`

This is the most general syntax that combines the different variants.

8.7.8 Conversion tactics applied to hypotheses

`conv_tactic in ident1 ... identn`

Applies the conversion tactic `conv_tactic` to the hypotheses `ident1, ..., identn`. The tactic `conv_tactic` is any of the conversion tactics listed in this section.

If `identi` is a local definition, then `identi` can be replaced by `(Type of identi)` to address not the body but the type of the local definition. Example: `unfold not in (Type of H1) (Type of H3)`.

Error messages:

1. No such hypothesis: `ident`.

8.8 Automation

8.8.1 `auto`

This tactic implements a Prolog-like resolution procedure to solve the current goal. It first tries to solve the goal using the `assumption` tactic, then it reduces the goal to an atomic one using `intros` and introduces the newly generated hypotheses as hints. Then it looks at the list of tactics associated to the head symbol of the goal and tries to apply one of them (starting from the tactics with lower cost). This process is recursively applied to the generated subgoals.

By default, `auto` only uses the hypotheses of the current goal and the hints of the database named `core`.

Variants:

1. `auto num`
Forces the search depth to be `num`. The maximal search depth is 5 by default.
2. `auto with ident1 ... identn`
Uses the hint databases `ident1 ... identn` in addition to the database `core`. See Section 8.9.1 for the list of pre-defined databases and the way to create or extend a database.
3. `auto with *`
Uses all existing hint databases. See Section 8.9.1
4. `auto using lemma1 , ... , lemman`
Uses `lemma1, ..., lemman` in addition to hints (can be combined with the `with ident` option). If `lemmai` is an inductive type, it is the collection of its constructors which is added as hints.
5. `info_auto`
Behaves like `auto` but shows the tactics it uses to solve the goal. This variant is very useful for getting a better understanding of automation, or to know what lemmas/assumptions were used.
6. `debug auto` Behaves like `auto` but shows the tactics it tries to solve the goal, including failing paths.
7. `[info_]auto [num] [using lemma1 , ... , lemman] [with ident1 ... identn]`
This is the most general form, combining the various options.

8. trivial

This tactic is a restriction of `auto` that is not recursive and tries only hints that cost 0. Typically it solves trivial equalities like $X = X$.

9. trivial with *ident*₁ ... *ident*_n

10. trivial with *

11. trivial using *lemma*₁ , ... , *lemma*_n

12. info_trivial

13. debug trivial

14. [info_]trivial [using *lemma*₁ , ... , *lemma*_n] [with *ident*₁ ... *ident*_n]

Remark: `auto` either solves completely the goal or else leaves it intact. `auto` and `trivial` never fail.

Remark: The following options enable printing of informative or debug information for the `auto` and `trivial` tactics:

```
Set Info Auto      Set Debug Auto      Set Info Trivial   Set Debug
Trivial
```

See also: Section 8.9.1

8.8.2 eauto

This tactic generalizes `auto`. While `auto` does not try resolution hints which would leave existential variables in the goal, `eauto` does try them (informally speaking, it uses simple `eapply` where `auto` uses simple `apply`). As a consequence, `eauto` can solve such a goal:

```
Coq < Hint Resolve ex_intro.
the hint: eapply ex_intro will only be used by eauto
Coq < Goal forall P:nat -> Prop, P 0 -> exists n, P n.
1 subgoal

=====
forall P : nat -> Prop, P 0 -> exists n : nat, P n

Coq < eauto.
No more subgoals.
```

Note that `ex_intro` should be declared as a hint.

Variants:

1. [info_]eauto [num] [using *lemma*₁ , ... , *lemma*_n] [with *ident*₁ ... *ident*_n]

The various options for `eauto` are the same as for `auto`.

Remark: `eauto` obeys the following options:

```
Set Info Eauto Set Debug Eauto
```

See also: Section 8.9.1

8.8.3 autounfold with $ident_1 \dots ident_n$

This tactic unfolds constants that were declared through a `Hint Unfold` in the given databases.

Variants:

1. `autounfold with $ident_1 \dots ident_n$ in clause`

Performs the unfolding in the given clause.

2. `autounfold with *`

Uses the unfold hints declared in all the hint databases.

8.8.4 autorewrite with $ident_1 \dots ident_n$

This tactic ⁴ carries out rewritings according the rewriting rule bases $ident_1 \dots ident_n$.

Each rewriting rule of a base $ident_i$ is applied to the main subgoal until it fails. Once all the rules have been processed, if the main subgoal has progressed (e.g., if it is distinct from the initial main goal) then the rules of this base are processed again. If the main subgoal has not progressed then the next base is processed. For the bases, the behavior is exactly similar to the processing of the rewriting rules.

The rewriting rule bases are built with the `Hint Rewrite` vernacular command.

Warning: This tactic may loop if you build non terminating rewriting systems.

Variant:

1. `autorewrite with $ident_1 \dots ident_n$ using tactic`

Performs, in the same way, all the rewritings of the bases $ident_1 \dots ident_n$ applying *tactic* to the main subgoal after each rewriting step.

2. `autorewrite with $ident_1 \dots ident_n$ in qualid`

Performs all the rewritings in hypothesis *qualid*.

3. `autorewrite with $ident_1 \dots ident_n$ in qualid using tactic`

Performs all the rewritings in hypothesis *qualid* applying *tactic* to the main subgoal after each rewriting step.

4. `autorewrite with $ident_1 \dots ident_n$ in clause`

Performs all the rewriting in the clause *clause*. The *clause* argument must not contain any `type of` or `value of`.

See also: Section 8.9.5 for feeding the database of lemmas used by `autorewrite`.

See also: Section 10.2 for examples showing the use of this tactic.

⁴The behavior of this tactic has much changed compared to the versions available in the previous distributions (V6). This may cause significant changes in your theories to obtain the same result. As a drawback of the re-engineering of the code, this tactic has also been completely revised to get a very compact and readable version.

8.8.5 `easy`

This tactic tries to solve the current goal by a number of standard closing steps. In particular, it tries to close the current goal using the closing tactics `trivial`, `reflexivity`, `symmetry`, `contradiction` and `inversion of hypothesis`. If this fails, it tries introducing variables and splitting and-hypotheses, using the closing tactics afterwards, and splitting the goal using `split` and `recursing`.

This tactic solves goals that belong to many common classes; in particular, many cases of unsatisfiable hypotheses, and simple equality goals are usually solved by this tactic.

Variant:

1. `now tactic`

Run `tactic` followed by `easy`. This is a notation for `tactic; easy`.

8.9 Controlling automation

8.9.1 The hints databases for `auto` and `eauto`

The hints for `auto` and `eauto` are stored in databases. Each database maps head symbols to a list of hints. One can use the command `Print Hint ident` to display the hints associated to the head symbol `ident` (see 8.9.4). Each hint has a cost that is a nonnegative integer, and an optional pattern. The hints with lower cost are tried first. A hint is tried by `auto` when the conclusion of the current goal matches its pattern or when it has no pattern.

Creating Hint databases

One can optionally declare a hint database using the command `Create HintDb`. If a hint is added to an unknown database, it will be automatically created.

```
Create HintDb ident [discriminated]
```

This command creates a new database named `ident`. The database is implemented by a Discrimination Tree (DT) that serves as an index of all the lemmas. The DT can use transparency information to decide if a constant should be indexed or not (c.f. 8.9.1), making the retrieval more efficient. The legacy implementation (the default one for new databases) uses the DT only on goals without existentials (i.e., `auto` goals), for non-Immediate hints and do not make use of transparency hints, putting more work on the unification that is run after retrieval (it keeps a list of the lemmas in case the DT is not used). The new implementation enabled by the `discriminated` option makes use of DTs in all cases and takes transparency information into account. However, the order in which hints are retrieved from the DT may differ from the order in which they were inserted, making this implementation observationally different from the legacy one.

The general command to add a hint to some databases `ident1, ..., identn` is

```
Hint hint_definition : ident1 ... identn
```

Variants:

1. `Hint hint_definition`

No database name is given: the hint is registered in the `core` database.

2. Local Hint *hint_definition* : *ident*₁ ... *ident*_{*n*}

This is used to declare hints that must not be exported to the other modules that require and import the current module. Inside a section, the option `Local` is useless since hints do not survive anyway to the closure of sections.

3. Local Hint *hint_definition*

Idem for the `core` database.

The *hint_definition* is one of the following expressions:

- `Resolve term [| [num] [pattern]]`

This command adds `simple apply term` to the hint list with the head symbol of the type of *term*. The cost of that hint is the number of subgoals generated by `simple apply term` or *num* if specified. The associated pattern is inferred from the conclusion of the type of *term* or the given *pattern* if specified.

In case the inferred type of *term* does not start with a product the tactic added in the hint list is `exact term`. In case this type can however be reduced to a type starting with a product, the tactic `simple apply term` is also stored in the hints list.

If the inferred type of *term* contains a dependent quantification on a variable which occurs only in the premisses of the type and not in its conclusion, no instance could be inferred for the variable by unification with the goal. In this case, the hint is added to the hint list of `eauto` (see 8.8.2) instead of the hint list of `auto` and a warning is printed. A typical example of a hint that is used only by `eauto` is a transitivity lemma.

Error messages:

1. *term* cannot be used as a hint

The head symbol of the type of *term* is a bound variable such that this tactic cannot be associated to a constant.

Variants:

1. `Resolve term1 ... termm`

Adds each `Resolve termi`.

2. `Resolve -> term`

Adds the left-to-right implication of an equivalence as a hint (informally the hint will be used as `apply <- term`, although as mentionned before, the tactic actually used is a restricted version of `apply`).

3. `Resolve <- term`

Adds the right-to-left implication of an equivalence as a hint.

- `Immediate term`

This command adds `simple apply term; trivial` to the hint list associated with the head symbol of the type of *ident* in the given database. This tactic will fail if all the subgoals generated by `simple apply term` are not solved immediately by the `trivial` tactic (which only tries tactics with cost 0).

This command is useful for theorems such as the symmetry of equality or $n+1 = m+1 \rightarrow n = m$ that we may like to introduce with a limited use in order to avoid useless proof-search.

The cost of this tactic (which never generates subgoals) is always 1, so that it is not used by `trivial` itself.

Error messages:

1. `term` cannot be used as a hint

Variants:

1. Immediate `term1 ... termm`
Adds each Immediate `termi`.

- Constructors `ident`

If `ident` is an inductive type, this command adds all its constructors as hints of type `Resolve`. Then, when the conclusion of current goal has the form `(ident ...)`, `auto` will try to apply each constructor.

Error messages:

1. `ident` is not an inductive type

Variants:

1. Constructors `ident1 ... identm`
Adds each Constructors `identi`.

- Unfold `qualid`

This adds the tactic `unfold qualid` to the hint list that will only be used when the head constant of the goal is `ident`. Its cost is 4.

Variants:

1. Unfold `ident1 ... identm`
Adds each Unfold `identi`.

- Transparent, Opaque `qualid`

This adds a transparency hint to the database, making `qualid` a transparent or opaque constant during resolution. This information is used during unification of the goal with any lemma in the database and inside the discrimination network to relax or constrain it in the case of discriminated databases.

Variants:

1. Transparent, Opaque `ident1 ... identm`
Declares each `identi` as a transparent or opaque constant.

- Extern *num* [*pattern*] => *tactic*

This hint type is to extend `auto` with tactics other than `apply` and `unfold`. For that, we must specify a cost, an optional pattern and a tactic to execute. Here is an example:

```
Hint Extern 4 (~(_ = _)) => discriminate.
```

Now, when the head of the goal is a disequality, `auto` will try `discriminate` if it does not manage to solve the goal with hints with a cost less than 4.

One can even use some sub-patterns of the pattern in the tactic script. A sub-pattern is a question mark followed by an identifier, like `?X1` or `?X2`. Here is an example:

```
Coq < Require Import List.

Coq < Hint Extern 5    ({?X1 = ?X2} + {?X1 <> ?X2}) =>
    generalize X1, X2; decide equality : eqdec.

Coq < Goal
    forall a b : list (nat * nat), {a = b} + {a <> b}.
1 subgoal

=====
forall a b : list (nat * nat), {a = b} + {a <> b}

Coq < Info 1 auto with eqdec.
<ltac_plugin::auto@0> "eqdec"
No more subgoals.
```

- Cut *regexp*

Warning: these hints currently only apply to typeclass proof search and the typeclasses `eauto` tactic (20.6.5).

This command can be used to cut the proof-search tree according to a regular expression matching paths to be cut. The grammar for regular expressions is the following. Beware, there is no operator precedence during parsing, one can check with `Print HintDb` to verify the current cut expression:

$e ::=$	<i>ident</i>	hint or instance identifier
	$-$	any hint
	$e e'$	disjunction
	ee'	sequence
	e^*	Kleene star
	<code>emp</code>	empty
	<code>eps</code>	epsilon
	(e)	

The `emp` regexp does not match any search path while `eps` matches the empty path. During proof search, the path of successive successful hints on a search branch is recorded, as a list of identifiers for the hints (note `Hint Extern`'s do not have an associated identifier). Before applying any hint *ident* the current path *p* extended with *ident* is matched against the current cut expression *c* associated to the hint database. If matching succeeds, the hint is *not* applied. The semantics of `Hint Cut e` is to set the cut expression to $c|e$, the initial cut expression being `emp`.

- Mode *qualid* (+ | ! | -)*

This sets an optional mode of use of the identifier *qualid*. When proof-search faces a goal that ends in an application of *qualid* to arguments $term_1 \dots term_n$, the mode tells if the hints associated to *qualid* can be applied or not. A mode specification is a list of n +, ! or - items that specify if an argument of the identifier is to be treated as an input (+), if its head only is an input (!) or an output (-) of the identifier. For a mode to match a list of arguments, input terms and input heads *must not* contain existential variables or be existential variables respectively, while outputs can be any term. Multiple modes can be declared for a single identifier, in that case only one mode needs to match the arguments for the hints to be applied.

The head of a term is understood here as the applicative head, or the match or projection scrutinee's head, recursively, casts being ignored.

Hint Mode is especially useful for typeclasses, when one does not want to support default instances and avoid ambiguity in general. Setting a parameter of a class as an input forces proof-search to be driven by that index of the class, with ! giving more flexibility by allowing existentials to still appear deeper in the index but not at its head.

Remark: One can use an `Extern` hint with no pattern to do pattern-matching on hypotheses using `match goal with` inside the tactic.

8.9.2 Hint databases defined in the COQ standard library

Several hint databases are defined in the COQ standard library. The actual content of a database is the collection of the hints declared to belong to this database in each of the various modules currently loaded. Especially, requiring new modules potentially extend a database. At COQ startup, only the `core` database is non empty and can be used.

`core` This special database is automatically used by `auto`, except when pseudo-database `nocore` is given to `auto`. The `core` database contains only basic lemmas about negation, conjunction, and so on from. Most of the hints in this database come from the `Init` and `Logic` directories.

`arith` This database contains all lemmas about Peano's arithmetic proved in the directories `Init` and `Arith`

`zarith` contains lemmas about binary signed integers from the directories `theories/ZArith`. When required, the module `Omega` also extends the database `zarith` with a high-cost hint that calls `omega` on equations and inequalities in `nat` or `Z`.

`bool` contains lemmas about booleans, mostly from directory `theories/Bool`.

`datatypes` is for lemmas about lists, streams and so on that are mainly proved in the `Lists` subdirectory.

`sets` contains lemmas about sets and relations from the directories `Sets` and `Relations`.

`typeclass_instances` contains all the type class instances declared in the environment, including those used for `setoid_rewrite`, from the `Classes` directory.

You are advised not to put your own hints in the `core` database, but use one or several databases specific to your development.

8.9.3 Remove Hints $term_1 \dots term_n : ident_1 \dots ident_m$

This command removes the hints associated to terms $term_1 \dots term_n$ in databases $ident_1 \dots ident_m$.

8.9.4 Print Hint

This command displays all hints that apply to the current goal. It fails if no proof is being edited, while the two variants can be used at every moment.

Variants:

1. Print Hint *ident*

This command displays only tactics associated with *ident* in the hints list. This is independent of the goal being edited, so this command will not fail if no goal is being edited.

2. Print Hint *

This command displays all declared hints.

3. Print HintDb *ident*

This command displays all hints from database *ident*.

8.9.5 Hint Rewrite $term_1 \dots term_n : ident_1 \dots ident_m$

This vernacular command adds the terms $term_1 \dots term_n$ (their types must be equalities) in the rewriting bases $ident_1, \dots, ident_m$ with the default orientation (left to right). Notice that the rewriting bases are distinct from the `auto` hint bases and that `auto` does not take them into account.

This command is synchronous with the section mechanism (see 2.4): when closing a section, all aliases created by `Hint Rewrite` in that section are lost. Conversely, when loading a module, all `Hint Rewrite` declarations at the global level of that module are loaded.

Variants:

1. Hint Rewrite $\rightarrow term_1 \dots term_n : ident_1 \dots ident_m$

This is strictly equivalent to the command above (we only make explicit the orientation which otherwise defaults to \rightarrow).

2. Hint Rewrite $\leftarrow term_1 \dots term_n : ident_1 \dots ident_m$

Adds the rewriting rules $term_1 \dots term_n$ with a right-to-left orientation in the bases $ident_1, \dots, ident_m$.

3. Hint Rewrite $term_1 \dots term_n$ using *tactic* : $ident_1 \dots ident_m$

When the rewriting rules $term_1 \dots term_n$ in $ident_1, \dots, ident_m$ will be used, the tactic *tactic* will be applied to the generated subgoals, the main subgoal excluded.

4. Print Rewrite HintDb *ident*

This command displays all rewrite hints contained in *ident*.

8.9.6 Hint locality

Hints provided by the `Hint` commands are erased when closing a section. Conversely, all hints of a module `A` that are not defined inside a section (and not defined with option `Local`) become available when the module `A` is imported (using e.g. `Require Import A.`).

As of today, hints only have a binary behavior regarding locality, as described above: either they disappear at the end of a section scope, or they remain global forever. This causes a scalability issue, because hints coming from an unrelated part of the code may badly influence another development. It can be mitigated to some extent thanks to the `Remove Hints` command (see 8.9.3), but this is a mere workaround and has some limitations (for instance, external hints cannot be removed).

A proper way to fix this issue is to bind the hints to their module scope, as for most of the other objects Coq uses. Hints should only be made available when the module they are defined in is imported, not just required. It is very difficult to change the historical behavior, as it would break a lot of scripts. We propose a smooth transitional path by providing the `Loose Hint Behavior` option which accepts three flags allowing for a fine-grained handling of non-imported hints.

Variants:

1. Set `Loose Hint Behavior "Lax"`

This is the default, and corresponds to the historical behavior, that is, hints defined outside of a section have a global scope.

2. Set `Loose Hint Behavior "Warn"`

When set, it outputs a warning when a non-imported hint is used. Note that this is an over-approximation, because a hint may be triggered by a run that will eventually fail and backtrack, resulting in the hint not being actually useful for the proof.

3. Set `Loose Hint Behavior "Strict"`

When set, it changes the behavior of an unloaded hint to a immediate fail tactic, allowing to emulate an import-scoped hint mechanism.

8.9.7 Setting implicit automation tactics

`Proof with tactic`

This command may be used to start a proof. It defines a default tactic to be used each time a tactic command `tactic1` is ended by “`. . .`”. In this case the tactic command typed by the user is equivalent to `tactic1;tactic`.

See also: `Proof .` in Section 7.1.4.

Variants:

1. `Proof with tactic using ident1 . . . identn`

Combines in a single line `Proof with` and `Proof using`, see 7.1.5

2. `Proof using ident1 . . . identn with tactic`

Combines in a single line `Proof with` and `Proof using`, see 7.1.5

Declare Implicit Tactic *tactic*

This command declares a tactic to be used to solve implicit arguments that COQ does not know how to solve by unification. It is used every time the term argument of a tactic has one of its holes not fully resolved.

Here is an example:

```
Coq < Parameter quo : nat -> forall n:nat, n<>0 -> nat.
quo is declared

Coq < Notation "x // y" := (quo x y _) (at level 40).

Coq < Declare Implicit Tactic assumption.

Coq < Goal forall n m, m<>0 -> { q:nat & { r | q * m + r = n } }.
1 subgoal

=====
forall n m : nat, m <> 0 -> {q : nat & {r : nat | q * m + r = n}}

Coq < intros.
1 subgoal

n, m : nat
H : m <> 0
=====
{q : nat & {r : nat | q * m + r = n}}

Coq < exists (n // m).
1 subgoal

n, m : nat
H : m <> 0
=====
{r : nat | n // m * m + r = n}
```

The tactic `exists (n // m)` did not fail. The hole was solved by `assumption` so that it behaved as `exists (quo n m H)`.

8.10 Decision procedures

8.10.1 `tauto`

This tactic implements a decision procedure for intuitionistic propositional calculus based on the contraction-free sequent calculi LJ^T* of Roy Dyckhoff [56]. Note that `tauto` succeeds on any instance of an intuitionistic tautological proposition. `tauto` unfolds negations and logical equivalence but does not unfold any other definition.

The following goal can be proved by `tauto` whereas `auto` would fail:

```
Coq < Goal forall (x:nat) (P:nat -> Prop), x = 0 \/ P x -> x <> 0 -> P x.
1 subgoal

=====
forall (x : nat) (P : nat -> Prop), x = 0 \/ P x -> x <> 0 -> P x
```

```

Coq < intros.
1 subgoal

  x : nat
  P : nat -> Prop
  H : x = 0 \ / P x
  H0 : x <> 0
  =====
  P x
Coq < tauto.
No more subgoals.

```

Moreover, if it has nothing else to do, `tauto` performs introductions. Therefore, the use of `intros` in the previous proof is unnecessary. `tauto` can for instance prove the following:

```

Coq < (* auto would fail *)
      Goal forall (A:Prop) (P:nat -> Prop),
        A \ / (forall x:nat, ~ A -> P x) -> forall x:nat, ~ A -> P x.
1 subgoal

  =====
  forall (A : Prop) (P : nat -> Prop),
  A \ / (forall x : nat, ~ A -> P x) -> forall x : nat, ~ A -> P x
Coq < tauto.
No more subgoals.

```

Remark: In contrast, `tauto` cannot solve the following goal

```

Coq < Goal forall (A:Prop) (P:nat -> Prop),
      A \ / (forall x:nat, ~ A -> P x) -> forall x:nat, ~ ~ (A \ / P x).

```

because $(\text{forall } x:\text{nat}, \sim A \rightarrow P\ x)$ cannot be treated as atomic and an instantiation of x is necessary.

Variants:

1. `dtauto`

While `tauto` recognizes inductively defined connectives isomorphic to the standard connective `and`, `prod`, `or`, `sum`, `False`, `Empty_set`, `unit`, `True`, `dtauto` recognizes also all inductive types with one constructors and no indices, i.e. record-style connectives.

8.10.2 intuition *tactic*

The tactic `intuition` takes advantage of the search-tree built by the decision procedure involved in the tactic `tauto`. It uses this information to generate a set of subgoals equivalent to the original one (but simpler than it) and applies the tactic *tactic* to them [113]. If this tactic fails on some goals then `intuition` fails. In fact, `tauto` is simply `intuition fail`.

For instance, the tactic `intuition auto` applied to the goal

```
(forall (x:nat), P x) /\ B -> (forall (y:nat), P y) /\ P O /\ B /\ P O
```

internally replaces it by the equivalent one:

```
(forall (x:nat), P x), B |- P O
```

and then uses `auto` which completes the proof.

Originally due to César Muñoz, these tactics (`tauto` and `intuition`) have been completely re-engineered by David Delahaye using mainly the tactic language (see Chapter 9). The code is now much shorter and a significant increase in performance has been noticed. The general behavior with respect to dependent types, unfolding and introductions has slightly changed to get clearer semantics. This may lead to some incompatibilities.

Variants:

1. `intuition`

Is equivalent to `intuition auto` with `*`.

2. `dintuition`

While `intuition` recognizes inductively defined connectives isomorphic to the standard connective `and`, `prod`, `or`, `sum`, `False`, `Empty_set`, `unit`, `True`, `dintuition` recognizes also all inductive types with one constructors and no indices, i.e. record-style connectives.

Some aspects of the tactic `intuition` can be controlled using options. To avoid that inner negations which do not need to be unfolded are unfolded, use:

```
Unset Intuition Negation Unfolding
```

To do that all negations of the goal are unfolded even inner ones (this is the default), use:

```
Set Intuition Negation Unfolding
```

To avoid that inner occurrence of `iff` which do not need to be unfolded are unfolded (this is the default), use:

```
Unset Intuition Iff Unfolding
```

To do that all negations of the goal are unfolded even inner ones (this is the default), use:

```
Set Intuition Iff Unfolding
```

8.10.3 `rtauto`

The `rtauto` tactic solves propositional tautologies similarly to what `tauto` does. The main difference is that the proof term is built using a reflection scheme applied to a sequent calculus proof of the goal. The search procedure is also implemented using a different technique.

Users should be aware that this difference may result in faster proof-search but slower proof-checking, and `rtauto` might not solve goals that `tauto` would be able to solve (e.g. goals involving universal quantifiers).

8.10.4 firstorder

The tactic `firstorder` is an *experimental* extension of `tauto` to first-order reasoning, written by Pierre Corbineau. It is not restricted to usual logical connectives but instead may reason about any first-order class inductive definition.

The default tactic used by `firstorder` when no rule applies is `auto` with `*`, it can be reset locally or globally using the `Set Firstorder Solver tactic vernacular` command and printed using `Print Firstorder Solver`.

Variants:

1. `firstorder tactic`
Tries to solve the goal with `tactic` when no logical rule may apply.
2. `firstorder using qualid1 , ... , qualidn`
Adds lemmas `qualid1 ... qualidn` to the proof-search environment. If `qualidi` refers to an inductive type, it is the collection of its constructors which are added to the proof-search environment.
3. `firstorder with ident1 ... identn`
Adds lemmas from `auto` hint bases `ident1 ... identn` to the proof-search environment.
4. `firstorder tactic using qualid1 , ... , qualidn with ident1 ... identn`
This combines the effects of the different variants of `firstorder`.

Proof-search is bounded by a depth parameter which can be set by typing the `Set Firstorder Depth n vernacular` command.

8.10.5 congruence

The tactic `congruence`, by Pierre Corbineau, implements the standard Nelson and Oppen congruence closure algorithm, which is a decision procedure for ground equalities with uninterpreted symbols. It also include the constructor theory (see 8.5.7 and 8.5.6). If the goal is a non-quantified equality, `congruence` tries to prove it with non-quantified equalities in the context. Otherwise it tries to infer a discriminable equality from those in the context. Alternatively, `congruence` tries to prove that a hypothesis is equal to the goal or to the negation of another hypothesis.

`congruence` is also able to take advantage of hypotheses stating quantified equalities, you have to provide a bound for the number of extra equalities generated that way. Please note that one of the members of the equality must contain all the quantified variables in order for `congruence` to match against it.

```
Coq < Theorem T:
      a=(f a) -> (g b (f a))=(f (f a)) -> (g a b)=(f (g b a)) -> (g a b)=a.
1 subgoal

=====
a = f a -> g b (f a) = f (f a) -> g a b = f (g b a) -> g a b = a

Coq < intros.
1 subgoal

H : a = f a
```

```

H0 : g b (f a) = f (f a)
H1 : g a b = f (g b a)
=====
g a b = a

Coq < congruence.
No more subgoals.

Coq < Theorem inj : f = pair a -> Some (f c) = Some (f d) -> c=d.
1 subgoal

=====
f = pair a -> Some (f c) = Some (f d) -> c = d

Coq < intros.
1 subgoal

H : f = pair a
H0 : Some (f c) = Some (f d)
=====
c = d

Coq < congruence.
No more subgoals.

```

Variants:

1. congruence *n*
Tries to add at most *n* instances of hypotheses stating quantified equalities to the problem in order to solve it. A bigger value of *n* does not make success slower, only failure. You might consider adding some lemmas as hypotheses using `assert` in order for congruence to use them.
2. congruence with *term*₁ ... *term*_{*n*}
Adds *term*₁ ... *term*_{*n*} to the pool of terms used by congruence. This helps in case you have partially applied constructors in your goal.

Error messages:

1. I don't know how to handle dependent equality
The decision procedure managed to find a proof of the goal or of a discriminable equality but this proof could not be built in COQ because of dependently-typed functions.
2. Goal is solvable by congruence but some arguments are missing.
Try "congruence with ...", replacing metavariables by arbitrary terms.
The decision procedure could solve the goal with the provision that additional arguments are supplied for some partially applied constructors. Any term of an appropriate type will allow the tactic to successfully solve the goal. Those additional arguments can be given to congruence by filling in the holes in the terms given in the error message, using the `with` variant described above.

Remark: congruence can be made to print debug information by setting the following option:

```
Set Congruence Verbose
```

8.11 Checking properties of terms

Each of the following tactics acts as the identity if the check succeeds, and results in an error otherwise.

8.11.1 `constr_eq term1 term2`

This tactic checks whether its arguments are equal modulo alpha conversion and casts.

Error message: `Not equal`

8.11.2 `unify term1 term2`

This tactic checks whether its arguments are unifiable, potentially instantiating existential variables.

Error message: `Not unifiable`

Variants:

1. `unify term1 term2 with ident`

Unification takes the transparency information defined in the hint database *ident* into account (see Section 8.9.1).

8.11.3 `is_evar term`

This tactic checks whether its argument is a current existential variable. Existential variables are uninstantiated variables generated by `eapply` (see Section 8.2.4) and some other tactics.

Error message: `Not an evar`

8.11.4 `has_evar term`

This tactic checks whether its argument has an existential variable as a subterm. Unlike `context` patterns combined with `is_evar`, this tactic scans all subterms, including those under binders.

Error message: `No evars`

8.11.5 `is_var term`

This tactic checks whether its argument is a variable or hypothesis in the current goal context or in the opened sections.

Error message: `Not a variable or hypothesis`

8.12 Equality

8.12.1 `f_equal`

This tactic applies to a goal of the form $f\ a_1 \dots a_n = f'\ a'_1 \dots a'_n$. Using `f_equal` on such a goal leads to subgoals $f = f'$ and $a_1 = a'_1$ and so on up to $a_n = a'_n$. Amongst these subgoals, the simple ones (e.g. provable by reflexivity or congruence) are automatically solved by `f_equal`.

8.12.2 `reflexivity`

This tactic applies to a goal that has the form $t=u$. It checks that t and u are convertible and then solves the goal. It is equivalent to apply `refl_equal`.

Error messages:

1. The conclusion is not a substitutive equation
2. Unable to unify ... with ...

8.12.3 `symmetry`

This tactic applies to a goal that has the form $t=u$ and changes it into $u=t$.

Variants:

1. `symmetry in ident`

If the statement of the hypothesis *ident* has the form $t=u$, the tactic changes it to $u=t$.

8.12.4 `transitivity term`

This tactic applies to a goal that has the form $t=u$ and transforms it into the two subgoals $t=term$ and $term=u$.

8.13 Equality and inductive sets

We describe in this section some special purpose tactics dealing with equality and inductive sets or types. These tactics use the equality `eq:forall (A:Type), A->A->Prop`, simply written with the infix symbol `=`.

8.13.1 `decide equality`

This tactic solves a goal of the form `forall x y:R, {x=y}+{~x=y}`, where R is an inductive type such that its constructors do not take proofs or functions as arguments, nor objects in dependent types. It solves goals of the form $\{x=y\}+{\sim x=y}$ as well.

8.13.2 `compare term1 term2`

This tactic compares two given objects $term_1$ and $term_2$ of an inductive datatype. If G is the current goal, it leaves the sub-goals $term_1=term_2 \rightarrow G$ and $\sim term_1=term_2 \rightarrow G$. The type of $term_1$ and $term_2$ must satisfy the same restrictions as in the tactic `decide equality`.

8.13.3 `simplify_eq term`

Let $term$ be the proof of a statement of conclusion $term_1=term_2$. If $term_1$ and $term_2$ are structurally different (in the sense described for the tactic `discriminate`), then the tactic `simplify_eq` behaves as `discriminate term`, otherwise it behaves as `injection term`.

Remark: If some quantified hypothesis of the goal is named *ident*, then `simplify_eq ident` first introduces the hypothesis in the local context using `intros until ident`.

Variants:

1. `simplify_eq num`

This does the same thing as `intros until num` then `simplify_eq ident` where *ident* is the identifier for the last introduced hypothesis.

2. `simplify_eq term` with *bindings_list*

This does the same as `simplify_eq term` but using the given bindings to instantiate parameters or hypotheses of *term*.

3. `esimplify_eq num`

`esimplify_eq term [with bindings_list]`

This works the same as `simplify_eq` but if the type of *term*, or the type of the hypothesis referred to by *num*, has uninstantiated parameters, these parameters are left as existential variables.

4. `simplify_eq`

If the current goal has form $t_1 <> t_2$, it behaves as `intro ident; simplify_eq ident`.

8.13.4 `dependent rewrite -> ident`

This tactic applies to any goal. If *ident* has type $(\text{existT } B \ a \ b) = (\text{existT } B \ a' \ b')$ in the local context (i.e. each term of the equality has a sigma type $\{a : A \ \& \ (B \ a)\}$) this tactic rewrites *a* into *a'* and *b* into *b'* in the current goal. This tactic works even if *B* is also a sigma type. This kind of equalities between dependent pairs may be derived by the injection and inversion tactics.

Variants:1. `dependent rewrite <- ident`

Analogous to `dependent rewrite ->` but uses the equality from right to left.

8.14 Inversion**8.14.1** `functional inversion ident`

`functional inversion` is a tactic that performs inversion on hypothesis *ident* of the form *qualid term₁...term_n = term* or *term = qualid term₁...term_n* where *qualid* must have been defined using `Function` (see Section 2.3). Note that this tactic is only available after a `Require Import FunInd`.

Error messages:1. Hypothesis *ident* must contain at least one `Function`2. Cannot find inversion information for hypothesis *ident*

This error may be raised when some inversion lemma failed to be generated by `Function`.

Variants:1. `functional inversion num`

This does the same thing as `intros until num` then `functional inversion ident` where *ident* is the identifier for the last introduced hypothesis.

2. functional inversion *ident qualid*
functional inversion *num qualid*

If the hypothesis *ident* (or *num*) has a type of the form $qualid_1 \text{ term}_1 \dots \text{term}_n = qualid_2 \text{ term}_{n+1} \dots \text{term}_{n+m}$ where $qualid_1$ and $qualid_2$ are valid candidates to functional inversion, this variant allows choosing which *qualid* is inverted.

8.14.2 quote *ident*

This kind of inversion has nothing to do with the tactic *inversion* above. This tactic does change (*ident* *t*), where *t* is a term built in order to ensure the convertibility. In other words, it does inversion of the function *ident*. This function must be a fixpoint on a simple recursive datatype: see 10.3 for the full details.

Error messages:

1. quote: not a simple fixpoint
Happens when *quote* is not able to perform inversion properly.

Variants:

1. quote *ident* [*ident*₁ ... *ident*_{*n*}]
All terms that are built only with *ident*₁ ... *ident*_{*n*} will be considered by *quote* as constants rather than variables.

8.15 Classical tactics

In order to ease the proving process, when the *Classical* module is loaded. A few more tactics are available. Make sure to load the module using the *Require Import* command.

8.15.1 classical_left and classical_right

The tactics *classical_left* and *classical_right* are the analog of the *left* and *right* but using classical logic. They can only be used for disjunctions. Use *classical_left* to prove the left part of the disjunction with the assumption that the negation of right part holds. Use *classical_right* to prove the right part of the disjunction with the assumption that the negation of left part holds.

8.16 Automating

8.16.1 btauto

The tactic *btauto* implements a reflexive solver for boolean tautologies. It solves goals of the form *t* = *u* where *t* and *u* are constructed over the following grammar:

$$t ::= x \mid \text{true} \mid \text{false} \mid \text{orb } t_1 \ t_2 \mid \text{andb } t_1 \ t_2 \mid \text{xorb } t_1 \ t_2 \mid \text{negb } t \mid \text{if } t_1 \text{ then } t_2 \text{ else } t_3$$

Whenever the formula supplied is not a tautology, it also provides a counter-example. Internally, it uses a system very similar to the one of the *ring* tactic.

8.16.2 `omega`

The tactic `omega`, due to Pierre Crégut, is an automatic decision procedure for Presburger arithmetic. It solves quantifier-free formulas built with \sim , \wedge , \vee , \rightarrow on top of equalities, inequalities and disequalities on both the type `nat` of natural numbers and `Z` of binary integers. This tactic must be loaded by the command `Require Import Omega`. See the additional documentation about `omega` (see Chapter 21).

8.16.3 `ring` and `ring_simplify term1 ... termn`

The `ring` tactic solves equations upon polynomial expressions of a ring (or semi-ring) structure. It proceeds by normalizing both hand sides of the equation (w.r.t. associativity, commutativity and distributivity, constant propagation) and comparing syntactically the results.

`ring_simplify` applies the normalization procedure described above to the terms given. The tactic then replaces all occurrences of the terms given in the conclusion of the goal by their normal forms. If no term is given, then the conclusion should be an equation and both hand sides are normalized.

See Chapter 25 for more information on the tactic and how to declare new ring structures. All declared field structures can be printed with the `Print Rings` command.

8.16.4 `field`, `field_simplify term1 ... termn`, and `field_simplify_eq`

The `field` tactic is built on the same ideas as `ring`: this is a reflexive tactic that solves or simplifies equations in a field structure. The main idea is to reduce a field expression (which is an extension of ring expressions with the inverse and division operations) to a fraction made of two polynomial expressions.

Tactic `field` is used to solve subgoals, whereas `field_simplify term1 ... termn` replaces the provided terms by their reduced fraction. `field_simplify_eq` applies when the conclusion is an equation: it simplifies both hand sides and multiplies so as to cancel denominators. So it produces an equation without division nor inverse.

All of these 3 tactics may generate a subgoal in order to prove that denominators are different from zero.

See Chapter 25 for more information on the tactic and how to declare new field structures. All declared field structures can be printed with the `Print Fields` command.

Example:

```
Coq < Require Import Reals.
Coq < Goal forall x y:R,
    (x * y > 0)%R ->
    (x * (1 / x + x / (x + y)))%R =
    ((- 1 / y) * y * (- x * (x / (x + y)) - 1))%R.

Coq < intros; field.
1 subgoal

  x, y : R
  H : (x * y > 0)%R
  =====
  (x + y)%R <> 0%R /\ y <> 0%R /\ x <> 0%R
```

See also: file `plugins/setoid_ring/RealField.v` for an example of instantiation, theory `theories/Reals` for many examples of use of `field`.

8.16.5 `fourier`

This tactic written by Loïc Pottier solves linear inequalities on real numbers using Fourier’s method [65]. This tactic must be loaded by `Require Import Fourier`.

Example:

```
Coq < Require Import Reals.
Coq < Require Import Fourier.
Coq < Goal forall x y:R, (x < y)%R -> (y + 1 >= x - 1)%R.

Coq < intros; fourier.
No more subgoals.
```

8.17 Non-logical tactics

8.17.1 `cycle num`

This tactic puts the *num* first goals at the end of the list of goals. If *num* is negative, it will put the last $|num|$ goals at the beginning of the list.

Example:

```
Coq < Parameter P : nat -> Prop.
Coq < Goal P 1 /\ P 2 /\ P 3 /\ P 4 /\ P 5.

Coq < repeat split.
5 subgoals

=====
P 1
subgoal 2 is:
P 2
subgoal 3 is:
P 3
subgoal 4 is:
P 4
subgoal 5 is:
P 5

Coq < all: cycle 2.
5 subgoals

=====
P 3
subgoal 2 is:
P 4
subgoal 3 is:
P 5
subgoal 4 is:
P 1
subgoal 5 is:
P 2
```



```

Coq < all: cycle -3.
5 subgoals

=====
P 5
subgoal 2 is:
P 1
subgoal 3 is:
P 2
subgoal 4 is:
P 3
subgoal 5 is:
P 4

```

8.17.2 swap num_1 num_2

This tactic switches the position of the goals of indices num_1 and num_2 . If either num_1 or num_2 is negative then goals are counted from the end of the focused goal list. Goals are indexed from 1, there is no goal with position 0.

Example:

```

Coq < Parameter P : nat -> Prop.
Coq < Goal P 1 /\ P 2 /\ P 3 /\ P 4 /\ P 5.
Coq < repeat split.
5 subgoals

=====
P 1
subgoal 2 is:
P 2
subgoal 3 is:
P 3
subgoal 4 is:
P 4
subgoal 5 is:
P 5

Coq < all: swap 1 3.
5 subgoals

=====
P 3
subgoal 2 is:
P 2
subgoal 3 is:
P 1
subgoal 4 is:
P 4
subgoal 5 is:
P 5

Coq < all: swap 1 -1.
5 subgoals

```

```

=====
P 5
subgoal 2 is:
P 2
subgoal 3 is:
P 1
subgoal 4 is:
P 4
subgoal 5 is:
P 3

```

8.17.3 revgoals

This tactics reverses the list of the focused goals.

Example:

```

Coq < Parameter P : nat -> Prop.
Coq < Goal P 1 /\ P 2 /\ P 3 /\ P 4 /\ P 5.
Coq < repeat split.
5 subgoals

```

```

=====
P 1
subgoal 2 is:
P 2
subgoal 3 is:
P 3
subgoal 4 is:
P 4
subgoal 5 is:
P 5

Coq < all: revgoals.
5 subgoals

```

```

=====
P 5
subgoal 2 is:
P 4
subgoal 3 is:
P 3
subgoal 4 is:
P 2
subgoal 5 is:
P 1

```

8.17.4 shelve

This tactic moves all goals under focus to a shelf. While on the shelf, goals will not be focused on. They can be solved by unification, or they can be called back into focus with the command `Unshelve` (Section 8.17.5).

Variants:1. `shelve_unifiable`

Shelves only the goals under focus that are mentioned in other goals. Goals that appear in the type of other goals can be solved by unification.

Example:

```
Coq < Goal exists n, n=0.
1 subgoal

=====
exists n : nat, n = 0
Coq < refine (ex_intro _ _ _).
1 focused subgoal
(shelved: 1)

=====
?Goal = 0
Coq < all:shelve_unifiable.
1 focused subgoal
(shelved: 1)

=====
?Goal = 0
Coq < reflexivity.
No more subgoals.
```

8.17.5 Unshelve

This command moves all the goals on the shelf (see Section 8.17.4) from the shelf into focus, by appending them to the end of the current list of focused goals.

8.17.6 give_up

This tactic removes the focused goals from the proof. They are not solved, and cannot be solved later in the proof. As the goals are not solved, the proof cannot be closed.

The `give_up` tactic can be used while editing a proof, to choose to write the proof script in a non-sequential order.

8.18 Simple tactic macros

A simple example has more value than a long explanation:

```
Coq < Ltac Solve := simpl; intros; auto.
Solve is defined

Coq < Ltac ElimBoolRewrite b H1 H2 :=
  elim b; [ intros; rewrite H1; eauto | intros; rewrite H2; eauto ].
ElimBoolRewrite is defined
```

The tactics macros are synchronous with the COQ section mechanism: a tactic definition is deleted from the current environment when you close the section (see also 2.4) where it was defined. If you want that a tactic macro defined in a module is usable in the modules that require it, you should put it outside of any section.

Chapter 9 gives examples of more complex user-defined tactics.

Chapter 9

The tactic language

This chapter gives a compact documentation of Ltac, the tactic language available in COQ. We start by giving the syntax, and next, we present the informal semantics. If you want to know more regarding this language and especially about its foundations, you can refer to [43]. Chapter 10 is devoted to giving examples of use of this language on small but also with non-trivial problems.

9.1 Syntax

The syntax of the tactic language is given Figures 9.1 and 9.2. See Chapter 1 for a description of the BNF metasyntax used in these grammar rules. Various already defined entries will be used in this chapter: entries *natural*, *integer*, *ident*, *qualid*, *term*, *cpattern* and *atomic_tactic* represent respectively the natural and integer numbers, the authorized identifiers and qualified names, COQ's terms and patterns and all the atomic tactics described in Chapter 8. The syntax of *cpattern* is the same as that of terms, but it is extended with pattern matching metavariables. In *cpattern*, a pattern-matching metavariable is represented with the syntax *?id* where *id* is an *ident*. The notation *_* can also be used to denote metavariable whose instance is irrelevant. In the notation *?id*, the identifier allows us to keep instantiations and to make constraints whereas *_* shows that we are not interested in what will be matched. On the right hand side of pattern-matching clauses, the named metavariable are used without the question mark prefix. There is also a special notation for second-order pattern-matching problems: in an applicative pattern of the form *@?id id₁ ... id_n*, the variable *id* matches any complex expression with (possible) dependencies in the variables *id₁ ... id_n* and returns a functional term of the form *fun id₁ ... id_n => term*.

The main entry of the grammar is *expr*. This language is used in proof mode but it can also be used in toplevel definitions as shown in Figure 9.3.

Remarks:

1. The infix tacticals “... || ...”, “... + ...”, and “... ; ...” are associative.
2. In *tacarg*, there is an overlap between *qualid* as a direct tactic argument and *qualid* as a particular case of *term*. The resolution is done by first looking for a reference of the tactic language and if it fails, for a reference to a term. To force the resolution as a reference of the tactic language, use the form *ltac : qualid*. To force the resolution as a reference to a term, use the syntax *(qualid)*.
3. As shown by the figure, tactical *||* binds more than the prefix tacticals *try*, *repeat*, *do* and *abstract* which themselves bind more than the postfix tactical “... ; [...]” which binds more than “... ; ...”.

For instance

```
try repeat tactic1 || tactic2; tactic3; [tactic31 | . . . | tactic3n] ; tactic4 .
```

is understood as

```
(try (repeat (tactic1 || tactic2))) ;  
( (tactic3; [tactic31 | . . . | tactic3n] ) ; tactic4 ) .
```

9.2 Semantics

Tactic expressions can only be applied in the context of a proof. The evaluation yields either a term, an integer or a tactic. Intermediary results can be terms or integers but the final result must be a tactic which is then applied to the focused goals.

There is a special case for `match goal` expressions of which the clauses evaluate to tactics. Such expressions can only be used as end result of a tactic expression (never as argument of a non recursive local definition or of an application).

The rest of this section explains the semantics of every construction of Ltac.

Sequence

A sequence is an expression of the following form:

```
expr1 ; expr2
```

The expression *expr*₁ is evaluated to *v*₁, which must be a tactic value. The tactic *v*₁ is applied to the current goal, possibly producing more goals. Then *expr*₂ is evaluated to produce *v*₂, which must be a tactic value. The tactic *v*₂ is applied to all the goals produced by the prior application. Sequence is associative.

Local application of tactics

Different tactics can be applied to the different goals using the following form:

```
[ > expr1 | ... | exprn ]
```

The expressions *expr*_{*i*} are evaluated to *v*_{*i*}, for *i* = 0, ..., *n* and all have to be tactics. The *v*_{*i*} is applied to the *i*-th goal, for *i* = 1, ..., *n*. It fails if the number of focused goals is not exactly *n*.

Variants:

1. If no tactic is given for the *i*-th goal, it behaves as if the tactic `idtac` were given. For instance, `[> | auto]` is a shortcut for `[> idtac | auto]`.
2. `[> expr1 | ... | expri | expr . . | expri+1+j | ... | exprn]`

In this variant, *expr* is used for each goal numbered from *i* + 1 to *i* + *j* (assuming *n* is the number of goals).

Note that `. .` is part of the syntax, while `...` is the meta-symbol used to describe a list of *expr* of arbitrary length. goals numbered from *i* + 1 to *i* + *j*.

```

expr      ::= expr ; expr
             |  [ > expr | ... | expr ]
             |  expr ; [ expr | ... | expr ]
             |  tacexpr3

tacexpr3 ::= do (natural | ident) tacexpr3
             |  progress tacexpr3
             |  repeat tacexpr3
             |  try tacexpr3
             |  once tacexpr3
             |  exactly_once tacexpr3
             |  timeout (natural | ident) tacexpr3
             |  time [string] tacexpr3
             |  only selector : tacexpr3
             |  tacexpr2

tacexpr2 ::= tacexpr1 || tacexpr3
             |  tacexpr1 + tacexpr3
             |  tryif tacexpr1 then tacexpr1 else tacexpr1
             |  tacexpr1

tacexpr1 ::= fun name ... name => atom
             |  let [rec] let_clause with ... with let_clause in atom
             |  match goal with context_rule | ... | context_rule end
             |  match reverse goal with context_rule | ... | context_rule end
             |  match expr with match_rule | ... | match_rule end
             |  lazy match goal with context_rule | ... | context_rule end
             |  lazy match reverse goal with context_rule | ... | context_rule end
             |  lazy match expr with match_rule | ... | match_rule end
             |  multimatch goal with context_rule | ... | context_rule end
             |  multimatch reverse goal with context_rule | ... | context_rule end
             |  multimatch expr with match_rule | ... | match_rule end
             |  abstract atom
             |  abstract atom using ident
             |  first [ expr | ... | expr ]
             |  solve [ expr | ... | expr ]
             |  idtac [message_token ... message_token]
             |  fail [natural] [message_token ... message_token]
             |  gfail [natural] [message_token ... message_token]
             |  fresh | fresh string | fresh qualid
             |  context ident [ term ]
             |  eval redexpr in term
             |  type of term
             |  external string string tacarg ... tacarg
             |  constr : term
             |  unconstr : term
             |  type_term term
             |  numgoals
             |  guard test
             |  atomic_tactic
             |  qualid tacarg ... tacarg
             |  atom

```

Figure 9.1: Syntax of the tactic language

```

atom          ::=  qualid
                |   ()
                |   integer
                |   ( expr )

message_token  ::=  string | ident | integer

tacarg        ::=  qualid
                |   ()
                |   ltac : atom
                |   term

let_clause    ::=  ident [name ... name] := expr

context_rule  ::=  context_hyp , ... , context_hyp | -cpattern => expr
                |   | - cpattern => expr
                |   _ => expr

context_hyp   ::=  name : cpattern
                |   name := cpattern [: cpattern]

match_rule    ::=  cpattern => expr
                |   context [ident] [ cpattern ] => expr
                |   appcontext [ident] [ cpattern ] => expr
                |   _ => expr

test          ::=  integer = integer
                |   integer < integer
                |   integer <= integer
                |   integer > integer
                |   integer >= integer

selector      ::=  [ident]
                |   integer
                |   (integer | integer - integer) , ... , (integer | integer - integer)

toplevel_selector ::= selector
                |   all
                |   par

```

Figure 9.2: Syntax of the tactic language (continued)

3. [> $expr_1$ | ... | $expr_i$ | . . | $expr_{i+1+j}$ | ... | $expr_n$]

In this variant, `idtac` is used for the goals numbered from $i + 1$ to $i + j$.

4. [> $expr$. .]

$\begin{aligned} \text{top} & ::= [\text{Local}] \text{Ltac } \text{ltac_def} \text{ with } \dots \text{ with } \text{ltac_def} \\ \\ \text{ltac_def} & ::= \text{ident } [\text{ident } \dots \text{ ident}] := \text{expr} \\ & \quad \quad \text{qualid } [\text{ident } \dots \text{ ident}] := \text{expr} \end{aligned}$

Figure 9.3: Tactic toplevel definitions

In this variant, the tactic *expr* is applied independently to each of the goals, rather than globally. In particular, if there are no goal, the tactic is not run at all. A tactic which expects multiple goals, such as `swap`, would act as if a single goal is focused.

5. *expr* ; [*expr*₁ | ... | *expr*_{*n*}]

This variant of local tactic application is paired with a sequence. In this variant, *n* must be the number of goals generated by the application of *expr* to each of the individual goals independently. All the above variants work in this form too. Formally, *expr* ; [...] is equivalent to

$$[\text{> } \text{expr} ; [\text{> } \dots] \dots]$$

Goal selectors

We can restrict the application of a tactic to a subset of the currently focused goals with:

$$\text{toplevel_selector} : \text{expr}$$

We can also use selectors as a tactical, which allows to use them nested in a tactic expression, by using the keyword `only`:

$$\text{only selector} : \text{expr}$$

When selecting several goals, the tactic *expr* is applied globally to all selected goals.

Variants:

1. [*ident*] : *expr*

In this variant, *expr* is applied locally to a goal previously named by the user (see 2.11).

2. *num* : *expr*

In this variant, *expr* is applied locally to the *num*-th goal.

3. *n*₁-*m*₁, ..., *n*_{*k*}-*m*_{*k*} : *expr*

In this variant, *expr* is applied globally to the subset of goals described by the given ranges. You can write a single *n* as a shortcut for *n*-*n* when specifying multiple ranges.

4. `all` : *expr*

In this variant, *expr* is applied to all focused goals. `all` : can only be used at the toplevel of a tactic expression.

5. `par : expr`

In this variant, *expr* is applied to all focused goals in parallel. The number of workers can be controlled via the command line option `-async-proofs-tac-j` taking as argument the desired number of workers. Limitations: `par :` only works on goals containing no existential variables and *expr* must either solve the goal completely or do nothing (i.e. it cannot make some progress). `par :` can only be used at the toplevel of a tactic expression.

Error message: `No such goal`

For loop

There is a for loop that repeats a tactic *num* times:

```
do num expr
```

expr is evaluated to *v* which must be a tactic value. This tactic value *v* is applied *num* times. Supposing *num* > 1, after the first application of *v*, *v* is applied, at least once, to the generated subgoals and so on. It fails if the application of *v* fails before the *num* applications have been completed.

Repeat loop

We have a repeat loop with:

```
repeat expr
```

expr is evaluated to *v*. If *v* denotes a tactic, this tactic is applied to each focused goal independently. If the application succeeds, the tactic is applied recursively to all the generated subgoals until it eventually fails. The recursion stops in a subgoal when the tactic has failed *to make progress*. The tactic `repeat expr` itself never fails.

Error catching

We can catch the tactic errors with:

```
try expr
```

expr is evaluated to *v* which must be a tactic value. The tactic value *v* is applied to each focused goal independently. If the application of *v* fails in a goal, it catches the error and leaves the goal unchanged. If the level of the exception is positive, then the exception is re-raised with its level decremented.

Detecting progress

We can check if a tactic made progress with:

```
progress expr
```

expr is evaluated to *v* which must be a tactic value. The tactic value *v* is applied to each focused subgoal independently. If the application of *v* to one of the focused subgoal produced subgoals equal to the initial goals (up to syntactical equality), then an error of level 0 is raised.

Error message: `Failed to progress`

Backtracking branching

We can branch with the following structure:

$$expr_1 + expr_2$$

$expr_1$ and $expr_2$ are evaluated to v_1 and v_2 which must be tactic values. The tactic value v_1 is applied to each focused goal independently and if it fails or a later tactic fails, then the proof backtracks to the current goal and v_2 is applied.

Tactics can be seen as having several successes. When a tactic fails it asks for more successes of the prior tactics. $expr_1 + expr_2$ has all the successes of v_1 followed by all the successes of v_2 . Algebraically, $(expr_1 + expr_2);expr_3 = (expr_1;expr_3) + (expr_2;expr_3)$.

Branching is left-associative.

First tactic to work

Backtracking branching may be too expensive. In this case we may restrict to a local, left biased, branching and consider the first tactic to work (i.e. which does not fail) among a panel of tactics:

$$\text{first } [expr_1 \mid \dots \mid expr_n]$$

$expr_i$ are evaluated to v_i and v_i must be tactic values, for $i = 1, \dots, n$. Supposing $n > 1$, it applies, in each focused goal independently, v_1 , if it works, it stops otherwise it tries to apply v_2 and so on. It fails when there is no applicable tactic. In other words, $\text{first } [expr_1 \mid \dots \mid expr_n]$ behaves, in each goal, as the the first v_i to have *at least* one success.

Error message: No applicable tactic

Variant: `first expr`

This is an Ltac alias that gives a primitive access to the `first` tactical as a Ltac definition without going through a parsing rule. It expects to be given a list of tactics through a `Tactic Notation`, allowing to write notations of the following form.

Example:

```
Tactic Notation "foo" tactic_list(tacs) := first tacs.
```

Left-biased branching

Yet another way of branching without backtracking is the following structure:

$$expr_1 \mid\mid expr_2$$

$expr_1$ and $expr_2$ are evaluated to v_1 and v_2 which must be tactic values. The tactic value v_1 is applied in each subgoal independently and if it fails *to progress* then v_2 is applied. $expr_1 \mid\mid expr_2$ is equivalent to $\text{first } [\text{progress } expr_1 \mid expr_2]$ (except that if it fails, it fails like v_2). Branching is left-associative.

Generalized biased branching

The tactic

```
tryif  $expr_1$  then  $expr_2$  else  $expr_3$ 
```

is a generalization of the biased-branching tactics above. The expression $expr_1$ is evaluated to v_1 , which is then applied to each subgoal independently. For each goal where v_1 succeeds at least once, $expr_2$ is evaluated to v_2 which is then applied collectively to the generated subgoals. The v_2 tactic can trigger backtracking points in v_1 : where v_1 succeeds at least once, $tryif\ expr_1\ then\ expr_2\ else\ expr_3$ is equivalent to $v_1; v_2$. In each of the goals where v_1 does not succeed at least once, $expr_3$ is evaluated in v_3 which is then applied to the goal.

Soft cut

Another way of restricting backtracking is to restrict a tactic to a single success *a posteriori*:

```
once  $expr$ 
```

$expr$ is evaluated to v which must be a tactic value. The tactic value v is applied but only its first success is used. If v fails, $once\ expr$ fails like v . If v has at least one success, $once\ expr$ succeeds once, but cannot produce more successes.

Checking the successes

Coq provides an experimental way to check that a tactic has *exactly one* success:

```
exactly_once  $expr$ 
```

$expr$ is evaluated to v which must be a tactic value. The tactic value v is applied if it has at most one success. If v fails, $exactly_once\ expr$ fails like v . If v has exactly one success, $exactly_once\ expr$ succeeds like v . If v has two or more successes, $exactly_once\ expr$ fails.

The experimental status of this tactic pertains to the fact if v performs side effects, they may occur in an unpredictable way. Indeed, normally v would only be executed up to the first success until backtracking is needed, however $exactly_once$ needs to look ahead to see whether a second success exists, and may run further effects immediately.

Error message: This tactic has more than one success

Solving

We may consider the first to solve (i.e. which generates no subgoal) among a panel of tactics:

```
solve [  $expr_1$  | ... |  $expr_n$  ]
```

$expr_i$ are evaluated to v_i and v_i must be tactic values, for $i = 1, \dots, n$. Supposing $n > 1$, it applies v_1 to each goal independently, if it doesn't solve the goal then it tries to apply v_2 and so on. It fails if there is no solving tactic.

Error message: Cannot solve the goal

Variant: `solve $expr$`

This is an Ltac alias that gives a primitive access to the `solve` tactical. See the `first` tactical for more information.

Identity

The constant `idtac` is the identity tactic: it leaves any goal unchanged but it appears in the proof script.

Variant: `idtac message_token ... message_token`

This prints the given tokens. Strings and integers are printed literally. If a (term) variable is given, its contents are printed.

Failing

The tactic `fail` is the always-failing tactic: it does not solve any goal. It is useful for defining other tacticals since it can be caught by `try`, `repeat`, `match goal`, or the branching tacticals. The `fail` tactic will, however, succeed if all the goals have already been solved.

Variants:

1. `fail n`
The number n is the failure level. If no level is specified, it defaults to 0. The level is used by `try`, `repeat`, `match goal` and the branching tacticals. If 0, it makes `match goal` considering the next clause (backtracking). If non zero, the current `match goal` block, `try`, `repeat`, or branching command is aborted and the level is decremented. In the case of `+`, a non-zero level skips the first backtrack point, even if the call to `fail n` is not enclosed in a `+` command, respecting the algebraic identity.
2. `fail message_token ... message_token`
The given tokens are used for printing the failure message.
3. `fail n message_token ... message_token`
This is a combination of the previous variants.
4. `gfail`
This variant fails even if there are no goals left.
5. `gfail message_token ... message_token`
`gfail n message_token ... message_token`
These variants fail with an error message or an error level even if there are no goals left. Be careful however if Coq terms have to be printed as part of the failure: term construction always forces the tactic into the goals, meaning that if there are no goals when it is evaluated, a tactic call like `let x:=H in fail 0 x` will succeed.

Error message: `Tactic Failure message (level n).`

Timeout

We can force a tactic to stop if it has not finished after a certain amount of time:

`timeout num expr`

`expr` is evaluated to v which must be a tactic value. The tactic value v is applied normally, except that it is interrupted after `num` seconds if it is still running. In this case the outcome is a failure.

Warning: For the moment, `timeout` is based on elapsed time in seconds, which is very machine-dependent: a script that works on a quick machine may fail on a slow one. The converse is even

possible if you combine a `timeout` with some other tacticals. This tactical is hence proposed only for convenience during debug or other development phases, we strongly advise you to not leave any `timeout` in final scripts. Note also that this tactical isn't available on the native Windows port of Coq.

Timing a tactic

A tactic execution can be timed:

```
time string expr
```

evaluates *expr* and displays the time the tactic expression ran, whether it fails or successes. In case of several successes, the time for each successive runs is displayed. Time is in seconds and is machine-dependent. The *string* argument is optional. When provided, it is used to identify this particular occurrence of `time`.

Local definitions

Local definitions can be done as follows:

```
let ident1 := expr1
with ident2 := expr2
...
with identn := exprn in
expr
```

each *expr_i* is evaluated to *v_i*, then, *expr* is evaluated by substituting *v_i* to each occurrence of *ident_i*, for *i* = 1, ..., *n*. There is no dependencies between the *expr_i* and the *ident_i*.

Local definitions can be recursive by using `let rec` instead of `let`. In this latter case, the definitions are evaluated lazily so that the `rec` keyword can be used also in non recursive cases so as to avoid the eager evaluation of local definitions.

Application

An application is an expression of the following form:

```
qualid tacarg1 ... tacargn
```

The reference *qualid* must be bound to some defined tactic definition expecting at least *n* arguments. The expressions *expr_i* are evaluated to *v_i*, for *i* = 1, ..., *n*.

Function construction

A parameterized tactic can be built anonymously (without resorting to local definitions) with:

```
fun ident1 ... identn => expr
```

Indeed, local definitions of functions are a syntactic sugar for binding a `fun` tactic to an identifier.

Pattern matching on terms

We can carry out pattern matching on terms with:

```

match expr with
  cpattern1 => expr1
  | cpattern2 => expr2
  ...
  | cpatternn => exprn
  | _ => exprn+1
end

```

The expression *expr* is evaluated and should yield a term which is matched against *cpattern*₁. The matching is non-linear: if a metavariable occurs more than once, it should match the same expression every time. It is first-order except on the variables of the form @?id that occur in head position of an application. For these variables, the matching is second-order and returns a functional term.

Alternatively, when a metavariable of the form ?id occurs under binders, say *x*₁, ..., *x*_{*n*} and the expression matches, the metavariable is instantiated by a term which can then be used in any context which also binds the variables *x*₁, ..., *x*_{*n*} with same types. This provides with a primitive form of matching under context which does not require manipulating a functional term.

If the matching with *cpattern*₁ succeeds, then *expr*₁ is evaluated into some value by substituting the pattern matching instantiations to the metavariables. If *expr*₁ evaluates to a tactic and the match expression is in position to be applied to a goal (e.g. it is not bound to a variable by a let in), then this tactic is applied. If the tactic succeeds, the list of resulting subgoals is the result of the match expression. If *expr*₁ does not evaluate to a tactic or if the match expression is not in position to be applied to a goal, then the result of the evaluation of *expr*₁ is the result of the match expression.

If the matching with *cpattern*₁ fails, or if it succeeds but the evaluation of *expr*₁ fails, or if the evaluation of *expr*₁ succeeds but returns a tactic in execution position whose execution fails, then *cpattern*₂ is used and so on. The pattern _ matches any term and shunts all remaining patterns if any. If all clauses fail (in particular, there is no pattern _) then a no-matching-clause error is raised.

Failures in subsequent tactics do not cause backtracking to select new branches or inside the right-hand side of the selected branch even if it has backtracking points.

Error messages:

1. No matching clauses for match
No pattern can be used and, in particular, there is no _ pattern.
2. Argument of match does not evaluate to a term
This happens when *expr* does not denote a term.

Variants:

1. Using `multimatch` instead of `match` will allow subsequent tactics to backtrack into a right-hand side tactic which has backtracking points left and trigger the selection of a new matching branch when all the backtracking points of the right-hand side have been consumed.

The syntax `match ...` is, in fact, a shorthand for `once multimatch`

2. Using `lazymatch` instead of `match` will perform the same pattern matching procedure but will commit to the first matching branch rather than trying a new matching if the right-hand side fails. If the right-hand side of the selected branch is a tactic with backtracking points, then subsequent failures cause this tactic to backtrack.
3. There is a special form of patterns to match a subterm against the pattern:

```
context ident [ cpattern ]
```

It matches any term with a subterm matching *cpattern*. If there is a match, the optional *ident* is assigned the “matched context”, i.e. the initial term where the matched subterm is replaced by a hole. The example below will show how to use such term contexts.

If the evaluation of the right-hand-side of a valid match fails, the next matching subterm is tried. If no further subterm matches, the next clause is tried. Matching subterms are considered top-bottom and from left to right (with respect to the raw printing obtained by setting option `Printing All`, see Section 2.9).

```
Coq < Ltac f x :=
  match x with
  | context f [S ?X] =>
    idtac X; (* To display the evaluation order *)
    assert (p := eq_refl 1 : X=1); (* To filter the case X=1 *)
    let x:= context f[0] in assert (x=0) (* To observe the context *)
  end.
f is defined

Coq < Goal True.
1 subgoal

=====
True

Coq < f (3+4) .
2
1
2 subgoals

p : 1 = 1
=====
1 + 4 = 0
subgoal 2 is:
True
```

4. For historical reasons, `context` used to consider n -ary applications such as $(f\ 1\ 2)$ as a whole, and not as a sequence of unary applications $((f\ 1)\ 2)$. Hence `context [f ?x]` would fail to find a matching subterm in $(f\ 1\ 2)$: if the pattern was a partial application, the matched subterms would have necessarily been applications with exactly the same number of arguments. As a workaround, one could use the following variant of `context`:

```
appcontext ident [ cpattern ]
```

This syntax is now deprecated, as `context` behaves as intended. The former behavior can be retrieved with the `Tactic Compat Context` flag.

Pattern matching on goals

We can make pattern matching on goals using the following expression:

```
match goal with
|  $hyp_{1,1}, \dots, hyp_{1,m_1}$  |  $-cpattern_1 \Rightarrow expr_1$ 
|  $hyp_{2,1}, \dots, hyp_{2,m_2}$  |  $-cpattern_2 \Rightarrow expr_2$ 
...
|  $hyp_{n,1}, \dots, hyp_{n,m_n}$  |  $-cpattern_n \Rightarrow expr_n$ 
| _  $\Rightarrow expr_{n+1}$ 
end
```

If each hypothesis pattern $hyp_{1,i}$, with $i = 1, \dots, m_1$ is matched (non-linear first-order unification) by an hypothesis of the goal and if $cpattern_1$ is matched by the conclusion of the goal, then $expr_1$ is evaluated to v_1 by substituting the pattern matching to the metavariables and the real hypothesis names bound to the possible hypothesis names occurring in the hypothesis patterns. If v_1 is a tactic value, then it is applied to the goal. If this application fails, then another combination of hypotheses is tried with the same proof context pattern. If there is no other combination of hypotheses then the second proof context pattern is tried and so on. If the next to last proof context pattern fails then $expr_{n+1}$ is evaluated to v_{n+1} and v_{n+1} is applied. Note also that matching against subterms (using the `context ident [cpattern]`) is available and is also subject to yielding several matchings.

Failures in subsequent tactics do not cause backtracking to select new branches or combinations of hypotheses, or inside the right-hand side of the selected branch even if it has backtracking points.

Error message: No matching clauses for match goal

No clause succeeds, i.e. all matching patterns, if any, fail at the application of the right-hand-side.

It is important to know that each hypothesis of the goal can be matched by at most one hypothesis pattern. The order of matching is the following: hypothesis patterns are examined from the right to the left (i.e. hyp_{i,m_i} before $hyp_{i,1}$). For each hypothesis pattern, the goal hypothesis are matched in order (fresher hypothesis first), but it possible to reverse this order (older first) with the `match reverse goal with variant`.

Variant:

Using `multimatch` instead of `match` will allow subsequent tactics to backtrack into a right-hand side tactic which has backtracking points left and trigger the selection of a new matching branch or combination of hypotheses when all the backtracking points of the right-hand side have been consumed.

The syntax `match [reverse] goal ...` is, in fact, a shorthand for `once multimatch [reverse] goal`

Using `lazymatch` instead of `match` will perform the same pattern matching procedure but will commit to the first matching branch with the first matching combination of hypotheses rather than trying a new matching if the right-hand side fails. If the right-hand side of the selected branch is a tactic with backtracking points, then subsequent failures cause this tactic to backtrack.

Filling a term context

The following expression is not a tactic in the sense that it does not produce subgoals but generates a term to be used in tactic expressions:

```
context ident [ expr ]
```

ident must denote a context variable bound by a `context` pattern of a `match` expression. This expression evaluates replaces the hole of the value of *ident* by the value of *expr*.

Error message: `not a context variable`

Generating fresh hypothesis names

Tactics sometimes have to generate new names for hypothesis. Letting the system decide a name with the `intro` tactic is not so good since it is very awkward to retrieve the name the system gave. The following expression returns an identifier:

```
fresh component ... component
```

It evaluates to an identifier unbound in the goal. This fresh identifier is obtained by concatenating the value of the *component*'s (each of them is, either an *qualid* which has to refer to a (unqualified) name, or directly a name denoted by a *string*). If the resulting name is already used, it is padded with a number so that it becomes fresh. If no component is given, the name is a fresh derivative of the name `H`.

Computing in a `constr`

Evaluation of a term can be performed with:

```
eval redexpr in term
```

where *redexpr* is a reduction tactic among `red`, `hnf`, `compute`, `simpl`, `cbv`, `lazy`, `unfold`, `fold`, `pattern`.

Recovering the type of a term

The following returns the type of *term*:

```
type of term
```

Manipulating untyped terms

The terms built in Ltac are well-typed by default. It may not be appropriate for building large terms using a recursive Ltac function: the term has to be entirely type checked at each step, resulting in potentially very slow behavior. It is possible to build untyped terms using Ltac with the syntax

```
uconstr : term
```

An untyped term, in Ltac, can contain references to hypotheses or to Ltac variables containing typed or untyped terms. An untyped term can be type-checked using the function `type_term` whose argument is parsed as an untyped term and returns a well-typed term which can be used in tactics.

```
type_term term
```

Untyped terms built using `uconstr` : can also be used as arguments to the `refine` tactic 8.2.3. In that case the untyped term is type checked against the conclusion of the goal, and the holes which are not solved by the typing procedure are turned into new subgoals.

Counting the goals

The number of goals under focus can be recovered using the `numgoals` function. Combined with the `guard` command below, it can be used to branch over the number of goals produced by previous tactics.

```
Coq < Ltac pr_numgoals := let n := numgoals in idtac "There are" n "goals".
Coq < Goal True /\ True /\ True.
Coq < split;[|split|.
Coq < all:pr_numgoals.
There are 3 goals
3 subgoals

=====
True
subgoal 2 is:
True
subgoal 3 is:
True
```

Testing boolean expressions

The `guard` tactic tests a boolean expression, and fails if the expression evaluates to false. If the expression evaluates to true, it succeeds without affecting the proof.

```
guard test
```

The accepted tests are simple integer comparisons.

```
Coq < Goal True /\ True /\ True.
Coq < split;[|split|.
Coq < all:let n:= numgoals in guard n<4.
3 subgoals

=====
True
subgoal 2 is:
True
subgoal 3 is:
True

Coq < Fail all:let n:= numgoals in guard n=2.
The command has indeed failed with message:
Ltac call to "guard (test)" failed.
Condition not satisfied: 3=2
3 subgoals

=====
True
subgoal 2 is:
True
subgoal 3 is:
True
```

Error messages:

1. Condition not satisfied

Proving a subgoal as a separate lemma

From the outside “`abstract expr`” is the same as `solve expr`. Internally it saves an auxiliary lemma called `ident_subproof n` where `ident` is the name of the current goal and n is chosen so that this is a fresh name. Such auxiliary lemma is inlined in the final proof term unless the proof is ended with “`Qed exporting`”. In such case the lemma is preserved. The syntax “`Qed exporting ident $_1$, ..., ident $_n$` ” is also supported. In such case the system checks that the names given by the user actually exist when the proof is ended.

This tactical is useful with tactics such as `omega` or `discriminate` that generate huge proof terms. With that tool the user can avoid the explosion at time of the `Save` command without having to cut manually the proof in smaller lemmas.

It may be useful to generate lemmas minimal w.r.t. the assumptions they depend on. This can be obtained thanks to the option below.

```
Set Shrink Abstract
```

Deprecated since 8.7

When `set` (default), all lemmas generated through `abstract expr` and `transparent_abstract expr` are quantified only over the variables that appear in the term constructed by `expr`.

Variants:

1. `abstract expr` using `ident`.
Give explicitly the name of the auxiliary lemma. Use this feature at your own risk; explicitly named and reused subterms don’t play well with asynchronous proofs.
2. `transparent_abstract expr`.
Save the subproof in a transparent lemma rather than an opaque one. Use this feature at your own risk; building computationally relevant terms with tactics is fragile.
3. `transparent_abstract expr` using `ident`.
Give explicitly the name of the auxiliary transparent lemma. Use this feature at your own risk; building computationally relevant terms with tactics is fragile, and explicitly named and reused subterms don’t play well with asynchronous proofs.

Error message: Proof is not complete

9.3 Tactic toplevel definitions

9.3.1 Defining \mathcal{L}_{tac} functions

Basically, \mathcal{L}_{tac} toplevel definitions are made as follows:

```
Ltac ident ident $_1$  ... ident $_n$  := expr
```

This defines a new \mathcal{L}_{tac} function that can be used in any tactic script or new \mathcal{L}_{tac} toplevel definition.

Remark: The preceding definition can equivalently be written:

```
Ltac ident := fun ident1 ... identn => expr
```

Recursive and mutual recursive function definitions are also possible with the syntax:

```
Ltac ident1 ident1,1 ... ident1,m1 := expr1
with ident2 ident2,1 ... ident2,m2 := expr2
...
with identn identn,1 ... identn,mn := exprn
```

It is also possible to *redefine* an existing user-defined tactic using the syntax:

```
Ltac qualid ident1 ... identn := expr
```

A previous definition of *qualid* must exist in the environment. The new definition will always be used instead of the old one and it goes across module boundaries.

If preceded by the keyword `Local` the tactic definition will not be exported outside the current module.

9.3.2 Printing \mathcal{L}_{tac} tactics

Defined \mathcal{L}_{tac} functions can be displayed using the command

```
Print Ltac qualid.
```

The command `Print Ltac Signatures` displays a list of all user-defined tactics, with their arguments.

9.4 Debugging \mathcal{L}_{tac} tactics

9.4.1 Info trace

It is possible to print the trace of the path eventually taken by an \mathcal{L}_{tac} script. That is, the list of executed tactics, discarding all the branches which have failed. To that end the `Info` command can be used with the following syntax.

```
Info num expr.
```

The number *num* is the unfolding level of tactics in the trace. At level 0, the trace contains a sequence of tactics in the actual script, at level 1, the trace will be the concatenation of the traces of these tactics, etc...

```
Coq < Ltac t x := exists x; reflexivity.
Coq < Goal exists n, n=0.
Coq < Info 0 t 1||t 0.
t <constr:(0)>
No more subgoals.
```

```
Coq < Undo.
Coq < Info 1 t 1||t 0.
exists 0; reflexivity
No more subgoals.
```

The trace produced by `Info` tries its best to be a reparsable \mathcal{L}_{tac} script, but this goal is not achievable in all generality. So some of the output traces will contain oddities.

As an additional help for debugging, the trace produced by `Info` contains (in comments) the messages produced by the `idtac` tacticals 9.2 at the right position in the script. In particular, the calls to `idtac` in branches which failed are not printed.

An alternative to the `Info` command is to use the `Info Level` option as follows:

```
Set Info Level num.
```

This will automatically print the same trace as `Info num` at each tactic call. The unfolding level can be overridden by a call to the `Info` command. And this option can be turned off with:

```
Unset Info Level num.
```

The current value for the `Info Level` option can be checked using the `Test Info Level` command.

9.4.2 Interactive debugger

The \mathcal{L}_{tac} interpreter comes with a step-by-step debugger. The debugger can be activated using the command

```
Set Ltac Debug.
```

and deactivated using the command

```
Unset Ltac Debug.
```

To know if the debugger is on, use the command `Test Ltac Debug`. When the debugger is activated, it stops at every step of the evaluation of the current \mathcal{L}_{tac} expression and it prints information on what it is doing. The debugger stops, prompting for a command which can be one of the following:

simple newline:	go to the next step
h:	get help
x:	exit current evaluation
s:	continue current evaluation without stopping
r n:	advance n steps further
r string:	advance up to the next call to “ <code>idtac string</code> ”

A non-interactive mode for the debugger is available via the command

```
Set Ltac Batch Debug.
```

This option has the effect of presenting a newline at every prompt, when the debugger is on. The debug log thus created, which does not require user input to generate when this option is set, can then be run through external tools such as `diff`.

9.4.3 Profiling \mathcal{L}_{tac} tactics

It is possible to measure the time spent in invocations of primitive tactics as well as tactics defined in \mathcal{L}_{tac} and their inner invocations. The primary use is the development of complex tactics, which can sometimes be so slow as to impede interactive usage. The reasons for the performance degradation can be intricate, like a slowly performing \mathcal{L}_{tac} match or a sub-tactic whose performance only degrades in certain situations. The profiler generates a call tree and indicates the time spent in a tactic depending its calling context. Thus it allows to locate the part of a tactic definition that contains the performance bug.

```
Set Ltac Profiling.
```

Enables the profiler

```
Unset Ltac Profiling.
```

Disables the profiler

```
Show Ltac Profile.
```

Prints the profile

```
Show Ltac Profile string.
```

Prints a profile for all tactics that start with *string*. Append a period (.) to the string if you only want exactly that name.

```
Reset Ltac Profile.
```

Resets the profile, that is, deletes all accumulated information. Note that backtracking across a Reset Ltac Profile will not restore the information.

```
Coq < Require Import Coq.omega.Omega.
```

```
Coq < Ltac mytauto := tauto.
```

```
Coq < Ltac tac := intros; repeat split; omega || mytauto.
```

```
Coq < Notation max x y := (x + (y - x)) (only parsing).
```

```
Coq < Goal forall x y z A B C D E F G H I J K L M N O P Q R S T U V W X Y Z,
      max x (max y z) = max (max x y) z /\ max x (max y z) = max (max x y) z
      /\ (A /\ B /\ C /\ D /\ E /\ F /\ G /\ H /\ I /\ J /\ K /\ L /\ M /\ N /\ O /\
      -> Z /\ Y /\ X /\ W /\ V /\ U /\ T /\ S /\ R /\ Q /\ P /\ O /\ N /\ M /\ I
```

```
Coq < Proof.
```

```
Coq < Set Ltac Profiling.
```

```
Coq < tac.
```

```
No more subgoals.
```

```
Coq < Show Ltac Profile.
```

```
total time:      6.371s
```

tactic		local	total	calls	max
-tac -----	0.1% 100.0%	1	6.371s		

```

-<Coq.Init.Tauto.with_uniform_flags> --    0.0%  65.8%    26    0.278s
-<Coq.Init.Tauto.tauto_gen> -----    0.0%  65.8%    26    0.278s
-<Coq.Init.Tauto.tauto_intuitionistic> -    0.0%  65.7%    26    0.278s
-t_tauto_intuit -----    1.7%  65.7%    26    0.278s
-<Coq.Init.Tauto.simplif> -----    43.0%  63.5%    26    0.272s
-omega -----    33.7%  33.7%    28    0.956s
-<Coq.Init.Tauto.is_conj> -----    11.5%  11.5%  28756    0.034s
-elim id -----    5.4%   5.4%   650    0.001s
tactic                                local  total  calls      max

```

```

-tac -----    0.1% 100.0%     1    6.371s
|  -<Coq.Init.Tauto.with_uniform_flags> -    0.0%  65.8%    26    0.278s
|  |<Coq.Init.Tauto.tauto_gen> -----    0.0%  65.8%    26    0.278s
|  |<Coq.Init.Tauto.tauto_intuitionistic>    0.0%  65.7%    26    0.278s
|  |t_tauto_intuit -----    1.7%  65.7%    26    0.278s
|  |<Coq.Init.Tauto.simplif> -----    43.0%  63.5%    26    0.272s
|  |  |<Coq.Init.Tauto.is_conj> -----    11.5%  11.5%  28756    0.034s
|  |  |elim id -----    5.4%   5.4%   650    0.001s
|  |omega -----    33.7%  33.7%    28    0.956s

```

Coq < Show Ltac Profile "omega".

total time: 6.371s

```

tactic                                local  total  calls      max
-----
-omega -----    33.7%  33.7%    28    0.956s
tactic                                local  total  calls      max
-----

```

Coq < Abort.

Coq < Unset Ltac Profiling.

The following two tactics behave like `idtac` but enable and disable the profiling. They allow you to exclude parts of a proof script from profiling.

```
start ltac profiling.
```

```
stop ltac profiling.
```

You can also pass the `-profile-ltac` command line option to `coqc`, which performs a `Set Ltac Profiling` at the beginning of each document, and a `Show Ltac Profile` at the end.

Note that the profiler currently does not handle backtracking into multi-success tactics, and issues a warning to this effect in many cases when such backtracking occurs.

Chapter 10

Detailed examples of tactics

This chapter presents detailed examples of certain tactics, to illustrate their behavior.

10.1 dependent induction

The tactics `dependent induction` and `dependent destruction` are another solution for inverting inductive predicate instances and potentially doing induction at the same time. It is based on the `BasicElim` tactic of Conor McBride which works by abstracting each argument of an inductive instance by a variable and constraining it by equalities afterwards. This way, the usual `induction` and `destruct` tactics can be applied to the abstracted instance and after simplification of the equalities we get the expected goals.

The abstracting tactic is called `generalize_eqs` and it takes as argument an hypothesis to generalize. It uses the `JMeq` datatype defined in `Coq.Logic.JMeq`, hence we need to require it before. For example, revisiting the first example of the inversion documentation above:

```
Coq < Require Import Coq.Logic.JMeq.

Coq < Goal forall n m:nat, Le (S n) m -> P n m.
Coq < intros n m H.
Coq < generalize_eqs H.
1 subgoal

  n, m, gen_x : nat
  H : Le gen_x m
  =====
  gen_x = S n -> P n m
```

The index `S n` gets abstracted by a variable here, but a corresponding equality is added under the abstract instance so that no information is actually lost. The goal is now almost amenable to do induction or case analysis. One should indeed first move `n` into the goal to strengthen it before doing induction, or `n` will be fixed in the inductive hypotheses (this does not matter for case analysis). As a rule of thumb, all the variables that appear inside constructors in the indices of the hypothesis should be generalized. This is exactly what the `generalize_eqs_vars` variant does:

```
Coq < generalize_eqs_vars H.
1 subgoal
```

```

m, gen_x : nat
H : Le gen_x m
=====
forall n : nat, gen_x = S n -> P n m
Coq < induction H.
2 subgoals

n : nat
=====
forall n0 : nat, 0 = S n0 -> P n0 n
subgoal 2 is:
forall n0 : nat, S n = S n0 -> P n0 (S m)

```

As the hypothesis itself did not appear in the goal, we did not need to use an heterogeneous equality to relate the new hypothesis to the old one (which just disappeared here). However, the tactic works just as well in this case, e.g.:

```

Coq < Goal forall n m (p : Le (S n) m), Q (S n) m p.
1 subgoal

=====
forall (n m : nat) (p : Le (S n) m), Q (S n) m p
Coq < intros n m p ; generalize_eqs_vars p.
1 subgoal

m, gen_x : nat
p : Le gen_x m
=====
forall (n : nat) (p0 : Le (S n) m),
gen_x = S n -> p ~ = p0 -> Q (S n) m p0

```

One drawback of this approach is that in the branches one will have to substitute the equalities back into the instance to get the right assumptions. Sometimes injection of constructors will also be needed to recover the needed equalities. Also, some subgoals should be directly solved because of inconsistent contexts arising from the constraints on indexes. The nice thing is that we can make a tactic based on discriminate, injection and variants of substitution to automatically do such simplifications (which may involve the K axiom). This is what the `simplify_dep_elim` tactic from `Coq.Program.Equality` does. For example, we might simplify the previous goals considerably:

```

Coq < induction p ; simplify_dep_elim.
1 subgoal

n, m : nat
p : Le n m
IHp : forall (n0 : nat) (p0 : Le (S n0) m),
      n = S n0 -> p ~ = p0 -> Q (S n0) m p0
=====
Q (S n) (S m) (LeS n m p)

```

The higher-order tactic `do_depind` defined in `Coq.Program.Equality` takes a tactic and combines the building blocks we have seen with it: generalizing by equalities calling the given tactic with

the generalized induction hypothesis as argument and cleaning the subgoals with respect to equalities. Its most important instantiations are `dependent induction` and `dependent destruction` that do induction or simply case analysis on the generalized hypothesis. For example we can redo what we've done manually with `dependent destruction`:

```
Coq < Require Import Coq.Program.Equality.
Coq < Lemma ex : forall n m:nat, Le (S n) m -> P n m.
Coq < intros n m H.
Coq < dependent destruction H.
1 subgoal

  n, m : nat
  H : Le n m
  =====
  P n (S m)
```

This gives essentially the same result as `inversion`. Now if the destructed hypothesis actually appeared in the goal, the tactic would still be able to invert it, contrary to `dependent inversion`. Consider the following example on vectors:

```
Coq < Require Import Coq.Program.Equality.
Coq < Set Implicit Arguments.
Coq < Variable A : Set.
Coq < Inductive vector : nat -> Type :=
  | vnil : vector 0
  | vcons : A -> forall n, vector n -> vector (S n).
Coq < Goal forall n, forall v : vector (S n),
  exists v' : vector n, exists a : A, v = vcons a v'.
Coq < intros n v.
Coq < dependent destruction v.
1 subgoal

  n : nat
  a : A
  v : vector n
  =====
  exists (v' : vector n) (a0 : A), vcons a v = vcons a0 v'
```

In this case, the `v` variable can be replaced in the goal by the generalized hypothesis only when it has a type of the form `vector (S n)`, that is only in the second case of the `destruct`. The first one is dismissed because `S n <> 0`.

10.1.1 A larger example

Let's see how the technique works with `induction` on inductive predicates on a real example. We will develop an example application to the theory of simply-typed lambda-calculus formalized in a dependently-typed style:

```

Coq < Inductive type : Type :=
  | base : type
  | arrow : type -> type -> type.
Coq < Notation " t -> t' " := (arrow t t') (at level 20, t' at next level).
Coq < Inductive ctx : Type :=
  | empty : ctx
  | snoc : ctx -> type -> ctx.
Coq < Notation " G , tau " := (snoc G tau) (at level 20, tau at next level).
Coq < Fixpoint conc (G D : ctx) : ctx :=
  match D with
  | empty => G
  | snoc D' x => snoc (conc G D') x
  end.
Coq < Notation " G ; D " := (conc G D) (at level 20).
Coq < Inductive term : ctx -> type -> Type :=
  | ax : forall G tau, term (G, tau) tau
  | weak : forall G tau,
    term G tau -> forall tau', term (G, tau') tau
  | abs : forall G tau tau',
    term (G , tau) tau' -> term G (tau -> tau')
  | app : forall G tau tau',
    term G (tau -> tau') -> term G tau -> term G tau'.

```

We have defined types and contexts which are snoc-lists of types. We also have a `conc` operation that concatenates two contexts. The `term` datatype represents in fact the possible typing derivations of the calculus, which are isomorphic to the well-typed terms, hence the name. A term is either an application of:

- the axiom rule to type a reference to the first variable in a context,
- the weakening rule to type an object in a larger context
- the abstraction or lambda rule to type a function
- the application to type an application of a function to an argument

Once we have this datatype we want to do proofs on it, like weakening:

```

Coq < Lemma weakening : forall G D tau, term (G ; D) tau ->
  forall tau', term (G , tau' ; D) tau.

```

The problem here is that we can't just use `induction` on the typing derivation because it will forget about the `G ; D` constraint appearing in the instance. A solution would be to rewrite the goal as:

```

Coq < Lemma weakening' : forall G' tau, term G' tau ->
  forall G D, (G ; D) = G' ->
  forall tau', term (G, tau' ; D) tau.

```

With this proper separation of the index from the instance and the right induction loading (putting `G` and `D` after the inducted-on hypothesis), the proof will go through, but it is a very tedious process. One is also forced to make a wrapper lemma to get back the more natural statement. The `dependent induction` tactic alleviates this trouble by doing all of this plumbing of generalizing and substituting back automatically. Indeed we can simply write:

```

Coq < Require Import Coq.Program.Tactics.
Coq < Lemma weakening : forall G D tau, term (G ; D) tau ->
    forall tau', term (G , tau' ; D) tau.
Coq < Proof with simpl in * ; simpl_depind ; auto.
Coq <   intros G D tau H. dependent induction H generalizing G D ; intros.

```

This call to `dependent induction` has an additional arguments which is a list of variables appearing in the instance that should be generalized in the goal, so that they can vary in the induction hypotheses. By default, all variables appearing inside constructors (except in a parameter position) of the instantiated hypothesis will be generalized automatically but one can always give the list explicitly.

```

Coq <   Show.
4 subgoals

  G0 : ctx
  tau : type
  G, D : ctx
  x : G0, tau = G; D
  tau' : type
  =====
  term ((G, tau'); D) tau
subgoal 2 is:
  term ((G, tau'0); D) tau
subgoal 3 is:
  term ((G, tau'0); D) (tau -> tau')
subgoal 4 is:
  term ((G, tau'0); D) tau'

```

The `simpl_depind` tactic includes an automatic tactic that tries to simplify equalities appearing at the beginning of induction hypotheses, generally using trivial applications of reflexivity. In cases where the equality is not between constructor forms though, one must help the automation by giving some arguments, using the `specialize` tactic for example.

```

Coq < destruct D... apply weak ; apply ax. apply ax.
Coq < destruct D...
Coq < Show.
4 subgoals

  G0 : ctx
  tau : type
  H : term G0 tau
  tau' : type
  IHterm : forall G D : ctx,
    G0 = G; D -> forall tau' : type, term ((G, tau'); D) tau
  tau'0 : type
  =====
  term ((G0, tau'), tau'0) tau
subgoal 2 is:
  term ((G, tau'0); D), t) tau
subgoal 3 is:

```

```

term ((G, tau'0); D) (tau -> tau')
subgoal 4 is:
term ((G, tau'0); D) tau'
Coq < specialize (IHterm G0 empty eq_refl).
4 subgoals

G0 : ctx
tau : type
H : term G0 tau
tau' : type
IHterm : forall tau' : type, term ((G0, tau')); empty) tau
tau'0 : type
=====
term ((G0, tau'), tau'0) tau
subgoal 2 is:
term (((G, tau'0); D), t) tau
subgoal 3 is:
term ((G, tau'0); D) (tau -> tau')
subgoal 4 is:
term ((G, tau'0); D) tau'

```

Once the induction hypothesis has been narrowed to the right equality, it can be used directly.

```

Coq < apply weak, IHterm.
3 subgoals

tau : type
G, D : ctx
IHterm : forall G0 D0 : ctx,
          G; D = G0; D0 ->
          forall tau' : type, term ((G0, tau')); D0) tau
H : term (G; D) tau
t, tau'0 : type
=====
term (((G, tau'0); D), t) tau
subgoal 2 is:
term ((G, tau'0); D) (tau -> tau')
subgoal 3 is:
term ((G, tau'0); D) tau'

```

If there is an easy first-order solution to these equations as in this subgoal, the `specialize_eqs` tactic can be used instead of giving explicit proof terms:

```

Coq < specialize_eqs IHterm.
Toplevel input, characters 2-23:
> specialize_eqs IHterm.
> ^^^^^^^^^^^^^^^^^^^^^^^^^^^
Error:
Ltac call to "specialize_eqs (var)" failed.
Specialization not allowed on dependent hypotheses

```

This concludes our example. **See also:** The induction [9](#), case [9](#) and inversion [8.14](#) tactics.

10.2 autorewrite

Here are two examples of `autorewrite` use. The first one (*Ackermann function*) shows actually a quite basic use where there is no conditional rewriting. The second one (*Mac Carthy function*) involves conditional rewritings and shows how to deal with them using the optional tactic of the `Hint Rewrite` command.

Example 1: Ackermann function

```
Coq < Reset Initial.
Coq < Require Import Arith.
Coq < Variable Ack :
      nat -> nat -> nat.
Coq < Axiom Ack0 :
      forall m:nat, Ack 0 m = S m.
Coq < Axiom Ack1 : forall n:nat, Ack (S n) 0 = Ack n 1.
Coq < Axiom Ack2 : forall n m:nat, Ack (S n) (S m) = Ack n (Ack (S n) m).

Coq < Hint Rewrite Ack0 Ack1 Ack2 : base0.
Coq < Lemma ResAck0 :
      Ack 3 2 = 29.
1 subgoal

=====
Ack 3 2 = 29

Coq < autorewrite with base0 using try reflexivity.
No more subgoals.
```

Example 2: Mac Carthy function

```
Coq < Require Import Omega.
Coq < Variable g :
      nat -> nat -> nat.
Coq < Axiom g0 :
      forall m:nat, g 0 m = m.
Coq < Axiom
      g1 :
      forall n m:nat,
        (n > 0) -> (m > 100) -> g n m = g (pred n) (m - 10).
Coq < Axiom
      g2 :
      forall n m:nat,
        (n > 0) -> (m <= 100) -> g n m = g (S n) (m + 11).

Coq < Hint Rewrite g0 g1 g2 using omega : base1.
Coq < Lemma Resg0 :
      g 1 110 = 100.
```

```

1 subgoal

=====
g 1 110 = 100

Coq < autorewrite with base1 using reflexivity || simpl.
No more subgoals.

Coq < Lemma Resg1 : g 1 95 = 91.
1 subgoal

=====
g 1 95 = 91

Coq < autorewrite with base1 using reflexivity || simpl.
No more subgoals.

```

10.3 quote

The tactic `quote` allows using Barendregt's so-called 2-level approach without writing any ML code. Suppose you have a language L of 'abstract terms' and a type A of 'concrete terms' and a function $f : L \rightarrow A$. If L is a simple inductive datatype and f a simple fixpoint, `quote f` will replace the head of current goal by a convertible term of the form $(f \ t)$. L must have a constructor of type: $A \rightarrow L$.

Here is an example:

```

Coq < Require Import Quote.

Coq < Parameters A B C : Prop.
A is declared
B is declared
C is declared

Coq < Inductive formula : Type :=
  | f_and : formula -> formula -> formula (* binary constructor *)
  | f_or : formula -> formula -> formula
  | f_not : formula -> formula (* unary constructor *)
  | f_true : formula (* 0-ary constructor *)
  | f_const : Prop -> formula (* constructor for constants *).
formula is defined
formula_rect is defined
formula_ind is defined
formula_rec is defined

Coq < Fixpoint interp_f (f:
                                formula) : Prop :=
  match f with
  | f_and f1 f2 => interp_f f1 /\ interp_f f2
  | f_or f1 f2 => interp_f f1 \/ interp_f f2
  | f_not f1 => ~ interp_f f1
  | f_true => True
  | f_const c => c
  end.
interp_f is defined
interp_f is recursively defined (decreasing on 1st argument)

```



```

Coq < Goal A /\ (A \/ True) /\ ~ B /\ (A <-> A) .
1 subgoal

=====
A /\ (A \/ True) /\ ~ B /\ (A <-> A)

Coq < quote interp_f.
1 subgoal

=====
interp_f
  (f_and (f_const A)
    (f_and (f_or (f_const A) f_true)
      (f_and (f_not (f_const B)) (f_const (A <-> A))))))

```

The algorithm to perform this inversion is: try to match the term with right-hand sides expression of `f`. If there is a match, apply the corresponding left-hand side and call yourself recursively on sub-terms. If there is no match, we are at a leaf: return the corresponding constructor (here `f_const`) applied to the term.

Error messages:

1. quote: not a simple fixpoint
Happens when quote is not able to perform inversion properly.

10.3.1 Introducing variables map

The normal use of `quote` is to make proofs by reflection: one defines a function `simplify : formula -> formula` and proves a theorem `simplify_ok: (f:formula) (interp_f (simplify f)) -> (interp_f f)`. Then, one can simplify formulas by doing:

```

quote interp_f.
apply simplify_ok.
compute.

```

But there is a problem with leafs: in the example above one cannot write a function that implements, for example, the logical simplifications $A \wedge A \rightarrow A$ or $A \wedge \neg A \rightarrow \text{False}$. This is because the `Prop` is impredicative.

It is better to use that type of formulas:

```

Coq < Inductive formula : Set :=
  | f_and : formula -> formula -> formula
  | f_or  : formula -> formula -> formula
  | f_not : formula -> formula
  | f_true : formula
  | f_atom : index -> formula.
formula is defined
formula_rect is defined
formula_ind is defined
formula_rec is defined

```

`index` is defined in module `quote`. Equality on that type is decidable so we are able to simplify $A \wedge A$ into A at the abstract level.

When there are variables, there are bindings, and `quote` provides also a type `(varmap A)` of bindings from index to any set `A`, and a function `varmap_find` to search in such maps. The interpretation function has now another argument, a variables map:

```
Coq < Fixpoint interp_f (vm:
                                varmap Prop) (f:formula) {struct f} : Prop :=
  match f with
  | f_and f1 f2 => interp_f vm f1 /\ interp_f vm f2
  | f_or f1 f2  => interp_f vm f1 \/ interp_f vm f2
  | f_not f1    => ~ interp_f vm f1
  | f_true      => True
  | f_atom i    => varmap_find True i vm
  end.
interp_f is defined
interp_f is recursively defined (decreasing on 2nd argument)
```

`quote` handles this second case properly:

```
Coq < Goal A /\ (B \/ A) /\ (A \/ ~ B).
1 subgoal

=====
A /\ (B \/ A) /\ (A \/ ~ B)

Coq < quote interp_f.
1 subgoal

=====
interp_f
(Node_vm B (Node_vm A (Empty_vm Prop) (Empty_vm Prop)) (Empty_vm Prop))
(f_and (f_atom (Left_idx End_idx))
  (f_and (f_or (f_atom End_idx) (f_atom (Left_idx End_idx)))
    (f_or (f_atom (Left_idx End_idx)) (f_not (f_atom End_idx)))))
```

It builds `vm` and `t` such that `(f vm t)` is convertible with the conclusion of current goal.

10.3.2 Combining variables and constants

One can have both variables and constants in abstracts terms; that is the case, for example, for the `ring` tactic (chapter 25). Then one must provide to `quote` a list of *constructors of constants*. For example, if the list is `[O S]` then closed natural numbers will be considered as constants and other terms as variables.

Example:

```
Coq < Inductive formula : Type :=
  | f_and : formula -> formula -> formula
  | f_or  : formula -> formula -> formula
  | f_not : formula -> formula
  | f_true : formula
  | f_const : Prop -> formula (* constructor for constants *)
  | f_atom : index -> formula.

Coq < Fixpoint interp_f
```

```

      (vm:          (* constructor for variables *)
      varmap Prop) (f:formula) {struct f} : Prop :=
      match f with
      | f_and f1 f2 => interp_f vm f1 /\ interp_f vm f2
      | f_or f1 f2 => interp_f vm f1 \/ interp_f vm f2
      | f_not f1 => ~ interp_f vm f1
      | f_true => True
      | f_const c => c
      | f_atom i => varmap_find True i vm
      end.

Coq < Goal
      A /\ (A \/ True) /\ ~ B /\ (C <-> C).

Coq < quote interp_f [ A B ].
1 subgoal

=====
interp_f (Node_vm (C <-> C) (Empty_vm Prop) (Empty_vm Prop))
  (f_and (f_const A)
    (f_and (f_or (f_const A) f_true)
      (f_and (f_not (f_const B)) (f_atom End_idx))))

Coq < Undo.
1 subgoal

=====
A /\ (A \/ True) /\ ~ B /\ (C <-> C)

Coq < quote interp_f [ B C iff ].
1 subgoal

=====
interp_f (Node_vm A (Empty_vm Prop) (Empty_vm Prop))
  (f_and (f_atom End_idx)
    (f_and (f_or (f_atom End_idx) f_true)
      (f_and (f_not (f_const B)) (f_const (C <-> C)))))

```

Warning: Since function inversion is undecidable in general case, don't expect miracles from it!

Variants:

1. quote *ident* in *term* using *tactic*

tactic must be a functional tactic (starting with `fun x =>`) and will be called with the quoted version of *term* according to *ident*.

2. quote *ident* [*ident*₁ ... *ident*_{*n*}] in *term* using *tactic*

Same as above, but will use *ident*₁, ..., *ident*_{*n*} to chose which subterms are constants (see above).

See also: comments of source file `plugins/quote/quote.ml`

See also: the `ring` tactic (Chapter 25)

```

Coq < Section Sort.
Coq < Variable A : Set.
Coq < Inductive permut : list A -> list A -> Prop :=
  | permut_refl    : forall l, permut l l
  | permut_cons    :
      forall a l0 l1, permut l0 l1 -> permut (a :: l0) (a :: l1)
  | permut_append  : forall a l, permut (a :: l) (l ++ a :: nil)
  | permut_trans   :
      forall l0 l1 l2, permut l0 l1 -> permut l1 l2 -> permut l0 l2.
Coq < End Sort.

```

Figure 10.1: Definition of the permutation predicate

10.4 Using the tactical language

10.4.1 About the cardinality of the set of natural numbers

A first example which shows how to use the pattern matching over the proof contexts is the proof that natural numbers have more than two elements. The proof of such a lemma can be done as follows:

```

Coq < Lemma card_nat :
  ~ (exists x : nat, exists y : nat, forall z:nat, x = z /\ y = z).

Coq < Proof.

Coq < red; intros (x, (y, Hy)).

Coq < elim (Hy 0); elim (Hy 1); elim (Hy 2); intros;
  match goal with
  | [_: (?a = ?b), _ : (?a = ?c) |- _] =>
    cut (b = c); [ discriminate | transitivity a; auto ]
  end.

Coq < Qed.

```

We can notice that all the (very similar) cases coming from the three eliminations (with three distinct natural numbers) are successfully solved by a `match goal` structure and, in particular, with only one pattern (use of non-linear matching).

10.4.2 Permutation on closed lists

Another more complex example is the problem of permutation on closed lists. The aim is to show that a closed list is a permutation of another one.

First, we define the permutation predicate as shown in table 10.1.

A more complex example is the problem of permutation on closed lists. The aim is to show that a closed list is a permutation of another one. First, we define the permutation predicate as shown on Figure 10.1.

Next, we can write naturally the tactic and the result can be seen on Figure 10.2. We can notice that we use two toplevel definitions `PermutProve` and `Permut`. The function to be called is `PermutProve` which computes the lengths of the two lists and calls `Permut` with the length if the two lists have the same length. `Permut` works as expected. If the two lists are equal, it concludes. Otherwise, if the lists have identical first elements, it applies `Permut` on the tail of the lists. Finally, if the

```

Coq < Ltac Permut n :=
  match goal with
  | |- (permut _ ?l ?l) => apply permut_refl
  | |- (permut _ (?a :: ?l1) (?a :: ?l2)) =>
    let newn := eval compute in (length l1) in
    (apply permut_cons; Permut newn)
  | |- (permut ?A (?a :: ?l1) ?l2) =>
    match eval compute in n with
    | 1 => fail
    | _ =>
      let l1' := constr:(l1 ++ a :: nil) in
      (apply (permut_trans A (a :: l1) l1' l2);
       [ apply permut_append | compute; Permut (pred n) ])
    end
  end.
Permut is defined

Coq < Ltac PermutProve :=
  match goal with
  | |- (permut _ ?l1 ?l2) =>
    match eval compute in (length l1 = length l2) with
    | (?n = ?n) => Permut n
    end
  end.
PermutProve is defined

```

Figure 10.2: Permutation tactic

lists have different first elements, it puts the first element of one of the lists (here the second one which appears in the `permut` predicate) at the end if that is possible, i.e., if the new first element has been at this place previously. To verify that all rotations have been done for a list, we use the length of the list as an argument for `Permut` and this length is decremented for each rotation down to, but not including, 1 because for a list of length n , we can make exactly $n - 1$ rotations to generate at most n distinct lists. Here, it must be noticed that we use the natural numbers of COQ for the rotation counter. On Figure 9.1, we can see that it is possible to use usual natural numbers but they are only used as arguments for primitive tactics and they cannot be handled, in particular, we cannot make computations with them. So, a natural choice is to use COQ data structures so that COQ makes the computations (reductions) by `eval compute in` and we can get the terms back by `match`.

With `PermutProve`, we can now prove lemmas as follows:

```

Coq < Lemma permut_ex1 :
  permut nat (1 :: 2 :: 3 :: nil) (3 :: 2 :: 1 :: nil).

Coq < Proof. PermutProve. Qed.

Coq < Lemma permut_ex2 :
  permut nat
    (0 :: 1 :: 2 :: 3 :: 4 :: 5 :: 6 :: 7 :: 8 :: 9 :: nil)
    (0 :: 2 :: 4 :: 6 :: 8 :: 9 :: 7 :: 5 :: 3 :: 1 :: nil).

Coq < Proof. PermutProve. Qed.

```

10.4.3 Deciding intuitionistic propositional logic

The pattern matching on goals allows a complete and so a powerful backtracking when returning tactic values. An interesting application is the problem of deciding intuitionistic propositional logic. Considering the contraction-free sequent calculi LJT^* of Roy Dyckhoff ([56]), it is quite natural to code such a tactic using the tactic language as shown on Figures 10.3 and 10.4. The tactic `Axioms` tries to conclude using usual axioms. The tactic `DSimplif` applies all the reversible rules of Dyckhoff’s system. Finally, the tactic `TautoProp` (the main tactic to be called) simplifies with `DSimplif`, tries to conclude with `Axioms` and tries several paths using the backtracking rules (one of the four Dyckhoff’s rules for the left implication to get rid of the contraction and the right or).

For example, with `TautoProp`, we can prove tautologies like those:

```
Coq < Lemma tauto_ex1 : forall A B:Prop, A /\ B -> A \/ B.
Coq < Proof. TautoProp. Qed.

Coq < Lemma tauto_ex2 :
      forall A B:Prop, (~ ~ B -> B) -> (A -> B) -> ~ ~ A -> B.
Coq < Proof. TautoProp. Qed.
```

10.4.4 Deciding type isomorphisms

A more tricky problem is to decide equalities between types and modulo isomorphisms. Here, we choose to use the isomorphisms of the simply typed λ -calculus with Cartesian product and *unit* type (see, for example, [45]). The axioms of this λ -calculus are given by table 10.5.

A more tricky problem is to decide equalities between types and modulo isomorphisms. Here, we choose to use the isomorphisms of the simply typed λ -calculus with Cartesian product and *unit* type (see, for example, [45]). The axioms of this λ -calculus are given on Figure 10.5.

The tactic to judge equalities modulo this axiomatization can be written as shown on Figures 10.6 and 10.7. The algorithm is quite simple. Types are reduced using axioms that can be oriented (this done by `MainSimplif`). The normal forms are sequences of Cartesian products without Cartesian product in the left component. These normal forms are then compared modulo permutation of the components (this is done by `CompareStruct`). The main tactic to be called and realizing this algorithm is `IsoProve`.

Here are examples of what can be solved by `IsoProve`.

```
Coq < Lemma isos_ex1 :
      forall A B:Set, A * unit * B = B * (unit * A).
Coq < Proof.
Coq < intros; IsoProve.
```

```
Coq < Ltac Axioms :=
      match goal with
      | |- True => trivial
      | _:False |- _ => elimtype False; assumption
      | _:?A |- ?A => auto
      end.
Axioms is defined
```

Figure 10.3: Deciding intuitionistic propositions (1)

```

Coq < Ltac DSimplif :=
  repeat
    (intros;
     match goal with
     | id:(~ _) |- _ => red in id
     | id:(_ /\ _) |- _ =>
       elim id; do 2 intro; clear id
     | id:(_ \/ _) |- _ =>
       elim id; intro; clear id
     | id:(?A /\ ?B -> ?C) |- _ =>
       cut (A -> B -> C);
       [ intro | intros; apply id; split; assumption ]
     | id:(?A \/ ?B -> ?C) |- _ =>
       cut (B -> C);
       [ cut (A -> C);
         [ intros; clear id
           | intro; apply id; left; assumption ]
         | intro; apply id; right; assumption ]
     | id0:(?A -> ?B),id1:?A |- _ =>
       cut B; [ intro; clear id0 | apply id0; assumption ]
     | |- (_ /\ _) => split
     | |- (~ _) => red
     end).
DSimplif is defined

Coq < Ltac TautoProp :=
  DSimplif;
  Axioms ||
  match goal with
  | id:(?A -> ?B) -> ?C |- _ =>
    cut (B -> C);
    [ intro; cut (A -> B);
      [ intro; cut C;
        [ intro; clear id | apply id; assumption ]
        | clear id ]
      | intro; apply id; intro; assumption ]; TautoProp
  | id:(~ ?A -> ?B) |- _ =>
    cut (False -> B);
    [ intro; cut (A -> False);
      [ intro; cut B;
        [ intro; clear id | apply id; assumption ]
        | clear id ]
      | intro; apply id; red; intro; assumption ]; TautoProp
  | |- (_ \/ _) => (left; TautoProp) || (right; TautoProp)
  end.
TautoProp is defined

```

Figure 10.4: Deciding intuitionistic propositions (2)

```
Coq < Qed.
```

```

Coq < Lemma isos_ex2 :
  forall A B C:Set,
    (A * unit -> B * (C * unit)) =

```

```

Coq < Open Scope type_scope.
Coq < Section Iso_axioms.
Coq < Variables A B C : Set.
Coq < Axiom Com : A * B = B * A.
Coq < Axiom Ass : A * (B * C) = A * B * C.
Coq < Axiom Cur : (A * B -> C) = (A -> B -> C).
Coq < Axiom Dis : (A -> B * C) = (A -> B) * (A -> C).
Coq < Axiom P_unit : A * unit = A.
Coq < Axiom AR_unit : (A -> unit) = unit.
Coq < Axiom AL_unit : (unit -> A) = A.
Coq < Lemma Cons : B = C -> A * B = A * C.
Coq < Proof.
Coq < intro Heq; rewrite Heq; reflexivity.
Coq < Qed.
Coq < End Iso_axioms.

```

Figure 10.5: Type isomorphism axioms

```

      (A * unit -> (C -> unit) * C) * (unit -> A -> B).
Coq < Proof.
Coq < intros; IsoProve.
Coq < Qed.

```



```

Coq < Ltac DSimplif trm :=
  match trm with
  | (?A * ?B * ?C) =>
    rewrite <- (Ass A B C); try MainSimplif
  | (?A * ?B -> ?C) =>
    rewrite (Cur A B C); try MainSimplif
  | (?A -> ?B * ?C) =>
    rewrite (Dis A B C); try MainSimplif
  | (?A * unit) =>
    rewrite (P_unit A); try MainSimplif
  | (unit * ?B) =>
    rewrite (Com unit B); try MainSimplif
  | (?A -> unit) =>
    rewrite (AR_unit A); try MainSimplif
  | (unit -> ?B) =>
    rewrite (AL_unit B); try MainSimplif
  | (?A * ?B) =>
    (DSimplif A; try MainSimplif) || (DSimplif B; try MainSimplif)
  | (?A -> ?B) =>
    (DSimplif A; try MainSimplif) || (DSimplif B; try MainSimplif)
  end
with MainSimplif :=
  match goal with
  | |- (?A = ?B) => try DSimplif A; try DSimplif B
  end.
DSimplif is defined
MainSimplif is defined

Coq < Ltac Length trm :=
  match trm with
  | (_ * ?B) => let succ := Length B in constr:(S succ)
  | _ => constr:(1)
  end.
Length is defined

Coq < Ltac assoc := repeat rewrite <- Ass.
assoc is defined

```

Figure 10.6: Type isomorphism tactic (1)

```

Coq < Ltac DoCompare n :=
  match goal with
  | [ |- (?A = ?A) ] => reflexivity
  | [ |- (?A * ?B = ?A * ?C) ] =>
    apply Cons; let newn := Length B in
      DoCompare newn
  | [ |- (?A * ?B = ?C) ] =>
    match eval compute in n with
    | 1 => fail
    | _ =>
      pattern (A * B) at 1; rewrite Com; assoc; DoCompare (pred n)
    end
  end.
DoCompare is defined

Coq < Ltac CompareStruct :=
  match goal with
  | [ |- (?A = ?B) ] =>
    let l1 := Length A
    with l2 := Length B in
      match eval compute in (l1 = l2) with
      | (?n = ?n) => DoCompare n
      end
  end.
CompareStruct is defined

Coq < Ltac IsoProve := MainSimplif; CompareStruct.
IsoProve is defined

```

Figure 10.7: Type isomorphism tactic (2)

Chapter 11

The SSReflect proof language

Georges Gonthier, Assia Mahboubi, Enrico Tassi

11.1 Introduction

This chapter describes a set of tactics known as SSREFLECT originally designed to provide support for the so-called *small scale reflection* proof methodology. Despite the original purpose this set of tactic is of general interest and is available in Coq starting from version 8.7.

SSREFLECT was developed independently of the tactics described in Chapter 8. Indeed the scope of the tactics part of SSREFLECT largely overlaps with the standard set of tactics. Eventually the overlap will be reduced in future releases of Coq.

Proofs written in SSREFLECT typically look quite different from the ones written using only tactics as per Chapter 8. We try to summarise here the most “visible” ones in order to help the reader already accustomed to the tactics described in Chapter 8 to read this chapter.

The first difference between the tactics described in this chapter and the tactics described in Chapter 8 is the way hypotheses are managed (we call this *bookkeeping*). In Chapter 8 the most common approach is to avoid moving explicitly hypotheses back and forth between the context and the conclusion of the goal. On the contrary in SSREFLECT all bookkeeping is performed on the conclusion of the goal, using for that purpose a couple of syntactic constructions behaving similar to tacticals (and often named as such in this chapter). The `:` tactical moves hypotheses from the context to the conclusion, while `=>` moves hypotheses from the conclusion to the context, and `in` moves back and forth an hypothesis from the context to the conclusion for the time of applying an action to it.

While naming hypotheses is commonly done by means of an `as` clause in the basic model of Chapter 8, it is here to `=>` that this task is devoted. As tactics leave new assumptions in the conclusion, and are often followed by `=>` to explicitly name them. While generalizing the goal is normally not explicitly needed in Chapter 8, it is an explicit operation performed by `:`.

Beside the difference of bookkeeping model, this chapter includes specific tactics which have no explicit counterpart in Chapter 8 such as tactics to mix forward steps and generalizations as `generally have` or `without loss`.

SSREFLECT adopts the point of view that rewriting, definition expansion and partial evaluation participate all to a same concept of rewriting a goal in a larger sense. As such, all these functionalities are provided by the `rewrite` tactic.

SSREFLECT includes a little language of patterns to select subterms in tactics or tacticals where it matters. Its most notable application is in the `rewrite` tactic, where patterns are used to specify where the rewriting step has to take place.

Finally, SSREFLECT supports so-called reflection steps, typically allowing to switch back and forth between the computational view and logical view of a concept.

To conclude it is worth mentioning that SSREFLECT tactics can be mixed with non SSREFLECT tactics in the same proof, or in the same Ltac expression. The few exceptions to this statement are described in section 11.2.2.

We follow the default color scheme of the SSREFLECT mode for ProofGeneral provided in the distribution:

`tactic` or `Command` or `keyword` or `tactical`

Closing tactics/tacticals like `exact` or `by` (see section 11.6.2) are in red.

Acknowledgments

The authors would like to thank Frédéric Blanqui, François Pottier and Laurence Rideau for their comments and suggestions.

11.2 Usage

11.2.1 Getting started

To be available, the tactics presented in this manual need the following minimal set of libraries to be loaded: `ssreflect.v`, `ssrfun.v` and `ssrbool.v`. Moreover, these tactics come with a methodology specific to the authors of `Ssreflect` and which requires a few options to be set in a different way than in their default way. All in all, this corresponds to working in the following context:

```
From Coq Require Import ssreflect ssrfun ssrbool.
Set Implicit Arguments.
Unset Strict Implicit.
Unset Printing Implicit Defensive.
```

11.2.2 Compatibility issues

Requiring the above modules creates an environment which is mostly compatible with the rest of COQ, up to a few discrepancies:

- New keywords (`is`) might clash with variable, constant, tactic or tactical names, or with quasi-keywords in tactic or vernacular notations.
- New tactic(al)s names (`last`, `done`, `have`, `suffices`, `suff`, `without loss`, `wlog`, `congr`, `unlock`) might clash with user tactic names.
- Identifiers with both leading and trailing `_`, such as `_x_`, are reserved by `SSREFLECT` and cannot appear in scripts.
- The extensions to the `rewrite` tactic are partly incompatible with those available in current versions of COQ; in particular: `rewrite .. in (type of k)` or `rewrite .. in *` or any other variant of `rewrite` will not work, and the `SSREFLECT` syntax and semantics for occurrence selection and rule chaining is different.
Use an explicit rewrite direction (`rewrite <- ...` or `rewrite -> ...`) to access the COQ `rewrite` tactic.
- New symbols (`///`, `/=`, `/// $=$`) might clash with adjacent existing symbols (e.g., `'/'`) instead of `'/'`). This can be avoided by inserting white spaces.
- New constant and theorem names might clash with the user theory. This can be avoided by not importing all of `SSREFLECT`:

```
From Coq Require ssreflect.
Import ssreflect.SsrSyntax.
```

Note that the full syntax of `SSREFLECT`'s `rewrite` and reserved identifiers are enabled only if the `ssreflect` module has been required and if `SsrSyntax` has been imported. Thus a file that requires (without importing) `ssreflect` and imports `SsrSyntax`, can be required and imported without automatically enabling `SSREFLECT`'s extended rewrite syntax and reserved identifiers.

- Some user notations (in particular, defining an infix `';`) might interfere with the "open term", parenthesis free, syntax of tactics such as `have`, `set` and `pose`.

- The generalization of `if` statements to non-Boolean conditions is turned off by `SSREFLECT`, because it is mostly subsumed by `Coercion` to `bool` of the `sumXXX` types (declared in `ssrfun.v`) and the `if term is pattern then term else term` construct (see 11.3.2). To use the generalized form, turn off the `SSREFLECT` Boolean `if` notation using the command:

```
Close Scope boolean_if_scope.
```

- The following two options can be unset to disable the incompatible `rewrite` syntax and allow reserved identifiers to appear in scripts.

```
Unset SsrRewrite.  
Unset SsrIdents.
```

11.3 Gallina extensions

Small-scale reflection makes an extensive use of the programming subset of Gallina, COQ's logical specification language. This subset is quite suited to the description of functions on representations, because it closely follows the well-established design of the ML programming language. The `SSREFLECT` extension provides three additions to Gallina, for pattern assignment, pattern testing, and polymorphism; these mitigate minor but annoying discrepancies between Gallina and ML.

11.3.1 Pattern assignment

The `SSREFLECT` extension provides the following construct for irrefutable pattern matching, that is, destructuring assignment:

```
let : pattern := term1 in term2
```

Note the colon ':' after the `let` keyword, which avoids any ambiguity with a function definition or COQ's basic destructuring `let`. The `let :` construct differs from the latter in that

- The pattern can be nested (deep pattern matching), in particular, this allows expression of the form:

```
let : exist (x, y) p_xy := Hp in ...
```

- The destructured constructor is explicitly given in the pattern, and is used for type inference, e.g.,

```
Let f u := let : (m, n) := u in m + n.
```

using a colon `let :`, infers `f : nat * nat -> nat`, whereas

```
Let f u := let (m, n) := u in m + n.
```

with a usual `let`, requires an extra type annotation.

The `let :` construct is just (more legible) notation for the primitive Gallina expression

```
match term1 with pattern => term2 end
```

The `SSREFLECT` destructuring assignment supports all the dependent match annotations; the full syntax is

```
let : pattern1 as ident in pattern2 := term1 return term2 in term3
```

where *pattern₂* is a *type* pattern and *term₁* and *term₂* are types.

When the *as* and *return* are both present, then *ident* is bound in both the type *term₂* and the expression *term₃*; variables in the optional type pattern *pattern₂* are bound only in the type *term₂*, and other variables in *pattern₁* are bound only in the expression *term₃*, however.

11.3.2 Pattern conditional

The following construct can be used for a refutable pattern matching, that is, pattern testing:

```
if term1 is pattern1 then term2 else term3
```

Although this construct is not strictly ML (it does exist in variants such as the pattern calculus or the ρ -calculus), it turns out to be very convenient for writing functions on representations, because most such functions manipulate simple datatypes such as Peano integers, options, lists, or binary trees, and the pattern conditional above is almost always the right construct for analyzing such simple types. For example, the *null* and *all* list function(al)s can be defined as follows:

```
Variable d: Set.
Fixpoint null (s : list d) := if s is nil then true else false.
Variable a : d -> bool.
Fixpoint all (s : list d) : bool :=
  if s is cons x s' then a x && all s' else true.
```

The pattern conditional also provides a notation for destructuring assignment with a refutable pattern, adapted to the pure functional setting of Gallina, which lacks a *Match_Failure* exception.

Like *let*: above, the *if...is* construct is just (more legible) notation for the primitive Gallina expression:

```
match term1 with pattern => term2 | _ => term2 end
```

Similarly, it will always be displayed as the expansion of this form in terms of primitive *match* expressions (where the default expression *term₃* may be replicated).

Explicit pattern testing also largely subsumes the generalization of the *if* construct to all binary datatypes; compare:

```
if term is inl _ then terml else termr
```

and:

```
if term then terml else termr
```

The latter appears to be marginally shorter, but it is quite ambiguous, and indeed often requires an explicit annotation *term* : {*_*} + {*_*} to type-check, which evens the character count.

Therefore, SSREFLECT restricts by default the condition of a plain *if* construct to the standard *bool* type; this avoids spurious type annotations, e.g., in:

```
Definition orb b1 b2 := if b1 then true else b2.
```

As pointed out in section 11.2.2, this restriction can be removed with the command:

```
Close Scope boolean_if_scope.
```

Like `let`: above, the `if term is pattern else term` construct supports the dependent `match` annotations:

```
if term1 is pattern1 as ident in pattern2 return term2 then term3 else term4
```

As in `let`: the variable `ident` (and those in the type pattern `pattern2`) are bound in `term2`; `ident` is also bound in `term3` (but not in `term4`), while the variables in `pattern1` are bound only in `term3`. Another variant allows to treat the else case first:

```
if term1 isn't pattern1 then term2 else term3
```

Note that `pattern1` eventually binds variables in `term3` and not `term2`.

11.3.3 Parametric polymorphism

Unlike ML, polymorphism in core Gallina is explicit: the type parameters of polymorphic functions must be declared explicitly, and supplied at each point of use. However, COQ provides two features to suppress redundant parameters:

- Sections are used to provide (possibly implicit) parameters for a set of definitions.
- Implicit arguments declarations are used to tell COQ to use type inference to deduce some parameters from the context at each point of call.

The combination of these features provides a fairly good emulation of ML-style polymorphism, but unfortunately this emulation breaks down for higher-order programming. Implicit arguments are indeed not inferred at all points of use, but only at points of call, leading to expressions such as

```
Definition all_null (s : list T) := all (@null T) s.
```

Unfortunately, such higher-order expressions are quite frequent in representation functions, especially those which use COQ's `Structures` to emulate Haskell type classes.

Therefore, SSREFLECT provides a variant of COQ's implicit argument declaration, which causes COQ to fill in some implicit parameters at each point of use, e.g., the above definition can be written:

```
Definition all_null (s : list d) := all null s.
```

Better yet, it can be omitted entirely, since `all_null s` isn't much of an improvement over `all null s`.

The syntax of the new declaration is

```
Prenex Implicits ident+.
```

Let us denote $c_1 \dots c_n$ the list of identifiers given to a `Prenex Implicits` command. The command checks that each c_i is the name of a functional constant, whose implicit arguments are prenex, i.e., the first $n_i > 0$ arguments of c_i are implicit; then it assigns `Maximal Implicit` status to these arguments.

As these prenex implicit arguments are ubiquitous and have often large display strings, it is strongly recommended to change the default display settings of COQ so that they are not printed (except after a `Set Printing All` command). All SSREFLECT library files thus start with the incantation


```
Set Implicit Arguments.
Unset Strict Implicit.
Unset Printing Implicit Defensive.
```

11.3.4 Anonymous arguments

When in a definition, the type of a certain argument is mandatory, but not its name, one usually use “arrow” abstractions for prenex arguments, or the $(_ : term)$ syntax for inner arguments. In SSREFLECT, the latter can be replaced by the open syntax ‘*of term*’ or (equivalently) ‘*&term*’, which are both syntactically equivalent to a $(_ : term)$ expression.

For instance, the usual two-contructor polymorphic type `list`, i.e. the one of the standard `List` library, can be defined by the following declaration:

```
Inductive list (A : Type) : Type := nil | cons of A & list A.
```

11.3.5 Wildcards

The terms passed as arguments to SSREFLECT tactics can contain *holes*, materialized by wildcards `_`. Since SSREFLECT allows a more powerful form of type inference for these arguments, it enhances the possibilities of using such wildcards. These holes are in particular used as a convenient shorthand for abstractions, especially in local definitions or type expressions.

Wildcards may be interpreted as abstractions (see for example sections 11.4.1 and 11.6.6), or their content can be inferred from the whole context of the goal (see for example section 11.4.2).

11.4 Definitions

11.4.1 Definitions

The `pose` tactic allows to add a defined constant to a proof context. SSREFLECT generalizes this tactic in several ways. In particular, the SSREFLECT `pose` tactic supports *open syntax*: the body of the definition does not need surrounding parentheses. For instance:

```
pose t := x + y.
```

is a valid tactic expression.

The `pose` tactic is also improved for the local definition of higher order terms. Local definitions of functions can use the same syntax as global ones. The tactic:

```
pose f x y := x + y.
```

adds to the context the defined constant:

```
f := fun x y : nat => x + y : nat -> nat -> nat
```

The SSREFLECT `pose` tactic also supports (co)fixpoints, by providing the local counterpart of the `Fixpoint` `f :=...` and `CoFixpoint` `f :=...` constructs. For instance, the following tactic:

```
pose fix f (x y : nat) {struct x} : nat :=
  if x is S p then S (f p y) else 0.
```

defines a local fixpoint `f`, which mimics the standard `plus` operation on natural numbers.

Similarly, local cofixpoints can be defined by a tactic of the form:

```
pose cofix f (arg : T) ...
```

The possibility to include wildcards in the body of the definitions offers a smooth way of defining local abstractions. The type of “holes” is guessed by type inference, and the holes are abstracted. For instance the tactic:

```
pose f := _ + 1.
```

is shorthand for:

```
pose f n := n + 1.
```

When the local definition of a function involves both arguments and holes, hole abstractions appear first. For instance, the tactic:

```
pose f x := x + _.
```

is shorthand for:

```
pose f n x := x + n.
```

The interaction of the `pose` tactic with the interpretation of implicit arguments results in a powerful and concise syntax for local definitions involving dependent types. For instance, the tactic:

```
pose f x y := (x, y).
```

adds to the context the local definition:

```
pose f (Tx Ty : Type) (x : Tx) (y : Ty) := (x, y).
```

The generalization of wildcards makes the use of the `pose` tactic resemble ML-like definitions of polymorphic functions.

11.4.2 Abbreviations

The SSREFLECT `set` tactic performs abbreviations: it introduces a defined constant for a subterm appearing in the goal and/or in the context.

SSREFLECT extends the `set` tactic by supplying:

- an open syntax, similarly to the `pose` tactic;
- a more aggressive matching algorithm;
- an improved interpretation of wildcards, taking advantage of the matching algorithm;
- an improved occurrence selection mechanism allowing to abstract only selected occurrences of a term.

The general syntax of this tactic is

$$\text{set } \textit{ident} [: \textit{term}_1] := [\textit{occ-switch}] \textit{term}_2$$

$$\textit{occ-switch} ::= \{ [+|-] \textit{natural}^* \}$$

where:

- *ident* is a fresh identifier chosen by the user.

- $term_1$ is an optional type annotation. The type annotation $term_1$ can be given in open syntax (no surrounding parentheses). If no *occ-switch* (described hereafter) is present, it is also the case for $term_2$. On the other hand, in presence of *occ-switch*, parentheses surrounding $term_2$ are mandatory.
- In the occurrence switch *occ-switch*, if the first element of the list is a *natural*, this element should be a number, and not an Ltac variable. The empty list $\{ \}$ is not interpreted as a valid occurrence switch.

The tactic:

```
set t := f _.
```

transforms the goal $f\ x + f\ x = f\ x$ into $t + t = t$, adding $t := f\ x$ to the context, and the tactic:

```
set t := {2} (f _).
```

transforms it into $f\ x + t = f\ x$, adding $t := f\ x$ to the context.

The type annotation $term_1$ may contain wildcards, which will be filled with the appropriate value by the matching process.

The tactic first tries to find a subterm of the goal matching $term_2$ (and its type $term_1$), and stops at the first subterm it finds. Then the occurrences of this subterm selected by the optional *occ-switch* are replaced by *ident* and a definition $ident := term$ is added to the context. If no *occ-switch* is present, then all the occurrences are abstracted.

Matching

The matching algorithm compares a pattern $term$ with a subterm of the goal by comparing their heads and then pairwise unifying their arguments (modulo conversion). Head symbols match under the following conditions:

- If the head of $term$ is a constant, then it should be syntactically equal to the head symbol of the subterm.
- If this head is a projection of a canonical structure, then canonical structure equations are used for the matching.
- If the head of $term$ is *not* a constant, the subterm should have the same structure (λ abstraction, `let...in` structure ...).
- If the head of $term$ is a hole, the subterm should have at least as many arguments as $term$. For instance the tactic:

```
set t := _ x.
```

transforms the goal $x + y = z$ into $t\ y = z$ and adds $t := plus\ x : nat \rightarrow nat$ to the context.

- In the special case where $term$ is of the form $(let\ f := t_0\ in\ f)\ t_1 \dots t_n$, then the pattern $term$ is treated as $(_t_1 \dots t_n)$. For each subterm in the goal having the form $(A\ u_1 \dots u_{n'})$ with $n' \geq n$, the matching algorithm successively tries to find the largest partial application $(A\ u_1 \dots u_{i'})$ convertible to the head t_0 of $term$. For instance the following tactic:

```
set t := (let g y z := y.+1 + z in g) 2.
```

transforms the goal

```
(let f x y z := x + y + z in f 1) 2 3 = 6.
```

into $t \ 3 = 6$ and adds the local definition of t to the context.

Moreover:

- Multiple holes in *term* are treated as independent placeholders. For instance, the tactic:

```
set t := _ + _.
```

transforms the goal $x + y = z$ into $t = z$ and pushes $t := x + y : \text{nat}$ in the context.

- The type of the subterm matched should fit the type (possibly casted by some type annotations) of the pattern *term*.
- The replacement of the subterm found by the instantiated pattern should not capture variables, hence the following script:

```
Goal forall x : nat, x + 1 = 0.
set u := _ + 1.
```

raises an error message, since x is bound in the goal.

- Typeclass inference should fill in any residual hole, but matching should never assign a value to a global existential variable.

Occurrence selection

SSREFLECT provides a generic syntax for the selection of occurrences by their position indexes. These *occurrence switches* are shared by all SSREFLECT tactics which require control on subterm selection like rewriting, generalization, ...

An *occurrence switch* can be:

- A list $\{ \text{natural}^* \}$ of occurrences affected by the tactic. For instance, the tactic:

```
set x := {1 3} (f 2).
```

transforms the goal $f \ 2 + f \ 8 = f \ 2 + f \ 2$ into $x + f \ 8 = f \ 2 + x$, and adds the abbreviation $x := f \ 2$ in the context. Notice that some occurrences of a given term may be hidden to the user, for example because of a notation. The vernacular **Set Printing All** command displays all these hidden occurrences and should be used to find the correct coding of the occurrences to be selected¹. For instance, the following script:

```
Notation "a < b" := (le (S a) b).
Goal forall x y, x < y -> S x < S y.
intros x y; set t := S x.
```

generates the goal $t \leq y \rightarrow t < S \ y$ since $x < y$ is now a notation for $S \ x \leq y$.

¹Unfortunately, even after a call to the Set Printing All command, some occurrences are still not displayed to the user, essentially the ones possibly hidden in the predicate of a dependent match structure.

- A list $\{natural^+\}$. This is equivalent to $\{natural^+\}$ but the list should start with a number, and not with an Ltac variable.
- A list $\{natural^*\}$ of occurrences *not* to be affected by the tactic. For instance, the tactic:

```
set x := {-2} (f 2) .
```

behaves like

```
set x := {1 3} (f 2) .
```

on the goal $f\ 2 + f\ 8 = f\ 2 + f\ 2$ which has three occurrences of the term $f\ 2$

- In particular, the switch $\{+\}$ selects *all* the occurrences. This switch is useful to turn off the default behavior of a tactic which automatically clears some assumptions (see section 11.5.3 for instance).
- The switch $\{-\}$ imposes that *no* occurrences of the term should be affected by the tactic. The tactic:

```
set x := {-} (f 2) .
```

leaves the goal unchanged and adds the definition $x := f\ 2$ to the context. This kind of tactic may be used to take advantage of the power of the matching algorithm in a local definition, instead of copying large terms by hand.

It is important to remember that matching *precedes* occurrence selection, hence the tactic:

```
set a := {2} ( _ + _ ) .
```

transforms the goal $x + y = x + y + z$ into $x + y = a + z$ and fails on the goal $(x + y) + (z + z) = z + z$ with the error message:

```
User error: only 1 < 2 occurrence of (x + y + (z + z))
```

11.4.3 Localization

It is possible to define an abbreviation for a term appearing in the context of a goal thanks to the *in* tactical.

A tactic of the form:

```
set x := term in fact1 . . . factn .
```

introduces a defined constant called x in the context, and folds it in the facts $fact_1 \dots fact_n$. The body of x is the first subterm matching *term* in $fact_1 \dots fact_n$.

A tactic of the form:

```
set x := term in fact1 . . . factn * .
```

matches *term* and then folds x similarly in $fact_1 \dots fact_n$, but also folds x in the goal.

A goal $x + t = 4$, whose context contains $Hx : x = 3$, is left unchanged by the tactic:

```
set z := 3 in Hx .
```

but the context is extended with the definition $z := 3$ and Hx becomes $Hx : x = z$. On the same goal and context, the tactic:

```
set z := 3 in Hx *.
```

will moreover change the goal into $x + t = S\ z$. Indeed, remember that 4 is just a notation for $(S\ 3)$.

The use of the `in` tactical is not limited to the localization of abbreviations: for a complete description of the `in` tactical, see section 11.5.1.

11.5 Basic tactics

A sizable fraction of proof scripts consists of steps that do not "prove" anything new, but instead perform menial bookkeeping tasks such as selecting the names of constants and assumptions or splitting conjuncts. Although they are logically trivial, bookkeeping steps are extremely important because they define the structure of the data-flow of a proof script. This is especially true for reflection-based proofs, which often involve large numbers of constants and assumptions. Good bookkeeping consists in always explicitly declaring (i.e., naming) all new constants and assumptions in the script, and systematically pruning irrelevant constants and assumptions in the context. This is essential in the context of an interactive development environment (IDE), because it facilitates navigating the proof, allowing to instantly "jump back" to the point at which a questionable assumption was added, and to find relevant assumptions by browsing the pruned context. While novice or casual COQ users may find the automatic name selection feature convenient, the usage of such a feature severely undermines the readability and maintainability of proof scripts, much like automatic variable declaration in programming languages. The SSREFLECT tactics are therefore designed to support precise bookkeeping and to eliminate name generation heuristics. The bookkeeping features of SSREFLECT are implemented as tacticals (or pseudo-tacticals), shared across most SSREFLECT tactics, and thus form the foundation of the SSREFLECT proof language.

11.5.1 Bookkeeping

During the course of a proof COQ always present the user with a *sequent* whose general form is

$$\begin{array}{c}
 c_i : T_i \\
 \dots \\
 d_j := e_j : T_j \\
 \dots \\
 F_k : P_k \\
 \dots \\
 \hline
 \text{forall } (x_\ell : T_\ell) \dots, \\
 \text{let } y_m := b_m \text{ in } \dots \text{ in} \\
 P_n \rightarrow \dots \rightarrow C
 \end{array}$$

The *goal* to be proved appears below the double line; above the line is the *context* of the sequent, a set of declarations of *constants* c_i , *defined constants* d_i , and *facts* F_k that can be used to prove the goal (usually, $T_i, T_j : \text{Type}$ and $P_k : \text{Prop}$). The various kinds of declarations can come in any order. The top part of the context consists of declarations produced by the `Section` commands `Variable`, `Let`, and `Hypothesis`. This *section context* is never affected by the SSREFLECT tactics: they only operate

on the lower part — the *proof context*. As in the figure above, the goal often decomposes into a series of (universally) quantified *variables* $(x_\ell : T_\ell)$, local *definitions* `let $y_m := b_m$ in`, and *assumptions* $P_n \rightarrow$, and a *conclusion* C (as in the context, variables, definitions, and assumptions can appear in any order). The conclusion is what actually needs to be proved — the rest of the goal can be seen as a part of the proof context that happens to be “below the line”.

However, although they are logically equivalent, there are fundamental differences between constants and facts on the one hand, and variables and assumptions on the others. Constants and facts are *unordered*, but *named* explicitly in the proof text; variables and assumptions are *ordered*, but *unnamed*: the display names of variables may change at any time because of α -conversion.

Similarly, basic deductive steps such as `apply` can only operate on the goal because the Gallina terms that control their action (e.g., the type of the lemma used by `apply`) only provide unnamed bound variables.² Since the proof script can only refer directly to the context, it must constantly shift declarations from the goal to the context and conversely in between deductive steps.

In SSREFLECT these moves are performed by two *tacticals* ‘`=>`’ and ‘`:`’, so that the bookkeeping required by a deductive step can be directly associated to that step, and that tactics in an SSREFLECT script correspond to actual logical steps in the proof rather than merely shuffle facts. Still, some isolated bookkeeping is unavoidable, such as naming variables and assumptions at the beginning of a proof. SSREFLECT provides a specific `move` tactic for this purpose.

Now `move` does essentially nothing: it is mostly a placeholder for ‘`=>`’ and ‘`:`’. The ‘`=>`’ tactical moves variables, local definitions, and assumptions to the context, while the ‘`:`’ tactical moves facts and constants to the goal. For example, the proof of³

```
Lemma subnK : forall m n, n <= m -> m - n + n = m.
```

might start with

```
move=> m n le_n_m.
```

where `move` does nothing, but `=> m n le_n_m` changes the variables and assumption of the goal in the constants `m n : nat` and the fact `le_n_m : n <= m`, thus exposing the conclusion `m - n + n = m`.

The ‘`:`’ tactical is the converse of ‘`=>`’: it removes facts and constants from the context by turning them into variables and assumptions. Thus

```
move: m le_n_m.
```

turns back `m` and `le_n_m` into a variable and an assumption, removing them from the proof context, and changing the goal to

```
forall m, n <= m -> m - n + n = m.
```

which can be proved by induction on `n` using `elim: n`.

Because they are tacticals, ‘`:`’ and ‘`=>`’ can be combined, as in

```
move: m le_n_m => p le_n_p.
```

simultaneously renames `m` and `le_n_m` into `p` and `le_n_p`, respectively, by first turning them into unnamed variables, then turning these variables back into constants and facts.

Furthermore, SSREFLECT redefines the basic COQ tactics `case`, `elim`, and `apply` so that they can take better advantage of ‘`:`’ and ‘`=>`’. In there SSREFLECT variants, these tactic operate on the first

²Thus scripts that depend on bound variable names, e.g., via `intros` or `with`, are inherently fragile.

³The name `subnK` reads as “right cancellation rule for `nat` subtraction”.

variable or constant of the goal and they do not use or change the proof context. The `:` tactical is used to operate on an element in the context. For instance the proof of `subnK` could continue with

```
elim: n.
```

instead of `elim n`; this has the advantage of removing `n` from the context. Better yet, this `elim` can be combined with previous `move` and with the branching version of the `=>` tactical (described in 11.5.4), to encapsulate the inductive step in a single command:

```
elim: n m le_n_m => [|n IHn] m => [_ | lt_n_m].
```

which breaks down the proof into two subgoals,

```
m - 0 + 0 = m
```

given `m : nat`, and

```
m - S n + S n = m
```

given `m n : nat`, `lt_n_m : S n <= m`, and

```
IHn : forall m, n <= m -> m - n + n = m.
```

The `:` and `=>` tacticals can be explained very simply if one views the goal as a stack of variables and assumptions piled on a conclusion:

- *tactic : a b c* pushes the context constants *a*, *b*, *c* as goal variables *before* performing *tactic*.
- *tactic => a b c* pops the top three goal variables as context constants *a*, *b*, *c*, *after* *tactic* has been performed.

These pushes and pops do not need to balance out as in the examples above, so

```
move: m le_n_m => p.
```

would rename `m` into `p`, but leave an extra assumption `n <= p` in the goal.

Basic tactics like `apply` and `elim` can also be used without the `:` tactical: for example we can directly start a proof of `subnK` by induction on the top variable `m` with

```
elim=> [|m IHm] n le_n.
```

The general form of the localization tactical `in` is also best explained in terms of the goal stack:

```
tactic in a H1 H2 *.
```

is basically equivalent to

```
move: a H1 H2; tactic => a H1 H2.
```

with two differences: the `in` tactical will preserve the body of `a` if `a` is a defined constant, and if the `*` is omitted it will use a temporary abbreviation to hide the statement of the goal from *tactic*.

The general form of the `in` tactical can be used directly with the `move`, `case` and `elim` tactics, so that one can write

```
elim: n => [|n IHn] in m le_n_m *.
```

instead of

```
elim: n m le_n_m => [|n IHn] m le_n_m.
```

This is quite useful for inductive proofs that involve many facts.

See section 11.6.5 for the general syntax and presentation of the `in` tactical.

11.5.2 The defective tactics

In this section we briefly present the three basic tactics performing context manipulations and the main backward chaining tool.

The `move` tactic.

The `move` tactic, in its defective form, behaves like the primitive `hnf` COQ tactic. For example, such a defective:

```
move .
```

exposes the first assumption in the goal, i.e. its changes the goal $\sim \text{False}$ into $\text{False} \rightarrow \text{False}$.

More precisely, the `move` tactic inspects the goal and does nothing (`idtac`) if an introduction step is possible, i.e. if the goal is a product or a `let...in`, and performs `hnf` otherwise.

Of course this tactic is most often used in combination with the bookkeeping tacticals (see section 11.5.4 and 11.5.3). These combinations mostly subsume the `intros`, `generalize`, `revert`, `rename`, `clear` and `pattern` tactics.

The `case` tactic.

The `case` tactic performs *primitive case analysis* on (co)inductive types; specifically, it destructs the top variable or assumption of the goal, exposing its constructor(s) and its arguments, as well as setting the value of its type family indices if it belongs to a type family (see section 11.5.6).

The SSREFLECT `case` tactic has a special behavior on equalities. If the top assumption of the goal is an equality, the `case` tactic “destructs” it as a set of equalities between the constructor arguments of its left and right hand sides, as per the tactic `injection`. For example, `case` changes the goal

```
(x, y) = (1, 2) -> G.
```

into

```
x = 1 -> y = 2 -> G.
```

Note also that the case of SSREFLECT performs `False` elimination, even if no branch is generated by this case operation. Hence the command:

```
case .
```

on a goal of the form $\text{False} \rightarrow G$ will succeed and prove the goal.

The `elim` tactic.

The `elim` tactic performs inductive elimination on inductive types. The defective:

```
elim.
```

tactic performs inductive elimination on a goal whose top assumption has an inductive type. For example on goal of the form:

```
forall n : nat, m <= n
```

in a context containing `m : nat`, the

```
elim.
```

tactic produces two goals,

```
m <= 0
```

on one hand and

```
forall n : nat, m <= n -> m <= S n
```

on the other hand.

The **apply** tactic.

The **apply** tactic is the main backward chaining tactic of the proof system. It takes as argument any *term* and applies it to the goal. Assumptions in the type of *term* that don't directly match the goal may generate one or more subgoals.

In fact the SSREFLECT tactic:

```
apply.
```

is a synonym for:

```
intro top; first [refine top | refine (top _) | refine (top _ _) |
...]; clear top.
```

where `top` is fresh name, and the sequence of **refine** tactics tries to catch the appropriate number of wildcards to be inserted. Note that this use of the **refine** tactic implies that the tactic tries to match the goal up to expansion of constants and evaluation of subterms.

SSREFLECT's **apply** has a special behaviour on goals containing existential metavariables of sort **Prop**. Consider the following example:

```
Goal (forall y, 1 < y -> y < 2 -> exists x : { n | n < 3 },
proj1_sig x > 0).
move=> y y_gt1 y_lt2; apply: (ex_intro _ (exist _ y _)).
by apply: gt_trans _ y_lt2.
by move=> y_lt3; apply: lt_trans y_gt1.
```

Note that the last `_` of the tactic **apply: (ex_intro _ (exist _ y _))** represents a proof that $y < 3$. Instead of generating the following goal

```
0 < (n:=3) (m:=y) ?54
```

the system tries to prove $y < 3$ calling the **trivial** tactic. If it succeeds, let's say because the context contains $H : y < 3$, then the system generates the following goal:

```
0 < proj1_sig (exist (fun n => n < 3) y H
```

Otherwise the missing proof is considered to be irrelevant, and is thus discharged generating the following goals:

```
y < 3
forall H : y < 3, proj1_sig (exist (fun n => n < 3) y H)
```

Last, the user can replace the **trivial** tactic by defining an Ltac expression named `ssrautoprop`.

11.5.3 Discharge

The general syntax of the discharging tactical `'.'` is:

```
tactic [ident] : d-item1 ... d-itemn [clear-switch]
```

where $n > 0$, and *d-item* and *clear-switch* are defined as

$$\begin{aligned}
 d\text{-item} &::= [\text{occ-switch} \mid \text{clear-switch}] \text{ term} \\
 \text{clear-switch} &::= \{ \text{ident}_1 \dots \text{ident}_m \}
 \end{aligned}$$

with the following requirements:

- *tactic* must be one of the four basic tactics described in 11.5.2, i.e., `move`, `case`, `elim` or `apply`, the `exact` tactic (section 11.6.2), the `congr` tactic (section 11.7.4), or the application of the *view* tactical ‘/’ (section 11.9.2) to one of `move`, `case`, or `elim`.
- The optional *ident* specifies *equation generation* (section 11.5.5), and is only allowed if *tactic* is `move`, `case` or `elim`, or the application of the *view* tactical ‘/’ (section 11.9.2) to `case` or `elim`.
- An *occ-switch* selects occurrences of *term*, as in 11.4.2; *occ-switch* is not allowed if *tactic* is `apply` or `exact`.
- A clear item *clear-switch* specifies facts and constants to be deleted from the proof context (as per the `clear` tactic).

The ‘:’ tactical first *discharges* all the *d-items*, right to left, and then performs *tactic*, i.e., for each *d-item*, starting with *d-item_n*:

1. The SSREFLECT matching algorithm described in section 11.4.2 is used to find occurrences of *term* in the goal, after filling any holes ‘_’ in *term*; however if *tactic* is `apply` or `exact` a different matching algorithm, described below, is used ⁴.
2. These occurrences are replaced by a new variable; in particular, if *term* is a fact, this adds an assumption to the goal.
3. If *term* is *exactly* the name of a constant or fact in the proof context, it is deleted from the context, unless there is an *occ-switch*.

Finally, *tactic* is performed just after *d-item₁* has been generalized — that is, between steps 2 and 3 for *d-item₁*. The names listed in the final *clear-switch* (if it is present) are cleared first, before *d-item_n* is discharged.

Switches affect the discharging of a *d-item* as follows:

- An *occ-switch* restricts generalization (step 2) to a specific subset of the occurrences of *term*, as per 11.4.2, and prevents clearing (step 3).
- All the names specified by a *clear-switch* are deleted from the context in step 3, possibly in addition to *term*.

For example, the tactic:

```
move: n {2}n (refl_equal n).
```

- first generalizes `(refl_equal n : n = n)`;
- then generalizes the second occurrence of `n`.

⁴Also, a slightly different variant may be used for the first *d-item* of `case` and `elim`; see section 11.5.6.

- finally generalizes all the other occurrences of n , and clears n from the proof context (assuming n is a proof constant).

Therefore this tactic changes any goal G into

```
forall n n0 : nat, n = n0 -> G.
```

where the name $n0$ is picked by the COQ display function, and assuming n appeared only in G .

Finally, note that a discharge operation generalizes defined constants as variables, and not as local definitions. To override this behavior, prefix the name of the local definition with a `@`, like in `move : @n`.

This is in contrast with the behavior of the `in` tactical (see section 11.6.5), which preserves local definitions by default.

Clear rules

The clear step will fail if *term* is a proof constant that appears in other facts; in that case either the facts should be cleared explicitly with a *clear-switch*, or the clear step should be disabled. The latter can be done by adding an *occ-switch* or simply by putting parentheses around *term*: both

```
move : (n) .
```

and

```
move : {+}n.
```

generalize n without clearing n from the proof context.

The clear step will also fail if the *clear-switch* contains a *ident* that is not in the *proof* context. Note that SSREFLECT never clears a section constant.

If *tactic* is `move` or `case` and an equation *ident* is given, then clear (step 3) for $d\text{-item}_1$ is suppressed (see section 11.5.5).

Matching for `apply` and `exact`

The matching algorithm for *d-items* of the SSREFLECT `apply` and `exact` tactics exploits the type of $d\text{-item}_1$ to interpret wildcards in the other *d-item* and to determine which occurrences of these should be generalized. Therefore, *occur switches* are not needed for `apply` and `exact`.

Indeed, the SSREFLECT tactic `apply : H x` is equivalent to

```
refine (@H _ ... _ x); clear H x
```

with an appropriate number of wildcards between H and x .

Note that this means that matching for `apply` and `exact` has much more context to interpret wildcards; in particular it can accommodate the ‘`_`’ *d-item*, which would always be rejected after ‘`move :` ’. For example, the tactic

```
apply : trans_equal (Hfg _) _.
```

transforms the goal $f\ a = g\ b$, whose context contains $(Hfg : forall\ x, f\ x = g\ x)$, into $g\ a = g\ b$. This tactic is equivalent (see section 11.5.1) to:

```
refine (trans_equal (Hfg _) _).
```

and this is a common idiom for applying transitivity on the left hand side of an equation.

The `abstract` tactic

The `abstract` tactic assigns an abstract constant previously introduced with the `[: name] intro` pattern (see section 11.5.4, page 319). In a goal like the following:

```
m : nat
abs : <hidden>
n : nat
=====
m < 5 + n
```

The tactic `abstract : abs n` first generalizes the goal with respect to `n` (that is not visible to the abstract constant `abs`) and then assigns `abs`. The resulting goal is:

```
m : nat
n : nat
=====
m < 5 + n
```

Once this subgoal is closed, all other goals having `abs` in their context see the type assigned to `abs`. In this case:

```
m : nat
abs : forall n, m < 5 + n
```

For a more detailed example the user should refer to section 11.6.6, page 329.

11.5.4 Introduction

The application of a tactic to a given goal can generate (quantified) variables, assumptions, or definitions, which the user may want to *introduce* as new facts, constants or defined constants, respectively. If the tactic splits the goal into several subgoals, each of them may require the introduction of different constants and facts. Furthermore it is very common to immediately decompose or rewrite with an assumption instead of adding it to the context, as the goal can often be simplified and even proved after this.

All these operations are performed by the introduction tactical ‘`=>`’, whose general syntax is

$$tactic \Rightarrow i\text{-}item_1 \dots i\text{-}item_n$$

where *tactic* can be any tactic, $n > 0$ and

$$\begin{aligned} i\text{-}item &::= i\text{-}pattern \mid s\text{-}item \mid clear\text{-}switch \mid /term \\ s\text{-}item &::= /= \mid // \mid // = \\ i\text{-}pattern &::= ident \mid _ \mid ? \mid * \mid [occ\text{-}switch] \rightarrow \mid [occ\text{-}switch] <- \mid \\ &\quad [i\text{-}item_1^* \mid \dots \mid i\text{-}item_m^*] \mid - \mid [: ident^+] \end{aligned}$$

The ‘`=>`’ tactical first executes *tactic*, then the *i-items*, left to right, i.e., starting from *i-item*₁. An *s-item* specifies a simplification operation; a *clear switch* specifies context pruning as in 11.5.3. The *i-patterns* can be seen as a variant of *intro patterns* 8.3.2: each performs an introduction operation, i.e., pops some variables or assumptions from the goal.

An *s-item* can simplify the set of subgoals or the subgoal themselves:

- `//` removes all the “trivial” subgoals that can be resolved by the `SSREFLECT` tactic `done` de-

scribed in 11.6.2, i.e., it executes `try done`.

- `/=` simplifies the goal by performing partial evaluation, as per the tactic `simpl`.⁵
- `//=` combines both kinds of simplification; it is equivalent to `/= //`, i.e., `simpl; try done`.

When an *s-item* bears a *clear-switch*, then the *clear-switch* is executed *after* the *s-item*, e.g., `{IHn} //` will solve some subgoals, possibly using the fact `IHn`, and will erase `IHn` from the context of the remaining subgoals.

The last entry in the *i-item* grammar rule, */term*, represents a view (see section 11.9). If *i-item*_{k+1} is a view *i-item*, the view is applied to the assumption in top position once *i-item*₁ . . . *i-item*_k have been performed.

The view is applied to the top assumption.

SSREFLECT supports the following *i-patterns*:

- *ident* pops the top variable, assumption, or local definition into a new constant, fact, or defined constant *ident*, respectively. Note that defined constants cannot be introduced when δ -expansion is required to expose the top variable or assumption.
- `?` pops the top variable into an anonymous constant or fact, whose name is picked by the tactic interpreter. SSREFLECT only generates names that cannot appear later in the user script.⁶
- `_` pops the top variable into an anonymous constant that will be deleted from the proof context of all the subgoals produced by the `=>` tactical. They should thus never be displayed, except in an error message if the constant is still actually used in the goal or context after the last *i-item* has been executed (*s-items* can erase goals or terms where the constant appears).
- `*` pops all the remaining apparent variables/assumptions as anonymous constants/facts. Unlike `?` and `move` the `*` *i-item* does not expand definitions in the goal to expose quantifiers, so it may be useful to repeat a `move=> *` tactic, e.g., on the goal

```
forall a b : bool, a <> b
```

a first `move=> *` adds only `_a_ : bool` and `_b_ : bool` to the context; it takes a second `move=> *` to add `_Hyp_ : _a_ = _b_`.

- `[occ-switch]->` (resp. `[occ-switch]<-`) pops the top assumption (which should be a rewritable proposition) into an anonymous fact, rewrites (resp. rewrites right to left) the goal with this fact (using the SSREFLECT `rewrite` tactic described in section 11.7, and honoring the optional occurrence selector), and finally deletes the anonymous fact from the context.
- `[i-item1* | . . . | i-itemm*]`, when it is the very *first i-pattern* after *tactic* `=>` tactical and *tactic* is not a `move`, is a *branching i-pattern*. It executes the sequence *i-item*_i^{*} on the *i*th subgoal produced by *tactic*. The execution of *tactic* should thus generate exactly *m* subgoals, unless the `[. . .]` *i-pattern* comes after an initial `//` or `//=` *s-item* that closes some of the goals produced by *tactic*, in which case exactly *m* subgoals should remain after the *s-item*, or we have the trivial branching *i-pattern* `[]`, which always does nothing, regardless of the number of remaining subgoals.

⁵Except `/=` does not expand the local definitions created by the SSREFLECT `in` tactical.

⁶SSREFLECT reserves all identifiers of the form “`_x_`”, which is used for such generated names.

- $[i\text{-item}_1^* \mid \dots \mid i\text{-item}_m^*]$, when it is *not* the first *i-pattern* or when *tactic* is a `move`, is a *destructing i-pattern*. It starts by destructing the top variable, using the `SSREFLECT case` tactic described in 11.5.2. It then behaves as the corresponding branching *i-pattern*, executing the sequence $i\text{-item}_i^*$ in the i^{th} subgoal generated by the case analysis; unless we have the trivial destructing *i-pattern* `[]`, the latter should generate exactly m subgoals, i.e., the top variable should have an inductive type with exactly m constructors.⁷ While it is good style to use the $i\text{-item}_i^*$ to pop the variables and assumptions corresponding to each constructor, this is not enforced by `SSREFLECT`.
- `-` does nothing, but counts as an intro pattern. It can also be used to force the interpretation of $[i\text{-item}_1^* \mid \dots \mid i\text{-item}_m^*]$ as a case analysis like in `move=> - [H1 H2]`. It can also be used to indicate explicitly the link between a view and a name like in `move=> /eqP-H1`. Last, it can serve as a separator between views. Section 11.9.9 explains in which respect the tactic `move=> /v1/v2` differs from the tactic `move=> /v1-/v2`.
- $[:ident^+]$ introduces in the context an abstract constant for each *ident*. Its type has to be fixed later on by using the `abstract` tactic (see page 317). Before then the type displayed is `<hidden>`.

Note that `SSREFLECT` does not support the syntax $(ipat, \dots, ipat)$ for destructing intro-patterns.

Clears are deferred until the end of the intro pattern. For example, given the goal:

```
x, y : nat
=====
0 < x = true -> (0 < x) && (y < 2) = true
```

the tactic `move=> {x}` -> successfully rewrites the goal and deletes `x` and the anonymous equation. The goal is thus turned into:

```
y : nat
=====
true && (y < 2) = true
```

If the cleared names are reused in the same intro pattern, a renaming is performed behind the scenes.

Facts mentioned in a clear switch must be valid names in the proof context (excluding the section context).

The rules for interpreting branching and destructing *i-pattern* are motivated by the fact that it would be pointless to have a branching pattern if *tactic* is a `move`, and in most of the remaining cases *tactic* is `case` or `elim`, which implies destruction. The rules above imply that

```
move=> [a b].
case=> [a b].
case=> a b.
```

are all equivalent, so which one to use is a matter of style; `move` should be used for casual decomposition, such as splitting a pair, and `case` should be used for actual decompositions, in particular for type families (see 11.5.6) and proof by contradiction.

The trivial branching *i-pattern* can be used to force the branching interpretation, e.g.,

```
case=> [] [a b] c.
move=> [[a b] c].
case; case=> a b c.
```

⁷More precisely, it should have a quantified inductive type with a assumptions and $m - a$ constructors.

are all equivalent.

11.5.5 Generation of equations

The generation of named equations option stores the definition of a new constant as an equation. The tactic:

```
move En: (size l) => n.
```

where l is a list, replaces `size l` by n in the goal and adds the fact $En : size\ l = n$ to the context. This is quite different from:

```
pose n := (size l).
```

which generates a definition $n := (size\ l)$. It is not possible to generalize or rewrite such a definition; on the other hand, it is automatically expanded during computation, whereas expanding the equation En requires explicit rewriting.

The use of this equation name generation option with a `case` or an `elim` tactic changes the status of the first *i-item*, in order to deal with the possible parameters of the constants introduced.

On the goal $a <> b$ where a, b are natural numbers, the tactic:

```
case E : a => [|n].
```

generates two subgoals. The equation $E : a = 0$ (resp. $E : a = S\ n$, and the constant $n : nat$) has been added to the context of the goal $0 <> b$ (resp. $S\ n <> b$).

If the user does not provide a branching *i-item* as first *i-item*, or if the *i-item* does not provide enough names for the arguments of a constructor, then the constants generated are introduced under fresh SSREFLECT names. For instance, on the goal $a <> b$, the tactic:

```
case E : a => H.
```

also generates two subgoals, both requiring a proof of `False`. The hypotheses $E : a = 0$ and $H : 0 = b$ (resp. $E : a = S\ _n_$ and $H : S\ _n_ = b$) have been added to the context of the first subgoal (resp. the second subgoal).

Combining the generation of named equations mechanism with the `case` tactic strengthens the power of a case analysis. On the other hand, when combined with the `elim` tactic, this feature is mostly useful for debug purposes, to trace the values of decomposed parameters and pinpoint failing branches.

11.5.6 Type families

When the top assumption of a goal has an inductive type, two specific operations are possible: the case analysis performed by the `case` tactic, and the application of an induction principle, performed by the `elim` tactic. When this top assumption has an inductive type, which is moreover an instance of a type family, COQ may need help from the user to specify which occurrences of the parameters of the type should be substituted.

A specific / switch indicates the type family parameters of the type of a *d-item* immediately following this / switch, using the syntax:

$$[\text{case} \mid \text{elim}] : d\text{-item}^+ / d\text{-item}^*$$

The *d-items* on the right side of the / switch are discharged as described in section 11.5.3. The case analysis or elimination will be done on the type of the top assumption after these discharge operations.

Every *d-item* preceding the `/` is interpreted as arguments of this type, which should be an instance of an inductive type family. These terms are not actually generalized, but rather selected for substitution. Occurrence switches can be used to restrict the substitution. If a *term* is left completely implicit (e.g. writing just `_`), then a pattern is inferred looking at the type of the top assumption. This allows for the compact syntax `case: {2}_ / eqP`, where `_` is interpreted as `(_ == _)`. Moreover if the *d-items* list is too short, it is padded with an initial sequence of `_` of the right length.

Here is a small example on lists. We define first a function which adds an element at the end of a given list.

```
Require Import List.

Section LastCases.
Variable A : Type.

Fixpoint add_last (a : A) (l : list A) : list A :=
match l with
| nil => a :: nil
| hd :: tl => hd :: (add_last a tl)
end.
```

Then we define an inductive predicate for case analysis on lists according to their last element:

```
Inductive last_spec : list A -> Type :=
| LastSeq0 : last_spec nil
| LastAdd s x : last_spec (add_last x s).

Theorem lastP : forall l : list A, last_spec l.
```

Applied to the goal:

```
Goal forall l : list A, (length l) * 2 = length (app l l).
```

the command:

```
move=> l; case: (lastP l).
```

generates two subgoals:

```
length nil * 2 = length (nil ++ nil)
```

and

```
forall (s : list A) (x : A),
length (add_last x s) * 2 = length (add_last x s ++ add_last x s)
```

both having `l : list A` in their context.

Applied to the same goal, the command:

```
move=> l; case: l / (lastP l).
```

generates the same subgoals but `l` has been cleared from both contexts.

Again applied to the same goal, the command:

```
move=> l; case: {1 3}l / (lastP l).
```

generates the subgoals $\text{length } l * 2 = \text{length } (\text{nil} ++ l)$ and $\text{forall } (s : \text{list } A) (x : A), \text{length } l * 2 = \text{length } (\text{add_last } x \ s ++ l)$ where the selected occurrences on the left of the $/$ switch have been substituted with l instead of being affected by the case analysis.

The equation name generation feature combined with a type family $/$ switch generates an equation for the *first* dependent d -item specified by the user. Again starting with the above goal, the command:

```
move=> l; case E: {l 3}l / (lastP l)=>[|s x].
```

adds $E : l = \text{nil}$ and $E : l = \text{add_last } x \ s$, respectively, to the context of the two subgoals it generates.

There must be at least one *d-item* to the left of the $/$ switch; this prevents any confusion with the view feature. However, the *d-items* to the right of the $/$ are optional, and if they are omitted the first assumption provides the instance of the type family.

The equation always refers to the first *d-item* in the actual tactic call, before any padding with initial $_s$. Thus, if an inductive type has two family parameters, it is possible to have SSREFLECT generate an equation for the second one by omitting the pattern for the first; note however that this will fail if the type of the second parameter depends on the value of the first parameter.

11.6 Control flow

11.6.1 Indentation and bullets

A linear development of COQ scripts gives little information on the structure of the proof. In addition, replaying a proof after some changes in the statement to be proved will usually not display information to distinguish between the various branches of case analysis for instance.

To help the user in this organization of the proof script at development time, SSREFLECT provides some bullets to highlight the structure of branching proofs. The available bullets are $-$, $+$ and $*$. Combined with tabulation, this lets us highlight four nested levels of branching; the most we have ever needed is three. Indeed, the use of “*simpl* and *closing*” switches, of terminators (see above section 11.6.2) and selectors (see section 11.6.3) is powerful enough to avoid most of the time more than two levels of indentation.

Here is a fragment of such a structured script:

```
case E1: (abezoutn _ _) => [| k1] [| k2]].
- rewrite !muln0 !gexpn0 mulg1 => H1.
  move/eqP: (sym_equal F0); rewrite -H1 orderg1 eqn_mull.
  by case/andP; move/eqP.
- rewrite muln0 gexpn0 mulg1 => H1.
  have F1: t %| t * S k2.+1 - 1.
    apply: (@dvdn_trans (orderg x)); first by rewrite F0; exact:
      dvdn_mull.
    rewrite orderg_dvd; apply/eqP; apply: (mulgI x).
    rewrite -{1}(gexpn1 x) mulg1 gexpn_add leq_add_sub //.
    by move: P1; case t.
  rewrite dvdn_subr in F1; last by exact: dvdn_mulr.
+ rewrite H1 F0 -{2}(muln1 (p ^ 1)); congr (_ * _).
  by apply/eqP; rewrite -dvdn1.
+ by move: P1; case: (t) => [| [| s1]].
```

```
- rewrite muln0 gexpn0 mullg => H1.
...
```

11.6.2 Terminators

To further structure scripts, SSREFLECT supplies *terminating* tacticals to explicitly close off tactics. When replaying scripts, we then have the nice property that an error immediately occurs when a closed tactic fails to prove its subgoal.

It is hence recommended practice that the proof of any subgoal should end with a tactic which *fails if it does not solve the current goal*, like `discriminate`, `contradiction` or `assumption`.

In fact, SSREFLECT provides a generic tactical which turns any tactic into a closing one (similar to `now`). Its general syntax is:

`by tactic .`

The Ltac expression:

`by [tactic1 | [tactic2 | ...]] .`

is equivalent to:

`[by tactic1 | by tactic2 | ...] .`

and this form should be preferred to the former.

In the script provided as example in section 11.6.1, the paragraph corresponding to each sub-case ends with a tactic line prefixed with a `by`, like in:

`by apply/eqP; rewrite -dvdn1.`

The `by` tactical is implemented using the user-defined, and extensible `done` tactic. This `done` tactic tries to solve the current goal by some trivial means and fails if it doesn't succeed. Indeed, the tactic expression:

`by tactic .`

is equivalent to:

`tactic; done .`

Conversely, the tactic

`by [] .`

is equivalent to:

`done .`

The default implementation of the `done` tactic, in the `ssreflect.v` file, is:

```

Ltac done :=
  trivial; hnf; intros; solve
  [ do ![solve [trivial | apply: sym_equal; trivial]
    | discriminate | contradiction | split]
  | case not_locked_false_eq_true; assumption
  | match goal with H : ~ _ |- _ => solve [case H; trivial] end ].

```

The lemma `not_locked_false_eq_true` is needed to discriminate *locked* boolean predicates (see section 11.7.3). The iterator tactical `do` is presented in section 11.6.4. This tactic can be customized by the user, for instance to include an `auto` tactic.

A natural and common way of closing a goal is to apply a lemma which is the `exact` one needed for the goal to be solved. The defective form of the tactic:

```
exact.
```

is equivalent to:

```
do [done | by move=> top; apply top].
```

where `top` is a fresh name affected to the top assumption of the goal. This applied form is supported by the `: discharge` tactical, and the tactic:

```
exact: MyLemma.
```

is equivalent to:

```
by apply: MyLemma.
```

(see section 11.5.3 for the documentation of the `apply:` combination).

Warning The list of tactics, possibly chained by semi-columns, that follows a `by` keyword is considered as a parenthesized block applied to the current goal. Hence for example if the tactic:

```
by rewrite my_lemma1.
```

succeeds, then the tactic:

```
by rewrite my_lemma1; apply my_lemma2.
```

usually fails since it is equivalent to:

```
by (rewrite my_lemma1; apply my_lemma2).
```

11.6.3 Selectors

When composing tactics, the two tacticals `first` and `last` let the user restrict the application of a tactic to only one of the subgoals generated by the previous tactic. This covers the frequent cases where a tactic generates two subgoals one of which can be easily disposed of.

This is an other powerful way of linearization of scripts, since it happens very often that a trivial subgoal can be solved in a less than one line tactic. For instance, the tactic:

```
tactic1; last by tactic2.
```

tries to solve the last subgoal generated by `tactic1` using the `tactic2`, and fails if it does not succeeds. Its analogous

```
tactic1; first by tactic2.
```

tries to solve the first subgoal generated by *tactic*₁ using the tactic *tactic*₂, and fails if it does not succeeds.

SSREFLECT also offers an extension of this facility, by supplying tactics to *permute* the subgoals generated by a tactic. The tactic:

tactic; *last first*.

inverts the order of the subgoals generated by *tactic*. It is equivalent to:

tactic; *first last*.

More generally, the tactic:

tactic; *lastnatural first*.

where *natural* is a COQ numeral, or and Ltac variable denoting a COQ numeral, having the value *k*. It rotates the *n* subgoals G_1, \dots, G_n generated by *tactic*. The first subgoal becomes G_{n+1-k} and the circular order of subgoals remains unchanged.

Conversely, the tactic:

tactic; *firstnatural last*.

rotates the *n* subgoals G_1, \dots, G_n generated by *tactic* in order that the first subgoal becomes G_k .

Finally, the tactics *last* and *first* combine with the branching syntax of Ltac: if the tactic *tactic*₀ generates *n* subgoals on a given goal, then the tactic

*tactic*₀; *lastnatural* [*tactic*₁ | ... | *tactic*_{*m*}] | | *tactic*_{*m*+1}.

where *natural* denotes the integer *k* as above, applies *tactic*₁ to the $n - k + 1$ -th goal, ... *tactic*_{*m*} to the $n - k + 2 - m$ -th goal and *tactic*_{*m*+1} to the others.

For instance, the script:

```
Inductive test : nat -> Prop :=
  C1 : forall n, test n | C2 : forall n, test n |
  C3 : forall n, test n | C4 : forall n, test n.

Goal forall n, test n -> True.
move=> n t; case: t; last 2 [move=> k | move=> l]; idtac.
```

creates a goal with four subgoals, the first and the last being *nat* -> True, the second and the third being True with respectively *k* : nat and *l* : nat in their context.

11.6.4 Iteration

SSREFLECT offers an accurate control on the repetition of tactics, thanks to the *do* tactical, whose general syntax is:

do [*mult*] [*tactic*₁ | ... | *tactic*_{*n*}]

where *mult* is a *multiplier*.

Brackets can only be omitted if a single tactic is given *and* a multiplier is present.

A tactic of the form:

do [*tactic*₁ | ... | *tactic*_{*n*}] .

is equivalent to the standard Ltac expression:

```
first [ tactic1 | ... | tacticn ] .
```

The optional multiplier *mult* specifies how many times the action of *tactic* should be repeated on the current subgoal.

There are four kinds of multipliers:

- *n*!: the step *tactic* is repeated exactly *n* times (where *n* is a positive integer argument).
- !: the step *tactic* is repeated as many times as possible, and done at least once.
- ?: the step *tactic* is repeated as many times as possible, optionally.
- *n*?: the step *tactic* is repeated up to *n* times, optionally.

For instance, the tactic:

```
tactic ; do 1?rewrite mult_comm.
```

rewrites at most one time the lemma `mult_com` in all the subgoals generated by *tactic*, whereas the tactic:

```
tactic ; do 2!rewrite mult_comm.
```

rewrites exactly two times the lemma `mult_com` in all the subgoals generated by *tactic*, and fails if this rewrite is not possible in some subgoal.

Note that the combination of multipliers and `rewrite` is so often used that multipliers are in fact integrated to the syntax of the SSREFLECT `rewrite` tactic, see section 11.7.

11.6.5 Localization

In sections 11.4.3 and 11.5.1, we have already presented the *localization* tactical `in`, whose general syntax is:

$$tactic \text{ in } ident^+ [*]$$

where $ident^+$ is a non empty list of fact names in the context. On the left side of `in`, *tactic* can be `move`, `case`, `elim`, `rewrite`, `set`, or any tactic formed with the general iteration tactical `do` (see section 11.6.4).

The operation described by *tactic* is performed in the facts listed in $ident^+$ and in the goal if a `*` ends the list.

The `in` tactical successively:

- generalizes the selected hypotheses, possibly “protecting” the goal if `*` is not present,
- performs *tactic*, on the obtained goal,
- reintroduces the generalized facts, under the same names.

This defective form of the `do` tactical is useful to avoid clashes between standard Ltac `in` and the SSREFLECT tactical `in`. For example, in the following script:

```

Ltac mytac H := rewrite H.

Goal forall x y, x = y -> y = 3 -> x + y = 6.
move=> x y H1 H2.
do [mytac H2] in H1 *.

```

the last tactic rewrites the hypothesis $H2 : y = 3$ both in $H1 : x = y$ and in the goal $x + y = 6$.

By default `in` keeps the body of local definitions. To erase the body of a local definition during the generalization phase, the name of the local definition must be written between parentheses, like in `rewrite H in H1 (def_n) H2`.

From SSREFLECT 1.5 the grammar for the `in` tactical has been extended to the following one:

$$\text{tactic } \text{in} [\text{clear-switch} \mid [@] \text{ident} \mid (\text{ident}) \mid ([@] \text{ident} := \text{c-pattern})]^+ [*]$$

In its simplest form the last option lets one rename hypotheses that can't be cleared (like section variables). For example $(y := x)$ generalizes over x and reintroduces the generalized variable under the name y (and does not clear x).

For a more precise description the $([@] \text{ident} := \text{c-pattern})$ item refer to the “Advanced generalization” paragraph at page 332.

11.6.6 Structure

Forward reasoning structures the script by explicitly specifying some assumptions to be added to the proof context. It is closely associated with the declarative style of proof, since an extensive use of these highlighted statements make the script closer to a (very detailed) text book proof.

Forward chaining tactics allow to state an intermediate lemma and start a piece of script dedicated to the proof of this statement. The use of closing tactics (see section 11.6.2) and of indentation makes syntactically explicit the portion of the script building the proof of the intermediate statement.

The **have** tactic.

The main SSREFLECT forward reasoning tactic is the `have` tactic. It can be use in two modes: one starts a new (sub)proof for an intermediate result in the main proof, and the other provides explicitly a proof term for this intermediate step.

In the first mode, the syntax of `have` in its defective form is:

```
have : term .
```

This tactic supports open syntax for *term*. Applied to a goal G , it generates a first subgoal requiring a proof of *term* in the context of G . The second generated subgoal is of the form $\text{term} \rightarrow G$, where *term* becomes the new top assumption, instead of being introduced with a fresh name. At the proof-term level, the `have` tactic creates a β redex, and introduces the lemma under a fresh name, automatically chosen.

Like in the case of the `pose` tactic (see section 11.4.1), the types of the holes are abstracted in *term*. For instance, the tactic:

```
have : _ * 0 = 0 .
```

is equivalent to:

```
have : forall n : nat, n * 0 = 0 .
```

The `have` tactic also enjoys the same abstraction mechanism as the `pose` tactic for the non-inferred implicit arguments. For instance, the tactic:

```
have: forall x y, (x, y) = (x, y + 0).
```

opens a new subgoal to prove that:

```
forall (T : Type) (x : T) (y : nat), (x, y) = (x, y + 0)
```

The behavior of the defective `have` tactic makes it possible to generalize it in the following general construction:

```
have i-item* [i-pattern] [s-item | binder+] [: term1] [:= term2 | by tactic]
```

Open syntax is supported for $term_1$ and $term_2$. For the description of *i-items* and clear switches see section 11.5.4. The first mode of the `have` tactic, which opens a sub-proof for an intermediate result, uses tactics of the form:

```
have clear-switch i-item : term by tactic.
```

which behave like:

```
have: term ; first by tactic.
```

```
move=> clear-switch i-item.
```

Note that the *clear-switch* precedes the *i-item*, which allows to reuse a name of the context, possibly used by the proof of the assumption, to introduce the new assumption itself.

The `by` feature is especially convenient when the proof script of the statement is very short, basically when it fits in one line like in:

```
have H : forall x y, x + y = y + x by move=> x y; rewrite addnC.
```

The possibility of using *i-items* supplies a very concise syntax for the further use of the intermediate step. For instance,

```
have -> : forall x, x * a = a.
```

on a goal G , opens a new subgoal asking for a proof of `forall x, x * a = a`, and a second subgoal in which the lemma `forall x, x * a = a` has been rewritten in the goal G . Note that in this last subgoal, the intermediate result does not appear in the context. Note that, thanks to the deferred execution of clears, the following idiom is supported (assuming x occurs in the goal only):

```
have {x} -> : x = y
```

An other frequent use of the intro patterns combined with `have` is the destruction of existential assumptions like in the tactic:

```
have [x Px]: exists x : nat, x > 0.
```

which opens a new subgoal asking for a proof of `exists x : nat, x > 0` and a second subgoal in which the witness is introduced under the name `x : nat`, and its property under the name `Px : x > 0`.

An alternative use of the `have` tactic is to provide the explicit proof term for the intermediate lemma, using tactics of the form:

`have [ident] := term.`

This tactic creates a new assumption of type the type of *term*. If the optional *ident* is present, this assumption is introduced under the name *ident*. Note that the body of the constant is lost for the user.

Again, non inferred implicit arguments and explicit holes are abstracted. For instance, the tactic:

```
have H := forall x, (x, x) = (x, x).
```

adds to the context `H : Type -> Prop`. This is a schematic example but the feature is specially useful when the proof term to give involves for instance a lemma with some hidden implicit arguments.

After the *i-pattern*, a list of binders is allowed. For example, if `Pos_to_P` is a lemma that proves that *P* holds for any positive, the following command:

```
have H x (y : nat) : 2 * x + y = x + x + y by auto.
```

will put in the context `H : forall x, 2 * x = x + x`. A proof term provided after `:=` can mention these bound variables (that are automatically introduced with the given names). Since the *i-pattern* can be omitted, to avoid ambiguity, bound variables can be surrounded with parentheses even if no type is specified:

```
have (x) : 2 * x = x + x by auto.
```

The *i-items* and *s-item* can be used to interpret the asserted hypothesis with views (see section 11.9) or simplify the resulting goals.

The `have` tactic also supports a `suff` modifier which allows for asserting that a given statement implies the current goal without copying the goal itself. For example, given a goal *G* the tactic `have suff H : P` results in the following two goals:

```
|- P -> G
H : P -> G |- G
```

Note that *H* is introduced in the second goal. The `suff` modifier is not compatible with the presence of a list of binders.

Generating `let in` context entries with `have`

Since SSREFLECT 1.5 the `have` tactic supports a “transparent” modifier to generate `let in` context entries: the `@` symbol in front of the context entry name. For example:

```
have @i : 'I_n by apply: (Sub m); auto.
```

generates the following two context entry:

```
i := Sub m proof_produced_by_auto : 'I_n
```

Note that the sub-term produced by `auto` is in general huge and uninteresting, and hence one may want to hide it.

For this purpose the `[: name]` intro pattern and the tactic `abstract` (see page 317) are provided. Example:

```
have [ :blurb] @i : 'I_n by apply: (Sub m); abstract: blurb; auto.
```

generates the following two context entries:

```
blurb : (m < n) (*1*)
i := Sub m blurb : 'I_n
```

The type of `blurb` can be cleaned up by its annotations by just simplifying it. The annotations are there for technical reasons only.

When intro patterns for abstract constants are used in conjunction with `have` and an explicit term, they must be used as follows:

```
have [:blurb] @i : 'I_n := Sub m blurb.
  by auto.
```

In this case the abstract constant `blurb` is assigned by using it in the term that follows `:=` and its corresponding goal is left to be solved. Goals corresponding to intro patterns for abstract constants are opened in the order in which the abstract constants are declared (not in the “order” in which they are used in the term).

Note that abstract constants do respect scopes. Hence, if a variable is declared after their introduction, it has to be properly generalized (i.e. explicitly passed to the abstract constant when one makes use of it). For example any of the following two lines:

```
have [:blurb] @i k : 'I_(n+k) by apply: (Sub m); abstract: blurb k;
  auto.
have [:blurb] @i k : 'I_(n+k) := apply: Sub m (blurb k); first by
  auto.
```

generates the following context:

```
blurb : (forall k, m < n+k) (*1*)
i := fun k => Sub m (blurb k) : forall k, 'I_(n+k)
```

Last, notice that the use of intro patterns for abstract constants is orthogonal to the transparent flag `@` for `have`.

The `have` tactic and type classes resolution

Since SSREFLECT 1.5 the `have` tactic behaves as follows with respect to type classes inference.

- `have foo : ty`. Full inference for `ty`. The first subgoal demands a proof of such instantiated statement.
- `have foo : ty := .` No inference for `ty`. Unresolved instances are quantified in `ty`. The first subgoal demands a proof of such quantified statement. Note that no proof term follows `:=`, hence two subgoals are generated.
- `have foo : ty := t`. No inference for `ty` and `t`.
- `have foo := t`. No inference for `t`. Unresolved instances are quantified in the (inferred) type of `t` and abstracted in `t`.

The behavior of SSREFLECT 1.4 and below (never resolve type classes) can be restored with the option `Set SsrHave NoTCResolution`.

Variants: the `suff` and `wlog` tactics.

As it is often the case in mathematical textbooks, forward reasoning may be used in slightly different variants. One of these variants is to show that the intermediate step L easily implies the initial goal

G . By easily we mean here that the proof of $L \Rightarrow G$ is shorter than the one of L itself. This kind of reasoning step usually starts with: “It suffices to show that ...”.

This is such a frequent way of reasoning that SSREFLECT has a variant of the `have` tactic called `suffices` (whose abridged name is `suff`). The `have` and `suff` tactics are equivalent and have the same syntax but:

- the order of the generated subgoals is inversed
- but the optional clear item is still performed in the *second* branch. This means that the tactic:

```
suff {H} H : forall x : nat, x >= 0.
```

fails if the context of the current goal indeed contains an assumption named H .

The rationale of this clearing policy is to make possible “trivial” refinements of an assumption, without changing its name in the main branch of the reasoning.

The `have` modifier can follow the `suff` tactic. For example, given a goal G the tactic `suff have H : P` results in the following two goals:

```
H : P |- G
|- (P -> G) -> G
```

Note that, in contrast with `have suff`, the name H has been introduced in the first goal.

Another useful construct is reduction, showing that a particular case is in fact general enough to prove a general property. This kind of reasoning step usually starts with: “Without loss of generality, we can suppose that ...”. Formally, this corresponds to the proof of a goal G by introducing a cut `wlog_statement -> G`. Hence the user shall provide a proof for both $(wlog_statement -> G) -> G$ and `wlog_statement -> G`. However, such cuts are usually rather painful to perform by hand, because the statement `wlog_statement` is tedious to write by hand, and sometimes even to read.

SSREFLECT implements this kind of reasoning step through the `without loss` tactic, whose short name is `wlog`. It offers support to describe the shape of the cut statements, by providing the simplifying hypothesis and by pointing at the elements of the initial goals which should be generalized. The general syntax of `without loss` is:

$$wlog [suff] [clear-switch] [i-item] : [ident_1 \dots ident_n] / term$$

where $ident_1 \dots ident_n$ are identifiers for constants in the context of the goal. Open syntax is supported for `term`.

In its defective form:

```
wlog: / term.
```

on a goal G , it creates two subgoals: a first one to prove the formula $(term -> G) -> G$ and a second one to prove the formula `term -> G`.

:browse confirm wa If the optional list $ident_1 \dots ident_n$ is present on the left side of `/`, these constants are generalized in the premise $(term -> G)$ of the first subgoal. By default the body of local definitions is erased. This behavior can be inhibited prefixing the name of the local definition with the `@` character.

In the second subgoal, the tactic:

```
move=> clear-switch i-item .
```

is performed if at least one of these optional switches is present in the `wlog` tactic.

The `wlog` tactic is specially useful when a symmetry argument simplifies a proof. Here is an example showing the beginning of the proof that quotient and remainder of natural number euclidean division are unique.

```
Lemma quo_rem_unicity: forall d q1 q2 r1 r2,
  q1*d + r1 = q2*d + r2 -> r1 < d -> r2 < d -> (q1, r1) = (q2, r2).
move=> d q1 q2 r1 r2.
wlog: q1 q2 r1 r2 / q1 <= q2.
  by case (le_gt_dec q1 q2)=> H; last symmetry; eauto with arith.
```

The `wlog suff` variant is simpler, since it cuts `wlog_statement` instead of `wlog_statement -> G`. It thus opens the goals `wlog_statement -> G` and `wlog_statement`.

In its simplest form the `generally have : ...` tactic is equivalent to `wlog suff : ...` followed by `last first`. When the `have` tactic is used with the `generally` (or `gen`) modifier it accepts an extra identifier followed by a comma before the usual intro pattern. The identifier will name the new hypothesis in its more general form, while the intro pattern will be used to process its instance. For example:

```
Lemma simple n (ngt0 : 0 < n) : P n.
gen have ltnV, /andP[nge0 neq0] : n ngt0 / (0 <= n) && (n != 0).
```

The first subgoal will be

```
n : nat
ngt0 : 0 < n
=====
(0 <= n) && (n != 0)
```

while the second one will be

```
n : nat
ltnV : forall n, 0 < n -> (0 <= n) && (n != 0)
nge0 : 0 <= n
neqn0 : n != 0
=====
P n
```

Advanced generalization The complete syntax for the items on the left hand side of the `/` separator is the following one:

$$\text{clear-switch} \mid [\text{@}] \text{ident} \mid ([\text{@}]/\text{ident} := \text{c-pattern})$$

Clear operations are intertwined with generalization operations. This helps in particular avoiding dependency issues while generalizing some facts.

If an *ident* is prefixed with the `@` prefix mark, then a let-in redex is created, which keeps track if its body (if any). The syntax $(\text{ident} := \text{c-pattern})$ allows to generalize an arbitrary term using a given name. Note that its simplest form $(x := y)$ is just a renaming of y into x . In particular, this can be useful in order to simulate the generalization of a section variable, otherwise not allowed. Indeed renaming does not require the original variable to be cleared.

The syntax `(@x := y)` generates a let-in abstraction but with the following caveat: `x` will not bind `y`, but its body, whenever `y` can be unfolded. This covers the case of both local and global definitions, as illustrated in the following example:

```
Section Test.
Variable x : nat.
Definition addx z := z + x.
Lemma test : x <= addx x.
wlog H : (y := x) (@twoy := addx x) / twoy = 2 * y.
```

The first subgoal is:

```
(forall y : nat, let twoy := y + y in twoy = 2 * y -> y <= twoy)
->
x <= addx x
```

To avoid unfolding the term captured by the pattern `add x` one can use the pattern `id (addx x)`, that would produce the following first subgoal:

```
(forall y : nat, let twoy := addx y in twoy = 2 * y -> y <= twoy)
->
x <= addx x
```

11.7 Rewriting

The generalized use of reflection implies that most of the intermediate results handled are properties of effectively computable functions. The most efficient mean of establishing such results are computation and simplification of expressions involving such functions, i.e., rewriting. SSREFLECT therefore includes an extended `rewrite` tactic, that unifies and combines most of the rewriting functionalities.

11.7.1 An extended `rewrite` tactic

The main features of the `rewrite` tactic are:

- It can perform an entire series of such operations in any subset of the goal and/or context;
- It allows to perform rewriting, simplifications, folding/unfolding of definitions, closing of goals;
- Several rewriting operations can be chained in a single tactic;
- Control over the occurrence at which rewriting is to be performed is significantly enhanced.

The general form of an SSREFLECT `rewrite` tactic is:

`rewrite` *rstep*⁺.

The combination of a `rewrite` tactic with the `in` tactical (see section 11.4.3) performs rewriting in both the context and the goal.

A `rewrite step` *rstep* has the general form:

[r-prefix] *r-item*

where:

$$\begin{aligned}
r\text{-prefix} &::= [-] [mult] [occ\text{-}switch \mid clear\text{-}switch] [[r\text{-}pattern]] \\
r\text{-pattern} &::= term \mid in [ident in] term \mid [term in \mid term as] ident in term \\
r\text{-item} &::= [/]term \mid s\text{-item}
\end{aligned}$$

An *r-prefix* contains annotations to qualify where and how the rewrite operation should be performed:

- The optional initial `-` indicates the direction of the rewriting of *r-item*: if present the direction is right-to-left and it is left-to-right otherwise.
- The multiplier *mult* (see section 11.6.4) specifies if and how the rewrite operation should be repeated.
- A rewrite operation matches the occurrences of a *rewrite pattern*, and replaces these occurrences by an other term, according to the given *r-item*. The optional *redex switch* `[r-pattern]`, which should always be surrounded by brackets, gives explicitly this rewrite pattern. In its simplest form, it is a regular term. If no explicit redex switch is present the rewrite pattern to be matched is inferred from the *r-item*.
- This optional *term*, or the *r-item*, may be preceded by an occurrence switch (see section 11.6.3) or a clear item (see section 11.5.3), these two possibilities being exclusive. An occurrence switch selects the occurrences of the rewrite pattern which should be affected by the rewrite operation.

An *r-item* can be:

- A *simplification r-item*, represented by a *s-item* (see section 11.5.4). Simplification operations are intertwined with the possible other rewrite operations specified by the list of *r-items*.
- A *folding/unfolding r-item*. The tactic:

```
rewrite /term
```

unfolds the head constant of *term* in every occurrence of the first matching of *term* in the goal. In particular, if `my_def` is a (local or global) defined constant, the tactic:

```
rewrite /my_def.
```

is analogous to:

```
unfold my_def.
```

Conversely:

```
rewrite -/my_def.
```

is equivalent to:

```
fold my_def.
```

When an *unfold r-item* is combined with a redex pattern, a conversion operation is performed. A tactic of the form:

```
rewrite -[term1] /term2.
```

is equivalent to:

`change term1 with term2 .`

If *term₂* is a single constant and *term₁* head symbol is not *term₂*, then the head symbol of *term₁* is repeatedly unfolded until *term₂* appears.

```
Definition double x := x + x.
Definition ddouble x := double (double x).
Lemma ex1 x : ddouble x = 4 * x.
rewrite [ddouble _]/double.
```

The resulting goal is:

```
double x + double x = 4 * x
```

Warning The SSREFLECT terms containing holes are *not* typed as abstractions in this context. Hence the following script:

```
Definition f := fun x y => x + y.
Goal forall x y, x + y = f y x.
move=> x y.
rewrite -[f y]/(y + _).
```

raises the error message

```
User error: fold pattern (y + _) does not match redex (f y)
```

but the script obtained by replacing the last line with:

```
rewrite -[f y x]/(y + _).
```

is valid.

- A term, which can be:
 - A term whose type has the form:

$$\text{forall } (x_1 : A_1) \dots (x_n : A_n), \text{eq } \text{term}_1 \text{ term}_2$$

where *eq* is the Leibniz equality or a registered setoid equality.

- A list of terms (t_1, \dots, t_n) , each t_i having a type of the form:

$$\text{forall } (x_1 : A_1) \dots (x_n : A_n), \text{eq } \text{term}_1 \text{ term}_2$$

where *eq* is the Leibniz equality or a registered setoid equality. The tactic:

```
rewrite r-prefix (t1, ..., tn) .
```

is equivalent to:

```
do [rewrite r-prefix t1 | ... | rewrite
r-prefix tn] .
```

- An anonymous rewrite lemma $(_ : term)$, where $term$ has again the form:

`foralll` $(x_1 : A_1) \dots (x_n : A_n), eq\ term_1\ term_2$

The tactic:

`rewrite` $(_ : term)$

is in fact synonym of:

`cutrewrite` $(term)$.

11.7.2 Remarks and examples

Rewrite redex selection

The general strategy of SSREFLECT is to grasp as many redexes as possible and to let the user select the ones to be rewritten thanks to the improved syntax for the control of rewriting.

This may be a source of incompatibilities between the two `rewrite` tactics.

In a rewrite tactic of the form:

`rewrite` *occ-switch* $[term_1]\ term_2$.

$term_1$ is the explicit rewrite redex and $term_2$ is the rewrite rule. This execution of this tactic unfolds as follows:

- First $term_1$ and $term_2$ are $\beta\iota$ normalized. Then $term_2$ is put in head normal form if the Leibniz equality constructor `eq` is not the head symbol. This may involve ζ reductions.
- Then, the matching algorithm (see section 11.4.2) determines the first subterm of the goal matching the rewrite pattern. The rewrite pattern is given by $term_1$, if an explicit redex pattern switch is provided, or by the type of $term_2$ otherwise. However, matching skips over matches that would lead to trivial rewrites. All the occurrences of this subterm in the goal are candidates for rewriting.
- Then only the occurrences coded by *occ-switch* (see again section 11.4.2) are finally selected for rewriting.
- The left hand side of $term_2$ is unified with the subterm found by the matching algorithm, and if this succeeds, all the selected occurrences in the goal are replaced by the right hand side of $term_2$.
- Finally the goal is $\beta\iota$ normalized.

In the case $term_2$ is a list of terms, the first top-down (in the goal) left-to-right (in the list) matching rule gets selected.

Chained rewrite steps

The possibility to chain rewrite operations in a single tactic makes scripts more compact and gathers in a single command line a bunch of surgical operations which would be described by a one sentence in a pen and paper proof.

Performing rewrite and simplification operations in a single tactic enhances significantly the concision of scripts. For instance the tactic:

`rewrite` /my_def {2} [f _] /= my_eq // =.

unfolds `my_def` in the goal, simplifies the second occurrence of the first subterm matching pattern `[f _]`, rewrites `my_eq`, simplifies the whole goal and closes trivial goals.

Here are some concrete examples of chained rewrite operations, in the proof of basic results on natural numbers arithmetic:

```
Lemma addnS : forall m n, m + n.+1 = (m + n).+1.
Proof. by move=> m n; elim: m. Qed.
```

```
Lemma addSnnS : forall m n, m.+1 + n = m + n.+1.
Proof. move=> *; rewrite addnS; apply addSn. Qed.
```

```
Lemma addnCA : forall m n p, m + (n + p) = n + (m + p).
Proof. by move=> m n; elim: m => [|m Hrec] p; rewrite ?addSnnS -?
  addnS. Qed.
```

```
Lemma addnC : forall m n, m + n = n + m.
Proof. by move=> m n; rewrite -{1}[n]addn0 addnCA addn0. Qed.
```

Note the use of the `?` switch for parallel rewrite operations in the proof of `addnCA`.

Explicit redex switches are matched first

If an *r-prefix* involves a *redex switch*, the first step is to find a subterm matching this redex pattern, independently from the left hand side `tl` of the equality the user wants to rewrite.

For instance, if `H : forall t u, t + u = u + t` is in the context of a goal `x + y = y + x`, the tactic:

```
rewrite [y + _]H.
```

transforms the goal into `x + y = x + y`.

Note that if this first pattern matching is not compatible with the *r-item*, the rewrite fails, even if the goal contains a correct redex matching both the redex switch and the left hand side of the equality. For instance, if `H : forall t u, t + u * 0 = t` is in the context of a goal `x + y * 4 + 2 * 0 = x + 2 * 0`, then tactic:

```
rewrite [x + _]H.
```

raises the error message:

```
User error: rewrite rule H doesn't match redex (x + y * 4)
```

while the tactic:

```
rewrite (H _ 2).
```

transforms the goal into `x + y * 4 = x + 2 * 0`.

Occurrence switches and redex switches

The tactic:

```
rewrite {2}[_ + y + 0](_ : forall z, z + 0 = z).
```

transforms the goal:

$$x + y + 0 = x + y + y + 0 + 0 + (x + y + 0)$$

into:

$$x + y + 0 = x + y + y + 0 + 0 + (x + y)$$

and generates a second subgoal:

```
forall z : nat, z + 0 = z
```

The second subgoal is generated by the use of an anonymous lemma in the rewrite tactic. The effect of the tactic on the initial goal is to rewrite this lemma at the second occurrence of the first matching $x + y + 0$ of the explicit rewrite `redex _ + y + 0`.

Occurrence selection and repetition

Occurrence selection has priority over repetition switches. This means the repetition of a rewrite tactic specified by a multiplier will perform matching each time an elementary rewrite operation is performed. Repeated rewrite tactics apply to every subgoal generated by the previous tactic, including the previous instances of the repetition. For example:

```
Goal forall x y z : nat, x + 1 = x + y + 1.
move=> x y z.
```

creates a goal $x + 1 = x + y + 1$, which is turned into $z = z$ by the additional tactic:

```
rewrite 2!(_ : _ + 1 = z).
```

In fact, this last tactic generates *three* subgoals, respectively $x + y + 1 = z$, $z = z$ and $x + 1 = z$. Indeed, the second rewrite operation specified with the $2!$ multiplier applies to the two subgoals generated by the first rewrite.

Multi-rule rewriting

The `rewrite` tactic can be provided a *tuple* of rewrite rules, or more generally a tree of such rules, since this tuple can feature arbitrary inner parentheses. We call *multirule* such a generalized rewrite rule. This feature is of special interest when it is combined with multiplier switches, which makes the `rewrite` tactic iterates the rewrite operations prescribed by the rules on the current goal. For instance, let us define two triples `multi1` and `multi2` as:

```
Variables (a b c : nat).
```

```
Hypothesis eqab : a = b.
```

```
Hypothesis eqac : a = c.
```

Executing the tactic:

```
rewrite (eqab, eqac)
```

on the goal:

```
=====
a = a
```

turns it into $b = b$, as rule `eqab` is the first to apply among the ones gathered in the tuple passed to the `rewrite` tactic. This multirule `(eqab, eqac)` is actually a COQ term and we can name it with a definition:

```
Definition multi1 := (eqab, eqac).
```

In this case, the tactic `rewrite multi1` is a synonym for `(eqab, eqac)`. More precisely, a multirule rewrites the first subterm to which one of the rules applies in a left-to-right traversal of the goal, with the first rule from the multirule tree in left-to-right order. Matching is performed according to the algorithm described in Section 11.4.2, but literal matches have priority. For instance if we add a definition and a new multirule to our context:

```
Definition d := a.
```

```
Hypotheses eqd0 : d = 0.
```

```
Definition multi2 := (eqab, eqd0).
```

then executing the tactic:

```
rewrite multi2.
```

on the goal:

```
=====
d = b
```

turns it into $0 = b$, as rule `eqd0` applies without unfolding the definition of `d`. For repeated rewrites the selection process is repeated anew. For instance, if we define:

```
Hypothesis eq_adda_b : forall x, x + a = b.
```

```
Hypothesis eq_adda_c : forall x, x + a = c.
```

```
Hypothesis eqb0 : b = 0.
```

```
Definition multi3 := (eq_adda_b, eq_adda_c, eqb0).
```

then executing the tactic:

```
rewrite 2!multi3.
```

on the goal:

```
=====
1 + a = 12 + a
```

turns it into $0 = 12 + a$: it uses `eq_adda_b` then `eqb0` on the left-hand side only. Now executing the tactic `rewrite multi3!` turns the same goal into $0 = 0$.

The grouping of rules inside a multirule does not affect the selection strategy but can make it easier to include one rule set in another or to (universally) quantify over the parameters of a subset of rules (as there is special code that will omit unnecessary quantifiers for rules that can be syntactically extracted). It is also possible to reverse the direction of a rule subset, using a special dedicated syntax: the tactic `rewrite (~ multi1)` is equivalent to `rewrite multi1_rev` with:

```
Hypothesis eqba : b = a.
```

```
Hypothesis eqca : c = a.
```

```
Definition multil_rev := (eqba, eqca).
```

except that the constants `eqba`, `eqab`, `multl_rev` have not been created.

Rewriting with multirules is useful to implement simplification or transformation procedures, to be applied on terms of small to medium size. For instance the library `ssrnat` provides two implementations for arithmetic operations on natural numbers: an elementary one and a tail recursive version, less inefficient but also less convenient for reasoning purposes. The library also provides one lemma per such operation, stating that both versions return the same values when applied to the same arguments:

```
Lemma addE : add =2 addn.
```

```
Lemma doubleE : double =1 doublen.
```

```
Lemma add_mulE n m s : add_mul n m s = addn (muln n m) s.
```

```
Lemma mulE : mul =2 muln.
```

```
Lemma mul_expE m n p : mul_exp m n p = muln (expn m n) p.
```

```
Lemma expE : exp =2 expn.
```

```
Lemma oddE : odd =1 oddn.
```

The operation on the left hand side of each lemma is the efficient version, and the corresponding naive implementation is on the right hand side. In order to reason conveniently on expressions involving the efficient operations, we gather all these rules in the definition `trecE`:

```
Definition trecE := (addE, (doubleE, oddE), (mulE, add_mulE, (expE,
  mul_expE))).
```

The tactic:

```
rewrite !trecE.
```

restores the naive versions of each operation in a goal involving the efficient ones, e.g. for the purpose of a correctness proof.

Wildcards vs abstractions

The `rewrite` tactic supports `r`-items containing holes. For example in the tactic (1):

```
rewrite (_ : _ * 0 = 0).
```

the term `_ * 0 = 0` is interpreted as `forall n : nat, n * 0 = 0`. Anyway this tactic is *not* equivalent to the tactic (2):

```
rewrite (_ : forall x, x * 0 = 0).
```

The tactic (1) transforms the goal $(y * 0) + y * (z * 0) = 0$ into $y * (z * 0) = 0$ and generates a new subgoal to prove the statement $y * 0 = 0$, which is the *instance* of the `forall x, x * 0 = 0` rewrite rule that has been used to perform the rewriting. On the other hand, tactic (2) performs the same rewriting on the current goal but generates a subgoal to prove `forall x, x * 0 = 0`.

When SSREFLECT `rewrite` fails on standard COQ `licit` rewrite

In a few cases, the SSREFLECT `rewrite` tactic fails rewriting some redexes which standard COQ successfully rewrites. There are two main cases:

- SSREFLECT never accepts to rewrite indeterminate patterns like:

```
Lemma foo : forall x : unit, x = tt.
```

SSREFLECT will however accept the $\eta\zeta$ expansion of this rule:

```
Lemma fubar : forall x : unit, (let u := x in u) = tt.
```

- In standard COQ, suppose that we work in the following context:

```
Variable g : nat -> nat.
Definition f := g.
```

then rewriting `H : forall x, f x = 0` in the goal `g 3 + g 3 = g 6` succeeds and transforms the goal into `0 + 0 = g 6`.

This rewriting is not possible in SSREFLECT because there is no occurrence of the head symbol `f` of the rewrite rule in the goal. Rewriting with `H` first requires unfolding the occurrences of `f` where the substitution is to be performed (here there is a single such occurrence), using tactic `rewrite /f` (for a global replacement of `f` by `g`) or `rewrite pattern/f`, for a finer selection.

Existential metavariables and rewriting

The `rewrite` tactic will not instantiate existing existential metavariables when matching a redex pattern.

If a rewrite rule generates a goal with new existential metavariables, these will be generalized as for `apply` (see page 314) and corresponding new goals will be generated. For example, consider the following script:

```
Lemma ex3 (x : 'I_2) y (le_1 : y < 1) (E : val x = y) : Some x =
  insub y.
rewrite insubT ?(leq_trans le_1)// => le_2.
```

Since `insubT` has the following type:

```
forall T P (sT : subType P) (x : T) (Px : P x), insub x = Some (Sub
  x Px)
```

and since the implicit argument corresponding to the `Px` abstraction is not supplied by the user, the resulting goal should be `Some x = Some (Sub y ?Px)`. Instead, SSREFLECT `rewrite` tactic generates the two following goals:

```
y < 2
forall Hyp0 : y < 2, Some x = Some (Sub y Hyp0)
```

The script closes the former with `?(leq_trans le_1)//`, then it introduces the new generalization naming it `le_2`.

```

x : 'I_2
y : nat
le_1 : y < 1
E : val x = y
le_2 : y < 2
=====
Some x = Some (Sub y le_2)

```

As a temporary limitation, this behavior is available only if the rewriting rule is stated using Leibniz equality (as opposed to setoid relations). It will be extended to other rewriting relations in the future.

11.7.3 Locking, unlocking

As program proofs tend to generate large goals, it is important to be able to control the partial evaluation performed by the simplification operations that are performed by the tactics. These evaluations can for example come from a `/=` simplification switch, or from rewrite steps which may expand large terms while performing conversion. We definitely want to avoid repeating large subterms of the goal in the proof script. We do this by “clamping down” selected function symbols in the goal, which prevents them from being considered in simplification or rewriting steps. This clamping is accomplished by using the occurrence switches (see section 11.4.2) together with “term tagging” operations.

SSREFLECT provides two levels of tagging.

The first one uses auxiliary definitions to introduce a provably equal copy of any term t . However this copy is (on purpose) *not convertible* to t in the COQ system⁸. The job is done by the following construction:

```

Lemma master_key : unit. Proof. exact tt. Qed.
Definition locked A := let: tt := master_key in fun x : A => x.
Lemma lock : forall A x, x = locked x :> A.

```

Note that the definition of `master_key` is explicitly opaque. The equation $t = \text{locked } t$ given by the `lock` lemma can be used for selective rewriting, blocking on the fly the reduction in the term t . For example the script:

```

Require Import List.
Variable A : Type.

Fixpoint my_has (p : A -> bool) (l : list A) {struct l} : bool :=
  match l with
  | nil => false
  | cons x l => p x || (my_has p l)
  end.

Goal forall a x y l, a x = true -> my_has a ( x :: y :: l ) = true.
move=> a x y l Hax.

```

where `||` denotes the boolean disjunction, results in a goal `my_has a (x :: y :: l) = true`. The tactic:

```

rewrite {2}[cons]lock /= -lock.

```

⁸This is an implementation feature: there is not such obstruction in the metatheory

turns it into `a x || my_has a (y :: l) = true`. Let us now start by reducing the initial goal without blocking reduction. The script:

```
Goal forall a x y l, a x = true -> my_has a ( x :: y :: l) = true.
move=> a x y l Hax /=.

```

creates a goal `(a x) || (a y) || (my_has a l) = true`. Now the tactic:

```
rewrite {1}[orb]lock orbC -lock.

```

where `orbC` states the commutativity of `orb`, changes the goal into

`(a x) || (my_has a l) || (a y) = true`: only the arguments of the second disjunction where permuted.

It is sometimes desirable to globally prevent a definition from being expanded by simplification; this is done by adding `locked` in the definition.

For instance, the function `fgraph_of_fun` maps a function whose domain and codomain are finite types to a concrete representation of its (finite) graph. Whatever implementation of this transformation we may use, we want it to be hidden to simplifications and tactics, to avoid the collapse of the graph object:

```
Definition fgraph_of_fun :=
  locked
  (fun (d1 : finType) (d2 : eqType) (f : d1 -> d2) => Fgraph (
    size_maps f _)).

```

We provide a special tactic `unlock` for unfolding such definitions while removing “locks”, e.g., the tactic:

```
unlock occ-switch fgraph_of_fun.

```

replaces the occurrence(s) of `fgraph_of_fun` coded by the `occ-switch` with `(Fgraph (size_maps _ _))` in the goal.

We found that it was usually preferable to prevent the expansion of some functions by the partial evaluation switch “`/=`”, unless this allowed the evaluation of a condition. This is possible thanks to an other mechanism of term tagging, resting on the following *Notation*:

```
Notation "'nosimpl' t" := (let: tt := tt in t).

```

The term `(nosimpl t)` simplifies to `t` *except* in a definition. More precisely, given:

```
Definition foo := (nosimpl bar).

```

the term `foo` (or `(foo t')`) will *not* be expanded by the `simpl` tactic unless it is in a forcing context (e.g., in `match foo t' with...end`, `foo t'` will be reduced if this allows `match` to be reduced). Note that `nosimpl bar` is simply notation for a term that reduces to `bar`; hence `unfold foo` will replace `foo` by `bar`, and `fold foo` will replace `bar` by `foo`.

Warning The `nosimpl` trick only works if no reduction is apparent in `t`; in particular, the declaration:

```
Definition foo x := nosimpl (bar x).

```

will usually not work. Anyway, the common practice is to tag only the function, and to use the following definition, which blocks the reduction as expected:

```
Definition foo x := nosimpl bar x.

```

A standard example making this technique shine is the case of arithmetic operations. We define for instance:

`Definition addn := nosimpl plus.`

The operation `addn` behaves exactly like `plus`, except that `(addn (S n) m)` will not simplify spontaneously to `(S (addn n m))` (the two terms, however, are inter-convertible). In addition, the unfolding step:

`rewrite /addn`

will replace `addn` directly with `plus`, so the `nosimpl` form is essentially invisible.

11.7.4 Congruence

Because of the way matching interferes with type families parameters, the tactic:

`apply: my_congr_property.`

will generally fail to perform congruence simplification, even on rather simple cases. We therefore provide a more robust alternative in which the function is supplied:

`congr [int] term`

This tactic:

- checks that the goal is a Leibniz equality
- matches both sides of this equality with “*term* applied to some arguments”, inferring the right number of arguments from the goal and the type of *term*. This may expand some definitions or fixpoints.
- generates the subgoals corresponding to pairwise equalities of the arguments present in the goal.

The goal can be a non dependent product $P \rightarrow Q$. In that case, the system asserts the equation $P = Q$, uses it to solve the goal, and calls the `congr` tactic on the remaining goal $P = Q$. This can be useful for instance to perform a transitivity step, like in the following situation:

```
x, y, z : nat
=====
x = y -> x = z
```

the tactic `congr` `(_ = _)` turns this goal into:

```
x, y, z : nat
=====
y = z
```

which can also be obtained starting from:

```
x, y, z : nat
h : x = y
=====
x = z
```

and using the tactic `congr` `(_ = _) : h`.

The optional *int* forces the number of arguments for which the tactic should generate equality proof obligations.

This tactic supports equalities between applications with dependent arguments. Yet dependent arguments should have exactly the same parameters on both sides, and these parameters should appear as first arguments.

The following script:

```
Definition f n := match n with 0 => plus | S _ => mult end.
Definition g (n m : nat) := plus.

Goal forall x y, f 0 x y = g 1 1 x y.
by move=> x y; congr plus.
Qed.
```

shows that the `congr` tactic matches `plus` with `f 0` on the left hand side and `g 1 1` on the right hand side, and solves the goal.

The script:

```
Goal forall n m, m <= n -> S m + (S n - S m) = S n.
move=> n m Hnm; congr S; rewrite -/plus.
```

generates the subgoal $m + (S n - S m) = n$. The tactic `rewrite -/plus` folds back the expansion of `plus` which was necessary for matching both sides of the equality with an application of `S`.

Like most `SSREFLECT` arguments, *term* can contain wildcards. The script:

```
Goal forall x y, x + (y * (y + x - x)) = x * 1 + (y + 0) * y.
move=> x y; congr ( _ + ( _ * _ ) ).
```

generates three subgoals, respectively $x = x * 1$, $y = y + 0$ and $y + x - x = y$.

11.8 Contextual patterns

The simple form of patterns used so far, *terms* possibly containing wild cards, often require an additional *occ-switch* to be specified. While this may work pretty fine for small goals, the use of polymorphic functions and dependent types may lead to an invisible duplication of functions arguments. These copies usually end up in types hidden by the implicit arguments machinery or by user defined notations. In these situations computing the right occurrence numbers is very tedious because they must be counted on the goal as printed after setting the `Printing All` flag. Moreover the resulting script is not really informative for the reader, since it refers to occurrence numbers he cannot easily see.

Contextual patterns mitigate these issues allowing to specify occurrences according to the context they occur in.

11.8.1 Syntax

The following table summarizes the full syntax of *c-pattern* and the corresponding subterm(s) identified by the pattern. In the third column we use s.m.r. for “the subterms matching the redex” specified in the second column.

<i>c-pattern</i>	redex	subterms affected
<i>term</i>	<i>term</i>	all occurrences of <i>term</i>
<i>ident in term</i>	subterm of <i>term</i> selected by <i>ident</i>	all the subterms identified by <i>ident</i> in all the occurrences of <i>term</i>
<i>term₁ in ident in term₂</i>	<i>term₁</i>	in all s.m.r. in all the subterms identified by <i>ident</i> in all the occurrences of <i>term₂</i>
<i>term₁ as ident in term₂</i>	<i>term₁</i>	in all the subterms identified by <i>ident</i> in all the occurrences of <i>term₂[term₁/ident]</i>

The `rewrite` tactic supports two more patterns obtained prefixing the first two with `in`. The intended meaning is that the pattern identifies all subterms of the specified context. The `rewrite` tactic will infer a pattern for the redex looking at the rule used for rewriting.

<i>r-pattern</i>	redex	subterms affected
<code>in term</code>	inferred from rule	in all s.m.r. in all occurrences of <i>term</i>
<code>in ident in term</code>	inferred from rule	in all s.m.r. in all the subterms identified by <i>ident</i> in all the occurrences of <i>term</i>

The first *c-pattern* is the simplest form matching any context but selecting a specific redex and has been described in the previous sections. We have seen so far that the possibility of selecting a redex using a term with holes is already a powerful mean of redex selection. Similarly, any *terms* provided by the user in the more complex forms of *c-patterns* presented in the tables above can contain holes.

For a quick glance at what can be expressed with the last *r-pattern* consider the goal $a = b$ and the tactic

```
rewrite [in X in _ = X] rule.
```

It rewrites all occurrences of the left hand side of `rule` inside `b` only (`a`, and the hidden type of the equality, are ignored). Note that the variant `rewrite [X in _ = X] rule` would have rewritten `b` exactly (i.e., it would only work if `b` and the left hand side of `rule` can be unified).

11.8.2 Matching contextual patterns

The *c-patterns* and *r-patterns* involving *terms* with holes are matched against the goal in order to find a closed instantiation. This matching proceeds as follows:

<i>c-pattern</i>	instantiation order and place for <i>term_i</i> and redex
<i>term</i>	<i>term</i> is matched against the goal, redex is unified with the instantiation of <i>term</i>
<i>ident in term</i>	<i>term</i> is matched against the goal, redex is unified with the subterm of the instantiation of <i>term</i> identified by <i>ident</i>
<i>term₁ in ident in term₂</i>	<i>term₂</i> is matched against the goal, <i>term₁</i> is matched against the subterm of the instantiation of <i>term₁</i> identified by <i>ident</i> , redex is unified with the instantiation of <i>term₁</i>
<i>term₁ as ident in term₂</i>	<i>term₂[term₁/ident]</i> is matched against the goal, redex is unified with the instantiation of <i>term₁</i>

In the following patterns, the redex is intended to be inferred from the rewrite rule.

<i>r-pattern</i>	instantiation order and place for $term_i$ and redex
<code>in ident in term</code>	<i>term</i> is matched against the goal, the redex is matched against the subterm of the instantiation of <i>term</i> identified by <i>ident</i>
<code>in term</code>	<i>term</i> is matched against the goal, redex is matched against the instantiation of <i>term</i>

11.8.3 Examples

Contextual pattern in `set` and the `:` tactical

As already mentioned in section 11.4.2 the `set` tactic takes as an argument a term in open syntax. This term is interpreted as the simplest form of *c-pattern*. To void confusion in the grammar, open syntax is supported only for the simplest form of patterns, while parentheses are required around more complex patterns.

```
set t := (X in _ = X) .
set t := (a + _ in X in _ = X) .
```

Given the goal $a + b + 1 = b + (a + 1)$ the first tactic captures $b + (a + 1)$, while the latter $a + 1$.

Since the user may define an infix notation for `in` the former tactic may result ambiguous. The disambiguation rule implemented is to prefer patterns over simple terms, but to interpret a pattern with double parentheses as a simple term. For example the following tactic would capture any occurrence of the term ‘`a in A`’.

```
set t := ((a in A)) .
```

Contextual pattern can also be used as arguments of the `:` tactical. For example:

```
elim: n (n in _ = n) (refl_equal n) .
```

Contextual patterns in `rewrite`

As a more comprehensive example consider the following goal:

$$(x.+1 + y) + f (x.+1 + y) (z + (x + y).+1) = 0$$

The tactic `rewrite [in f _ _] addSn` turns it into:

$$(x.+1 + y) + f (x + y).+1 (z + (x + y).+1) = 0$$

since the simplification rule `addSn` is applied only under the `f` symbol. Then we simplify also the first addition and expand `0` into `0+0`.

```
rewrite addSn -[X in _ = X] addn0 .
```

obtaining:

$$(x + y).+1 + f (x + y).+1 (z + (x + y).+1) = 0 + 0$$

Note that the right hand side of `addn0` is undetermined, but the `rewrite` pattern specifies the redex explicitly. The right hand side of `addn0` is unified with the term identified by `X`, `0` here.

The following pattern does not specify a redex, since it identifies an entire region, hence the `rewrite` rule has to be instantiated explicitly. Thus the tactic:

```
rewrite -{2}[in X in _ = X] (addn0 0) .
```

changes the goal as follows:

$$(x + y).+1 + f (x + y).+1 (z + (x + y).+1) = 0 + (0 + 0)$$

The following tactic is quite tricky:

```
rewrite [_.+1 in X in f _ X] (addnC x.+1) .
```

and the resulting goals is:

$$(x + y).+1 + f (x + y).+1 (z + (y + x.+1)) = 0 + (0 + 0)$$

The explicit redex $_.+1$ is important since its head constant S differs from the head constant inferred from $(\text{addnC } x.+1)$ (that is addn , denoted $+$ here). Moreover, the pattern $f _ X$ is important to rule out the first occurrence of $(x + y).+1$. Last, only the subterms of $f _ X$ identified by X are rewritten, thus the first argument of f is skipped too. Also note the pattern $_.+1$ is interpreted in the context identified by X , thus it gets instantiated to $(y + x).+1$ and not $(x + y).+1$.

The last rewrite pattern allows to specify exactly the shape of the term identified by X , that is thus unified with the left hand side of the rewrite rule.

```
rewrite [x.+1 + y as X in f X _] addnC.
```

The resulting goal is:

$$(x + y).+1 + f (y + x.+1) (z + (y + x.+1)) = 0 + (0 + 0)$$

11.8.4 Patterns for recurrent contexts

The user can define shortcuts for recurrent contexts corresponding to the *ident in term* part. The notation scope identified with `%pattern` provides a special notation ‘ $(X \text{ in } t)$ ’ the user must adopt to define context shortcuts.

The following example is taken from `ssreflect.v` where the LHS and RHS shortcuts are defined.

```
Notation RHS := (X in _ = X) %pattern.
```

```
Notation LHS := (X in X = _) %pattern.
```

Shortcuts defined this way can be freely used in place of the trailing *ident in term* part of any contextual pattern. Some examples follow:

```
set rhs := RHS.
rewrite [in RHS] rule.
case: (a + _ in RHS) .
```

11.9 Views and reflection

The bookkeeping facilities presented in section 11.5 are crafted to ease simultaneous introductions and generalizations of facts and casing, naming ... operations. It also a common practice to make a stack operation immediately followed by an *interpretation* of the fact being pushed, that is, to apply a lemma to this fact before passing it to a tactic for decomposition, application and so on.

SSREFLECT provides a convenient, unified syntax to combine these interpretation operations with the proof stack operations. This *view mechanism* relies on the combination of the `/ view` switch with bookkeeping tactics and tacticals.

11.9.1 Interpreting eliminations

The view syntax combined with the `elim` tactic specifies an elimination scheme to be used instead of the default, generated, one. Hence the `SSREFLECT` tactic:

```
elim/V.
```

is a synonym for:

```
intro top; elim top using V; clear top.
```

where `top` is a fresh name and `V` any second-order lemma.

Since an elimination view supports the two bookkeeping tacticals of discharge and introduction (see section 11.5), the `SSREFLECT` tactic:

```
elim/V: x => y.
```

is a synonym for:

```
elim x using V; clear x; intro y.
```

where `x` is a variable in the context, `y` a fresh name and `V` any second order lemma; `SSREFLECT` relaxes the syntactic restrictions of the COQ `elim`. The first pattern following `:` can be a `_` wildcard if the conclusion of the view `V` specifies a pattern for its last argument (e.g., if `V` is a functional induction lemma generated by the `Function` command).

The elimination view mechanism is compatible with the equation name generation (see section 11.5.5).

The following script illustrate a toy example of this feature. Let us define a function adding an element at the end of a list:

```
Require Import List.
```

```
Variable d : Type.
```

```
Fixpoint add_last(s : list d) (z : d) {struct s} : list d :=
  match s with
  | nil => z :: nil
  | cons x s' => cons x (add_last s' z)
  end.
```

One can define an alternative, reversed, induction principle on inductively defined lists, by proving the following lemma:

```
Lemma last_ind_list : forall (P : list d -> Type),
P nil ->
(forall (s : list d) (x : d), P s -> P (add_last s x)) -> forall s
: list d, P s.
```

Then the combination of elimination views with equation names result in a concise syntax for reasoning inductively using the user defined elimination scheme. The script:

```
Goal forall (x : d) (l : list d), l = l.
move=> x l.
elim/last_ind_list E : l=> [| u v]; last first.
```

generates two subgoals: the first one to prove $\text{nil} = \text{nil}$ in a context featuring $E : l = \text{nil}$ and the second to prove $\text{add_last } u \ v = \text{add_last } u \ v$, in a context containing $E : l = \text{add_last } u \ v$.

User provided eliminators (potentially generated with the `Function COQ's` command) can be combined with the type family switches described in section 11.5.6. Consider an eliminator `foo_ind` of type:

```
foo_ind : forall..., forall x : T, P p1...pm
```

and consider the tactic

```
elim/foo_ind: e1.../ e_n
```

The `elim/` tactic distinguishes two cases:

truncated eliminator when x does not occur in $P \ p_1 \dots p_m$ and the type of e_n unifies with T and e_n is not `_`. In that case, e_n is passed to the eliminator as the last argument (x in `foo_ind`) and $e_{n-1} \dots e_1$ are used as patterns to select in the goal the occurrences that will be bound by the predicate P , thus it must be possible to unify the sub-term of the goal matched by e_{n-1} with p_m , the one matched by e_{n-2} with p_{m-1} and so on.

regular eliminator in all the other cases. Here it must be possible to unify the term matched by e_n with p_m , the one matched by e_{n-1} with p_{m-1} and so on. Note that standard eliminators have the shape `...forall x, P...x`, thus e_n is the pattern identifying the eliminated term, as expected.

As explained in section 11.5.6, the initial prefix of e_i can be omitted.

Here an example of a regular, but non trivial, eliminator:

```
Function plus (m n : nat) {struct n} : nat :=
  match n with 0 => m | S p => S (plus m p) end.
```

The type of `plus_ind` is

```
plus_ind : forall (m : nat) (P : nat -> nat -> Prop),
  (forall n : nat, n = 0 -> P 0 m) ->
  (forall n p : nat, n = p.+1 -> P p (plus m p) -> P p.+1 (plus m p)
    .+1) ->
  forall n : nat, P n (plus m n)
```

Consider the following goal

```
Lemma exF x y z: plus (plus x y) z = plus x (plus y z).
```

The following tactics are all valid and perform the same elimination on that goal.

```
elim/plus_ind: z / (plus _ z).
elim/plus_ind: {z}(plus _ z).
elim/plus_ind: {z}_.
elim/plus_ind: z / _.
```

In the two latter examples, being the user provided pattern a wildcard, the pattern inferred from the type of the eliminator is used instead. For both cases it is `(plus _ _)` and matches the subterm `plus (plus x y) z` thus instantiating the latter `_` with `z`. Note that the tactic `elim/plus_ind: y / _` would have resulted in an error, since y and z do not unify but the type of the eliminator requires the second argument of P to be the same as the second argument of `plus` in the second argument of P .

Here an example of a truncated eliminator. Consider the goal

```

p : nat_eqType
n : nat
n_gt0 : 0 < n
pr_p : prime p
=====
p %| \prod_(i <- prime_decomp n | i \in prime_decomp n) i.1 ^ i.2
->
  exists2 x : nat * nat, x \in prime_decomp n & p = x.1

```

and the tactic

```
elim/big_prop: _ => [| u v IHu IHv | [q e] /=].
```

where the type of the eliminator is

```

big_prop: forall (R : Type) (Pb : R -> Type) (idx : R) (op1 : R -> R
-> R),
  Pb idx ->
  (forall x y : R, Pb x -> Pb y -> Pb (op1 x y)) ->
  forall (I : Type) (r : seq I) (P : pred I) (F : I -> R),
  (forall i : I, P i -> Pb (F i)) ->
  Pb (\big[op1/idx]_(i <- r | P i) F i)

```

Since the pattern for the argument of Pb is not specified, the inferred one is used instead: `(\big[_/_]_(i <- _ | _ i) _ i)`, and after the introductions, the following goals are generated.

```

subgoal 1 is:
p %| 1 -> exists2 x : nat * nat, x \in prime_decomp n & p = x.1
subgoal 2 is:
p %| u * v -> exists2 x : nat * nat, x \in prime_decomp n & p = x.1
subgoal 3 is:
(q, e) \in prime_decomp n -> p %| q ^ e ->
  exists2 x : nat * nat, x \in prime_decomp n & p = x.1

```

Note that the pattern matching algorithm instantiated all the variables occurring in the pattern.

11.9.2 Interpreting assumptions

Interpreting an assumption in the context of a proof is applying it a correspondence lemma before generalizing, and/or decomposing it. For instance, with the extensive use of boolean reflection (see section 11.9.4), it is quite frequent to need to decompose the logical interpretation of (the boolean expression of) a fact, rather than the fact itself. This can be achieved by a combination of `move : _ => _` switches, like in the following script, where `||` is a notation for the boolean disjunction:

```

Variables P Q : bool -> Prop.
Hypothesis P2Q : forall a b, P (a || b) -> Q a.

Goal forall a, P (a || a) -> True.
move=> a HPa; move: {HPa} (P2Q _ _ HPa) => HQa.

```

which transforms the hypothesis $HP_n : P \ n$ which has been introduced from the initial statement into $HQ_n : Q \ n$. This operation is so common that the tactic shell has specific syntax for it. The following scripts:

```
Goal forall a, P (a || a) -> True.
move=> a HPa; move/P2Q: HPa => HQa.
```

or more directly:

```
Goal forall a, P (a || a) -> True.
move=> a; move/P2Q=> HQa.
```

are equivalent to the former one. The former script shows how to interpret a fact (already in the context), thanks to the discharge tactical (see section 11.5.3) and the latter, how to interpret the top assumption of a goal. Note that the number of wildcards to be inserted to find the correct application of the view lemma to the hypothesis has been automatically inferred.

The view mechanism is compatible with the `case` tactic and with the equation name generation mechanism (see section 11.5.5):

```
Variables P Q: bool -> Prop.
Hypothesis Q2P : forall a b, Q (a || b) -> P a \/ P b.

Goal forall a b, Q (a || b) -> True.
move=> a b; case/Q2P=> [HPa | HPb].
```

creates two new subgoals whose contexts no more contain $HQ : Q \ (a \ || \ b)$ but respectively $HPa : P \ a$ and $HPb : P \ b$. This view tactic performs:

```
move=> a b HQ; case: {HQ} (Q2P _ _ HQ) => [HPa | HPb].
```

The term on the right of the `/` view switch is called a *view lemma*. Any SSREFLECT term coercing to a product type can be used as a view lemma.

The examples we have given so far explicitly provide the direction of the translation to be performed. In fact, view lemmas need not to be oriented. The view mechanism is able to detect which application is relevant for the current goal. For instance, the script:

```
Variables P Q: bool -> Prop.
Hypothesis PQequiv : forall a b, P (a || b) <-> Q a.

Goal forall a b, P (a || b) -> True.
move=> a b; move/PQequiv=> HQab.
```

has the same behavior as the first example above.

The view mechanism can insert automatically a *view hint* to transform the double implication into the expected simple implication. The last script is in fact equivalent to:

```
Goal forall a b, P (a || b) -> True.
move=> a b; move/(iffLR (PQequiv _ _)).
```

where:

```
Lemma iffLR : forall P Q, (P <-> Q) -> P -> Q.
```


Specializing assumptions

The special case when the *head symbol* of the view lemma is a wildcard is used to interpret an assumption by *specializing* it. The view mechanism hence offers the possibility to apply a higher-order assumption to some given arguments.

For example, the script:

```
Goal forall z, (forall x y, x + y = z -> z = x) -> z = 0.
move=> z; move/(_ 0 z).
```

changes the goal into:

```
(0 + z = z -> z = 0) -> z = 0
```

11.9.3 Interpreting goals

In a similar way, it is also often convenient to interpret a goal by changing it into an equivalent proposition. The view mechanism of SSREFLECT has a special syntax `apply/` for combining simultaneous goal interpretation operations and bookkeeping steps in a single tactic.

With the hypotheses of section 11.9.2, the following script, where `~~` denotes the boolean negation:

```
Goal forall a, P ((~~ a) || a).
move=> a; apply/PQequiv.
```

transforms the goal into $\mathcal{Q} \ (\sim\sim a)$, and is equivalent to:

```
Goal forall a, P ((~~ a) || a).
move=> a; apply: (iffRL (PQequiv _ _)).
```

where `iffRL` is the analogous of `iffRL` for the converse implication.

Any SSREFLECT term whose type coerces to a double implication can be used as a view for goal interpretation.

Note that the goal interpretation view mechanism supports both `apply` and `exact` tactics. As expected, a goal interpretation view command `exact/term` should solve the current goal or it will fail.

Warning Goal interpretation view tactics are *not* compatible with the bookkeeping tactical `=>` since this would be redundant with the `apply: term => _` construction.

11.9.4 Boolean reflection

In the Calculus of Inductive Construction, there is an obvious distinction between logical propositions and boolean values. On the one hand, logical propositions are objects of *sort Prop* which is the carrier of intuitionistic reasoning. Logical connectives in *Prop* are *types*, which give precise information on the structure of their proofs; this information is automatically exploited by COQ tactics. For example, COQ knows that a proof of $A \ \backslash / \ B$ is either a proof of A or a proof of B . The tactics `left` and `right` change the goal $A \ \backslash / \ B$ to A and B , respectively; dually, the tactic `case` reduces the goal $A \ \backslash / \ B \Rightarrow G$ to two subgoals $A \Rightarrow G$ and $B \Rightarrow G$.

On the other hand, `bool` is an inductive *datatype* with two constructors `true` and `false`. Logical connectives on `bool` are *computable functions*, defined by their truth tables, using case analysis:

```
Definition (b1 || b2) := if b1 then true else b2.
```

Properties of such connectives are also established using case analysis: the tactic `by case: b` solves the goal

```
b || ~~ b = true
```

by replacing `b` first by `true` and then by `false`; in either case, the resulting subgoal reduces by computation to the trivial `true = true`.

Thus, `Prop` and `bool` are truly complementary: the former supports robust natural deduction, the latter allows brute-force evaluation. SSREFLECT supplies a generic mechanism to have the best of the two worlds and move freely from a propositional version of a decidable predicate to its boolean version.

First, booleans are injected into propositions using the coercion mechanism:

```
Coercion is_true (b : bool) := b = true.
```

This allows any boolean formula `b` to be used in a context where COQ would expect a proposition, e.g., after `Lemma...:.` It is then interpreted as `(is_true b)`, i.e., the proposition `b = true`. Coercions are elided by the pretty-printer, so they are essentially transparent to the user.

11.9.5 The `reflect` predicate

To get all the benefits of the boolean reflection, it is in fact convenient to introduce the following inductive predicate `reflect` to relate propositions and booleans:

```
Inductive reflect (P : Prop) : bool -> Type :=
| Reflect_true : P => reflect P true
| Reflect_false : ~P => reflect P false.
```

The statement `(reflect P b)` asserts that `(is_true b)` and `P` are logically equivalent propositions.

For instance, the following lemma:

```
Lemma andP : forall b1 b2, reflect (b1 /\ b2) (b1 && b2).
```

relates the boolean conjunction `&&` to the logical one `/\`. Note that in `andP`, `b1` and `b2` are two boolean variables and the proposition `b1 /\ b2` hides two coercions. The conjunction of `b1` and `b2` can then be viewed as `b1 /\ b2` or as `b1 && b2`.

Expressing logical equivalences through this family of inductive types makes possible to take benefit from *rewritable equations* associated to the case analysis of COQ's inductive types.

Since the equivalence predicate is defined in COQ as:

```
Definition iff (A B : Prop) := (A -> B) /\ (B -> A).
```

where `/\` is a notation for `and`:

```
Inductive and (A B : Prop) : Prop :=
conj : A -> B -> and A B
```

This make case analysis very different according to the way an equivalence property has been defined.

For instance, if we have proved the lemma:

```
Lemma andE : forall b1 b2, (b1 /\ b2) <-> (b1 && b2).
```

let us compare the respective behaviours of `andE` and `andP` on a goal:

```
Goal forall b1 b2, if (b1 && b2) then b1 else ~~(b1 || b2).
```

The command:

```
move=> b1 b2; case (@andE b1 b2).
```

generates a single subgoal:

```
(b1 && b2 -> b1 /\ b2) -> (b1 /\ b2 -> b1 && b2) ->
  if b1 && b2 then b1 else ~~ (b1 || b2)
```

while the command:

```
move=> b1 b2; case (@andP b1 b2) .
```

generates two subgoals, respectively $b1 \wedge b2 \rightarrow b1$ and $\sim (b1 \wedge b2) \rightarrow \sim (b1 \vee b2)$.

Expressing reflection relation through the `reflect` predicate is hence a very convenient way to deal with classical reasoning, by case analysis. Using the `reflect` predicate allows moreover to program rich specifications inside its two constructors, which will be automatically taken into account during destruction. This formalisation style gives far more efficient specifications than quantified (double) implications.

A naming convention in SSREFLECT is to postfix the name of view lemmas with `P`. For example, `orP` relates `||` and `\/`, `negP` relates `~~` and `~`.

The view mechanism is compatible with `reflect` predicates.

For example, the script

```
Goal forall a b : bool, a -> b -> a /\ b.
move=> a b Ha Hb; apply/andP.
```

changes the goal $a \wedge b$ to $a \&\& b$ (see section 11.9.3).

Conversely, the script

```
Goal forall a b : bool, a /\ b -> a.
move=> a b; move/andP.
```

changes the goal $a \wedge b \rightarrow a$ into $a \&\& b \rightarrow a$ (see section 11.9.2).

The same tactics can also be used to perform the converse operation, changing a boolean conjunction into a logical one. The view mechanism guesses the direction of the transformation to be used i.e., the constructor of the `reflect` predicate which should be chosen.

11.9.6 General mechanism for interpreting goals and assumptions

Specializing assumptions

The SSREFLECT tactic:

```
move/(_ term1 ... termn)
```

is equivalent to the tactic:

```
intro top; generalize (top term1 ... termn); clear top.
```

where `top` is a fresh name for introducing the top assumption of the current goal.

Interpreting assumptions

The general form of an assumption view tactic is:

$$[\text{move} \mid \text{case}] / \text{term}_0$$

The term term_0 , called the *view lemma* can be:

- a (term coercible to a) function;

- a (possibly quantified) implication;
- a (possibly quantified) double implication;
- a (possibly quantified) instance of the `reflect` predicate (see section 11.9.5).

Let `top` be the top assumption in the goal.

There are three steps in the behaviour of an assumption view tactic:

- It first introduces `top`.
- If the type of `term0` is neither a double implication nor an instance of the `reflect` predicate, then the tactic automatically generalises a term of the form:

$$(term_0 \ term_1 \ \dots \ term_n)$$

where the terms `term1 ... termn` instantiate the possible quantified variables of `term0`, in order for `(term0 term1 ... termn top)` to be well typed.

- If the type of `term0` is an equivalence, or an instance of the `reflect` predicate, it generalises a term of the form:

$$(term_{vh} (term_0 \ term_1 \ \dots \ term_n))$$

where the term `termvh` inserted is called an *assumption interpretation view hint*.

- It finally clears `top`.

For a `case`/`term0` tactic, the generalisation step is replaced by a case analysis step.

View hints are declared by the user (see section 11.9.8) and are stored in the *Hint View* database. The proof engine automatically detects from the shape of the top assumption `top` and of the view lemma `term0` provided to the tactic the appropriate view hint in the database to be inserted.

If `term0` is a double implication, then the view hint *A* will be one of the defined view hints for implication. These hints are by default the ones present in the file `ssreflect.v`:

```
Lemma iffLR : forall P Q, (P <-> Q) -> P -> Q.
```

which transforms a double implication into the left-to-right one, or:

```
Lemma iffRL : forall P Q, (P <-> Q) -> Q -> P.
```

which produces the converse implication. In both cases, the two first *Prop* arguments are implicit.

If `term0` is an instance of the `reflect` predicate, then *A* will be one of the defined view hints for the `reflect` predicate, which are by default the ones present in the file `ssrbool.v`. These hints are not only used for choosing the appropriate direction of the translation, but they also allow complex transformation, involving negations. For instance the hint:

```
Lemma introN : forall (P : Prop) (b : bool), reflect P b -> ~ P ->
  ~~ b.
```

makes the following script:

```
Goal forall a b : bool, a -> b -> ~~ (a && b).
move=> a b Ha Hb. apply/andP.
```

transforms the goal into $\sim (a \ / \ b)$. In fact⁹ this last script does not exactly use the hint `introN`, but the more general hint:

```
Lemma introNTF : forall (P : Prop) (b c : bool),
  reflect P b -> (if c then ~ P else P) -> ~~ b = c
```

The lemma `introN` is an instantiation of `introNTF` using `c := true`.

Note that views, being part of *i-pattern*, can be used to interpret assertions too. For example the following script asserts `a && b` but actually used its propositional interpretation.

```
Lemma test (a b : bool) (pab : b && a) : b.
have /andP [pa ->] : (a && b) by rewrite andbC.
```

Interpreting goals

A goal interpretation view tactic of the form:

`apply / term0`

applied to a goal `top` is interpreted in the following way:

- If the type of `term0` is not an instance of the `reflect` predicate, nor an equivalence, then the term `term0` is applied to the current goal `top`, possibly inserting implicit arguments.
- If the type of `term0` is an instance of the `reflect` predicate or an equivalence, then a *goal interpretation view hint* can possibly be inserted, which corresponds to the application of a term $(term_{vh}(term_0_.._))$ to the current goal, possibly inserting implicit arguments.

Like assumption interpretation view hints, goal interpretation ones are user defined lemmas stored (see section 11.9.8) in the `Hint View` database bridging the possible gap between the type of `term0` and the type of the goal.

11.9.7 Interpreting equivalences

Equivalent boolean propositions are simply *equal* boolean terms. A special construction helps the user to prove boolean equalities by considering them as logical double implications (between their coerced versions), while performing at the same time logical operations on both sides.

The syntax of double views is:

`apply / terml / termr`

The term `terml` is the view lemma applied to the left hand side of the equality, `termr` is the one applied to the right hand side.

In this context, the identity view:

```
Lemma idP : reflect b1 b1.
```

is useful, for example the tactic:

```
apply / idP / idP .
```

⁹The current state of the proof shall be displayed by the `Show Proof` command of COQ proof mode.

transforms the goal $\sim\sim (b1 \mid\mid b2) = b3$ into two subgoals, respectively $\sim\sim (b1 \mid\mid b2) \rightarrow b3$ and $b3 \rightarrow \sim\sim (b1 \mid\mid b2)$.

The same goal can be decomposed in several ways, and the user may choose the most convenient interpretation. For instance, the tactic:

```
apply/norP/idP.
```

applied on the same goal $\sim\sim (b1 \mid\mid b2) = b3$ generates the subgoals $\sim\sim b1 /\ \sim\sim b2 \rightarrow b3$ and $b3 \rightarrow \sim\sim b1 /\ \sim\sim b2$.

11.9.8 Declaring new Hint Views

The database of hints for the view mechanism is extensible via a dedicated vernacular command. As library `ssrbool.v` already declares a corpus of hints, this feature is probably useful only for users who define their own logical connectives. Users can declare their own hints following the syntax used in `ssrbool.v`:

```
Hint View for tactic / ident [| natural]
```

where $tactic \in \{\text{move}, \text{apply}\}$, $ident$ is the name of the lemma to be declared as a hint, and $natural$ a natural number. If `move` is used as $tactic$, the hint is declared for assumption interpretation tactics, `apply` declares hints for goal interpretations. Goal interpretation view hints are declared for both simple views and left hand side views. The optional natural number $natural$ is the number of implicit arguments to be considered for the declared hint view lemma `name_of_the_lemma`.

The command:

```
Hint View for apply// ident [| natural].
```

with a double slash `//`, declares hint views for right hand sides of double views. See the files `ssreflect.v` and `ssrbool.v` for examples.

11.9.9 Multiple views

The hypotheses and the goal can be interpreted applying multiple views in sequence. Both `move` and `apply` can be followed by an arbitrary number of `/termi`. The main difference between the following two tactics

```
apply/v1/v2/v3.
apply/v1; apply/v2; apply/v3.
```

is that the former applies all the views to the principal goal. Applying a view with hypotheses generates new goals, and the second line would apply the view `v2` to all the goals generated by `apply/v1`. Note that the NO-OP intro pattern `-` can be used to separate two views, making the two following examples equivalent:

```
move=> /v1; move=> /v2.
move=> /v1-/v2.
```

The tactic `move` can be used together with the `in` tactical to pass a given hypothesis to a lemma. For example, if $P2Q : P \rightarrow Q$ and $Q2R : Q \rightarrow R$, the following tactic turns the hypothesis $p : P$ into $P : R$.

```
move/P2Q/Q2R in p.
```

If the list of views is of length two, **Hint Views** for interpreting equivalences are indeed taken into account, otherwise only single **Hint Views** are used.

11.10 SSREFLECT searching tool

SSREFLECT proposes an extension of the **Search** command. Its syntax is:

```
Search [pattern] [[-] [string[%key] | pattern]]* [in [[-] name ]+]
```

where *name* is the name of an open module. This command search returns the list of lemmas:

- whose *conclusion* contains a subterm matching the optional first *pattern*. A `-` reverses the test, producing the list of lemmas whose conclusion does not contain any subterm matching the pattern;
- whose name contains the given string. A `-` prefix reverses the test, producing the list of lemmas whose name does not contain the string. A string that contains symbols or is followed by a scope key, is interpreted as the constant whose notation involves that string (e.g., `+` for `addn`), if this is unambiguous; otherwise the diagnostic includes the output of the **Locate** vernacular command.
- whose statement, including assumptions and types, contains a subterm matching the next patterns. If a pattern is prefixed by `-`, the test is reversed;
- contained in the given list of modules, except the ones in the modules prefixed by a `-`.

Note that:

- As for regular terms, patterns can feature scope indications. For instance, the command:

```
Search _ ( _ + _ ) %N.
```

lists all the lemmas whose statement (conclusion or hypotheses) involve an application of the binary operation denoted by the infix `+` symbol in the `N` scope (which is SSREFLECT scope for natural numbers).

- Patterns with holes should be surrounded by parentheses.
- Search always volunteers the expansion of the notation, avoiding the need to execute **Locate** independently. Moreover, a string fragment looks for any notation that contains fragment as a substring. If the `ssrbool` library is imported, the command:

```
Search "~~".
```

answers :

```
"~~" is part of notation (~~ _)
In bool_scope, (~~ b) denotes negb b
negbT forall b : bool, b = false -> ~~ b
contra forall c b : bool, (c -> b) -> ~~ b -> ~~ c
introN forall (P : Prop) (b : bool), reflect P b -> ~ P -> ~~ b
```

- A diagnostic is issued if there are different matching notations; it is an error if all matches are partial.
- Similarly, a diagnostic warns about multiple interpretations, and signals an error if there is no default one.
- The command `Search in M.` is a way of obtaining the complete signature of the module `M`.
- Strings and pattern indications can be interleaved, but the first indication has a special status if it is a pattern, and only filters the conclusion of lemmas:

– The command :

```
Search (_ =1 _) "bij".
```

lists all the lemmas whose conclusion features a `'=1'` and whose name contains the string `bij`.

– The command :

```
Search "bij" (_ =1 _).
```

lists all the lemmas whose statement, including hypotheses, features a `'=1'` and whose name contains the string `bij`.

11.11 Synopsis and Index

Parameters

<i>d-tactic</i>	one of the <code>elim</code> , <code>case</code> , <code>congr</code> , <code>apply</code> , <code>exact</code> and <code>move</code> SSREFLECT tactics
<i>fix-body</i>	standard COQ <code>fix_body</code>
<i>ident</i>	standard COQ identifier
<i>int</i>	integer literal
<i>key</i>	notation scope
<i>name</i>	module name
<i>natural</i>	<code>int</code> or Ltac variable denoting a standard COQ numeral ^a
<i>pattern</i>	synonym for <i>term</i>
<i>string</i>	standard COQ string
<i>tactic</i>	standard COQ tactic or SSREFLECT tactic
<i>term</i>	Gallina term, possibly containing wildcards

^aThe name of this Ltac variable should not be the name of a tactic which can be followed by a bracket `[`, like `do`, `have`,...

Items and switches

<i>binder</i>	<code>ident (ident [: term])</code>	binder	p. 305
<i>clear-switch</i>	<code>{ ident⁺ }</code>	clear switch	p. 314
<i>c-pattern</i>	<code>[term in term as] ident in term</code>	context pattern	p. 345

<i>d-item</i>	$[occ\text{-}switch \mid clear\text{-}switch] [term \mid (c\text{-}pattern)]$	discharge item	p. 314
<i>gen-item</i>	$[@]ident \mid (ident) \mid ([@]ident := c\text{-}pattern)$	generalization item	p. 327
<i>i-pattern</i>	$ident \mid _ \mid ? \mid * \mid [occ\text{-}switch] \text{-} > \mid [occ\text{-}switch] < \text{-} \mid$ $[i\text{-}item^* \mid \dots \mid i\text{-}item^*] \mid \text{-} \mid [: ident^+]$	intro pattern	p. 317
<i>i-item</i>	$clear\text{-}switch \mid s\text{-}item \mid i\text{-}pattern \mid / term$	intro item	p. 317
<i>int-mult</i>	$[natural] mult\text{-}mark$	multiplier	p. 325
<i>occ-switch</i>	$\{ [+ \mid -] natural^* \}$	occur. switch	p. 308
<i>mult</i>	$[natural] mult\text{-}mark$	multiplier	p. 325
<i>mult-mark</i>	$? \mid !$	multiplier mark	p. 325
<i>r-item</i>	$[/] term \mid s\text{-}item$	rewrite item	p. 333
<i>r-prefix</i>	$[-] [int\text{-}mult] [occ\text{-}switch \mid clear\text{-}switch] [[r\text{-}pattern]]$	rewrite prefix	p. 333
<i>r-pattern</i>	$term \mid c\text{-}pattern \mid in [ident in] term$	rewrite pattern	p. 333
<i>r-step</i>	$[r\text{-}prefix] r\text{-}item$	rewrite step	p. 333
<i>s-item</i>	$/ = \mid // \mid // =$	simplify switch	p. 317

Tactics

Note: `without` `loss` and `suffices` are synonyms for `wlog` and `suff` respectively.

<code>move</code>	<code>idtac</code> or <code>hnf</code>	p. 310
<code>apply</code>	application	p. 313
<code>exact</code>		
<code>abstract</code>		p. 317, 329
<code>elim</code>	induction	p. 313
<code>case</code>	case analysis	p. 313
<code>rewrite rstep⁺</code>	rewrite	p. 333
<code>have i-item* [i-pattern] [s-item binder⁺] [: term] := term</code>	forward	p. 327
<code>have i-item* [i-pattern] [s-item binder⁺] : term [by tactic]</code>	chaining	
<code>have suff [clear-switch] [i-pattern] [: term] := term</code>		
<code>have suff [clear-switch] [i-pattern] : term [by tactic]</code>		
<code>gen have [ident,] [i-pattern] : gen-item⁺ / term [by tactic]</code>		

<code>wlog [suff] [i-item] : [gen-item clear-switch]* / term</code>	specializing	p. 327
<code>suff i-item* [i-pattern] [binder⁺] : term [by tactic]</code>	backchaining	p. 327
<code>suff [have] [clear-switch] [i-pattern] : term [by tactic]</code>		
<code>pose ident := term</code>	local definition	p. 305
<code>pose ident binder⁺ := term</code>	local function definition	
<code>pose fix fix-body</code>	local fix definition	
<code>pose cofix fix-body</code>	local cofix definition	
<code>set ident [: term] := [occ-switch] [term] (c-pattern)</code>	abbreviation	p. 306
<code>unlock [r-prefix] ident*</code>	unlock	p. 342
<code>congr [natural] term</code>	congruence	p. 344

Tacticals

<code>d-tactic [ident] : d-item⁺ [clear-switch]</code>	discharge	p. 314
<code>tactic => i-item⁺</code>	introduction	p. 317
<code>tactic in [gen-item clear-switch]⁺ [*]</code>	localization	p. 326
<code>do [mult] [tactic ... tactic]</code>	iteration	p. 325
<code>do mult tactic</code>		
<code>tactic ; first [natural] [tactic ... tactic]</code>	selector	p. 324
<code>tactic ; last [natural] [tactic ... tactic]</code>		
<code>tactic ; first [natural] last</code>	subgoals	p. 324
<code>tactic ; last [natural] first</code>	rotation	
<code>by [tactic ... tactic]</code>	closing	p. 323
<code>by []</code>		
<code>by tactic</code>		

Commands

<code>Hint View for [move apply] / ident [natural]</code>	view hint declaration	p. 358
<code>Hint View for apply // ident [natural]</code>	right hand side double view hint declaration	p. 358
<code>Prenex Implicits ident⁺</code>	prenex implicits decl.	p. 304

Part III

User extensions

Chapter 12

Syntax extensions and interpretation scopes

In this chapter, we introduce advanced commands to modify the way COQ parses and prints objects, i.e. the translations between the concrete and internal representations of terms and commands. The main commands are `Notation` and `Infix` which are described in section 12.1. It also happens that the same symbolic notation is expected in different contexts. To achieve this form of overloading, COQ offers a notion of interpretation scope. This is described in Section 12.2.

Remark: The commands `Grammar`, `Syntax` and `Distfix` which were present for a while in COQ are no longer available from COQ version 8.0. The underlying AST structure is also no longer available. The functionalities of the command `Syntactic Definition` are still available; see Section 12.3.

12.1 Notations

12.1.1 Basic notations

A *notation* is a symbolic abbreviation denoting some term or term pattern.

A typical notation is the use of the infix symbol `/\` to denote the logical conjunction (and). Such a notation is declared by

```
Coq < Notation "A /\ B" := (and A B).
```

The expression `(and A B)` is the abbreviated term and the string `"A /\ B"` (called a *notation*) tells how it is symbolically written.

A notation is always surrounded by double quotes (except when the abbreviation is a single identifier; see 12.3). The notation is composed of *tokens* separated by spaces. Identifiers in the string (such as `A` and `B`) are the *parameters* of the notation. They must occur at least once each in the denoted term. The other elements of the string (such as `/\`) are the *symbols*.

An identifier can be used as a symbol but it must be surrounded by simple quotes to avoid the confusion with a parameter. Similarly, every symbol of at least 3 characters and starting with a simple quote must be quoted (then it starts by two single quotes). Here is an example.

```
Coq < Notation "'IF' c1 'then' c2 'else' c3" := (IF_then_else c1 c2 c3).
```

A notation binds a syntactic expression to a term. Unless the parser and pretty-printer of COQ already know how to deal with the syntactic expression (see 12.1.7), explicit precedences and associativity rules have to be given.

Remark: The right-hand side of a notation is interpreted at the time the notation is given. In particular, implicit arguments (see Section 2.7), coercions (see Section 2.8), etc. are resolved at the time of the declaration of the notation.

12.1.2 Precedences and associativity

Mixing different symbolic notations in the same text may cause serious parsing ambiguity. To deal with the ambiguity of notations, COQ uses precedence levels ranging from 0 to 100 (plus one extra level numbered 200) and associativity rules.

Consider for example the new notation

```
Coq < Notation "A \/ B" := (or A B).
```

Clearly, an expression such as `forall A:Prop, True /\ A \/ A \/ False` is ambiguous. To tell the COQ parser how to interpret the expression, a priority between the symbols `/\` and `/\` has to be given. Assume for instance that we want conjunction to bind more than disjunction. This is expressed by assigning a precedence level to each notation, knowing that a lower level binds more than a higher level. Hence the level for disjunction must be higher than the level for conjunction.

Since connectives are not tight articulation points of a text, it is reasonable to choose levels not so far from the highest level which is 100, for example 85 for disjunction and 80 for conjunction¹.

Similarly, an associativity is needed to decide whether `True /\ False /\ False` defaults to `True /\ (False /\ False)` (right associativity) or to `(True /\ False) /\ False` (left associativity). We may even consider that the expression is not well-formed and that parentheses are mandatory (this is a “no associativity”)². We don’t know of a special convention of the associativity of disjunction and conjunction, so let’s apply for instance a right associativity (which is the choice of COQ).

Precedence levels and associativity rules of notations have to be given between parentheses in a list of modifiers that the `Notation` command understands. Here is how the previous examples refine.

```
Coq < Notation "A /\ B" := (and A B) (at level 80, right associativity).
Coq < Notation "A \/ B" := (or A B) (at level 85, right associativity).
```

By default, a notation is considered non associative, but the precedence level is mandatory (except for special cases whose level is canonical). The level is either a number or the phrase `next level` whose meaning is obvious. The list of levels already assigned is on Figure 3.1.

12.1.3 Complex notations

Notations can be made from arbitrarily complex symbols. One can for instance define prefix notations.

```
Coq < Notation "~ x" := (not x) (at level 75, right associativity).
```

¹ which are the levels effectively chosen in the current implementation of COQ

² COQ accepts notations declared as no associative but the parser on which COQ is built, namely CAMLP5, currently does not implement the no-associativity and replaces it by a left associativity; hence it is the same for COQ: no-associativity is in fact left associativity

One can also define notations for incomplete terms, with the hole expected to be inferred at typing time.

```
Coq < Notation "x = y" := (@eq _ x y) (at level 70, no associativity).
```

One can define *closed* notations whose both sides are symbols. In this case, the default precedence level for inner subexpression is 200.

```
Coq < Notation "( x , y )" := (@pair _ _ x y) (at level 0).
```

One can also define notations for binders.

```
Coq < Notation "{ x : A | P }" := (sig A (fun x => P)) (at level 0).
```

In the last case though, there is a conflict with the notation for type casts. This last notation, as shown by the command `Print Grammar constr` is at level 100. To avoid `x : A` being parsed as a type cast, it is necessary to put `x` at a level below 100, typically 99. Hence, a correct definition is

```
Coq < Notation "{ x : A | P }" := (sig A (fun x => P)) (at level 0, x at level 99).
```

See the next section for more about factorization.

12.1.4 Simple factorization rules

COQ extensible parsing is performed by `Camlp5` which is essentially a LL1 parser. Hence, some care has to be taken not to hide already existing rules by new rules. Some simple left factorization work has to be done. Here is an example.

```
Coq < Notation "x < y"      := (lt x y) (at level 70).
Coq < Notation "x < y < z" := (x < y /\ y < z) (at level 70).
```

In order to factorize the left part of the rules, the subexpression referred by `y` has to be at the same level in both rules. However the default behavior puts `y` at the next level below 70 in the first rule (no associativity is the default), and at the level 200 in the second rule (level 200 is the default for inner expressions). To fix this, we need to force the parsing level of `y`, as follows.

```
Coq < Notation "x < y"      := (lt x y) (at level 70).
Coq < Notation "x < y < z" := (x < y /\ y < z) (at level 70, y at next level).
```

For the sake of factorization with COQ predefined rules, simple rules have to be observed for notations starting with a symbol: e.g. rules starting with “{” or “(” should be put at level 0. The list of COQ predefined notations can be found in Chapter 3.

The command to display the current state of the COQ term parser is

```
Print Grammar constr.
```

Variant:

```
Print Grammar pattern.
```

This displays the state of the subparser of patterns (the parser used in the grammar of the `match with constructions`).

12.1.5 Displaying symbolic notations

The command `Notation` has an effect both on the COQ parser and on the COQ printer. For example:

```
Coq < Check (and True True).
True /\ True
      : Prop
```

However, printing, especially pretty-printing, requires more care than parsing. We may want specific indentations, line breaks, alignment if on several lines, etc.

The default printing of notations is very rudimentary. For printing a notation, a *formatting box* is opened in such a way that if the notation and its arguments cannot fit on a single line, a line break is inserted before the symbols of the notation and the arguments on the next lines are aligned with the argument on the first line.

A first, simple control that a user can have on the printing of a notation is the insertion of spaces at some places of the notation. This is performed by adding extra spaces between the symbols and parameters: each extra space (other than the single space needed to separate the components) is interpreted as a space to be inserted by the printer. Here is an example showing how to add spaces around the bar of the notation.

```
Coq < Notation "{ x : A | P }" := (sig (fun x : A => P))
      (at level 0, x at level 99).

Coq < Check (sig (fun x : nat => x=x)).
{{x : nat | x = x}}
      : Set
```

The second, more powerful control on printing is by using the `format` modifier. Here is an example

```
Coq < Notation "'If' c1 'then' c2 'else' c3" := (IF_then_else c1 c2 c3)
      (at level 200, right associativity, format
      "'[v ' 'If' c1 '/' '[' 'then' c2 ']' '/' '[' 'else' c3 ']' ']'").
Identifier 'If' now a keyword
```

A *format* is an extension of the string denoting the notation with the possible following elements delimited by single quotes:

- extra spaces are translated into simple spaces
- tokens of the form `' / '` are translated into breaking point, in case a line break occurs, an indentation of the number of spaces after the `“/”` is applied (2 spaces in the given example)
- token of the form `' / / '` force writing on a new line
- well-bracketed pairs of tokens of the form `' [' and '] '` are translated into printing boxes; in case a line break occurs, an extra indentation of the number of spaces given after the `“[”` is applied (4 spaces in the example)
- well-bracketed pairs of tokens of the form `' [hv ' and '] '` are translated into horizontal-or-vertical printing boxes; if the content of the box does not fit on a single line, then every breaking point forces a newline and an extra indentation of the number of spaces given after the `“[”` is applied at the beginning of each newline (3 spaces in the example)

- well-bracketed pairs of tokens of the form `' [v ' and '] '` are translated into vertical printing boxes; every breaking point forces a newline, even if the line is large enough to display the whole content of the box, and an extra indentation of the number of spaces given after the `"["` is applied at the beginning of each newline

Notations do not survive the end of sections. No typing of the denoted expression is performed at definition time. Type-checking is done only at the time of use of the notation.

```
Coq < Check
      (IF_then_else (IF_then_else True False True)
        (IF_then_else True False True)
        (IF_then_else True False True)).
If If True
  then False
  else True
then If True
  then False
  else True
else If True
  then False
  else True
: Prop
```

Remark: Sometimes, a notation is expected only for the parser. To do so, the option *only parsing* is allowed in the list of modifiers of `Notation`.

Conversely, the *only printing* can be used to declare that a notation should only be used for printing and should not declare a parsing rule. In particular, such notations do not modify the parser.

12.1.6 The `Infix` command

The `Infix` command is a shortening for declaring notations of infix symbols. Its syntax is

```
Infix "symbol" := qualid ( modifier , ... , modifier ).
```

and it is equivalent to

```
Notation "x symbol y" := ( qualid x y ) ( modifier , ... , modifier ).
```

where `x` and `y` are fresh names distinct from *qualid*. Here is an example.

```
Coq < Infix "/" := and (at level 80, right associativity).
```

12.1.7 Reserving notations

A given notation may be used in different contexts. COQ expects all uses of the notation to be defined at the same precedence and with the same associativity. To avoid giving the precedence and associativity every time, it is possible to declare a parsing rule in advance without giving its interpretation. Here is an example from the initial state of COQ.

```
Coq < Reserved Notation "x = y" (at level 70, no associativity).
```

Reserving a notation is also useful for simultaneously defining an inductive type or a recursive constant and a notation for it.

Remark: The notations mentioned on Figure 3.1 are reserved. Hence their precedence and associativity cannot be changed.

12.1.8 Simultaneous definition of terms and notations

Thanks to reserved notations, the inductive, co-inductive, recursive and corecursive definitions can benefit of customized notations. To do this, insert a `where` notation clause after the definition of the (co)inductive type or (co)recursive term (or after the definition of each of them in case of mutual definitions). The exact syntax is given on Figure 12.1. Here are examples:

```
Coq < Inductive and (A B:Prop) : Prop := conj : A -> B -> A /\ B
      where "A /\ B" := (and A B).
```

```
Coq < Fixpoint plus (n m:nat) {struct n} : nat :=
      match n with
      | 0 => m
      | S p => S (p+m)
      end
      where "n + m" := (plus n m).
```

12.1.9 Displaying informations about notations

To deactivate the printing of all notations, use the command

```
Unset Printing Notations.
```

To reactivate it, use the command

```
Set Printing Notations.
```

The default is to use notations for printing terms wherever possible.

See also: `Set Printing All` in Section 2.9.

12.1.10 Locating notations

To know to which notations a given symbol belongs to, use the command

```
Locate symbol
```

where *symbol* is any (composite) symbol surrounded by double quotes. To locate a particular notation, use a string where the variables of the notation are replaced by “_” and where possible single quotes inserted around identifiers or tokens starting with a single quote are dropped.

Example:

```
Coq < Locate "exists".
Notation
"'exists' x .. y , p" := ex (fun y => .. (ex (fun y => p)) ..)
: type_scope (default interpretation)
"'exists' ! x .. y , p" := ex
```

<i>sentence</i>	<code>::=</code>	<code>[Local] Notation <i>string</i> := <i>term</i> [<i>modifiers</i>] [:<i>scope</i>].</code> <code>[Local] Infix <i>string</i> := <i>qualid</i> [<i>modifiers</i>] [:<i>scope</i>].</code> <code>[Local] Reserved Notation <i>string</i> [<i>modifiers</i>].</code> <code>Inductive <i>ind_body</i> [<i>decl_notation</i>] with ... with <i>ind_body</i> [<i>decl_notation</i>].</code> <code>CoInductive <i>ind_body</i> [<i>decl_notation</i>] with ... with <i>ind_body</i> [<i>decl_notation</i>].</code> <code>Fixpoint <i>fix_body</i> [<i>decl_notation</i>] with ... with <i>fix_body</i> [<i>decl_notation</i>].</code> <code>CoFixpoint <i>cofix_body</i> [<i>decl_notation</i>] with ... with <i>cofix_body</i> [<i>decl_notation</i>].</code>
<i>decl_notation</i>	<code>::=</code>	<code>[where <i>string</i> := <i>term</i> [:<i>scope</i>] and ... and <i>string</i> := <i>term</i> [:<i>scope</i>]].</code>
<i>modifiers</i>	<code>::=</code>	<code><i>ident</i> , ... , <i>ident</i> at level <i>natural</i></code> <code><i>ident</i> , ... , <i>ident</i> at next level</code> <code>at level <i>natural</i></code> <code>left associativity</code> <code>right associativity</code> <code>no associativity</code> <code><i>ident</i> <i>ident</i></code> <code><i>ident</i> binder</code> <code><i>ident</i> closed binder</code> <code><i>ident</i> global</code> <code><i>ident</i> bigint</code> <code>only parsing</code> <code>only printing</code> <code>format <i>string</i></code>

Figure 12.1: Syntax of the variants of Notation

```

                (unique
                  (fun y =>
                    .. (ex (unique (fun y => p))) ..))
: type_scope (default interpretation)
Coq < Locate "exists _ .. _ , _".
Notation
"'exists' x .. y , p" := ex (fun y => .. (ex (fun y => p)) ..)
: type_scope (default interpretation)

```

See also: Section 6.3.10.

12.1.11 Notations and simple binders

Notations can be defined for binders as in the example:

```
Coq < Notation "{ x : A | P }" := (sig (fun x : A => P)) (at level 0).
```

The binding variables in the left-hand-side that occur as a parameter of the notation naturally bind all their occurrences appearing in their respective scope after instantiation of the parameters of the notation.

Contrastingly, the binding variables that are not a parameter of the notation do not capture the variables of same name that could appear in their scope after instantiation of the notation. E.g., for the notation

```
Coq < Notation "'exists_different' n" := (exists p:nat, p<>n) (at level 200).
```

the next command fails because p does not bind in the instance of n .

```
Coq < Fail Check (exists_different p).
The command has indeed failed with message:
The reference p was not found in the current environment.
```

Remark: Binding variables must not necessarily be parsed using the `ident` entry. For factorization purposes, they can be said to be parsed at another level (e.g. x in $\{ x : A \mid P \}$ must be parsed at level 99 to be factorized with the notation $\{ A \} + \{ B \}$ for which A can be any term). However, even if parsed as a term, this term must at the end be effectively a single identifier.

12.1.12 Notations with recursive patterns

A mechanism is provided for declaring elementary notations with recursive patterns. The basic example is:

```
Coq < Notation "[ x ; .. ; y ]" := (cons x .. (cons y nil) ..).
```

On the right-hand side, an extra construction of the form $.. t ..$ can be used. Notice that $..$ is part of the COQ syntax and it must not be confused with the three-dots notation $...$ used in this manual to denote a sequence of arbitrary size.

On the left-hand side, the part “ $x s .. s y$ ” of the notation parses any number of time (but at least one time) a sequence of expressions separated by the sequence of tokens s (in the example, s is just “;”).

In the right-hand side, the term enclosed within $..$ must be a pattern with two holes of the form $\phi([]_E, []_I)$ where the first hole is occupied either by x or by y and the second hole is occupied by an arbitrary term t called the *terminating* expression of the recursive notation. The subterm $.. \phi(x, t) ..$ (or $.. \phi(y, t) ..$) must itself occur at second position of the same pattern where the first hole is occupied by the other variable, y or x . Otherwise said, the right-hand side must contain a subterm of the form either $\phi(x, .. \phi(y, t) ..)$ or $\phi(y, .. \phi(x, t) ..)$. The pattern ϕ is the *iterator* of the recursive notation and, of course, the name x and y can be chosen arbitrarily.

The parsing phase produces a list of expressions which are used to fill in order the first hole of the iterating pattern which is repeatedly nested as many times as the length of the list, the second hole being the nesting point. In the innermost occurrence of the nested iterating pattern, the second hole is finally filled with the terminating expression.

In the example above, the iterator $\phi([]_E, []_I)$ is `cons []E []I` and the terminating expression is `nil`. Here are other examples:

```
Coq < Notation "( x , y , .. , z )" := (pair .. (pair x y) .. z) (at level 0).
Coq < Notation "[| t * ( x , y , .. , z ) ; ( a , b , .. , c ) * u |]" :=
  (pair (pair .. (pair (pair t x) (pair t y)) .. (pair t z))
    (pair .. (pair (pair a u) (pair b u)) .. (pair c u)))
  (t at level 39).
```

Recursive patterns can occur several times on the right-hand side. Here is an example:

```
Coq < Notation "> a , .. , b <" :=
  (cons a .. (cons b nil) .., cons b .. (cons a nil) ..).
```

Notations with recursive patterns can be reserved like standard notations, they can also be declared within interpretation scopes (see section 12.2).

12.1.13 Notations with recursive patterns involving binders

Recursive notations can also be used with binders. The basic example is:

```
Coq < Notation "'exists' x .. y , p" := (ex (fun x => .. (ex (fun y => p)) ..))
      (at level 200, x binder, y binder, right associativity).
```

The principle is the same as in Section 12.1.12 except that in the iterator $\phi([]_E, []_I)$, the first hole is a placeholder occurring at the position of the binding variable of a `fun` or a `forall`.

To specify that the part “ $x \dots y$ ” of the notation parses a sequence of binders, x and y must be marked as `binder` in the list of modifiers of the notation. Then, the list of binders produced at the parsing phase are used to fill in the first hole of the iterating pattern which is repeatedly nested as many times as the number of binders generated. If ever the generalization operator ``` (see Section 2.7.19) is used in the binding list, the added binders are taken into account too.

Binders parsing exist in two flavors. If x and y are marked as `binder`, then a sequence such as `a b c : T` will be accepted and interpreted as the sequence of binders $(a:T) (b:T) (c:T)$. For instance, in the notation above, the syntax `exists a b : nat, a = b` is provided.

The variables x and y can also be marked as `closed binder` in which case only well-bracketed binders of the form $(a b c:T)$ or $\{a b c:T\}$ etc. are accepted.

With closed binders, the recursive sequence in the left-hand side can be of the general form $x s \dots s y$ where s is an arbitrary sequence of tokens. With open binders though, s has to be empty. Here is an example of recursive notation with closed binders:

```
Coq < Notation "'mylet' f x .. y := t 'in' u" :=
      (let f := fun x => .. (fun y => t) .. in u)
      (at level 200, x closed binder, y closed binder, right associativity).
```

A recursive pattern for binders can be used in position of a recursive pattern for terms. Here is an example:

```
Coq < Notation "'FUNAPP' x .. y , f" :=
      (fun x => .. (fun y => (.. (f x) ..) y) ..)
      (at level 200, x binder, y binder, right associativity).
```

12.1.14 Summary

Syntax of notations The different syntactic variants of the command `Notation` are given on Figure 12.1. The optional `:scope` is described in the Section 12.2.

Remark: No typing of the denoted expression is performed at definition time. Type-checking is done only at the time of use of the notation.

Remark: Many examples of `Notation` may be found in the files composing the initial state of COQ (see directory `$COQLIB/theories/Init`).

Remark: The notation “ $\{ x \}$ ” has a special status in such a way that complex notations of the form “ $x + \{ y \}$ ” or “ $x * \{ y \}$ ” can be nested with correct precedences. Especially, every notation involving a pattern of the form “ $\{ x \}$ ” is parsed as a notation where the pattern “ $\{ x \}$ ” has been simply replaced by “ x ” and the curly brackets are parsed separately. E.g. “ $y + \{ z \}$ ” is not parsed as a term of the given form but as a term of the form “ $y + z$ ” where z has been parsed using the rule

parsing " $\{ x \}$ ". Especially, level and precedences for a rule including patterns of the form " $\{ x \}$ " are relative not to the textual notation but to the notation where the curly brackets have been removed (e.g. the level and the associativity given to some notation, say " $\{ y \} \& \{ z \}$ " in fact applies to the underlying " $\{ x \}$ "-free rule which is " $y \& z$ ").

Persistence of notations Notations do not survive the end of sections. They survive modules unless the command `Local Notation` is used instead of `Notation`.

12.2 Interpretation scopes

An *interpretation scope* is a set of notations for terms with their interpretation. Interpretation scopes provide a weak, purely syntactical form of notation overloading: the same notation, for instance the infix symbol $+$, can be used to denote distinct definitions of the additive operator. Depending on which interpretation scope is currently open, the interpretation is different. Interpretation scopes can include an interpretation for numerals and strings. However, this is only made possible at the OCAML level.

See Figure 12.1 for the syntax of notations including the possibility to declare them in a given scope. Here is a typical example which declares the notation for conjunction in the scope `type_scope`.

```
Notation "A /\ B" := (and A B) : type_scope.
```

Remark: A notation not defined in a scope is called a *lonely* notation.

12.2.1 Global interpretation rules for notations

At any time, the interpretation of a notation for term is done within a *stack* of interpretation scopes and lonely notations. In case a notation has several interpretations, the actual interpretation is the one defined by (or in) the more recently declared (or open) lonely notation (or interpretation scope) which defines this notation. Typically if a given notation is defined in some scope *scope* but has also an interpretation not assigned to a scope, then, if *scope* is open before the lonely interpretation is declared, then the lonely interpretation is used (and this is the case even if the interpretation of the notation in *scope* is given after the lonely interpretation: otherwise said, only the order of lonely interpretations and opening of scopes matters, and not the declaration of interpretations within a scope).

The initial state of COQ declares three interpretation scopes and no lonely notations. These scopes, in opening order, are `core_scope`, `type_scope` and `nat_scope`.

The command to add a scope to the interpretation scope stack is

```
Open Scope scope.
```

It is also possible to remove a scope from the interpretation scope stack by using the command

```
Close Scope scope.
```

Notice that this command does not only cancel the last `Open Scope scope` but all the invocation of it.

Remark: `Open Scope` and `Close Scope` do not survive the end of sections where they occur. When defined outside of a section, they are exported to the modules that import the module where they occur.

Variants:

1. Local Open Scope *scope*.

2. Local Close Scope *scope*.

These variants are not exported to the modules that import the module where they occur, even if outside a section.

3. Global Open Scope *scope*.

4. Global Close Scope *scope*.

These variants survive sections. They behave as if `Global` were absent when not inside a section.

12.2.2 Local interpretation rules for notations

In addition to the global rules of interpretation of notations, some ways to change the interpretation of subterms are available.

Local opening of an interpretation scope

It is possible to locally extend the interpretation scope stack using the syntax $(term)\%key$ (or simply $term\%key$ for atomic terms), where *key* is a special identifier called *delimiting key* and bound to a given scope.

In such a situation, the term *term*, and all its subterms, are interpreted in the scope stack extended with the scope bound to *key*.

To bind a delimiting key to a scope, use the command

```
Delimit Scope scope with ident
```

To remove a delimiting key of a scope, use the command

```
Undelimit Scope scope
```

Binding arguments of a constant to an interpretation scope

It is possible to set in advance that some arguments of a given constant have to be interpreted in a given scope. The command is

```
Arguments qualid name%scope ... name%scope
```

where the list is a prefix of the list of the arguments of *qualid* eventually annotated with their *scope*. Grouping round parentheses can be used to decorate multiple arguments with the same scope. *scope* can be either a scope name or its delimiting key. For example the following command puts the first two arguments of `plus_fct` in the scope delimited by the key `F` (`Rfun_scope`) and the last argument in the scope delimited by the key `R` (`R_scope`).

```
Coq < Arguments plus_fct (f1 f2)%F x%R.
```

The `Arguments` command accepts scopes decoration to all grouping parentheses. In the following example arguments `A` and `B` are marked as maximally inserted implicit arguments and are put into the `type_scope` scope.

```
Coq < Arguments respectful {A B}%type (R R')%signature _ _.
```

When interpreting a term, if some of the arguments of *qualid* are built from a notation, then this notation is interpreted in the scope stack extended by the scope bound (if any) to this argument. The effect of the scope is limited to the argument itself. It does not propagate to subterms but the subterms that, after interpretation of the notation, turn to be themselves arguments of a reference are interpreted accordingly to the arguments scopes bound to this reference.

Arguments scopes can be cleared with the following command:

```
Arguments qualid : clear scopes
```

Variants:

1. Global Arguments *qualid name%scope ... name%scope*

This behaves like Arguments *qualid name%scope ... name%scope* but survives when a section is closed instead of stopping working at section closing. Without the Global modifier, the effect of the command stops when the section it belongs to ends.

2. Local Arguments *qualid name%scope ... name%scope*

This behaves like Arguments *qualid name%scope ... name%scope* but does not survive modules and files. Without the Local modifier, the effect of the command is visible from within other modules or files.

See also: The command to show the scopes bound to the arguments of a function is described in Section 2.

Binding types of arguments to an interpretation scope

When an interpretation scope is naturally associated to a type (e.g. the scope of operations on the natural numbers), it may be convenient to bind it to this type. When a scope *scope* is bound to a type *type*, any new function defined later on gets its arguments of type *type* interpreted by default in scope *scope* (this default behavior can however be overwritten by explicitly using the command Arguments).

Whether the argument of a function has some type *type* is determined statically. For instance, if *f* is a polymorphic function of type $\text{forall } X:\text{Type}, X \rightarrow X$ and type *t* is bound to a scope *scope*, then *a* of type *t* in *f t a* is not recognized as an argument to be interpreted in scope *scope*.

More generally, any coercion *class* (see Chapter 18) can be bound to an interpretation scope. The command to do it is

```
Bind Scope scope with class
```

Example:

```
Coq < Parameter U : Set.
U is declared

Coq < Bind Scope U_scope with U.

Coq < Parameter Uplus : U -> U -> U.
Uplus is declared

Coq < Parameter P : forall T:Set, T -> U -> Prop.
```



```

P is declared

Coq < Parameter f : forall T:Set, T -> U.
f is declared

Coq < Infix "+" := Uplus : U_scope.

Coq < Unset Printing Notations.

Coq < Open Scope nat_scope. (* Define + on the nat as the default for + *)

Coq < Check (fun x y1 y2 z t => P _ (x + t) ((f _ (y1 + y2) + z))).
fun (x y1 y2 : nat) (z : U) (t : nat) =>
P nat (Nat.add x t) (Uplus (f nat (Nat.add y1 y2)) z)
: forall (_ : nat) (_ : nat) (_ : nat) (_ : U) (_ : nat), Prop

```

Remark: The scopes `type_scope` and `function_scope` also have a local effect on interpretation. See the next section.

See also: The command to show the scopes bound to the arguments of a function is described in Section 2.

Remark: In notations, the subterms matching the identifiers of the notations are interpreted in the scope in which the identifiers occurred at the time of the declaration of the notation. Here is an example:

```

Coq < Parameter g : bool -> bool.
g is declared

Coq < Notation "@@" := true (only parsing) : bool_scope.
Setting notation at level 0.

Coq < Notation "@@" := false (only parsing): mybool_scope.

Coq < (* Defining a notation while the argument of g is bound to bool_scope *)
      Bind Scope bool_scope with bool.

Coq < Notation "# x #" := (g x) (at level 40).

Coq < Check # @@ #.
g true
: bool

Coq < (* Rebinding the argument of g to mybool_scope has no effect on the notation *)
      Arguments g _%mybool_scope.

Coq < Check # @@ #.
g true
: bool

Coq < (* But we can force the scope *)
      Delimit Scope mybool_scope with mybool.

Coq < Check # @@%mybool #.
g false
: bool

```

12.2.3 The `type_scope` interpretation scope

The scope `type_scope` has a special status. It is a primitive interpretation scope which is temporarily activated each time a subterm of an expression is expected to be a type. It is delimited by the key `type`,

and bound to the coercion class `Sortclass`. It is also used in certain situations where an expression is statically known to be a type, including the conclusion and the type of hypotheses within an `Ltac` goal match (see Section 9.2) the statement of a theorem, the type of a definition, the type of a binder, the domain and codomain of implication, the codomain of products, and more generally any type argument of a declared or defined constant.

12.2.4 The `function_scope` interpretation scope

The scope `function_scope` also has a special status. It is temporarily activated each time the argument of a global reference is recognized to be a `Funclass` instance, i.e., of type `forall x:A, B or A -> B`.

12.2.5 Interpretation scopes used in the standard library of COQ

We give an overview of the scopes used in the standard library of COQ. For a complete list of notations in each scope, use the commands `Print Scopes` or `Print Scope scope`.

`type_scope`

This scope includes infix `*` for product types and infix `+` for sum types. It is delimited by key `type`, and bound to the coercion class `Sortclass`, as described at 12.2.2.

`nat_scope`

This scope includes the standard arithmetical operators and relations on type `nat`. Positive numerals in this scope are mapped to their canonical representant built from `O` and `S`. The scope is delimited by key `nat`, and bound to the type `nat` (see 12.2.2).

`N_scope`

This scope includes the standard arithmetical operators and relations on type `N` (binary natural numbers). It is delimited by key `N` and comes with an interpretation for numerals as closed term of type `N`.

`Z_scope`

This scope includes the standard arithmetical operators and relations on type `Z` (binary integer numbers). It is delimited by key `Z` and comes with an interpretation for numerals as closed term of type `Z`.

`positive_scope`

This scope includes the standard arithmetical operators and relations on type `positive` (binary strictly positive numbers). It is delimited by key `positive` and comes with an interpretation for numerals as closed term of type `positive`.

`Q_scope`

This scope includes the standard arithmetical operators and relations on type `Q` (rational numbers defined as fractions of an integer and a strictly positive integer modulo the equality of the numerator-denominator cross-product). As for numerals, only 0 and 1 have an interpretation in scope `Q_scope` (their interpretations are $\frac{0}{1}$ and $\frac{1}{1}$ respectively).

`Qc_scope`

This scope includes the standard arithmetical operators and relations on the type `Qc` of rational numbers defined as the type of irreducible fractions of an integer and a strictly positive integer.

`real_scope`

This scope includes the standard arithmetical operators and relations on type `R` (axiomatic real numbers). It is delimited by key `R` and comes with an interpretation for numerals using the `IZR` morphism from binary integer numbers to `R`.

`bool_scope`

This scope includes notations for the boolean operators. It is delimited by key `bool`, and bound to the type `bool` (see 12.2.2).

`list_scope`

This scope includes notations for the list operators. It is delimited by key `list`, and bound to the type `list` (see 12.2.2).

`function_scope`

This scope is delimited by the key `function`, and bound to the coercion class `FuncClass`, as described at 12.2.2.

`core_scope`

This scope includes the notation for pairs. It is delimited by key `core`.

`string_scope`

This scope includes notation for strings as elements of the type `string`. Special characters and escaping follow COQ conventions on strings (see Section 1.1). Especially, there is no convention to visualize non printable characters of a string. The file `String.v` shows an example that contains quotes, a newline and a beep (i.e. the ASCII character of code 7).

`char_scope`

This scope includes interpretation for all strings of the form `"c"` where `c` is an ASCII character, or of the form `"nnn"` where `nnn` is a three-digits number (possibly with leading 0's), or of the form `"\""`. Their respective denotations are the ASCII code of `c`, the decimal ASCII code `nnn`, or the ASCII code of the character `"` (i.e. the ASCII code 34), all of them being represented in the type `ascii`.

12.2.6 Displaying informations about scopes

`Print Visibility`

This displays the current stack of notations in scopes and lonely notations that is used to interpret a notation. The top of the stack is displayed last. Notations in scopes whose interpretation is hidden by

the same notation in a more recently open scope are not displayed. Hence each notation is displayed only once.

Variant:

Print Visibility *scope*

This displays the current stack of notations in scopes and lonely notations assuming that *scope* is pushed on top of the stack. This is useful to know how a subterm locally occurring in the scope of *scope* is interpreted.

Print Scope *scope*

This displays all the notations defined in interpretation scope *scope*. It also displays the delimiting key if any and the class to which the scope is bound, if any.

Print Scopes

This displays all the notations, delimiting keys and corresponding class of all the existing interpretation scopes. It also displays the lonely notations.

12.3 Abbreviations

An *abbreviation* is a name, possibly applied to arguments, that denotes a (presumably) more complex expression. Here are examples:

```
Coq < Notation Nlist := (list nat).

Coq < Check 1 :: 2 :: 3 :: nil.
[1; 2; 3]
      : Nlist

Coq < Notation reflexive R := (forall x, R x x).

Coq < Check forall A:Prop, A <-> A.
reflexive iff
      : Prop

Coq < Check reflexive iff.
reflexive iff
      : Prop
```

An abbreviation expects no precedence nor associativity, since it follows the usual syntax of application. Abbreviations are used as much as possible by the COQ printers unless the modifier (only parsing) is given.

Abbreviations are bound to an absolute name as an ordinary definition is, and they can be referred by qualified names too.

Abbreviations are syntactic in the sense that they are bound to expressions which are not typed at the time of the definition of the abbreviation but at the time it is used. Especially, abbreviations can be bound to terms with holes (i.e. with “_”). The general syntax for abbreviations is

$$[\text{Local}] \text{Notation } \textit{ident} [\textit{ident} \textit{ident} \dots \textit{ident} \textit{ident}] := \textit{term} [(\text{only parsing})].$$

Example:

```

Coq < Definition explicit_id (A:Set) (a:A) := a.
explicit_id is defined

Coq < Notation id := (explicit_id _).

Coq < Check (id 0).
id 0
    : nat

```

Abbreviations do not survive the end of sections. No typing of the denoted expression is performed at definition time. Type-checking is done only at the time of use of the abbreviation.

12.4 Tactic Notations

Tactic notations allow to customize the syntax of the tactics of the tactic language³. Tactic notations obey the following syntax

<i>sentence</i>	::=	[Local] Tactic Notation [<i>tactic_level</i>] [<i>prod_item</i> ... <i>prod_item</i>] := <i>tactic</i> .
<i>prod_item</i>	::=	<i>string</i> <i>tactic_argument_type</i> (<i>ident</i>)
<i>tactic_level</i>	::=	(at level <i>natural</i>)
<i>tactic_argument_type</i>	::=	ident simple_intropattern reference hyp hyp_list ne_hyp_list constr uconstr constr_list ne_constr_list integer integer_list ne_integer_list int_or_var int_or_var_list ne_int_or_var_list tactic tactic _n (for $0 \leq n \leq 5$)

A tactic notation `Tactic Notation tactic_level [prod_item ... prod_item] := tactic` extends the parser and pretty-printer of tactics with a new rule made of the list of production items. It then evaluates into the tactic expression *tactic*. For simple tactics, it is recommended to use a terminal symbol, i.e. a *string*, for the first production item. The tactic level indicates the parsing precedence of the tactic notation. This information is particularly relevant for notations of tacticals. Levels 0 to 5 are available (default is 0). To know the parsing precedences of the existing tacticals, use the command `Print Grammar tactic`.

Each type of tactic argument has a specific semantic regarding how it is parsed and how it is interpreted. The semantic is described in the following table. The last command gives examples of tactics which use the corresponding kind of argument.

³Tactic notations are just a simplification of the `Grammar tactic simple_tactic` command that existed in versions prior to version 8.0.

Tactic argument type	parsed as	interpreted as	as in tactic
<code>ident</code>	identifier	a user-given name	<code>intro</code>
<code>simple_intropattern</code>	<code>intro_pattern</code>	an <code>intro_pattern</code>	<code>intros</code>
<code>hyp</code>	identifier	an hypothesis defined in context	<code>clear</code>
<code>reference</code>	qualified identifier	a global reference of term	<code>unfold</code>
<code>constr</code>	term	a term	<code>exact</code>
<code>uconstr</code>	term	an untyped term	<code>refine</code>
<code>integer</code>	integer	an integer	
<code>int_or_var</code>	identifier or integer	an integer	<code>do</code>
<code>tactic</code>	tactic at level 5	a tactic	
<code>tacticlev n</code>	tactic at level n	a tactic	
<code>entry_list</code>	list of <code>entry</code>	a list of how <code>entry</code> is interpreted	
<code>ne_entry_list</code>	non-empty list of <code>entry</code>	a list of how <code>entry</code> is interpreted	

Remark: In order to be bound in tactic definitions, each syntactic entry for argument type must include the case of simple \mathcal{L}_{tac} identifier as part of what it parses. This is naturally the case for `ident`, `simple_intropattern`, `reference`, `constr`, ... but not for `integer`. This is the reason for introducing a special entry `int_or_var` which evaluates to integers only but which syntactically includes identifiers in order to be usable in tactic definitions.

Remark: The `entry_list` and `ne_entry_list` entries can be used in primitive tactics or in other notations at places where a list of the underlying entry can be used: `entry` is either `constr`, `hyp`, `integer` or `int_or_var`.

Tactic notations do not survive the end of sections. They survive modules unless the command `Local Tactic Notation` is used instead of `Tactic Notation`.

Chapter 13

Proof schemes

13.1 Generation of induction principles with `Scheme`

The `Scheme` command is a high-level tool for generating automatically (possibly mutual) induction principles for given types and sorts. Its syntax follows the schema:

```
Scheme ident1 := Induction for ident'1 Sort sort1
with
...
with identm := Induction for ident'm Sort sortm
```

where *ident'*₁ ... *ident'*_{*m*} are different inductive type identifiers belonging to the same package of mutual inductive definitions. This command generates *ident*₁ ... *ident*_{*m*} to be mutually recursive definitions. Each term *ident*_{*i*} proves a general principle of mutual induction for objects in type *term*_{*i*}.

Variants:

1. Scheme *ident*₁ := Minimality for *ident'*₁ Sort *sort*₁
with
...
with *ident*_{*m*} := Minimality for *ident'*_{*m*} Sort *sort*_{*m*}

Same as before but defines a non-dependent elimination principle more natural in case of inductively defined relations.

2. Scheme Equality for *ident*₁

Tries to generate a Boolean equality and a proof of the decidability of the usual equality. If *ident*_{*i*} involves some other inductive types, their equality has to be defined first.

3. Scheme Induction for *ident*₁ Sort *sort*₁
with
...
with Induction for *ident*_{*m*} Sort *sort*_{*m*}

If you do not provide the name of the schemes, they will be automatically computed from the sorts involved (works also with Minimality).

Example 1: Induction scheme for *tree* and *forest*

The definition of principle of mutual induction for `tree` and `forest` over the sort `Set` is defined by the command:

```

Coq < Inductive tree : Set :=
  node : A -> forest -> tree
  with forest : Set :=
    | leaf : B -> forest
    | cons : tree -> forest -> forest.

Coq < Scheme tree_forest_rec := Induction for tree Sort Set
  with forest_tree_rec := Induction for forest Sort Set.

```

You may now look at the type of `tree_forest_rec`:

```

Coq < Check tree_forest_rec.
tree_forest_rec
  : forall (P : tree -> Set) (P0 : forest -> Set),
    (forall (a : A) (f : forest), P0 f -> P (node a f)) ->
    (forall b : B, P0 (leaf b)) ->
    (forall t : tree, P t -> forall f1 : forest, P0 f1 -> P0 (cons t f1)) ->
    forall t : tree, P t

```

This principle involves two different predicates for trees and forests; it also has three premises each one corresponding to a constructor of one of the inductive definitions.

The principle `forest_tree_rec` shares exactly the same premises, only the conclusion now refers to the property of forests.

```

Coq < Check forest_tree_rec.
forest_tree_rec
  : forall (P : tree -> Set) (P0 : forest -> Set),
    (forall (a : A) (f : forest), P0 f -> P (node a f)) ->
    (forall b : B, P0 (leaf b)) ->
    (forall t : tree, P t -> forall f1 : forest, P0 f1 -> P0 (cons t f1)) ->
    forall f2 : forest, P0 f2

```

Example 2: Predicates odd and even on naturals

Let odd and even be inductively defined as:

```

Coq < Inductive odd : nat -> Prop :=
  oddS : forall n:nat, even n -> odd (S n)
  with even : nat -> Prop :=
    | even0 : even 0
    | evenS : forall n:nat, odd n -> even (S n).

```

The following command generates a powerful elimination principle:

```

Coq < Scheme odd_even := Minimality for odd Sort Prop
  with even_odd := Minimality for even Sort Prop.
even_odd is defined
odd_even is defined
odd_even, even_odd are recursively defined

```

The type of `odd_even` for instance will be:


```

Coq < Check odd_even.
odd_even
  : forall P P0 : nat -> Prop,
    (forall n : nat, even n -> P0 n -> P (S n)) ->
    P0 0 ->
    (forall n : nat, odd n -> P n -> P0 (S n)) ->
    forall n : nat, odd n -> P n

```

The type of `even_odd` shares the same premises but the conclusion is $(n : \text{nat}) (\text{even } n) \rightarrow (Q \ n)$.

13.1.1 Automatic declaration of schemes

It is possible to deactivate the automatic declaration of the induction principles when defining a new inductive type with the `Unset Elimination Schemes` command. It may be reactivated at any time with `Set Elimination Schemes`.

The types declared with the keywords `Variant` (see 1.3.3) and `Record` (see 2.1) do not have an automatic declaration of the induction principles. It can be activated with the command `Set Nonrecursive Elimination Schemes`. It can be deactivated again with `Unset Nonrecursive Elimination Schemes`. `Record Elimination Schemes` is a deprecated alias of `Nonrecursive Elimination Schemes`.

In addition, the `Case Analysis Schemes` flag governs the generation of case analysis lemmas for inductive types, i.e. corresponding to the pattern-matching term alone and without `fixpoint`.

You can also activate the automatic declaration of those Boolean equalities (see the second variant of `Scheme`) with respectively the commands `Set Boolean Equality Schemes` and `Set Decidable Equality Schemes`. However you have to be careful with this option since `COQ` may now reject well-defined inductive types because it cannot compute a Boolean equality for them.

The `Rewriting Schemes` flag governs generation of equality related schemes such as `congruence`.

13.1.2 Combined Scheme

The `Combined Scheme` command is a tool for combining induction principles generated by the `Scheme` command. Its syntax follows the schema :

```
Combined Scheme ident0 from ident1, ..., identn
```

where $\text{ident}_1 \dots \text{ident}_n$ are different inductive principles that must belong to the same package of mutual inductive principle definitions. This command generates ident_0 to be the conjunction of the principles: it is built from the common premises of the principles and concluded by the conjunction of their conclusions.

Example: We can define the induction principles for trees and forests using:

```

Coq < Scheme tree_forest_ind := Induction for tree Sort Prop
      with forest_tree_ind := Induction for forest Sort Prop.
forest_tree_ind is defined
tree_forest_ind is defined
tree_forest_ind, forest_tree_ind are recursively defined

```

Then we can build the combined induction principle which gives the conjunction of the conclusions of each individual principle:

```
Coq < Combined Scheme tree_forest_mutind from tree_forest_ind, forest_tree_ind.
tree_forest_mutind is defined
tree_forest_mutind is recursively defined
```

The type of `tree_forest_mutrec` will be:

```
Coq < Check tree_forest_mutind.
tree_forest_mutind
  : forall (P : tree -> Prop) (P0 : forest -> Prop),
    (forall (a : A) (f : forest), P0 f -> P (node a f)) ->
    (forall b : B, P0 (leaf b)) ->
    (forall t : tree, P t -> forall f1 : forest, P0 f1 -> P0 (cons t f1)) ->
    (forall t : tree, P t) /\ (forall f2 : forest, P0 f2)
```

13.2 Generation of induction principles with Functional Scheme

The `Functional Scheme` command is a high-level experimental tool for generating automatically induction principles corresponding to (possibly mutually recursive) functions. First, it must be made available via `Require Import FunInd`. Its syntax then follows the schema:

```
Functional Scheme ident1 := Induction for ident'1 Sort sort1
with
...
with identm := Induction for ident'm Sort sortm
```

where *ident*'₁ ... *ident*'_{*m*} are different mutually defined function names (they must be in the same order as when they were defined). This command generates the induction principles *ident*₁ ... *ident*_{*m*}, following the recursive structure and case analyses of the functions *ident*'₁ ... *ident*'_{*m*}.

Remark: There is a difference between obtaining an induction scheme by using `Functional Scheme` on a function defined by `Function` or not. Indeed `Function` generally produces smaller principles, closer to the definition written by the user.

Example 1: Induction scheme for `div2`

We define the function `div2` as follows:

```
Coq < Require Import Arith.
Coq < Fixpoint div2 (n:nat) : nat :=
  match n with
  | 0 => 0
  | S 0 => 0
  | S (S n') => S (div2 n')
  end.
```

The definition of a principle of induction corresponding to the recursive structure of `div2` is defined by the command:

```
Coq < Functional Scheme div2_ind := Induction for div2 Sort Prop.
div2_equation is defined
div2_ind is defined
```

You may now look at the type of `div2_ind`:

```
Coq < Check div2_ind.
div2_ind
  : forall P : nat -> nat -> Prop,
    (forall n : nat, n = 0 -> P 0 0) ->
    (forall n n0 : nat, n = S n0 -> n0 = 0 -> P 1 0) ->
    (forall n n0 : nat,
      n = S n0 ->
      forall n' : nat,
        n0 = S n' -> P n' (div2 n') -> P (S (S n')) (S (div2 n')))) ->
    forall n : nat, P n (div2 n)
```

We can now prove the following lemma using this principle:

```
Coq < Lemma div2_le' : forall n:nat, div2 n <= n.
Coq < intro n.
Coq < pattern n , (div2 n).
Coq < apply div2_ind; intros.
3 subgoals

  n, n0 : nat
  e : n0 = 0
  =====
  0 <= 0
subgoal 2 is:
  0 <= 1
subgoal 3 is:
  S (div2 n') <= S (S n')

Coq < auto with arith.
Coq < auto with arith.
Coq < simpl; auto with arith.
Coq < Qed.
```

We can use directly the functional induction (8.5.5) tactic instead of the pattern/apply trick:

```
Coq < Reset div2_le'.
Coq < Lemma div2_le : forall n:nat, div2 n <= n.
Coq < intro n.

Coq < functional induction (div2 n).
3 subgoals

  =====
  0 <= 0
subgoal 2 is:
  0 <= 1
subgoal 3 is:
  S (div2 n') <= S (S n')
```

```

Coq < auto with arith.
Coq < auto with arith.
Coq < auto with arith.
Coq < Qed.

```

Remark: There is a difference between obtaining an induction scheme for a function by using `Function` (see Section 2.3) and by using `Functional Scheme` after a normal definition using `Fixpoint` or `Definition`. See 2.3 for details.

Example 2: *Induction scheme for tree_size*

We define trees by the following mutual inductive type:

```

Coq < Variable A : Set.
Coq < Inductive tree : Set :=
    node : A -> forest -> tree
  with forest : Set :=
    | empty : forest
    | cons : tree -> forest -> forest.

```

We define the function `tree_size` that computes the size of a tree or a forest. Note that we use `Function` which generally produces better principles.

```

Coq < Require Import FunInd.
Coq < Function tree_size (t:tree) : nat :=
  match t with
  | node A f => S (forest_size f)
  end
  with forest_size (f:forest) : nat :=
  match f with
  | empty => 0
  | cons t f' => (tree_size t + forest_size f')
  end.

```

Remark: `Function` generates itself non mutual induction principles `tree_size_ind` and `forest_size_ind`:

```

Coq < Check tree_size_ind.
tree_size_ind
  : forall P : tree -> nat -> Prop,
    (forall (t : tree) (A : A) (f : forest),
      t = node A f -> P (node A f) (S (forest_size f))) ->
    forall t : tree, P t (tree_size t)

```

The definition of mutual induction principles following the recursive structure of `tree_size` and `forest_size` is defined by the command:

```

Coq < Functional Scheme tree_size_ind2 := Induction for tree_size Sort Prop
  with forest_size_ind2 := Induction for forest_size Sort Prop.

```

You may now look at the type of `tree_size_ind2`:

```

Coq < Check tree_size_ind2.
tree_size_ind2
  : forall (P : tree -> nat -> Prop) (P0 : forest -> nat -> Prop),
    (forall (t : tree) (A : A) (f : forest),
      t = node A f ->
        P0 f (forest_size f) -> P (node A f) (S (forest_size f))) ->
    (forall f0 : forest, f0 = empty -> P0 empty 0) ->
    (forall (f1 : forest) (t : tree) (f' : forest),
      f1 = cons t f' ->
        P t (tree_size t) ->
        P0 f' (forest_size f') ->
        P0 (cons t f') (tree_size t + forest_size f')) ->
    forall t : tree, P t (tree_size t)

```

13.3 Generation of inversion principles with Derive Inversion

The syntax of `Derive Inversion` follows the schema:

`Derive Inversion ident with forall ($\vec{x}:\vec{T}$), $I \vec{t}$ Sort sort`

This command generates an inversion principle for the `inversion ... using tactic`. Let I be an inductive predicate and \vec{x} the variables occurring in \vec{t} . This command generates and stocks the inversion lemma for the sort *sort* corresponding to the instance $\forall(\vec{x}:\vec{T}), I \vec{t}$ with the name *ident* in the **global** environment. When applied, it is equivalent to having inverted the instance with the tactic `inversion`.

Variants:

1. `Derive Inversion_clear ident with forall ($\vec{x}:\vec{T}$), $I \vec{t}$ Sort sort`
 When applied, it is equivalent to having inverted the instance with the tactic `inversion` replaced by the tactic `inversion_clear`.
2. `Derive Dependent Inversion ident with forall ($\vec{x}:\vec{T}$), $I \vec{t}$ Sort sort`
 When applied, it is equivalent to having inverted the instance with the tactic `dependent inversion`.
3. `Derive Dependent Inversion_clear ident with forall ($\vec{x}:\vec{T}$), $I \vec{t}$ Sort sort`
 When applied, it is equivalent to having inverted the instance with the tactic `dependent inversion_clear`.

Example:

Let us consider the relation `Le` over natural numbers and the following variable:

```

Coq < Inductive Le : nat -> nat -> Set :=
  | Le0 : forall n:nat, Le 0 n
  | LeS : forall n m:nat, Le n m -> Le (S n) (S m).
Coq < Variable P : nat -> nat -> Prop.

```

To generate the inversion lemma for the instance `(Le (S n) m)` and the sort `Prop`, we do:

```
Coq < Derive Inversion_clear leminv with (forall n m:nat, Le (S n) m) Sort Prop.
```

```
Coq < Check leminv.
```

```
leminv
      : forall (n m : nat) (P : nat -> nat -> Prop),
        (forall m0 : nat, Le n m0 -> P n (S m0)) -> Le (S n) m -> P n m
```

Then we can use the proven inversion lemma:

```
Coq < Show.
```

```
1 subgoal
```

```

n, m : nat
H : Le (S n) m
=====
P n m
```

```
Coq < inversion H using leminv.
```

```
1 subgoal
```

```

n, m : nat
H : Le (S n) m
=====
forall m0 : nat, Le n m0 -> P n (S m0)
```

Part IV

Practical tools

Chapter 14

The COQ commands

There are three COQ commands:

- `coqtop`: the COQ toplevel (interactive mode);
- `coqcc`: the COQ compiler (batch compilation);
- `coqchk`: the COQ checker (validation of compiled libraries).

The options are (basically) the same for the first two commands, and roughly described below. You can also look at the man pages of `coqtop` and `coqcc` for more details.

14.1 Interactive use (`coqtop`)

In the interactive mode, also known as the COQ toplevel, the user can develop his theories and proofs step by step. The COQ toplevel is run by the command `coqtop`.

They are two different binary images of COQ: the byte-code one and the native-code one (if OCAML provides a native-code compiler for your platform, which is supposed in the following). By default, `coqtop` executes the native-code version; run `coqtop.byte` to get the byte-code version.

The byte-code toplevel is based on an OCAML toplevel (to allow the dynamic link of tactics). You can switch to the OCAML toplevel with the command `Drop.`, and come back to the COQ toplevel with the command `Coqloop.loop();;`.

14.2 Batch compilation (`coqcc`)

The `coqcc` command takes a name *file* as argument. Then it looks for a vernacular file named *file.v*, and tries to compile it into a *file.vo* file (See [6.5](#)).

Warning: The name *file* should be a regular COQ identifier, as defined in [Section 1.1](#). It should contain only letters, digits or underscores (`_`). For instance, `/bar/foo/toto.v` is valid, but `/bar/foo/to-to.v` is invalid.

14.3 Customization at launch time

14.3.1 By resource file

When COQ is launched, with either `coqtop` or `coqc`, the resource file `$XDG_CONFIG_HOME/coq/coqrc.xxx` is loaded, where `$XDG_CONFIG_HOME` is the configuration directory of the user (by default its home directory `/.config` and `xxx` is the version number (e.g. 8.3). If this file is not found, then the file `$XDG_CONFIG_HOME/coqrc` is searched. You can also specify an arbitrary name for the resource file (see option `-init-file` below).

This file may contain, for instance, `Add LoadPath` commands to add directories to the load path of COQ. It is possible to skip the loading of the resource file with the option `-q`.

14.3.2 By environment variables

Load path can be specified to the COQ system by setting up `$COQPATH` environment variable. It is a list of directories separated by `:` (`;` on windows). COQ will also honor `$XDG_DATA_HOME` and `$XDG_DATA_DIRS` (see Section 2.6.3).

Some COQ commands call other COQ commands. In this case, they look for the commands in directory specified by `$COQBIN`. If this variable is not set, they look for the commands in the executable path.

The `$COQ_COLORS` environment variable can be used to specify the set of colors used by `coqtop` to highlight its output. It uses the same syntax as the `$LS_COLORS` variable from GNU's `ls`, that is, a colon-separated list of assignments of the form `name=attr1;...;attrn` where `name` is the name of the corresponding highlight tag and `attri` is an ANSI escape code. The list of highlight tags can be retrieved with the `-list-tags` command-line option of `coqtop`.

14.3.3 By command line options

The following command-line options are recognized by the commands `coqc` and `coqtop`, unless stated otherwise:

`-I directory`, `-include directory`

Add physical path *directory* to the OCAML loadpath.

See also: Section 2.6.1 and the command `Declare ML Module` Section 6.5.

`-Q directory dirpath`

Add physical path *directory* to the list of directories where COQ looks for a file and bind it to the logical directory *dirpath*. The subdirectory structure of *directory* is recursively available from COQ using absolute names (extending the *dirpath* prefix) (see Section 2.6.2).

Note that only those subdirectories and files which obey the lexical conventions of what is an *ident* (see Section 1.1) are taken into account. Conversely, the underlying file systems or operating systems may be more restrictive than COQ. While Linux's ext4 file system supports any COQ recursive layout (within the limit of 255 bytes per file name), the default on NTFS (Windows) or HFS+ (MacOS X) file systems is on the contrary to disallow two files differing only in the case in the same directory.

See also: Section 2.6.1.

-R *directory dirpath*

Do as **-Q *directory dirpath*** but make the subdirectory structure of *directory* recursively visible so that the recursive contents of physical *directory* is available from COQ using short or partially qualified names.

See also: Section 2.6.1.

-top *dirpath*

Set the toplevel module name to *dirpath* instead of `Top`. Not valid for `coqc` as the toplevel module name is inferred from the name of the output file.

-exclude-dir *directory*

Exclude any subdirectory named *directory* while processing options such as **-R** and **-Q**. By default, only the conventional version control management directories named `CVS` and `_darcs` are excluded.

-nois

Start from an empty state instead of loading the `Init.Prelude` module.

-init-file *file*

Load *file* as the resource file instead of loading the default resource file from the standard configuration directories.

-q

Do not to load the default resource file.

-load-ml-source *file*

Load the OCAML source file *file*.

-load-ml-object *file*

Load the OCAML object file *file*.

-l *file*, -load-vernac-source *file*

Load and execute the COQ script from *file.v*.

-lv *file*, -load-vernac-source-verbose *file*

Load and execute the COQ script from *file.v*. Output its content on the standard input as it is executed.

-load-vernac-object *dirpath*

Load COQ compiled library *dirpath*. This is equivalent to running `Require dirpath`.

-require *dirpath*

Load COQ compiled library *dirpath* and import it. This is equivalent to running `Require Import dirpath`.

-batch

Exit just after argument parsing. Available for `coqtop` only.

`-compile file.v`

Compile file *file.v* into *file.vo*. This options imply `-batch` (exit just after argument parsing). It is available only for `coqtop`, as this behavior is the purpose of `coqc`.

`-compile-verbose file.v`

Same as `-compile` but also output the content of *file.v* as it is compiled.

`-verbose`

Output the content of the input file as it is compiled. This option is available for `coqc` only; it is the counterpart of `-compile-verbose`.

`-w (all|none|w1,...,wn)`

Configure the display of warnings. This option expects `all`, `none` or a comma-separated list of warning names or categories (see Section 6.9.3).

`-with-geoproof (yes|no)`

Enable or not special functions for Geoproof within COQIDE (default is yes).

`-color (on|off|auto)`

Enable or not the coloring of output of `coqtop`. Default is `auto`, meaning that `coqtop` dynamically decides, depending on whether the output channel supports ANSI escape sequences.

`-beautify`

Pretty-print each command to *file.beautified* when compiling *file.v*, in order to get old-fashioned syntax/definitions/notations.

`-emacs, -ide-slave`

Start a special toplevel to communicate with a specific IDE.

`-impredicative-set`

Change the logical theory of COQ by declaring the sort `Set` impredicative. Warning: this is known to be inconsistent with some standard axioms of classical mathematics such as the functional axiom of choice or the principle of description.

`-type-in-type`

Collapse the universe hierarchy of COQ. Warning: this makes the logic inconsistent.

`-compat version`

Attempt to maintain some backward-compatibility with a previous version.

`-dump-glob file`

Dump references for global names in file *file* (to be used by `coqdoc`, see 15.4). By default, if *file.v* is being compiled, *file.glob* is used.

`-no-glob`

Disable the dumping of references for global names.

`-image` *file*

Set the binary image to be used by `coqc` to be *file* instead of the standard one. Not of general use.

`-bindir` *directory*

Set the directory containing COQ binaries to be used by `coqc`. It is equivalent to doing `export COQBIN=directory` before launching `coqc`.

`-where`

Print the location of COQ's standard library and exit.

`-config`

Print the locations of COQ's binaries, dependencies, and libraries, then exit.

`-filteropts`

Print the list of command line arguments that `coqtop` has recognized as options and exit.

`-v`

Print COQ's version and exit.

`-list-tags`

Print the highlight tags known by COQ as well as their currently associated color and exit.

`-h, -help`

Print a short usage and exit.

14.4 Compiled libraries checker (`coqchk`)

The `coqchk` command takes a list of library paths as argument. The corresponding compiled libraries (.vo files) are searched in the path, recursively processing the libraries they depend on. The content of all these libraries is then type-checked. The effect of `coqchk` is only to return with normal exit code in case of success, and with positive exit code if an error has been found. Error messages are not deemed to help the user understand what is wrong. In the current version, it does not modify the compiled libraries to mark them as successfully checked.

Note that non-logical information is not checked. By logical information, we mean the type and optional body associated to names. It excludes for instance anything related to the concrete syntax of objects (customized syntax rules, association between short and long names), implicit arguments, etc.

This tool can be used for several purposes. One is to check that a compiled library provided by a third-party has not been forged and that loading it cannot introduce inconsistencies.¹ Another point is to get an even higher level of security. Since `coqtop` can be extended with custom tactics, possibly ill-typed code, it cannot be guaranteed that the produced compiled libraries are correct. `coqchk` is a standalone verifier, and thus it cannot be tainted by such malicious code.

Command-line options `-I`, `-R`, `-where` and `-impredicative-set` are supported by `coqchk` and have the same meaning as for `coqtop`. Extra options are:

¹Ill-formed non-logical information might for instance bind `Coq.Init.Logic.True` to short name `False`, so apparently `False` is inhabited, but using fully qualified names, `Coq.Init.Logic.False` will always refer to the absurd proposition, what we guarantee is that there is no proof of this latter constant.

`-norec module`

Check *module* but do not check its dependencies.

`-admit module`

Do not check *module* and any of its dependencies, unless explicitly required.

`-o`

At exit, print a summary about the context. List the names of all assumptions and variables (constants without body).

`-silent`

Do not write progress information in standard output.

Environment variable `$COQLIB` can be set to override the location of the standard library.

The algorithm for deciding which modules are checked or admitted is the following: assuming that `coqchk` is called with argument M , option `-norec` N , and `-admit` A . Let us write \overline{S} the set of reflexive transitive dependencies of set S . Then:

- Modules $C = \overline{M} \setminus \overline{A} \cup M \cup N$ are loaded and type-checked before being added to the context.
- And $\overline{M} \cup \overline{N} \setminus C$ is the set of modules that are loaded and added to the context without type-checking. Basic integrity checks (checksums) are nonetheless performed.

As a rule of thumb, the `-admit` can be used to tell that some libraries have already been checked. So `coqchk A B` can be split in `coqchk A && coqchk B -admit A` without type-checking any definition twice. Of course, the latter is slightly slower since it makes more disk access. It is also less secure since an attacker might have replaced the compiled library A after it has been read by the first command, but before it has been read by the second command.

Chapter 15

Utilities

The distribution provides utilities to simplify some tedious works beside proof development, tactics writing or documentation.

15.1 Building a toplevel extended with user tactics

The native-code version of COQ cannot dynamically load user tactics using OCAML code. It is possible to build a toplevel of COQ, with OCAML code statically linked, with the tool `coqmktop`.

For example, one can build a native-code COQ toplevel extended with a tactic which source is in `tactic.ml` with the command

```
% coqmktop -opt -o mytop.out tactic.cmx
```

where `tactic.ml` has been compiled with the native-code compiler `ocamlopt`. This command generates an executable called `mytop.out`. To use this executable to compile your COQ files, use `coqc -image mytop.out`.

A basic example is the native-code version of COQ (`coqtop.opt`), which can be generated by `coqmktop -opt -o coqopt.opt`.

Application: how to use the OCAML debugger with Coq. One useful application of `coqmktop` is to build a COQ toplevel in order to debug your tactics with the OCAML debugger. You need to have configured and compiled COQ for debugging (see the file `INSTALL` included in the distribution). Then, you must compile the Caml modules of your tactic with the option `-g` (with the bytecode compiler) and build a stand-alone bytecode toplevel with the following command:

```
% coqmktop -g -o coq-debug <your .cmo files>
```

To launch the OCAML debugger with the image you need to execute it in an environment which correctly sets the `COQLIB` variable. Moreover, you have to indicate the directories in which `ocamldebug` should search for Caml modules.

A possible solution is to use a wrapper around `ocamldebug` which detects the executables containing the word `coq`. In this case, the debugger is called with the required additional arguments. In other cases, the debugger is simply called without additional arguments. Such a wrapper can be found in the `dev/` subdirectory of the sources.

15.2 Building a COQ project with `coq_makefile`

The majority of COQ projects are very similar: a collection of `.v` files and eventually some `.ml` ones (a COQ plugin). The main piece of metadata needed in order to build the project are the command line options to `coqc` (e.g. `-R`, `-I`, **See also:** Section 14.3.3). Collecting the list of files and options is the job of the `_CoqProject` file.

A simple example of a `_CoqProject` file follows:

```
-R theories/ MyCode
theories/foo.v
theories/bar.v
-I src/
src/baz.ml4
src/bazaux.ml
src/qux_plugin.mlpack
```

Currently, both COQIDE and Proof General (version ≥ 4.3 pre) understand `_CoqProject` files and invoke COQ with the desired options.

The `coq_makefile` utility can be used to set up a build infrastructure for the COQ project based on makefiles. The recommended way of invoking `coq_makefile` is the following one:

```
coq_makefile -f _CoqProject -o CoqMakefile
```

Such command generates the following files:

`CoqMakefile` is a generic makefile for GNU Make that provides targets to build the project (both `.v` and `.ml*` files), to install it system-wide in the `coq-contrib` directory (i.e. where COQ is installed) as well as to invoke `coqdoc` to generate html documentation.

`CoqMakefile.conf` contains make variables assignments that reflect the contents of the `_CoqProject` file as well as the path relevant to COQ.

An optional file `CoqMakefile.local` can be provided by the user in order to extend `CoqMakefile`. In particular one can declare custom actions to be performed before or after the build process. Similarly one can customize the install target or even provide new targets. Extension points are documented in paragraph 15.2.

The extensions of the files listed in `_CoqProject` is used in order to decide how to build them. In particular:

- COQ files must use the `.v` extension
- OCAML files must use the `.ml` or `.mli` extension
- OCAML files that require pre processing for syntax extensions (like `VERNAC EXTEND`) must use the `.ml4` extension
- In order to generate a plugin one has to list all OCAML modules (i.e. “Baz” for “baz.ml”) in a `.mlpack` file (or `.mllib` file).

The use of `.mlpack` files has to be preferred over `.mllib` files, since it results in a “packed” plugin: All auxiliary modules (as `Baz` and `Bazaux`) are hidden inside the plugin’s “name space” (`Qux_plugin`). This reduces the chances of begin unable to load two distinct plugins because of a clash in their auxiliary module names.

CoqMakefile.local The optional file `CoqMakefile.local` is included by the generated file `CoqMakefile`. Such can contain two kinds of directives.

Variable assignment to the variables listed in the `Parameters` section of the generated makefile. Here we describe only few of them.

CAMLPKGS can be used to specify third party findlib packages, and is passed to the OCaml compiler on building or linking of modules. Eg: `-package yojson`.

CAMLFLAGS can be used to specify additional flags to the OCaml compiler, like `-bin-annot` or `-w...`.

COQC, COQDEP, COQDOC can be set in order to use alternative binaries (e.g. wrappers)

COQ_SRC_SUBDIRS can be extended by including other paths in which `*.cm*` files are searched. For example `COQ_SRC_SUBDIRS+=user-contrib/Unicoq` lets you build a plugin containing OCaml code that depends on the OCaml code of `Unicoq`.

Rule extension The following makefile rules can be extended. For example

```
pre-all::
    echo "This line is print before making the all target"
install-extra::
    cp ThisExtraFile /there/it/goes
```

pre-all:: run before the `all` target. One can use this to configure the project, or initialize sub modules or check dependencies are met.

post-all:: run after the `all` target. One can use this to run a test suite, or compile extracted code.

install-extra:: run after `install`. One can use this to install extra files.

install-doc:: One can use this to install extra doc.

uninstall::

uninstall-doc::

clean::

cleanall::

archclean::

merlin-hook:: One can append lines to the generated `.merlin` file extending this target.

Timing targets and performance testing The generated `Makefile` supports the generation of two kinds of timing data: per-file build-times, and per-line times for an individual file.

The following targets and `Makefile` variables allow collection of per-file timing data:

- **TIMED=1** — passing this variable will cause `make` to emit a line describing the user-space build-time and peak memory usage for each file built.

Note: On Mac OS, this works best if you've installed `gnu-time`.

Example: For example, the output of `make TIMED=1` may look like this:

```
COQDEP Fast.v
COQDEP Slow.v
COQC Slow.v
Slow (user: 0.34 mem: 395448 ko)
COQC Fast.v
Fast (user: 0.01 mem: 45184 ko)
```

- `pretty-timed` — this target stores the output of `make TIMED=1` into `time-of-build.log`, and displays a table of the times, sorted from slowest to fastest, which is also stored in `time-of-build-pretty.log`. If you want to construct the log for targets other than the default one, you can pass them via the variable `TGTS`, e.g., `make pretty-timed TGTS="a.vo b.vo"`.

Note: This target requires `python` to build the table.

Note: This target will *append* to the timing log; if you want a fresh start, you must remove the file `time-of-build.log` or run `make cleanall`.

Example: For example, the output of `make pretty-timed` may look like this:

```
COQDEP Fast.v
COQDEP Slow.v
COQC Slow.v
Slow (user: 0.36 mem: 393912 ko)
COQC Fast.v
Fast (user: 0.05 mem: 45992 ko)
Time      | File Name
-----
0m00.41s | Total
-----
0m00.36s | Slow
0m00.05s | Fast
```

- `print-pretty-timed-diff` — this target builds a table of timing changes between two compilations; run `make make-pretty-timed-before` to build the log of the “before” times, and run `make make-pretty-timed-after` to build the log of the “after” times. The table is printed on the command line, and stored in `time-of-build-both.log`. This target is most useful for profiling the difference between two commits to a repo.

Note: This target requires `python` to build the table.

Note: The `make-pretty-timed-before` and `make-pretty-timed-after` targets will *append* to the timing log; if you want a fresh start, you must remove the files `time-of-build-before.log` and `time-of-build-after.log` or run `make cleanall` *before* building either the “before” or “after” targets.

Note: The table will be sorted first by absolute time differences rounded towards zero to a whole-number of seconds, then by times in the “after” column, and finally lexicographically by file name. This will put the biggest changes in either direction first, and will prefer sorting by build-time over subsecond changes in build time (which are frequently noise); lexicographic sorting forces an order on files which take effectively no time to compile.

Example: For example, the output table from `make print-pretty-timed-diff` may look like this:

After	File Name	Before	Change	% Change
0m00.39s	Total	0m00.35s	+0m00.03s	+11.42%
0m00.37s	Slow	0m00.01s	+0m00.36s	+3600.00%
0m00.02s	Fast	0m00.34s	-0m00.32s	-94.11%

The following targets and Makefile variables allow collection of per-line timing data:

- `TIMING=1` — passing this variable will cause `make` to use `coqc -time` to write to a `.v.timing` file for each `.v` file compiled, which contains line-by-line timing information.

Example: For example, running `make all TIMING=1` may result in a file like this:

```
Chars 0 - 26 [Require~Coq.ZArith.BinInt.] 0.157 secs (0.128u,0.028s)
Chars 27 - 68 [Declare~Reduction~comp~::~~vm_c...] 0. secs (0.u,0.s)
Chars 69 - 162 [Definition~foo0~::~~Eval~comp~i...] 0.153 secs (0.136u,0.019s)
Chars 163 - 208 [Definition~foo1~::~~Eval~comp~i...] 0.239 secs (0.236u,0.s)
```

- `print-pretty-single-time-diff BEFORE=path/to/file.v.before-timing AFTER=path/to/file.v.after-timing` — this target will make a sorted table of the per-line timing differences between the timing logs in the `BEFORE` and `AFTER` files, display it, and save it to the file specified by the `TIME_OF_PRETTY_BUILD_FILE` variable, which defaults to `time-of-build-pretty.log`.

To generate the `.v.before-timing` or `.v.after-timing` files, you should pass `TIMING=before` or `TIMING=after` rather than `TIMING=1`.

Note: The sorting used here is the same as in the `print-pretty-timed-diff` target.

Note: This target requires `python` to build the table.

Example: For example, running `print-pretty-single-time-diff` might give a table like this:

After	Code	Before
0m00.50s	Total	0m04.17s
0m00.145s	Chars 069 - 162 [Definition~foo0~::~~Eval~comp~i...]	0m00.192s
0m00.126s	Chars 000 - 026 [Require~Coq.ZArith.BinInt.]	0m00.143s
N/A	Chars 027 - 068 [Declare~Reduction~comp~::~~nati...]	0m00.s
0m00.s	Chars 027 - 068 [Declare~Reduction~comp~::~~vm_c...]	N/A
0m00.231s	Chars 163 - 208 [Definition~foo1~::~~Eval~comp~i...]	0m03.836s

- `all.timing.diff, path/to/file.v.timing.diff` — The `path/to/file.v.timing.diff` target will make a `.v.timing.diff` file for the corresponding `.v` file, with a table as would be generated by the

`print-pretty-single-time-diff` target; it depends on having already made the corresponding `.v.before-timing` and `.v.after-timing` files, which can be made by passing `TIMING=before` and `TIMING=after`. The `all.timing.diff` target will make such timing difference files for all of the `.v` files that the Makefile knows about. It will fail if some `.v.before-timing` or `.v.after-timing` files don't exist.

Note: This target requires `python` to build the table.

Reusing/extending the generated Makefile Including the generated makefile with an `include` directive is discouraged. The contents of this file, including variable names and status of rules shall change in the future. Users are advised to include `Makefile.conf` or call a target of the generated Makefile as in `make -f Makefile target` from another Makefile.

One way to get access to all targets of the generated `CoqMakefile` is to have a generic target for invoking unknown targets. For example:

```
# KNOWNTARGETS will not be passed along to CoqMakefile
KNOWNTARGETS := CoqMakefile extra-stuff extra-stuff2
# KNOWNFILES will not get implicit targets from the final rule, and so
# depending on them won't invoke the submake
# Warning: These files get declared as PHONY, so any targets depending
# on them always get rebuilt
KNOWNFILES    := Makefile _CoqProject

.DEFAULT_GOAL := invoke-coqmakefile

CoqMakefile: Makefile _CoqProject
$(COQBIN)coq_makefile -f _CoqProject -o CoqMakefile

invoke-coqmakefile: CoqMakefile
$(MAKE) --no-print-directory -f CoqMakefile $(filter-out $(KNOWNTARGETS),$(MAKECMDLINE))

.PHONY: invoke-coqmakefile $(KNOWNFILES)

#####
##                               Your targets here                               ##
#####

# This should be the last rule, to handle any targets not declared above
%: invoke-coqmakefile
@true
```

Building a subset of the targets with `-j` To build, say, two targets `foo.vo` and `bar.vo` in parallel one can use `make only TGTS="foo.vo bar.vo" -j`.

Note that `make foo.vo bar.vo -j` has a different meaning for the `make` utility, in particular it may build a shared prerequisite twice.

Notes for users of `coq_makefile` with version < 8.7

- Support for “sub-directory” is deprecated. To perform actions before or after the build (like invoking `make` on a subdirectory) one can hook in `pre-all` and `post-all` extension points
- `-extra-phony` and `-extra` are deprecated. To provide additional target (`.PHONY` or not) please use `CoqMakefile.local`

15.3 Modules dependencies

In order to compute modules dependencies (so to use `make`), COQ comes with an appropriate tool, `coqdep`.

`coqdep` computes inter-module dependencies for COQ and OCAML programs, and prints the dependencies on the standard output in a format readable by `make`. When a directory is given as argument, it is recursively looked at.

Dependencies of COQ modules are computed by looking at `Require` commands (`Require`, `Require Export`, `Require Import`, but also at the command `Declare ML Module`.

Dependencies of OCAML modules are computed by looking at `open` commands and the dot notation `module.value`. However, this is done approximately and you are advised to use `ocamldep` instead for the OCAML modules dependencies.

See the man page of `coqdep` for more details and options.

The build infrastructure generated by `coq_makefile` uses `coqdep` to automatically compute the dependencies among the files part of the project.

15.4 Documenting COQ files with `coqdoc`

`coqdoc` is a documentation tool for the proof assistant COQ, similar to `javadoc` or `ocamldoc`. The task of `coqdoc` is

1. to produce a nice \LaTeX and/or HTML document from the COQ sources, readable for a human and not only for the proof assistant;
2. to help the user navigating in his own (or third-party) sources.

15.4.1 Principles

Documentation is inserted into COQ files as *special comments*. Thus your files will compile as usual, whether you use `coqdoc` or not. `coqdoc` presupposes that the given COQ files are well-formed (at least lexically). Documentation starts with `(**`, followed by a space, and ends with the pending `*)`. The documentation format is inspired by Todd A. Coram’s *Almost Free Text (AFT)* tool: it is mainly ASCII text with some syntax-light controls, described below. `coqdoc` is robust: it shouldn’t fail, whatever the input is. But remember: “garbage in, garbage out”.

COQ material inside documentation. COQ material is quoted between the delimiters `[` and `]`. Square brackets may be nested, the inner ones being understood as being part of the quoted code (thus you can quote a term like `fun x => u` by writing `[fun x => u]`). Inside quotations, the code is pretty-printed in the same way as it is in code parts.

Pre-formatted vernacular is enclosed by `[[` and `]]`. The former must be followed by a newline and the latter must follow a newline.

Pretty-printing. `coqdoc` uses different faces for identifiers and keywords. The pretty-printing of COQ tokens (identifiers or symbols) can be controlled using one of the following commands:

```
(** printing token %... $\text{\LaTeX}$ ...% #...HTML...# *)
```

or

```
(** printing token $... $\text{\LaTeX}$  math...$ #...HTML...# *)
```

It gives the \LaTeX and HTML texts to be produced for the given COQ token. One of the \LaTeX or HTML text may be omitted, causing the default pretty-printing to be used for this token.

The printing for one token can be removed with

```
(** remove printing token *)
```

Initially, the pretty-printing table contains the following mapping:

\rightarrow	\rightarrow	$<-$	\leftarrow	$*$	\times
\leq	\leq	\geq	\geq	\Rightarrow	\Rightarrow
$<>$	\neq	$<->$	\leftrightarrow	\vdash	\vdash
\setminus	\vee	\wedge	\wedge	\sim	\neg

Any of these can be overwritten or suppressed using the `printing` commands.

Important note: the recognition of tokens is done by a (ocaml)lex automaton and thus applies the longest-match rule. For instance, $\rightarrow\sim$ is recognized as a single token, where COQ sees two tokens. It is the responsibility of the user to insert space between tokens *or* to give pretty-printing rules for the possible combinations, e.g.

```
(** printing  $\rightarrow\sim$  %\ensuremath{\rightarrow\sim}% *)
```

Sections. Sections are introduced by 1 to 4 leading stars (i.e. at the beginning of the line) followed by a space. One star is a section, two stars a sub-section, etc. The section title is given on the remaining of the line. Example:

```
(** * Well-founded relations
```

```
    In this section, we introduce... *)
```

Lists. List items are introduced by a leading dash. `coqdoc` uses whitespace to determine the depth of a new list item and which text belongs in which list items. A list ends when a line of text starts at or before the level of indenting of the list's dash. A list item's dash must always be the first non-space character on its line (so, in particular, a list can not begin on the first line of a comment - start it on the second line instead).

Example:

```
We go by induction on [n]:
```

```
- If [n] is 0...
```

```
- If [n] is [S n'] we require...
```

```
    two paragraphs of reasoning, and two subcases:
```

```
    - In the first case...
```

```
    - In the second case...
```

```
So the theorem holds.
```

Rules. More than 4 leading dashes produce a horizontal rule.

Emphasis. Text can be italicized by placing it in underscores. A non-identifier character must precede the leading underscore and follow the trailing underscore, so that uses of underscores in names aren't mistaken for emphasis. Usually, these are spaces or punctuation.

This sentence contains some `_emphasized text_`.

Escaping to \LaTeX and HTML. Pure \LaTeX or HTML material can be inserted using the following escape sequences:

- `$...LaTeX stuff...$` inserts some \LaTeX material in math mode. Simply discarded in HTML output.
- `%...LaTeX stuff...%` inserts some \LaTeX material. Simply discarded in HTML output.
- `#...HTML stuff...#` inserts some HTML material. Simply discarded in \LaTeX output.

Note: to simply output the characters `$`, `%` and `#` and escaping their escaping role, these characters must be doubled.

Verbatim. Verbatim material is introduced by a leading `<<` and closed by `>>` at the beginning of a line. Example:

```
Here is the corresponding caml code:
<<
  let rec fact n =
    if n <= 1 then 1 else n * fact (n-1)
  >>
```

Hyperlinks. Hyperlinks can be inserted into the HTML output, so that any identifier is linked to the place of its definition.

`coqc file.v` automatically dumps localization information in `file.glob` or appends it to a file specified using option `--dump-glob file`. Take care of erasing this global file, if any, when starting the whole compilation process.

Then invoke `coqdoc` or `coqdoc --glob-from file` to tell `coqdoc` to look for name resolutions into the file `file` (it will look in `file.glob` by default).

Identifiers from the COQ standard library are linked to the COQ web site at <http://coq.inria.fr/library/>. This behavior can be changed using command line options `--no-externals` and `--coqlib`; see below.

Hiding / Showing parts of the source. Some parts of the source can be hidden using command line options `-g` and `-l` (see below), or using such comments:

```
(* begin hide *)
some Coq material
(* end hide *)
```

Conversely, some parts of the source which would be hidden can be shown using such comments:

```
(* begin show *)
some Coq material
(* end show *)
```

The latter cannot be used around some inner parts of a proof, but can be used around a whole proof.

15.4.2 Usage

`coqdoc` is invoked on a shell command line as follows:

```
coqdoc < options and files >
```

Any command line argument which is not an option is considered to be a file (even if it starts with a `-`). COQ files are identified by the suffixes `.v` and `.g` and \LaTeX files by the suffix `.tex`.

HTML output

This is the default output. One HTML file is created for each COQ file given on the command line, together with a file `index.html` (unless option `-no-index` is passed). The HTML pages use a style sheet named `style.css`. Such a file is distributed with `coqdoc`.

\LaTeX output

A single \LaTeX file is created, on standard output. It can be redirected to a file with option `-o`. The order of files on the command line is kept in the final document. \LaTeX files given on the command line are copied ‘as is’ in the final document. DVI and PostScript can be produced directly with the options `-dvi` and `-ps` respectively.

\TeX macs output

To translate the input files to \TeX macs format, to be used by the \TeX macs Coq interface.

Command line options

Overall options

`--html`

Select a HTML output.

`--latex`

Select a \LaTeX output.

`--dvi`

Select a DVI output.

`--ps`

Select a PostScript output.

`--texmacs`

Select a \TeX macs output.

--stdout

Write output to stdout.

-o file, --output file

Redirect the output into the file '*file*' (meaningless with `-html`).

-d dir, --directory dir

Output files into directory '*dir*' instead of current directory (option `-d` does not change the filename specified with option `-o`, if any).

--body-only

Suppress the header and trailer of the final document. Thus, you can insert the resulting document into a larger one.

-p string, --preamble string

Insert some material in the \LaTeX preamble, right before `\begin{document}` (meaningless with `-html`).

--vernac-file file, --tex-file file

Considers the file '*file*' respectively as a `.v` (or `.g`) file or a `.tex` file.

--files-from file

Read file names to process in file '*file*' as if they were given on the command line. Useful for program sources split up into several directories.

-q, --quiet

Be quiet. Do not print anything except errors.

-h, --help

Give a short summary of the options and exit.

-v, --version

Print the version and exit.

Index options Default behavior is to build an index, for the HTML output only, into `index.html`.

--no-index

Do not output the index.

--multi-index

Generate one page for each category and each letter in the index, together with a top page `index.html`.

--index string

Make the filename of the index *string* instead of "index". Useful since "index.html" is special.

Table of contents option

-toc, --table-of-contents

Insert a table of contents. For a \LaTeX output, it inserts a `\tableofcontents` at the beginning of the document. For a HTML output, it builds a table of contents into `toc.html`.

--toc-depth *int*

Only include headers up to depth *int* in the table of contents.

Hyperlinks options

--glob-from *file*

Make references using COQ globalizations from file *file*. (Such globalizations are obtained with COQ option `-dump-glob`).

--no-externals

Do not insert links to the COQ standard library.

--external *url coqdir*

Use given URL for linking references whose name starts with prefix *coqdir*.

--coqlib *url*

Set base URL for the COQ standard library (default is <http://coq.inria.fr/library/>). This is equivalent to `--external url Coq`.

-R *dir coqdir*

Map physical directory *dir* to COQ logical directory *coqdir* (similarly to COQ option `-R`).

Note: option `-R` only has effect on the files *following* it on the command line, so you will probably need to put this option first.

Title options

-s , --short

Do not insert titles for the files. The default behavior is to insert a title like “Library Foo” for each file.

--lib-name *string*

Print “*string* Foo” instead of “Library Foo” in titles. For example “Chapter” and “Module” are reasonable choices.

--no-lib-name

Print just “Foo” instead of “Library Foo” in titles.

--lib-subtitles

Look for library subtitles. When enabled, the beginning of each file is checked for a comment of the form:

```
(** * ModuleName : text *)
```

where `ModuleName` must be the name of the file. If it is present, the `text` is used as a subtitle for the module in appropriate places.

-t *string*, --title *string*

Set the document title.

Contents options

-g, --gallina

Do not print proofs.

-l, --light

Light mode. Suppress proofs (as with `-g`) and the following commands:

- `[Recursive]Tactic Definition`
- `Hint / Hints`
- `Require`
- `Transparent / Opaque`
- `Implicit Argument / Implicits`
- `Section / Variable / Hypothesis / End`

The behavior of options `-g` and `-l` can be locally overridden using the `(* begin show *) ... (* end show *)` environment (see above).

There are a few options to drive the parsing of comments:

--parse-comments

Parses regular comments delimited by `(*` and `*)` as well. They are typeset inline.

--plain-comments

Do not interpret comments, simply copy them as plain-text.

--interpolate

Use the globalization information to typeset identifiers appearing in COQ escapings inside comments.

Language options Default behavior is to assume ASCII 7 bits input files.

-latin1, --latin1

Select ISO-8859-1 input files. It is equivalent to `--inputenc latin1 --charset iso-8859-1`.

-utf8, --utf8

Set `--inputenc utf8x` for \LaTeX output and `--charset utf-8` for HTML output. Also use Unicode replacements for a couple of standard plain ASCII notations such as \rightarrow for `->` and \forall for `forall`. \LaTeX UTF-8 support can be found at <http://www.ctan.org/pkg/unicode>.

For the interpretation of Unicode characters by \LaTeX , extra packages which `coqdoc` does not provide by default might be required, such as `textgreek` for some Greek letters or `stmaryrd` for some mathematical symbols. If a Unicode character is missing an interpretation in the `utf8x` input encoding, add `\DeclareUnicodeCharacter{code}{latex-interpretation}`. Packages and declarations can be added with option `-p`.

--inputenc *string*

Give a \LaTeX input encoding, as an option to \LaTeX package `inputenc`.

--charset *string*

Specify the HTML character set, to be inserted in the HTML header.

15.4.3 The `coqdoc` \LaTeX style file

In case you choose to produce a document without the default \LaTeX preamble (by using option `--no-preamble`), then you must insert into your own preamble the command

```
\usepackage{coqdoc}
```

The package optionally takes the argument `[color]` to typeset identifiers with colors (this requires the `xcolor` package).

Then you may alter the rendering of the document by redefining some macros:

`coqdockw`, `coqdocid`, ...

The one-argument macros for typesetting keywords and identifiers. Defaults are sans-serif for keywords and italic for identifiers.

For example, if you would like a slanted font for keywords, you may insert

```
\renewcommand{\coqdockw}[1]{\textsl{#1}}
```

anywhere between `\usepackage{coqdoc}` and `\begin{document}`.

`coqdocmodule`

One-argument macro for typesetting the title of a `.v` file. Default is

```
\newcommand{\coqdocmodule}[1]{\section*{Module #1}}
```

and you may redefine it using `\renewcommand`.

15.5 Embedded COQ phrases inside \LaTeX documents

When writing a documentation about a proof development, one may want to insert COQ phrases inside a \LaTeX document, possibly together with the corresponding answers of the system. We provide a mechanical way to process such COQ phrases embedded in \LaTeX files: the `coq-tex` filter. This filter extracts Coq phrases embedded in LaTeX files, evaluates them, and insert the outcome of the evaluation after each phrase.

Starting with a file `file.tex` containing COQ phrases, the `coq-tex` filter produces a file named `file.v.tex` with the COQ outcome.

There are options to produce the COQ parts in smaller font, italic, between horizontal rules, etc. See the man page of `coq-tex` for more details.

Remark. This Reference Manual and the Tutorial have been completely produced with `coq-tex`.

15.6 COQ and GNU EMACS

15.6.1 The COQ Emacs mode

COQ comes with a Major mode for GNU EMACS, `gallina.el`. This mode provides syntax highlighting and also a rudimentary indentation facility in the style of the Caml GNU EMACS mode.

Add the following lines to your `.emacs` file:

```
(setq auto-mode-alist (cons '("\\.v$" . coq-mode) auto-mode-alist))
(autoload 'coq-mode "gallina" "Major mode for editing Coq vernacular." t)
```

The COQ major mode is triggered by visiting a file with extension `.v`, or manually with the command `M-x coq-mode`. It gives you the correct syntax table for the COQ language, and also a rudimentary indentation facility:

- pressing TAB at the beginning of a line indents the line like the line above;
- extra TABs increase the indentation level (by 2 spaces by default);
- M-TAB decreases the indentation level.

An inferior mode to run COQ under Emacs, by Marco Maggesi, is also included in the distribution, in file `coq-inferior.el`. Instructions to use it are contained in this file.

15.6.2 PROOF GENERAL

PROOF GENERAL is a generic interface for proof assistants based on Emacs. The main idea is that the COQ commands you are editing are sent to a COQ toplevel running behind Emacs and the answers of the system automatically inserted into other Emacs buffers. Thus you don't need to copy-paste the COQ material from your files to the COQ toplevel or conversely from the COQ toplevel to some files.

PROOF GENERAL is developed and distributed independently of the system COQ. It is freely available at <https://proofgeneral.github.io/>.

15.7 Module specification

Given a COQ vernacular file, the `gallina` filter extracts its specification (inductive types declarations, definitions, type of lemmas and theorems), removing the proofs parts of the file. The COQ file *file.v* gives birth to the specification file *file.g* (where the suffix `.g` stands for GALLINA).

See the man page of `gallina` for more details and options.

15.8 Man pages

There are man pages for the commands `coqdep`, `gallina` and `coq-tex`. Man pages are installed at installation time (see installation instructions in file `INSTALL`, step 6).

Chapter 16

CoQ Integrated Development Environment

The CoQ Integrated Development Environment is a graphical tool, to be used as a user-friendly replacement to `coqtop`. Its main purpose is to allow the user to navigate forward and backward into a CoQ vernacular file, executing corresponding commands or undoing them respectively.

COQIDE is run by typing the command `coqide` on the command line. Without argument, the main screen is displayed with an “unnamed buffer”, and with a file name as argument, another buffer displaying the contents of that file. Additionally, `coqide` accepts the same options as `coqtop`, given in Chapter 14, the ones having obviously no meaning for COQIDE being ignored. Additionally, `coqide` accepts the option `-enable-geoproof` to enable the support for *GeoProof*¹.

A sample COQIDE main screen, while navigating into a file `Fermat.v`, is shown on Figure 16.1. At the top is a menu bar, and a tool bar below it. The large window on the left is displaying the various *script buffers*. The upper right window is the *goal window*, where goals to prove are displayed. The lower right window is the *message window*, where various messages resulting from commands are displayed. At the bottom is the status bar.

16.1 Managing files and buffers, basic edition

In the script window, you may open arbitrarily many buffers to edit. The *File* menu allows you to open files or create some, save them, print or export them into various formats. Among all these buffers, there is always one which is the current *running buffer*, whose name is displayed on a green background, which is the one where Coq commands are currently executed.

Buffers may be edited as in any text editor, and classical basic editing commands (Copy/Paste, ...) are available in the *Edit* menu. COQIDE offers only basic editing commands, so if you need more complex editing commands, you may launch your favorite text editor on the current buffer, using the *Edit/External Editor* menu.

¹*GeoProof* is dynamic geometry software which can be used in conjunction with COQIDE to interactively build a Coq statement corresponding to a geometric figure. More information about *GeoProof* can be found here: <http://home.gna.org/geoproof/>



Figure 16.1: COQIDE main screen

16.2 Interactive navigation into COQ scripts

The running buffer is the one where navigation takes place. The toolbar proposes five basic commands for this. The first one, represented by a down arrow icon, is for going forward executing one command. If that command is successful, the part of the script that has been executed is displayed on a green background. If that command fails, the error message is displayed in the message window, and the location of the error is emphasized by a red underline.

On Figure 16.1, the running buffer is `Fermat.v`, all commands until the `Theorem` have been already executed, and the user tried to go forward executing `Induction n`. That command failed because no such tactic exist (tactics are now in lowercase...), and the wrong word is underlined.

Notice that the green part of the running buffer is not editable. If you ever want to modify something you have to go backward using the up arrow tool, or even better, put the cursor where you want to go back and use the `goto` button. Unlike with `coqtop`, you should never use `Undo` to go backward.

Two additional tool buttons exist, one to go directly to the end and one to go back to the beginning. If you try to go to the end, or in general to run several commands using the `goto` button, the execution will stop whenever an error is found.

If you ever try to execute a command which happens to run during a long time, and would like to abort it before its termination, you may use the interrupt button (the white cross on a red circle).

Finally, notice that these navigation buttons are also available in the menu, where their keyboard shortcuts are given.

16.3 Try tactics automatically

The menu `Try Tactics` provides some features for automatically trying to solve the current goal using simple tactics. If such a tactic succeeds in solving the goal, then its text is automatically inserted into the script. There is finally a combination of these tactics, called the *proof wizard* which will try each of them in turn. This wizard is also available as a tool button (the light bulb). The set of tactics tried by the wizard is customizable in the preferences.

These tactics are general ones, in particular they do not refer to particular hypotheses. You may also try specific tactics related to the goal or one of the hypotheses, by clicking with the right mouse button on the goal or the considered hypothesis. This is the “contextual menu on goals” feature, that may be disabled in the preferences if undesirable.

16.4 Proof folding

As your script grows bigger and bigger, it might be useful to hide the proofs of your theorems and lemmas.

This feature is toggled via the `Hide` entry of the `Navigation` menu. The proof shall be enclosed between `Proof.` and `Qed.`, both with their final dots. The proof that shall be hidden or revealed is the first one whose beginning statement (such as `Theorem`) precedes the insertion cursor.

16.5 Vernacular commands, templates

The `Templates` menu allows using shortcuts to insert vernacular commands. This is a nice way to proceed if you are not sure of the spelling of the command you want.

Moreover, this menu offers some *templates* which will automatic insert a complex command like `Fixpoint` with a convenient shape for its arguments.

16.6 Queries

We call *query* any vernacular command that do not change the current state, such as `Check`, `Search`, etc. Those commands are of course useless during compilation of a file, hence should not be included in scripts. To run such commands without writing them in the script, COQIDE offers another input window called the *query window*. This window can be displayed on demand, either by using the `Window` menu, or directly using shortcuts given in the `Queries` menu. Indeed, with COQIDE the simplest way to perform a `Search` on some identifier is to select it using the mouse, and pressing `F2`. This will both make appear the query window and run the `Search` in it, displaying the result. Shortcuts `F3` and `F4` are for `Check` and `Print` respectively. Figure 16.2 displays the query window after selection of the word “mult” in the script windows, and pressing `F4` to print its definition.

16.7 Compilation

The `Compile` menu offers direct commands to:

- compile the current buffer
- run a compilation using `make`



Figure 16.2: COQIDE: the query window

- go to the last compilation error
- create a makefile using `coq_makefile`.

16.8 Customizations

You may customize your environment using menu `Edit/Preferences`. A new window will be displayed, with several customization sections presented as a notebook.

The first section is for selecting the text font used for scripts, goal and message windows.

The second section is devoted to file management: you may configure automatic saving of files, by periodically saving the contents into files named `#f#` for each opened file `f`. You may also activate the *revert* feature: in case a opened file is modified on the disk by a third party, COQIDE may read it again for you. Note that in the case you edited that same file, you will be prompt to choose to either discard your changes or not. The File charset encoding choice is described below in Section 16.9.3

The Externals section allows customizing the external commands for compilation, printing, web browsing. In the browser command, you may use `%s` to denote the URL to open, for example: `mozilla -remote "OpenURL(%s)"`.

The Tactics Wizard section allows defining the set of tactics that should be tried, in sequence, to solve the current goal.

The last section is for miscellaneous boolean settings, such as the “contextual menu on goals” feature presented in Section 16.3.

Notice that these settings are saved in the file `.coqiderc` of your home directory.

A gtk2 accelerator keymap is saved under the name `.coqide.keys`. It is not recommended to edit this file manually: to modify a given menu shortcut, go to the corresponding menu item without releasing the mouse button, press the key you want for the new shortcut, and release the mouse button afterwards. If your system does not allow it, you may still edit this configuration file by hand, but this is more involved.

16.9 Using Unicode symbols

COQIDE is based on GTK+ and inherits from it support for Unicode in its text windows. Consequently a large set of symbols is available for notations.

16.9.1 Displaying Unicode symbols

You just need to define suitable notations as described in Chapter 12. For example, to use the mathematical symbols \forall and \exists , you may define

```
Notation "∀ x : t, P" :=
  (forall x:t, P) (at level 200, x ident).
Notation "∃ x : t, P" :=
  (exists x:t, P) (at level 200, x ident).
```

There exists a small set of such notations already defined, in the file `utf8.v` of COQ library, so you may enable them just by `Require utf8` inside COQIDE, or equivalently, by starting COQIDE with `coqide -l utf8`.

However, there are some issues when using such Unicode symbols: you of course need to use a character font which supports them. In the Fonts section of the preferences, the Preview line displays some Unicode symbols, so you could figure out if the selected font is OK. Related to this, one thing you may need to do is choose whether GTK+ should use antialiased fonts or not, by setting the environment variable `GDK_USE_XFT` to 1 or 0 respectively.

16.9.2 Defining an input method for non ASCII symbols

To input a Unicode symbol, a general method provided by GTK+ is to simultaneously press the Control, Shift and “u” keys, release, then type the hexadecimal code of the symbol required, for example 2200 for the \forall symbol. A list of symbol codes is available at <http://www.unicode.org>.

An alternative method which does not require to know the hexadecimal code of the character is to use an Input Method Editor. On POSIX systems (Linux distributions, BSD variants and MacOS X), you can use `uim` version 1.6 or later which provides a \LaTeX -style input method.

To configure `uim`, execute `uim-pref-gtk` as your regular user. In the "Global Settings" group set the default Input Method to "ELatin" (don't forget to tick the checkbox "Specify default IM"). In the "ELatin" group set the layout to "TeX", and remember the content of the "[ELatin] on" field (by default Control- \backslash). You can now execute CoqIDE with the following commands (assuming you use a Bourne-style shell):

```
$ export GTK_IM_MODULE=uim
$ coqide
```

Activate the ELatin Input Method with Control- \backslash , then type the sequence `"\Gamma"`. You will see the sequence being replaced by Γ as soon as you type the second "a".

16.9.3 Character encoding for saved files

In the `Files` section of the preferences, the encoding option is related to the way files are saved.

If you have no need to exchange files with non UTF-8 aware applications, it is better to choose the UTF-8 encoding, since it guarantees that your files will be read again without problems. (This is because when COQIDE reads a file, it tries to automatically detect its character encoding.)

If you choose something else than UTF-8, then missing characters will be written encoded by `\x{ }` or `\x{ }` where each dot is an hexadecimal digit: the number between braces is the hexadecimal Unicode index for the missing character.

Part V

Addendum to the Reference Manual

Presentation of the Addendum

Here you will find several pieces of additional documentation for the COQ Reference Manual. Each of these chapters is concentrated on a particular topic, that should interest only a fraction of the COQ users: that's the reason why they are apart from the Reference Manual.

Extended pattern-matching This chapter details the use of generalized pattern-matching. It is contributed by Cristina Cornes and Hugo Herbelin.

Implicit coercions This chapter details the use of the coercion mechanism. It is contributed by Amokrane Saïbi.

Program extraction This chapter explains how to extract in practice ML files from F_ω terms. It is contributed by Jean-Christophe Filliâtre and Pierre Letouzey.

Program This chapter explains the use of the `Program` vernacular which allows the development of certified programs in COQ. It is contributed by Matthieu Sozeau and replaces the previous `Program` tactic by Catherine Parent.

omega `omega`, written by Pierre Crégut, solves a whole class of arithmetic problems.

The `ring` tactic This is a tactic to do AC rewriting. This chapter explains how to use it and how it works. The chapter is contributed by Patrick Loiseleur.

The `Setoid_replace` tactic This is a tactic to do rewriting on types equipped with specific (only partially substitutive) equality. The chapter is contributed by Clément Renard.

Calling external provers This chapter describes several tactics which call external provers.

Contents

The SSReflect proof language	299
Extended pattern-matching	427
Patterns	427
About patterns of parametric types	430
Matching objects of dependent types	432
Understanding dependencies in patterns	432
When the elimination predicate must be provided	433
Using pattern matching to write proofs	434
Pattern-matching on inductive objects involving local definitions	435

Pattern-matching and coercions	436
When does the expansion strategy fail ?	436
Implicit Coercions	439
General Presentation	439
Classes	439
Coercions	440
Identity Coercions	440
Inheritance Graph	441
Declaration of Coercions	441
Coercion <i>qualid</i> : <i>class</i> ₁ \rightarrow <i>class</i> ₂	441
Identity Coercion <i>ident</i> : <i>class</i> ₁ \rightarrow <i>class</i> ₂	442
Displaying Available Coercions	443
Print Classes.	443
Print Coercions.	443
Print Graph.	443
Print Coercion Paths <i>class</i> ₁ <i>class</i> ₂	443
Activating the Printing of Coercions	443
Set Printing Coercions.	443
Add Printing Coercion <i>qualid</i>	443
Classes as Records	444
Coercions and Sections	444
Coercions and Modules	444
Examples	444
Canonical Structures	449
Type Classes	459
Class and Instance declarations	459
Binding classes	460
Parameterized Instances	461
Sections and contexts	461
Building hierarchies	462
typeclasses eauto	465
autoapply <i>term</i> with <i>ident</i>	465
Omega: a solver of quantifier-free problems in Presburger Arithmetic	469
Description of omega	469
Arithmetical goals recognized by omega	469
Messages from omega	470
Technical data	471
Overview of the tactic	471
Overview of the <i>OMEGA</i> decision procedure	471
Bugs	472

Micromega: tactics for solving arithmetic goals over ordered rings	473
Short description of the tactics	473
<i>Positivstellensatz</i> refutations	474
lra: a decision procedure for linear real and rational arithmetic	474
lia: a tactic for linear integer arithmetic	475
nra: a proof procedure for non-linear arithmetic	476
nia: a proof procedure for non-linear integer arithmetic	476
psatz: a proof procedure for non-linear arithmetic	476
Extraction of programs in Objective Caml and Haskell	477
Generating ML code	477
Extraction options	478
Setting the target language	478
Inlining and optimizations	479
Extra elimination of useless arguments	480
Realizing axioms	480
Avoiding conflicts with existing filenames	482
Differences between COQ and ML type systems	483
Some examples	484
A detailed example: Euclidean division	484
Extraction's horror museum	485
Users' Contributions	486
PROGRAM	487
Elaborating programs	487
The <code>ring</code> and <code>field</code> tactic families	493
What does this tactic do?	493
The variables map	494
Is it automatic?	494
Concrete usage in COQ	494
Adding a ring structure	496
How does it work?	499
Dealing with fields	500
Adding a new field structure	501
History of <code>ring</code>	502
Discussion	503
Nsatz: tactics for proving equalities in integral domains	505
Using the basic tactic <code>nsatz</code>	505
More about <code>nsatz</code>	505
Generalized rewriting	507
Introduction to generalized rewriting	508
Commands and tactics	513
Extensions	515
Strategies for rewriting	517

Asynchronous and Parallel Proof Processing	519
Polymorphic Universes	523
General Presentation	523
Polymorphic, Monomorphic	525
Cumulative, NonCumulative	525
Global and local universes	528
Conversion and unification	528
Minimization	528
Explicit Universes	529
Miscellaneous extensions	531
Program derivation	531

Chapter 17

Extended pattern-matching

Cristina Cornes and Hugo Herbelin

This section describes the full form of pattern-matching in COQ terms.

17.1 Patterns

The full syntax of `match` is presented in Figures 1.1 and 1.2. Identifiers in patterns are either constructor names or variables. Any identifier that is not the constructor of an inductive or co-inductive type is considered to be a variable. A variable name cannot occur more than once in a given pattern. It is recommended to start variable names by a lowercase letter.

If a pattern has the form $(c \vec{x})$ where c is a constructor symbol and \vec{x} is a linear vector of (distinct) variables, it is called *simple*: it is the kind of pattern recognized by the basic version of `match`. On the opposite, if it is a variable x or has the form $(c \vec{p})$ with p not only made of variables, the pattern is called *nested*.

A variable pattern matches any value, and the identifier is bound to that value. The pattern “`_`” (called “don’t care” or “wildcard” symbol) also matches any value, but does not bind anything. It may occur an arbitrary number of times in a pattern. Alias patterns written $(pattern \text{ as } identifier)$ are also accepted. This pattern matches the same values as *pattern* does and *identifier* is bound to the matched value. A pattern of the form $pattern | pattern$ is called disjunctive. A list of patterns separated with commas is also considered as a pattern and is called *multiple pattern*. However multiple patterns can only occur at the root of pattern-matching equations. Disjunctions of *multiple pattern* are allowed though.

Since extended `match` expressions are compiled into the primitive ones, the expressiveness of the theory remains the same. Once the stage of parsing has finished only simple patterns remain. Re-nesting of pattern is performed at printing time. An easy way to see the result of the expansion is to toggle off the nesting performed at printing (use here `Set Printing Matching`), then by printing the term with `Print` if the term is a constant, or using the command `Check`.

The extended `match` still accepts an optional *elimination predicate* given after the keyword `return`. Given a pattern matching expression, if all the right-hand-sides of $=>$ (*rhs* in short) have the same type, then this type can be sometimes synthesized, and so we can omit the `return` part. Otherwise the predicate after `return` has to be provided, like for the basic `match`.

Let us illustrate through examples the different aspects of extended pattern matching. Consider for example the function that computes the maximum of two natural numbers. We can write it in primitive syntax by:

```
Coq < Fixpoint max (n m:nat) {struct m} : nat :=
  match n with
  | O => m
  | S n' => match m with
            | O => S n'
            | S m' => S (max n' m')
          end
  end.
max is defined
max is recursively defined (decreasing on 2nd argument)
```

Multiple patterns Using multiple patterns in the definition of `max` lets us write:

```
Coq < Fixpoint max (n m:nat) {struct m} : nat :=
  match n, m with
  | O, _ => m
  | S n', O => S n'
  | S n', S m' => S (max n' m')
  end.
max is defined
max is recursively defined (decreasing on 2nd argument)
```

which will be compiled into the previous form.

The pattern-matching compilation strategy examines patterns from left to right. A `match` expression is generated **only** when there is at least one constructor in the column of patterns. E.g. the following example does not build a `match` expression.

```
Coq < Check (fun x:nat => match x return nat with
                        | y => y
                      end).

fun x : nat => x
  : nat -> nat
```

Aliasing subpatterns We can also use “as *ident*” to associate a name to a sub-pattern:

```
Coq < Fixpoint max (n m:nat) {struct n} : nat :=
  match n, m with
  | O, _ => m
  | S n' as p, O => p
  | S n', S m' => S (max n' m')
  end.
max is defined
max is recursively defined (decreasing on 1st argument)
```

Nested patterns Here is now an example of nested patterns:

```

Coq < Fixpoint even (n:nat) : bool :=
  match n with
  | 0 => true
  | S 0 => false
  | S (S n') => even n'
  end.
even is defined
even is recursively defined (decreasing on 1st argument)

```

This is compiled into:

```

Coq < Unset Printing Matching.
Coq < Print even.
even =
fix even (n : nat) : bool :=
  match n with
  | 0 => true
  | S n0 => match n0 with
    | 0 => false
    | S n' => even n'
    end
  end
  : nat -> bool
Argument scope is [nat_scope]

```

In the previous examples patterns do not conflict with, but sometimes it is comfortable to write patterns that admit a non trivial superposition. Consider the boolean function `leq` that given two natural numbers yields `true` if the first one is less or equal than the second one and `false` otherwise. We can write it as follows:

```

Coq < Fixpoint leq (n m:nat) {struct m} : bool :=
  match n, m with
  | 0, x => true
  | x, 0 => false
  | S n, S m => leq n m
  end.
leq is defined
leq is recursively defined (decreasing on 2nd argument)

```

Note that the first and the second multiple pattern superpose because the couple of values `0 0` matches both. Thus, what is the result of the function on those values? To eliminate ambiguity we use the *textual priority rule*: we consider patterns ordered from top to bottom, then a value is matched by the pattern at the *i*th row if and only if it is not matched by some pattern of a previous row. Thus in the example, `0 0` is matched by the first pattern, and so `(leq 0 0)` yields `true`.

Another way to write this function is:

```

Coq < Fixpoint leq (n m:nat) {struct m} : bool :=
  match n, m with
  | 0, x => true
  | S n, S m => leq n m
  | _, _ => false
  end.

```

lef is defined
lef is recursively defined (decreasing on 2nd argument)

Here the last pattern superposes with the first two. Because of the priority rule, the last pattern will be used only for values that do not match neither the first nor the second one.

Terms with useless patterns are not accepted by the system. Here is an example:

```
Coq < Fail Check (fun x:nat =>
    match x with
    | 0 => true
    | S _ => false
    | x => true
    end).
```

The command has indeed failed with message:
This clause is redundant.

Disjunctive patterns Multiple patterns that share the same right-hand-side can be factorized using the notation *mult_pattern* | ... | *mult_pattern*. For instance, *max* can be rewritten as follows:

```
Coq < Fixpoint max (n m:nat) {struct m} : nat :=
    match n, m with
    | S n', S m' => S (max n' m')
    | 0, p | p, 0 => p
    end.
```

max is defined
max is recursively defined (decreasing on 2nd argument)

Similarly, factorization of (non necessary multiple) patterns that share the same variables is possible by using the notation *pattern* | ... | *pattern*. Here is an example:

```
Coq < Definition filter_2_4 (n:nat) : nat :=
    match n with
    | 2 as m | 4 as m => m
    | _ => 0
    end.
```

filter_2_4 is defined

Here is another example using disjunctive subpatterns.

```
Coq < Definition filter_some_square_corners (p:nat*nat) : nat*nat :=
    match p with
    | ((2 as m | 4 as m), (3 as n | 5 as n)) => (m,n)
    | _ => (0,0)
    end.
```

filter_some_square_corners is defined

17.2 About patterns of parametric types

Parameters in patterns When matching objects of a parametric type, parameters do not bind in patterns. They must be substituted by “_”. Consider for example the type of polymorphic lists:

```
Coq < Inductive List (A:Set) : Set :=
  | nil : List A
  | cons : A -> List A -> List A.
List is defined
List_rect is defined
List_ind is defined
List_rec is defined
```

We can check the function *tail*:

```
Coq < Check
  (fun l:List nat =>
    match l with
    | nil _ => nil nat
    | cons _ _ l' => l'
    end).
fun l : List nat =>
match l with
| nil _ => nil nat
| cons _ _ l' => l'
end
      : List nat -> List nat
```

When we use parameters in patterns there is an error message:

```
Coq < Fail Check
  (fun l:List nat =>
    match l with
    | nil A => nil nat
    | cons A _ l' => l'
    end).
The command has indeed failed with message:
The parameters do not bind in patterns; they must be replaced by '_'.
```

The option `Set Asymmetric Patterns` (off by default) removes parameters from constructors in patterns:

```
Coq < Set Asymmetric Patterns.
Coq < Check (fun l:List nat =>
  match l with
  | nil => nil
  | cons _ l' => l'
  end)
Unset Asymmetric Patterns.
Toplevel input, characters 111-119:
> Unset Asymmetric Patterns.
> ^^^^^^^
Error: The reference Patterns was not found in the current environment.
```

Implicit arguments in patterns By default, implicit arguments are omitted in patterns. So we write:

```
Coq < Arguments nil [A].
Coq < Arguments cons [A] _ _.
```

```

Coq < Check
  (fun l:List nat =>
    match l with
    | nil => nil
    | cons _ l' => l'
    end).
fun l : List nat => match l with
  | nil => nil
  | cons _ l' => l'
  end
: List nat -> List nat

```

But the possibility to use all the arguments is given by “@” implicit explicitations (as for terms 2.7.11).

```

Coq < Check
  (fun l:List nat =>
    match l with
    | @nil _ => @nil nat
    | @cons _ _ l' => l'
    end).
fun l : List nat => match l with
  | nil => nil
  | cons _ l' => l'
  end
: List nat -> List nat

```

17.3 Matching objects of dependent types

The previous examples illustrate pattern matching on objects of non-dependent types, but we can also use the expansion strategy to destructure objects of dependent type. Consider the type `listn` of lists of a certain length:

```

Coq < Inductive listn : nat -> Set :=
  | niln : listn 0
  | consn : forall n:nat, nat -> listn n -> listn (S n).
listn is defined
listn_rect is defined
listn_ind is defined
listn_rec is defined

```

17.3.1 Understanding dependencies in patterns

We can define the function `length` over `listn` by:

```

Coq < Definition length (n:nat) (l:listn n) := n.
length is defined

```

Just for illustrating pattern matching, we can define it by case analysis:


```

Coq < Definition length (n:nat) (l:listn n) :=
  match l with
  | niln => 0
  | consn n _ _ => S n
  end.
length is defined

```

We can understand the meaning of this definition using the same notions of usual pattern matching.

17.3.2 When the elimination predicate must be provided

Dependent pattern matching The examples given so far do not need an explicit elimination predicate because all the rhs have the same type and the strategy succeeds to synthesize it. Unfortunately when dealing with dependent patterns it often happens that we need to write cases where the type of the rhs are different instances of the elimination predicate. The function `concat` for `listn` is an example where the branches have different type and we need to provide the elimination predicate:

```

Coq < Fixpoint concat (n:nat) (l:listn n) (m:nat) (l':listn m) {struct l} :
  listn (n + m) :=
  match l in listn n return listn (n + m) with
  | niln => l'
  | consn n' a y => consn (n' + m) a (concat n' y m l')
  end.
concat is defined
concat is recursively defined (decreasing on 2nd argument)

```

The elimination predicate is `fun (n:nat) (l:listn n) => listn (n+m)`. In general if m has type $(I q_1 \dots q_r t_1 \dots t_s)$ where q_1, \dots, q_r are parameters, the elimination predicate should be of the form `fun y1 ... ys x : (I q1 ... qr y1 ... ys) => Q`.

In the concrete syntax, it should be written :

`match m as x in (I _ ... _ y1 ... ys) return Q with ... end`

The variables which appear in the `in` and `as` clause are new and bounded in the property Q in the return clause. The parameters of the inductive definitions should not be mentioned and are replaced by `_`.

Multiple dependent pattern matching Recall that a list of patterns is also a pattern. So, when we destructure several terms at the same time and the branches have different types we need to provide the elimination predicate for this multiple pattern. It is done using the same scheme, each term may be associated to an `as` and `in` clause in order to introduce a dependent product.

For example, an equivalent definition for `concat` (even though the matching on the second term is trivial) would have been:

```

Coq < Fixpoint concat (n:nat) (l:listn n) (m:nat) (l':listn m) {struct l} :
  listn (n + m) :=
  match l in listn n, l' return listn (n + m) with
  | niln, x => x
  | consn n' a y, x => consn (n' + m) a (concat n' y m x)
  end.
concat is defined
concat is recursively defined (decreasing on 2nd argument)

```

Even without real matching over the second term, this construction can be used to keep types linked. If a and b are two `listn` of the same length, by writing

```
Coq < Check (fun n (a b: listn n) => match a,b with
      | niln,b0 => tt
      | consn n' a y, bS => tt
    end).
fun (n : nat) (a _ : listn n) =>
match a with
| niln => tt
| consn n' _ _ => tt
end
      : forall n : nat, listn n -> listn n -> unit
```

I have a copy of b in type `listn 0` resp `listn (S n')`.

Patterns in in If the type of the matched term is more precise than an inductive applied to variables, arguments of the inductive in the `in` branch can be more complicated patterns than a variable.

Moreover, constructors whose type do not follow the same pattern will become impossible branches. In an impossible branch, you can answer anything but `False_rect unit` has the advantage to be subterm of anything.

To be concrete: the `tail` function can be written:

```
Coq < Definition tail n (v: listn (S n)) :=
      match v in listn (S m) return listn m with
      | niln => False_rect unit
      | consn n' a y => y
    end.
tail is defined
```

and `tail n v` will be subterm of v .

17.4 Using pattern matching to write proofs

In all the previous examples the elimination predicate does not depend on the object(s) matched. But it may depend and the typical case is when we write a proof by induction or a function that yields an object of dependent type. An example of proof using `match` in given in Section 8.2.3.

For example, we can write the function `buildlist` that given a natural number n builds a list of length n containing zeros as follows:

```
Coq < Fixpoint buildlist (n:nat) : listn n :=
      match n return listn n with
      | 0 => niln
      | S n => consn n 0 (buildlist n)
    end.
buildlist is defined
buildlist is recursively defined (decreasing on 1st argument)
```

We can also use multiple patterns. Consider the following definition of the predicate less-equal `Le`:

```
Coq < Inductive LE : nat -> nat -> Prop :=
  | LEO : forall n:nat, LE 0 n
  | LES : forall n m:nat, LE n m -> LE (S n) (S m).
LE is defined
LE_ind is defined
```

We can use multiple patterns to write the proof of the lemma `forall (n m:nat), (LE n m) \/(LE m n)`:

```
Coq < Fixpoint dec (n m:nat) {struct n} : LE n m \/(LE m n) :=
  match n, m return LE n m \/(LE m n) with
  | 0, x => or_introl (LE x 0) (LEO x)
  | x, 0 => or_intror (LE x 0) (LEO x)
  | S n as n', S m as m' =>
    match dec n m with
    | or_introl h => or_introl (LE m' n') (LES n m h)
    | or_intror h => or_intror (LE n' m') (LES m n h)
    end
  end.
dec is defined
dec is recursively defined (decreasing on 1st argument)
```

In the example of `dec`, the first match is dependent while the second is not.

The user can also use `match` in combination with the tactic `refine` (see Section 8.2.3) to build incomplete proofs beginning with a `match` construction.

17.5 Pattern-matching on inductive objects involving local definitions

If local definitions occur in the type of a constructor, then there are two ways to match on this constructor. Either the local definitions are skipped and matching is done only on the true arguments of the constructors, or the bindings for local definitions can also be caught in the matching.

Example.

```
Coq < Inductive list : nat -> Set :=
  | nil : list 0
  | cons : forall n:nat, let m := (2 * n) in list m -> list (S (S m)).
```

In the next example, the local definition is not caught.

```
Coq < Fixpoint length n (l:list n) {struct l} : nat :=
  match l with
  | nil => 0
  | cons n l0 => S (length (2 * n) l0)
  end.
length is defined
length is recursively defined (decreasing on 2nd argument)
```

But in this example, it is.

```
Coq < Fixpoint length' n (l:list n) {struct l} : nat :=
  match l with
```

```

      | nil => 0
      | @cons _ m l0 => S (length' m l0)
    end.
length' is defined
length' is recursively defined (decreasing on 2nd argument)

```

Remark: for a given matching clause, either none of the local definitions or all of them can be caught.

Remark: you can only catch `let` bindings in mode where you bind all variables and so you have to use `@` syntax.

Remark: this feature is incoherent with the fact that parameters cannot be caught and consequently is somehow hidden. For example, there is no mention of it in error messages.

17.6 Pattern-matching and coercions

If a mismatch occurs between the expected type of a pattern and its actual type, a coercion made from constructors is sought. If such a coercion can be found, it is automatically inserted around the pattern.

Example:

```

Coq < Inductive I : Set :=
      | C1 : nat -> I
      | C2 : I -> I.
I is defined
I_rect is defined
I_ind is defined
I_rec is defined

Coq < Coercion C1 : nat -> I.
C1 is now a coercion

Coq < Check (fun x => match x with
                      | C2 0 => 0
                      | _ => 0
                      end).

fun x : I =>
match x with
| C1 _ => 0
| C2 (C1 0) => 0
| C2 (C1 (S _)) => 0
| C2 (C2 _) => 0
end
      : I -> nat

```

17.7 When does the expansion strategy fail ?

The strategy works very like in ML languages when treating patterns of non-dependent type. But there are new cases of failure that are due to the presence of dependencies.

The error messages of the current implementation may be sometimes confusing. When the tactic fails because patterns are somehow incorrect then error messages refer to the initial expression. But the strategy may succeed to build an expression whose sub-expressions are well typed when the whole

expression is not. In this situation the message makes reference to the expanded expression. We encourage users, when they have patterns with the same outer constructor in different equations, to name the variable patterns in the same positions with the same name. E.g. to write $(\text{cons } n \ 0 \ x) \Rightarrow e1$ and $(\text{cons } n \ _ \ x) \Rightarrow e2$ instead of $(\text{cons } n \ 0 \ x) \Rightarrow e1$ and $(\text{cons } n' \ _ \ x') \Rightarrow e2$. This helps to maintain certain name correspondence between the generated expression and the original.

Here is a summary of the error messages corresponding to each situation:

Error messages:

1. The constructor *ident* expects *num* arguments

The variable *ident* is bound several times in pattern *term*

Found a constructor of inductive type *term* while a constructor of *term* is expected

Patterns are incorrect (because constructors are not applied to the correct number of the arguments, because they are not linear or they are wrongly typed).

2. Non exhaustive pattern-matching

The pattern matching is not exhaustive.

3. The elimination predicate *term* should be of arity *num* (for non dependent case) or *num* (for dependent case)

The elimination predicate provided to match has not the expected arity.

4. Unable to infer a match predicate

Either there is a type incompatibility or the problem involves dependencies

There is a type mismatch between the different branches. The user should provide an elimination predicate.

Chapter 18

Implicit Coercions

Amokrane Saïbi

18.1 General Presentation

This section describes the inheritance mechanism of COQ. In COQ with inheritance, we are not interested in adding any expressive power to our theory, but only convenience. Given a term, possibly not typable, we are interested in the problem of determining if it can be well typed modulo insertion of appropriate coercions. We allow to write:

- $f\ a$ where $f : forall\ x : A, B$ and $a : A'$ when A' can be seen in some sense as a subtype of A .
- $x : A$ when A is not a type, but can be seen in a certain sense as a type: set, group, category etc.
- $f\ a$ when f is not a function, but can be seen in a certain sense as a function: bijection, functor, any structure morphism etc.

18.2 Classes

A class with n parameters is any defined name with a type $forall\ (x_1 : A_1)..(x_n : A_n), s$ where s is a sort. Thus a class with parameters is considered as a single class and not as a family of classes. An object of a class C is any term of type $C\ t_1..t_n$. In addition to these user-classes, we have two abstract classes:

- `Sortclass`, the class of sorts; its objects are the terms whose type is a sort.
- `Funclass`, the class of functions; its objects are all the terms with a functional type, i.e. of form $forall\ x : A, B$.

Formally, the syntax of a classes is defined on Figure 18.1.

$class$	$::=$	$qualid$
		<code>Sortclass</code>
		<code>Funclass</code>

Figure 18.1: Syntax of classes

18.3 Coercions

A name f can be declared as a coercion between a source user-class C with n parameters and a target class D if one of these conditions holds:

- D is a user-class, then the type of f must have the form $forall (x_1 : A_1)..(x_n : A_n)(y : C\ x_1..x_n), D\ u_1..u_m$ where m is the number of parameters of D .
- D is `Funclass`, then the type of f must have the form $forall (x_1 : A_1)..(x_n : A_n)(y : C\ x_1..x_n)(x : A), B$.
- D is `Sortclass`, then the type of f must have the form $forall (x_1 : A_1)..(x_n : A_n)(y : C\ x_1..x_n), s$ with s a sort.

We then write $f : C \multimap D$. The restriction on the type of coercions is called *the uniform inheritance condition*. Remark that the abstract classes `Funclass` and `Sortclass` cannot be source classes.

To coerce an object $t : C\ t_1..t_n$ of C towards D , we have to apply the coercion f to it; the obtained term $f\ t_1..t_n\ t$ is then an object of D .

18.4 Identity Coercions

Identity coercions are special cases of coercions used to go around the uniform inheritance condition. Let C and D be two classes with respectively n and m parameters and $f : forall (x_1 : T_1)..(x_k : T_k)(y : C\ u_1..u_n), D\ v_1..v_m$ a function which does not verify the uniform inheritance condition. To declare f as coercion, one has first to declare a subclass C' of C :

$$C' := fun (x_1 : T_1)..(x_k : T_k) => C\ u_1..u_n$$

We then define an *identity coercion* between C' and C :

$$Id_C'_C := fun (x_1 : T_1)..(x_k : T_k)(y : C'\ x_1..x_k) => (y : C\ u_1..u_n)$$

We can now declare f as coercion from C' to D , since we can “cast” its type as $forall (x_1 : T_1)..(x_k : T_k)(y : C'\ x_1..x_k), D\ v_1..v_m$.

The identity coercions have a special status: to coerce an object $t : C'\ t_1..t_k$ of C' towards C , we does not have to insert explicitly $Id_C'_C$ since $Id_C'_C\ t_1..t_k\ t$ is convertible with t . However we “rewrite” the type of t to become an object of C ; in this case, it becomes $C\ u_1^*..u_k^*$ where each u_i^* is the result of the substitution in u_i of the variables x_j by t_j .

18.5 Inheritance Graph

Coercions form an inheritance graph with classes as nodes. We call *coercion path* an ordered list of coercions between two nodes of the graph. A class C is said to be a subclass of D if there is a coercion path in the graph from C to D ; we also say that C inherits from D . Our mechanism supports multiple inheritance since a class may inherit from several classes, contrary to simple inheritance where a class inherits from at most one class. However there must be at most one path between two classes. If this is not the case, only the *oldest* one is valid and the others are ignored. So the order of declaration of coercions is important.

We extend notations for coercions to coercion paths. For instance $[f_1; \dots; f_k] : C \multimap D$ is the coercion path composed by the coercions $f_1..f_k$. The application of a coercion path to a term consists of the successive application of its coercions.

18.6 Declaration of Coercions

18.6.1 Coercion *qualid* : $class_1 \multimap class_2$.

Declares the construction denoted by *qualid* as a coercion between $class_1$ and $class_2$.

Error messages:

1. *qualid* not declared
2. *qualid* is already a coercion
3. Funclass cannot be a source class
4. Sortclass cannot be a source class
5. *qualid* is not a function
6. Cannot find the source class of *qualid*
7. Cannot recognize $class_1$ as a source class of *qualid*
8. *qualid* does not respect the uniform inheritance condition
9. Found target class $class$ instead of $class_2$

When the coercion *qualid* is added to the inheritance graph, non valid coercion paths are ignored; they are signaled by a warning.

Warning :

1. Ambiguous paths: $[f_1^1; \dots; f_{n_1}^1] : C_1 \multimap D_1$
 \dots
 $[f_1^m; \dots; f_{n_m}^m] : C_m \multimap D_m$

Variants:

1. Local Coercion *qualid* : $class_1 \multimap class_2$.
 Declares the construction denoted by *qualid* as a coercion local to the current section.

2. `Coercion ident := term`
This defines *ident* just like `Definition ident := term`, and then declares *ident* as a coercion between its source and its target.
3. `Coercion ident := term : type`
This defines *ident* just like `Definition ident : type := term`, and then declares *ident* as a coercion between its source and its target.
4. `Local Coercion ident := term`
This defines *ident* just like `Let ident := term`, and then declares *ident* as a coercion between its source and its target.
5. Assumptions can be declared as coercions at declaration time. This extends the grammar of assumptions from Figure 1.3 as follows:

```

assumption      ::= assumption_keyword assums .

assums          ::= simple_assums
                  | ( simple_assums ) ... ( simple_assums )

simple_assums    ::= ident ... ident :[>] term

```

If the extra `>` is present before the type of some assumptions, these assumptions are declared as coercions.

6. Constructors of inductive types can be declared as coercions at definition time of the inductive type. This extends and modifies the grammar of inductive types from Figure 1.3 as follows:

```

inductive       ::= Inductive ind_body with ... with ind_body .
                  | CoInductive ind_body with ... with ind_body .

ind_body       ::= ident [binders] : term :=
                  [[|] constructor | ... | constructor]

constructor    ::= ident [binders] [:>] term

```

Especially, if the extra `>` is present in a constructor declaration, this constructor is declared as a coercion.

18.6.2 Identity Coercion `ident : class1 >-> class2`.

We check that *class₁* is a constant with a value of the form $\text{fun } (x_1 : T_1) .. (x_n : T_n) => (\text{class}_2 \ t_1 .. t_m)$ where *m* is the number of parameters of *class₂*. Then we define an identity function with the type $\text{forall } (x_1 : T_1) .. (x_n : T_n) (y : \text{class}_1 \ x_1 .. x_n), \text{class}_2 \ t_1 .. t_m$, and we declare it as an identity coercion between *class₁* and *class₂*.

Error messages:

1. *class₁* must be a transparent constant

Variants:

1. Local Identity Coercion `ident : ident1 >-> ident2.`
Idem but locally to the current section.
2. SubClass `ident := type.`
If `type` is a class `ident'` applied to some arguments then `ident` is defined and an identity coercion of name `Id_ident_ident'` is declared. Otherwise said, this is an abbreviation for
Definition `ident := type.`
followed by
Identity Coercion `Id_ident_ident' : ident >-> ident'.`
3. Local SubClass `ident := type.`
Same as before but locally to the current section.

18.7 Displaying Available Coercions

18.7.1 Print Classes.

Print the list of declared classes in the current context.

18.7.2 Print Coercions.

Print the list of declared coercions in the current context.

18.7.3 Print Graph.

Print the list of valid coercion paths in the current context.

18.7.4 Print Coercion Paths `class1 class2.`

Print the list of valid coercion paths from `class1` to `class2`.

18.8 Activating the Printing of Coercions

18.8.1 Set Printing Coercions.

This command forces all the coercions to be printed. Conversely, to skip the printing of coercions, use `Unset Printing Coercions`. By default, coercions are not printed.

18.8.2 Add Printing Coercion `qualid`.

This command forces coercion denoted by `qualid` to be printed. To skip the printing of coercion `qualid`, use `Remove Printing Coercion qualid`. By default, a coercion is never printed.

18.9 Classes as Records

We allow the definition of *Structures with Inheritance* (or classes as records) by extending the existing `Record` macro (see Section 2.1). Its new syntax is:

```
Record [>] ident [binders] : sort := [ident0] {
  ident1 [:|:>] term1 ;
  ...
  identn [:|:>] termn } .
```

The identifier *ident* is the name of the defined record and *sort* is its type. The identifier *ident₀* is the name of its constructor. The identifiers *ident₁*, ..., *ident_n* are the names of its fields and *term₁*, ..., *term_n* their respective types. The alternative `[:|:>]` is “:” or “:>”. If *ident_i :> term_i*, then *ident_i* is automatically declared as coercion from *ident* to the class of *term_i*. Remark that *ident_i* always verifies the uniform inheritance condition. If the optional “>” before *ident* is present, then *ident₀* (or the default name `Build_ident` if *ident₀* is omitted) is automatically declared as a coercion from the class of *term_n* to *ident* (this may fail if the uniform inheritance condition is not satisfied).

Remark: The keyword `Structure` is a synonym of `Record`.

18.10 Coercions and Sections

The inheritance mechanism is compatible with the section mechanism. The global classes and coercions defined inside a section are redefined after its closing, using their new value and new type. The classes and coercions which are local to the section are simply forgotten. Coercions with a local source class or a local target class, and coercions which do not verify the uniform inheritance condition any longer are also forgotten.

18.11 Coercions and Modules

From Coq version 8.3, the coercions present in a module are activated only when the module is explicitly imported. Formerly, the coercions were activated as soon as the module was required, whatever it was imported or not.

To recover the behavior of the versions of Coq prior to 8.3, use the following command:

```
Set Automatic Coercions Import.
```

To cancel the effect of the option, use instead:

```
Unset Automatic Coercions Import.
```

18.12 Examples

There are three situations:

- *f a* is ill-typed where *f* : forall *x* : *A*, *B* and *a* : *A'*. If there is a coercion path between *A'* and *A*, *f a* is transformed into *f a'* where *a'* is the result of the application of this coercion path to *a*.

We first give an example of coercion between atomic inductive types

```

Coq < Definition bool_in_nat (b:bool) := if b then 0 else 1.
bool_in_nat is defined

Coq < Coercion bool_in_nat : bool -> nat.
bool_in_nat is now a coercion

Coq < Check (0 = true).
0 = true
      : Prop

Coq < Set Printing Coercions.

Coq < Check (0 = true).
0 = bool_in_nat true
      : Prop

```

Warning: “Check true=0.” fails. This is “normal” behaviour of coercions. To validate true=0, the coercion is searched from nat to bool. There is none.

We give an example of coercion between classes with parameters.

```

Coq < Parameters
      (C : nat -> Set) (D : nat -> bool -> Set) (E : bool -> Set).
C is declared
D is declared
E is declared

Coq < Parameter f : forall n:nat, C n -> D (S n) true.
f is declared

Coq < Coercion f : C -> D.
f is now a coercion

Coq < Parameter g : forall (n:nat) (b:bool), D n b -> E b.
g is declared

Coq < Coercion g : D -> E.
g is now a coercion

Coq < Parameter c : C 0.
c is declared

Coq < Parameter T : E true -> nat.
T is declared

Coq < Check (T c).
T c
      : nat

Coq < Set Printing Coercions.

Coq < Check (T c).
T (g 1 true (f 0 c))
      : nat

```

We give now an example using identity coercions.

```

Coq < Definition D' (b:bool) := D 1 b.
D' is defined

Coq < Identity Coercion IdD'D : D' >-> D.

Coq < Print IdD'D.
IdD'D =
(fun (b : bool) (x : D' b) => x) : forall b : bool, D' b -> D 1 b
      : forall b : bool, D' b -> D 1 b
Argument scopes are [bool_scope _]
IdD'D is a coercion

Coq < Parameter d' : D' true.
d' is declared

Coq < Check (T d').
T d'
      : nat

Coq < Set Printing Coercions.

Coq < Check (T d').
T (g 1 true d')
      : nat

```

In the case of functional arguments, we use the monotonic rule of sub-typing. Approximatively, to coerce $t : \text{forall } x : A, B$ towards $\text{forall } x : A', B'$, one have to coerce A' towards A and B towards B' . An example is given below:

```

Coq < Parameters (A B : Set) (h : A -> B).
A is declared
B is declared
h is declared

Coq < Coercion h : A >-> B.
h is now a coercion

Coq < Parameter U : (A -> E true) -> nat.
U is declared

Coq < Parameter t : B -> C 0.
t is declared

Coq < Check (U t).
U (fun x : A => t x)
      : nat

Coq < Set Printing Coercions.

Coq < Check (U t).
U (fun x : A => g 1 true (f 0 (t (h x))))
      : nat

```

Remark the changes in the result following the modification of the previous example.

```

Coq < Parameter U' : (C 0 -> B) -> nat.
U' is declared

```

```

Coq < Parameter t' : E true -> A.
t' is declared

Coq < Check (U' t').
U' (fun x : C 0 => t' x)
    : nat

Coq < Set Printing Coercions.

Coq < Check (U' t').
U' (fun x : C 0 => h (t' (g 1 true (f 0 x))))
    : nat

```

- An assumption $x : A$ when A is not a type, is ill-typed. It is replaced by $x : A'$ where A' is the result of the application to A of the coercion path between the class of A and `Sortclass` if it exists. This case occurs in the abstraction $\text{fun } x : A \Rightarrow t$, universal quantification $\text{forall } x : A, B$, global variables and parameters of (co-)inductive definitions and functions. In $\text{forall } x : A, B$, such a coercion path may be applied to B also if necessary.

```

Coq < Parameter Graph : Type.
Graph is declared

Coq < Parameter Node : Graph -> Type.
Node is declared

Coq < Coercion Node : Graph -> Sortclass.
Node is now a coercion

Coq < Parameter G : Graph.
G is declared

Coq < Parameter Arrows : G -> G -> Type.
Arrows is declared

Coq < Check Arrows.
Arrows
    : G -> G -> Type

Coq < Parameter fg : G -> G.
fg is declared

Coq < Check fg.
fg
    : G -> G

Coq < Set Printing Coercions.

Coq < Check fg.
fg
    : Node G -> Node G

```

- $f a$ is ill-typed because $f : A$ is not a function. The term f is replaced by the term obtained by applying to f the coercion path between A and `Funclass` if it exists.

```

Coq < Parameter bij : Set -> Set -> Set.
bij is declared

Coq < Parameter ap : forall A B:Set, bij A B -> A -> B.

```

```

ap is declared

Coq < Coercion ap : bij >-> Funclass.
ap is now a coercion

Coq < Parameter b : bij nat nat.
b is declared

Coq < Check (b 0) .
b 0
      : nat

Coq < Set Printing Coercions.

Coq < Check (b 0) .
ap nat nat b 0
      : nat

```

Let us see the resulting graph of this session.

```

Coq < Print Graph.
[bool_in_nat] : bool >-> nat
[f] : C >-> D
[f; g] : C >-> E
[g] : D >-> E
[IdD'D] : D' >-> D
[IdD'D; g] : D' >-> E
[h] : A >-> B
[Node] : Graph >-> Sortclass
[ap] : bij >-> Funclass

```


Chapter 19

Canonical Structures

Assia Mahboubi and Enrico Tassi

This chapter explains the basics of Canonical Structure and how they can be used to overload notations and build a hierarchy of algebraic structures. The examples are taken from [103]. We invite the interested reader to refer to this paper for all the details that are omitted here for brevity. The interested reader shall also find in [76] a detailed description of another, complementary, use of Canonical Structures: advanced proof search. This latter papers also presents many techniques one can employ to tune the inference of Canonical Structures.

19.1 Notation overloading

We build an infix notation `==` for a comparison predicate. Such notation will be overloaded, and its meaning will depend on the types of the terms that are compared.

```
Coq < Module EQ.
Interactive Module EQ started

Coq <   Record class (T : Type) := Class { cmp : T -> T -> Prop }.
class is defined
cmp is defined

Coq <   Structure type := Pack { obj : Type; class_of : class obj }.
type is defined
obj is defined
class_of is defined

Coq <   Definition op (e : type) : obj e -> obj e -> Prop :=
      let 'Pack _ (Class _ the_cmp) := e in the_cmp.
op is defined

Coq <   Check op.
op
      : forall e : type, obj e -> obj e -> Prop

Coq <   Arguments op {e} x y : simpl never.

Coq <   Arguments Class {T} cmp.
```

```

Coq < Module theory.
Interactive Module theory started

Coq <      Notation "x == y" := (op x y) (at level 70).

Coq < End theory.
Module theory is defined

Coq < End EQ.
Module EQ is defined

```

We use Coq modules as name spaces. This allows us to follow the same pattern and naming convention for the rest of the chapter. The base name space contains the definitions of the algebraic structure. To keep the example small, the algebraic structure `EQ.type` we are defining is very simplistic, and characterizes terms on which a binary relation is defined, without requiring such relation to validate any property. The inner `theory` module contains the overloaded notation `==` and will eventually contain lemmas holding on all the instances of the algebraic structure (in this case there are no lemmas).

Note that in practice the user may want to declare `EQ.obj` as a coercion, but we will not do that here.

The following line tests that, when we assume a type `e` that is in the `EQ` class, then we can relate two of its objects with `==`.

```

Coq < Import EQ.theory.

Coq < Check forall (e : EQ.type) (a b : EQ.obj e), a == b.
forall (e : EQ.type) (a b : EQ.obj e), a == b
      : Prop

```

Still, no concrete type is in the `EQ` class. We amend that by equipping `nat` with a comparison relation.

```

Coq < Fail Check 3 == 3.
The command has indeed failed with message:
The term "3" has type "nat" while it is expected to have type
"EQ.obj ?e".

Coq < Definition nat_eq (x y : nat) := nat_compare x y = Eq.
nat_eq is defined

Coq < Definition nat_EQcl : EQ.class nat := EQ.Class nat_eq.
nat_EQcl is defined

Coq < Canonical Structure nat_EQty : EQ.type := EQ.Pack nat nat_EQcl.
nat_EQty is defined

Coq < Check 3 == 3.
3 == 3
      : Prop

Coq < Eval compute in 3 == 4.
= Lt = Eq
      : Prop

```

This last test shows that Coq is now not only able to typecheck `3==3`, but also that the infix relation was bound to the `nat_eq` relation. This relation is selected whenever `==` is used on terms of type `nat`. This can be read in the line declaring the canonical structure `nat_EQty`, where the first argument to

`Pack` is the key and its second argument a group of canonical values associated to the key. In this case we associate to `nat` only one canonical value (since its class, `nat_EQcl` has just one member). The use of the projection `op` requires its argument to be in the class `EQ`, and uses such a member (function) to actually compare its arguments.

Similarly, we could equip any other type with a comparison relation, and use the `==` notation on terms of this type.

19.1.1 Derived Canonical Structures

We know how to use `==` on base types, like `nat`, `bool`, `Z`. Here we show how to deal with type constructors, i.e. how to make the following example work:

```
Coq < Fail Check forall (e : EQ.type) (a b : EQ.obj e), (a,b) == (a,b).
The command has indeed failed with message:
In environment
e : EQ.type
a : EQ.obj e
b : EQ.obj e
The term "(a, b)" has type "(EQ.obj e * EQ.obj e)%type"
while it is expected to have type "EQ.obj ?e".
```

The error message is telling that Coq has no idea on how to compare pairs of objects. The following construction is telling Coq exactly how to do that.

```
Coq < Definition pair_eq (e1 e2 : EQ.type) (x y : EQ.obj e1 * EQ.obj e2) :=
  fst x == fst y /\ snd x == snd y.
pair_eq is defined

Coq < Definition pair_EQcl e1 e2 := EQ.Class (pair_eq e1 e2).
pair_EQcl is defined

Coq < Canonical Structure pair_EQty (e1 e2 : EQ.type) : EQ.type :=
  EQ.Pack (EQ.obj e1 * EQ.obj e2) (pair_EQcl e1 e2).
pair_EQty is defined

Coq < Check forall (e : EQ.type) (a b : EQ.obj e), (a,b) == (a,b).
forall (e : EQ.type) (a b : EQ.obj e), (a, b) == (a, b)
: Prop

Coq < Check forall n m : nat, (3,4) == (n,m).
forall n m : nat, (3, 4) == (n, m)
: Prop
```

Thanks to the `pair_EQty` declaration, Coq is able to build a comparison relation for pairs whenever it is able to build a comparison relation for each component of the pair. The declaration associates to the key `*` (the type constructor of pairs) the canonical comparison relation `pair_eq` whenever the type constructor `*` is applied to two types being themselves in the `EQ` class.

19.2 Hierarchy of structures

To get to an interesting example we need another base class to be available. We choose the class of types that are equipped with an order relation, to which we associate the infix `<=` notation.

```

Coq < Module LE.
Interactive Module LE started

Coq <   Record class T := Class { cmp : T -> T -> Prop }.
class is defined
cmp is defined

Coq <   Structure type := Pack { obj : Type; class_of : class obj }.
type is defined
obj is defined
class_of is defined

Coq <   Definition op (e : type) : obj e -> obj e -> Prop :=
      let 'Pack _ (Class _ f) := e in f.
op is defined

Coq <   Arguments op {_} x y : simpl never.

Coq <   Arguments Class {T} cmp.

Coq <   Module theory.
Interactive Module theory started

Coq <   Notation "x <= y" := (op x y) (at level 70).

Coq <   End theory.
Module theory is defined

Coq < End LE.
Module LE is defined

```

As before we register a canonical LE class for nat.

```

Coq < Import LE.theory.

Coq < Definition nat_le x y := nat_compare x y <> Gt.
nat_le is defined

Coq < Definition nat_LEcl : LE.class nat := LE.Class nat_le.
nat_LEcl is defined

Coq < Canonical Structure nat_LEty : LE.type := LE.Pack nat nat_LEcl.
nat_LEty is defined

```

And we enable Coq to relate pair of terms with <=.

```

Coq < Definition pair_le e1 e2 (x y : LE.obj e1 * LE.obj e2) :=
      fst x <= fst y /\ snd x <= snd y.
pair_le is defined

Coq < Definition pair_LEcl e1 e2 := LE.Class (pair_le e1 e2).
pair_LEcl is defined

Coq < Canonical Structure pair_LEty (e1 e2 : LE.type) : LE.type :=
      LE.Pack (LE.obj e1 * LE.obj e2) (pair_LEcl e1 e2).
pair_LEty is defined

Coq < Check (3,4,5) <= (3,4,5).
(3, 4, 5) <= (3, 4, 5)
      : Prop

```

At the current stage we can use `==` and `<=` on concrete types, like tuples of natural numbers, but we can't develop an algebraic theory over the types that are equipped with both relations.

```
Coq < Check 2 <= 3 /\ 2 == 2.
2 <= 3 /\ 2 == 2
      : Prop
```

```
Coq < Fail Check forall (e : EQ.type) (x y : EQ.obj e), x <= y -> y <= x -> x == y.
The command has indeed failed with message:
In environment
e : EQ.type
x : EQ.obj e
y : EQ.obj e
The term "x" has type "EQ.obj e" while it is expected to have type
"LE.obj ?e".
```

```
Coq < Fail Check forall (e : LE.type) (x y : LE.obj e), x <= y -> y <= x -> x == y.
The command has indeed failed with message:
In environment
e : LE.type
x : LE.obj e
y : LE.obj e
The term "x" has type "LE.obj e" while it is expected to have type
"EQ.obj ?e".
```

We need to define a new class that inherits from both `EQ` and `LE`.

```
Coq < Module LEQ.
Interactive Module LEQ started

Coq <   Record mixin (e : EQ.type) (le : EQ.obj e -> EQ.obj e -> Prop) :=
      Mixin { compat : forall x y : EQ.obj e, le x y /\ le y x <-> x == y }.
mixin is defined
compat is defined

Coq <   Record class T := Class {
      EQ_class      : EQ.class T;
      LE_class      : LE.class T;
      extra : mixin (EQ.Pack T EQ_class) (LE.cmp T LE_class) }.
class is defined
EQ_class is defined
LE_class is defined
extra is defined

Coq <   Structure type := _Pack { obj : Type; class_of : class obj }.
type is defined
obj is defined
class_of is defined

Coq <   Arguments Mixin {e le} _.
Coq <   Arguments Class {T} _ _ _.
```

The `mixin` component of the `LEQ` class contains all the extra content we are adding to `EQ` and `LE`. In particular it contains the requirement that the two relations we are combining are compatible.

Unfortunately there is still an obstacle to developing the algebraic theory of this new class.

```

Coq < Module theory.
Interactive Module theory started

Coq < Fail Check forall (le : type) (n m : obj le), n <= m -> n <= m -> n == m.
The command has indeed failed with message:
In environment
le : type
n : obj le
m : obj le
The term "n" has type "obj le" while it is expected to have type
"LE.obj ?e".

```

The problem is that the two classes `LE` and `LEQ` are not yet related by a subclass relation. In other words Coq does not see that an object of the `LEQ` class is also an object of the `LE` class.

The following two constructions tell Coq how to canonically build the `LE.type` and `EQ.type` structure given an `LEQ.type` structure on the same type.

```

Coq < Definition to_EQ (e : type) : EQ.type :=
      EQ.Pack (obj e) (EQ_class _ (class_of e)).
to_EQ is defined

Coq < Canonical Structure to_EQ.

Coq < Definition to_LE (e : type) : LE.type :=
      LE.Pack (obj e) (LE_class _ (class_of e)).
to_LE is defined

Coq < Canonical Structure to_LE.

```

We can now formulate our first theorem on the objects of the `LEQ` structure.

```

Coq < Lemma lele_eq (e : type) (x y : obj e) : x <= y -> y <= x -> x == y.
1 subgoal

  e : type
  x, y : obj e
  =====
  x <= y -> y <= x -> x == y

Coq < now intros; apply (compat _ _ (extra _ (class_of e)) x y); split. Qed.
No more subgoals.
lele_eq is defined

Coq < Arguments lele_eq {e} x y _ .

Coq < End theory.
Module theory is defined

Coq < End LEQ.
Module LEQ is defined

Coq < Import LEQ.theory.

Coq < Check lele_eq.
lele_eq
      : forall x y : LEQ.obj ?e, x <= y -> y <= x -> x == y
where
?e : [ |- LEQ.type]

```

Of course one would like to apply results proved in the algebraic setting to any concrete instance of the algebraic structure.

```
Coq < Example test_algebraic (n m : nat) : n <= m -> m <= n -> n == m.
1 subgoal

  n, m : nat
  =====
  n <= m -> m <= n -> n == m

Coq < Fail apply (lele_eq n m). Abort.
The command has indeed failed with message:
In environment
n, m : nat
The term "n" has type "nat" while it is expected to have type
"LEQ.obj ?e".
1 subgoal

  n, m : nat
  =====
  n <= m -> m <= n -> n == m

Coq < Example test_algebraic2 (l1 l2 : LEQ.type) (n m : LEQ.obj l1 * LEQ.obj l2) :
      n <= m -> m <= n -> n == m.
1 subgoal

  l1, l2 : LEQ.type
  n, m : LEQ.obj l1 * LEQ.obj l2
  =====
  n <= m -> m <= n -> n == m

Coq < Fail apply (lele_eq n m). Abort.
The command has indeed failed with message:
In environment
l1, l2 : LEQ.type
n, m : LEQ.obj l1 * LEQ.obj l2
The term "n" has type "(LEQ.obj l1 * LEQ.obj l2)%type"
while it is expected to have type "LEQ.obj ?e".
1 subgoal

  l1, l2 : LEQ.type
  n, m : LEQ.obj l1 * LEQ.obj l2
  =====
  n <= m -> m <= n -> n == m
```

Again one has to tell Coq that the type `nat` is in the `LEQ` class, and how the type constructor `*` interacts with the `LEQ` class. In the following proofs are omitted for brevity.

```
Coq < Lemma nat_LEQ_compat (n m : nat) : n <= m /\ m <= n <-> n == m.
1 subgoal

  n, m : nat
  =====
  n <= m /\ m <= n <-> n == m
```

```

Coq < Definition nat_LEQmx := LEQ.Mixin nat_LEQ_compat.
nat_LEQmx is defined

Coq < Lemma pair_LEQ_compat (l1 l2 : LEQ.type) (n m : LEQ.obj l1 * LEQ.obj l2) :
  n <= m /\ m <= n <-> n == m.
1 subgoal

  l1, l2 : LEQ.type
  n, m : LEQ.obj l1 * LEQ.obj l2
  =====
  n <= m /\ m <= n <-> n == m

Coq < Definition pair_LEQmx l1 l2 := LEQ.Mixin (pair_LEQ_compat l1 l2).
pair_LEQmx is defined

```

The following script registers an LEQ class for nat and for the type constructor *. It also tests that they work as expected.

Unfortunately, these declarations are very verbose. In the following subsection we show how to make these declaration more compact.

```

Coq < Module Add_instance_attempt.
Interactive Module Add_instance_attempt started

Coq < Canonical Structure nat_LEQty : LEQ.type :=
  LEQ._Pack nat (LEQ.Class nat_EQcl nat_LEcl nat_LEQmx).
nat_LEQty is defined

Coq < Canonical Structure pair_LEQty (l1 l2 : LEQ.type) : LEQ.type :=
  LEQ._Pack (LEQ.obj l1 * LEQ.obj l2)
    (LEQ.Class
      (EQ.class_of (pair_EQty (to_EQ l1) (to_EQ l2)))
      (LE.class_of (pair_LEty (to_LE l1) (to_LE l2)))
      (pair_LEQmx l1 l2)).
pair_LEQty is defined
Toplevel input, characters 2-263:
> Canonical Structure pair_LEQty (l1 l2 : LEQ.type) : LEQ.type :=
>   LEQ._Pack (LEQ.obj l1 * LEQ.obj l2)
>   (LEQ.Class
>     (EQ.class_of (pair_EQty (to_EQ l1) (to_EQ l2)))
>     (LE.class_of (pair_LEty (to_LE l1) (to_LE l2)))
>     (pair_LEQmx l1 l2)).
Warning: Ignoring canonical projection to LEQ.Class by LEQ.class_of in
pair_LEQty: redundant with nat_LEQty
[redundant-canonical-projection,typechecker]

Coq < Example test_algebraic (n m : nat) : n <= m -> m <= n -> n == m.
1 subgoal

  n, m : nat
  =====
  n <= m -> m <= n -> n == m

Coq < now apply (lele_eq n m). Qed.
No more subgoals.
test_algebraic is defined

```



```

Coq < Example test_algebraic2 (n m : nat * nat) : n <= m -> m <= n -> n == m.
1 subgoal

  n, m : nat * nat
  =====
  n <= m -> m <= n -> n == m

Coq < now apply (lele_eq n m). Qed.
No more subgoals.
test_algebraic2 is defined

Coq < End Add_instance_attempt.
Module Add_instance_attempt is defined

```

Note that no direct proof of $n \leq m \rightarrow m \leq n \rightarrow n = m$ is provided by the user for n and m of type `nat * nat`. What the user provides is a proof of this statement for n and m of type `nat` and a proof that the pair constructor preserves this property. The combination of these two facts is a simple form of proof search that Coq performs automatically while inferring canonical structures.

19.2.1 Compact declaration of Canonical Structures

We need some infrastructure for that.

```

Coq < Require Import Strings.String.

Coq < Module infrastructure.
Interactive Module infrastructure started

Coq < Inductive phantom {T : Type} (t : T) : Type := Phantom.
phantom is defined
phantom_rect is defined
phantom_ind is defined
phantom_rec is defined

Coq < Definition unify {T1 T2} (t1 : T1) (t2 : T2) (s : option string) :=
  phantom t1 -> phantom t2.
unify is defined

Coq < Definition id {T} {t : T} (x : phantom t) := x.
id is defined

Coq < Notation "[find v | t1 ~ t2 ] p" := (fun v (_ : unify t1 t2 None) => p)
  (at level 50, v ident, only parsing).

Coq < Notation "[find v | t1 ~ t2 | s ] p" := (fun v (_ : unify t1 t2 (Some s)) => p)
  (at level 50, v ident, only parsing).

Coq < Notation "'Error : t : s'" := (unify _ t (Some s))
  (at level 50, format "'Error' : t : s").

Coq < Open Scope string_scope.

Coq < End infrastructure.
Module infrastructure is defined

```

To explain the notation `[find v | t1 ~ t2]` let us pick one of its instances: `[find e | EQ.obj e ~ T | "is not an EQ.type"]`. It should be read as: “find a class `e` such that its objects have type `T` or fail with message “`T is not an EQ.type`””.

The other utilities are used to ask Coq to solve a specific unification problem, that will in turn require the inference of some canonical structures. They are explained in more details in [103].

We now have all we need to create a compact “packager” to declare instances of the LEQ class.

```
Coq < Import infrastructure.

Coq < Definition packager T e0 le0 (m0 : LEQ.mixin e0 le0) :=
  [find e   | EQ.obj e ~ T           | "is not an EQ.type" ]
  [find o   | LE.obj o ~ T           | "is not an LE.type" ]
  [find ce  | EQ.class_of e ~ ce     |
  [find co  | LE.class_of o ~ co     |
  [find m   | m ~ m0                 | "is not the right mixin" ]
  LEQ._Pack T (LEQ.Class ce co m).
packager is defined

Coq < Notation Pack T m := (packager T _ _ m _ id _ id _ id _ id).
```

The object Pack takes a type T (the key) and a mixin m. It infers all the other pieces of the class LEQ and declares them as canonical values associated to the T key. All in all, the only new piece of information we add in the LEQ class is the mixin, all the rest is already canonical for T and hence can be inferred by Coq.

Pack is a notation, hence it is not type checked at the time of its declaration. It will be type checked when it is used, and in that case T is going to be a concrete type. The odd arguments _ and id we pass to the packager represent respectively the classes to be inferred (like e, o, etc) and a token (id) to force their inference. Again, for all the details the reader can refer to [103].

The declaration of canonical instances can now be way more compact:

```
Coq < Canonical Structure nat_LEQty := Eval hnf in Pack nat nat_LEQmx.
nat_LEQty is defined

Coq < Canonical Structure pair_LEQty (l1 l2 : LEQ.type) :=
  Eval hnf in Pack (LEQ.obj l1 * LEQ.obj l2) (pair_LEQmx l1 l2).
pair_LEQty is defined
Toplevel input, characters 0-117:
> Canonical Structure pair_LEQty (l1 l2 : LEQ.type) :=
>   Eval hnf in Pack (LEQ.obj l1 * LEQ.obj l2) (pair_LEQmx l1 l2).
Warning: Ignoring canonical projection to LEQ.Class by LEQ.class_of in
pair_LEQty: redundant with nat_LEQty
[redundant-canonical-projection,typechecker]
```

Error messages are also quite intelligible (if one skips to the end of the message).

```
Coq < Fail Canonical Structure err := Eval hnf in Pack bool nat_LEQmx.
The command has indeed failed with message:
The term "id" has type "phantom (EQ.obj ?e) -> phantom (EQ.obj ?e)"
while it is expected to have type
  "'Error : bool : "is not an EQ.type"".
```

Chapter 20

Type Classes

Matthieu Sozeau

This chapter presents a quick reference of the commands related to type classes. For an actual introduction to type classes, there is a description of the system [136] and the literature on type classes in HASKELL which also applies.

20.1 Class and Instance declarations

The syntax for class and instance declarations is the same as record syntax of COQ:

$$\text{Class Id } (\alpha_1 : \tau_1) \cdots (\alpha_n : \tau_n) [: \text{sort}] := \{$$
$$\begin{array}{l} \mathbf{f}_1 \quad : \quad \text{type}_1; \\ \vdots \\ \mathbf{f}_m \quad : \quad \text{type}_m \}. \end{array}$$
$$\text{Instance } \textit{ident} : \text{Id } \textit{term}_1 \cdots \textit{term}_n := \{$$
$$\begin{array}{l} \mathbf{f}_1 \quad := \quad \textit{term}_{f_1}; \\ \vdots \\ \mathbf{f}_m \quad := \quad \textit{term}_{f_m} \}. \end{array}$$

The $\overrightarrow{\alpha_i : \tau_i}$ variables are called the *parameters* of the class and the $\overrightarrow{f_k : \text{type}_k}$ are called the *methods*. Each class definition gives rise to a corresponding record declaration and each instance is a regular definition whose name is given by *ident* and type is an instantiation of the record type.

We'll use the following example class in the rest of the chapter:

```
Coq < Class EqDec (A : Type) := {
  eqb : A -> A -> bool ;
  eqb_leibniz : forall x y, eqb x y = true -> x = y }.
```

This class implements a boolean equality test which is compatible with Leibniz equality on some type. An example implementation is:

```
Coq < Instance unit_EqDec : EqDec unit :=
  { eqb x y := true ;
    eqb_leibniz x y H :=
      match x, y return x = y with tt, tt => eq_refl tt end }.
```

If one does not give all the members in the Instance declaration, Coq enters the proof-mode and the user is asked to build inhabitants of the remaining fields, e.g.:

```
Coq < Instance eq_bool : EqDec bool :=
  { eqb x y := if x then y else negb y }.

Coq < Proof. intros x y H.
1 subgoal

=====
forall x y : bool, (if x then y else negb y) = true -> x = y
1 subgoal

x, y : bool
H : (if x then y else negb y) = true
=====
x = y

Coq < destruct x ; destruct y ; (discriminate || reflexivity).
No more subgoals.

Coq < Defined.
eq_bool is defined
```

One has to take care that the transparency of every field is determined by the transparency of the Instance proof. One can use alternatively the Program Instance variant which has richer facilities for dealing with obligations.

20.2 Binding classes

Once a type class is declared, one can use it in class binders:

```
Coq < Definition negb {A} {eqa : EqDec A} (x y : A) := negb (eqb x y).
negb is defined
```

When one calls a class method, a constraint is generated that is satisfied only in contexts where the appropriate instances can be found. In the example above, a constraint `EqDec A` is generated and satisfied by `eqa : EqDec A`. In case no satisfying constraint can be found, an error is raised:

```
Coq < Fail Definition negb' (A : Type) (x y : A) := negb (eqb x y).
The command has indeed failed with message:
Unable to satisfy the following constraints:
In environment:
A : Type
x, y : A
?EqDec : "EqDec A"
```

The algorithm used to solve constraints is a variant of the `eauto` tactic that does proof search with a set of lemmas (the instances). It will use local hypotheses as well as declared lemmas in the `typeclass_instances` database. Hence the example can also be written:

```
Coq < Definition negb' A (eqa : EqDec A) (x y : A) := negb (eqb x y).
negb' is defined
```

However, the generalizing binders should be used instead as they have particular support for type classes:

- They automatically set the maximally implicit status for type class arguments, making derived functions as easy to use as class methods. In the example above, `A` and `eqa` should be set maximally implicit.
- They support implicit quantification on partially applied type classes (§2.7.19). Any argument not given as part of a type class binder will be automatically generalized.
- They also support implicit quantification on superclasses (§20.5.1)

Following the previous example, one can write:

```
Coq < Definition negb_impl `{eqa : EqDec A} (x y : A) := negb (eqb x y).
negb_impl is defined
```

Here `A` is implicitly generalized, and the resulting function is equivalent to the one above.

20.3 Parameterized Instances

One can declare parameterized instances as in HASKELL simply by giving the constraints as a binding context before the instance, e.g.:

```
Coq < Instance prod_eqb `(EA : EqDec A, EB : EqDec B) : EqDec (A * B) :=
  { eqb x y := match x, y with
    | (la, ra), (lb, rb) => andb (eqb la lb) (eqb ra rb)
  end }.
```

These instances are used just as well as lemmas in the instance hint database.

20.4 Sections and contexts

To ease the parametrization of developments by type classes, we provide a new way to introduce variables into section contexts, compatible with the implicit argument mechanism. The new command works similarly to the `Variables vernacular` (see 1.3.1), except it accepts any binding context as argument. For example:

```
Coq < Section EqDec_defs.
Coq <   Context `{EA : EqDec A}.
A is declared
EA is declared
```

```
Coq < Global Instance option_eqb : EqDec (option A) :=
  { eqb x y := match x, y with
    | Some x, Some y => eqb x y
    | None, None => true
    | _, _ => false
  end }.
```

```
Coq < End EqDec_defs.
```

```
Coq < About option_eqb.
option_eqb : forall A : Type, EqDec A -> EqDec (option A)
Arguments A, EA are implicit and maximally inserted
Argument scopes are [type_scope _]
option_eqb is transparent
Expands to: Constant Top.option_eqb
```

Here the `Global` modifier redeclares the instance at the end of the section, once it has been generalized by the context variables it uses.

20.5 Building hierarchies

20.5.1 Superclasses

One can also parameterize classes by other classes, generating a hierarchy of classes and superclasses. In the same way, we give the superclasses as a binding context:

```
Coq < Class Ord `(E : EqDec A) :=
  { le : A -> A -> bool }.
```

Contrary to `HASKELL`, we have no special syntax for superclasses, but this declaration is morally equivalent to:

```
Class `(E : EqDec A) => Ord A :=
  { le : A -> A -> bool }.
```

This declaration means that any instance of the `Ord` class must have an instance of `EqDec`. The parameters of the subclass contain at least all the parameters of its superclasses in their order of appearance (here `A` is the only one). As we have seen, `Ord` is encoded as a record type with two parameters: a type `A` and an `E` of type `EqDec A`. However, one can still use it as if it had a single parameter inside generalizing binders: the generalization of superclasses will be done automatically.

```
Coq < Definition le_eqb `{Ord A} (x y : A) := andb (le x y) (le y x).
```

In some cases, to be able to specify sharing of structures, one may want to give explicitly the superclasses. It is possible to do it directly in regular binders, and using the `!` modifier in class binders. For example:

```
Coq < Definition lt `{eqa : EqDec A, ! Ord eqa} (x y : A) :=
  andb (le x y) (neqb x y).
```

The `!` modifier switches the way a binder is parsed back to the regular interpretation of `Coq`. In particular, it uses the implicit arguments mechanism if available, as shown in the example.

20.5.2 Substructures

Substructures are components of a class which are instances of a class themselves. They often arise when using classes for logical properties, e.g.:

```
Coq < Class Reflexive (A : Type) (R : relation A) :=
  reflexivity : forall x, R x x.

Coq < Class Transitive (A : Type) (R : relation A) :=
  transitivity : forall x y z, R x y -> R y z -> R x z.
```

This declares singleton classes for reflexive and transitive relations, (see [1](#) for an explanation). These may be used as part of other classes:

```
Coq < Class PreOrder (A : Type) (R : relation A) :=
  { PreOrder_Reflexive :> Reflexive A R ;
    PreOrder_Transitive :> Transitive A R }.
```

The syntax `:>` indicates that each `PreOrder` can be seen as a `Reflexive` relation. So each time a reflexive relation is needed, a preorder can be used instead. This is very similar to the coercion mechanism of `Structure` declarations. The implementation simply declares each projection as an instance.

One can also declare existing objects or structure projections using the `Existing Instance` command to achieve the same effect.

20.6 Summary of the commands

20.6.1 `Class ident binder1 ... bindern : sort := { field1 ; ... ; fieldk }.`

The `Class` command is used to declare a type class with parameters `binder1` to `bindern` and fields `field1` to `fieldk`.

Variants:

1. `Class ident binder1 ... bindern : sort := ident1 : type1.` This variant declares a *singleton* class whose only method is `ident1`. This singleton class is a so-called definitional class, represented simply as a definition `ident binder1 ... bindern := type1` and whose instances are themselves objects of this type. Definitional classes are not wrapped inside records, and the trivial projection of an instance of such a class is convertible to the instance itself. This can be useful to make instances of existing objects easily and to reduce proof size by not inserting useless projections. The class constant itself is declared rigid during resolution so that the class abstraction is maintained.
2. `Existing Class ident.` This variant declares a class a posteriori from a constant or inductive definition. No methods or instances are defined.

20.6.2 Instance *ident* *binder*₁ ... *binder*_{*n*} : Class *t*₁ ... *t*_{*n*} [| *priority*] := { field₁ := *b*₁ ; ... ; field_{*i*} := *b*_{*i*} }

The Instance command is used to declare a type class instance named *ident* of the class *Class* with parameters *t*₁ to *t*_{*n*} and fields *b*₁ to *b*_{*i*}, where each field must be a declared field of the class. Missing fields must be filled in interactive proof mode.

An arbitrary context of the form *binder*₁ ... *binder*_{*n*} can be put after the name of the instance and before the colon to declare a parameterized instance. An optional *priority* can be declared, 0 being the highest priority as for auto hints. If the priority is not specified, it defaults to *n*, the number of binders of the instance.

Variants:

1. Instance *ident* *binder*₁ ... *binder*_{*n*} : forall *binder*_{*n*+1} ... *binder*_{*m*}, Class *t*₁ ... *t*_{*n*} [| *priority*] := *term* This syntax is used for declaration of singleton class instances or for directly giving an explicit term of type forall *binder*_{*n*+1} ... *binder*_{*m*}, Class *t*₁ ... *t*_{*n*}. One need not even mention the unique field name for singleton classes.
2. Global Instance One can use the Global modifier on instances declared in a section so that their generalization is automatically redeclared after the section is closed.
3. Program Instance Switches the type-checking to PROGRAM (chapter 24) and uses the obligation mechanism to manage missing fields.
4. Declare Instance In a Module Type, this command states that a corresponding concrete instance should exist in any implementation of this Module Type. This is similar to the distinction between Parameter vs. Definition, or between Declare Module and Module.

Besides the Class and Instance vernacular commands, there are a few other commands related to type classes.

20.6.3 Existing Instance *ident* [| *priority*]

This commands adds an arbitrary constant whose type ends with an applied type class to the instance database with an optional priority. It can be used for redeclaring instances at the end of sections, or declaring structure projections as instances. This is almost equivalent to Hint Resolve *ident* : typeclass_instances.

Variants:

1. Existing Instances *ident*₁ ... *ident*_{*n*} [| *priority*] With this command, several existing instances can be declared at once.

20.6.4 Context *binder*₁ ... *binder*_{*n*}

Declares variables according to the given binding context, which might use implicit generalization (see 20.4).

20.6.5 `typeclasses eauto`

The `typeclasses eauto` tactic uses a different resolution engine than `eauto` and `auto`. The main differences are the following:

- Contrary to `eauto` and `auto`, the resolution is done entirely in the new proof engine (as of Coq v8.6), meaning that backtracking is available among dependent subgoals, and shelving goals is supported. `typeclasses eauto` is a multi-goal tactic. It analyses the dependencies between subgoals to avoid backtracking on subgoals that are entirely independent.
- When called with no arguments, `typeclasses eauto` uses the `typeclass_instances` database by default (instead of `core`). Dependent subgoals are automatically shelved, and shelved goals can remain after resolution ends (following the behavior of COQ 8.5).

Note: As of Coq 8.6, `all:once (typeclasses eauto)` faithfully mimicks what happens during typeclass resolution when it is called during refinement/type-inference, except that *only* declared class subgoals are considered at the start of resolution during type inference, while “all” can select non-class subgoals as well. It might move to `all:typeclasses eauto` in future versions when the refinement engine will be able to backtrack.

- When called with specific databases (e.g. `with`), `typeclasses eauto` allows shelved goals to remain at any point during search and treat typeclasses goals like any other.
- The transparency information of databases is used consistently for all hints declared in them. It is always used when calling the unifier. When considering the local hypotheses, we use the transparent state of the first hint database given. Using an empty database (created with `Create HintDb` for example) with unfoldable variables and constants as the first argument of `typeclasses eauto` hence makes resolution with the local hypotheses use full conversion during unification.

Variants:

1. `typeclasses eauto [num]` *Warning:* The semantics for the limit `num` is different than for `auto`. By default, if no limit is given the search is unbounded. Contrary to `auto`, introduction steps (`intro`) are counted, which might result in larger limits being necessary when searching with `typeclasses eauto` than `auto`.
2. `typeclasses eauto with ident1 ... identn`. This variant runs resolution with the given hint databases. It treats typeclass subgoals the same as other subgoals (no shelving of non-typeclass goals in particular).

20.6.6 `autoapply term with ident`

The tactic `autoapply` applies a term using the transparency information of the hint database `ident`, and does *no* typeclass resolution. This can be used in `Hint Extern's` for typeclass instances (in `hint db typeclass_instances`) to allow backtracking on the typeclass subgoals created by the lemma application, rather than doing type class resolution locally at the hint application time.

20.6.7 `Typeclasses Transparent, Opaque ident1 ... identn`

This commands defines the transparency of `ident1 ... identn` during type class resolution. It is useful when some constants prevent some unifications and make resolution fail. It is also useful to declare

constants which should never be unfolded during proof-search, like fixpoints or anything which does not look like an abbreviation. This can additionally speed up proof search as the typeclass map can be indexed by such rigid constants (see 8.9.1). By default, all constants and local variables are considered transparent. One should take care not to make opaque any constant that is used to abbreviate a type, like relation `A := A -> A -> Prop`.

This is equivalent to `Hint Transparent, Opaque ident : typeclass_instances`.

20.6.8 Set Typeclasses Dependency Order

This option (on by default since 8.6) respects the dependency order between subgoals, meaning that subgoals which are depended on by other subgoals come first, while the non-dependent subgoals were put before the dependent ones previously (Coq v8.5 and below). This can result in quite different performance behaviors of proof search.

20.6.9 Set Typeclasses Filtered Unification

This option, available since Coq 8.6 and off by default, switches the hint application procedure to a filter-then-unify strategy. To apply a hint, we first check that the goal *matches* syntactically the inferred or specified pattern of the hint, and only then try to *unify* the goal with the conclusion of the hint. This can drastically improve performance by calling unification less often, matching syntactic patterns being very quick. This also provides more control on the triggering of instances. For example, forcing a constant to explicitly appear in the pattern will make it never apply on a goal where there is a hole in that place.

20.6.10 Set Typeclasses Legacy Resolution

Deprecated since 8.7

This option (off by default) uses the 8.5 implementation of resolution. Use for compatibility purposes only (porting and debugging).

20.6.11 Set Typeclasses Module Eta

Deprecated since 8.7

This option allows eta-conversion for functions and records during unification of type-classes. This option is unsupported since 8.6 with `Typeclasses Filtered Unification` set, but still affects the default unification strategy, and the one used in `Legacy Resolution` mode. It is *unset* by default. If `Typeclasses Filtered Unification` is set, this has no effect and unification will find solutions up-to eta conversion. Note however that syntactic pattern-matching is not up-to eta.

20.6.12 Set Typeclasses Limit Intros

This option (on by default) controls the ability to apply hints while avoiding (functional) eta-expansions in the generated proof term. It does so by allowing hints that conclude in a product to apply to a goal with a matching product directly, avoiding an introduction. *Warning:* this can be expensive as it requires rebuilding hint clauses dynamically, and does not benefit from the invertibility status of the product introduction rule, resulting in potentially more expensive proof-search (i.e. more useless backtracking).

20.6.13 `Set Typeclass Resolution After Apply`*Deprecated since 8.6*

This option (off by default in Coq 8.6 and 8.5) controls the resolution of typeclass subgoals generated by the `apply` tactic.

20.6.14 `Set Typeclass Resolution For Conversion`

This option (on by default) controls the use of typeclass resolution when a unification problem cannot be solved during elaboration/type-inference. With this option on, when a unification fails, typeclass resolution is tried before launching unification once again.

20.6.15 `Set Typeclasses Strict Resolution`

Typeclass declarations introduced when this option is set have a stricter resolution behavior (the option is off by default). When looking for unifications of a goal with an instance of this class, we “freeze” all the existentials appearing in the goals, meaning that they are considered rigid during unification and cannot be instantiated.

20.6.16 `Set Typeclasses Unique Solutions`

When a typeclass resolution is launched we ensure that it has a single solution or fail. This ensures that the resolution is canonical, but can make proof search much more expensive.

20.6.17 `Set Typeclasses Unique Instances`

Typeclass declarations introduced when this option is set have a more efficient resolution behavior (the option is off by default). When a solution to the typeclass goal of this class is found, we never backtrack on it, assuming that it is canonical.

20.6.18 `Typeclasses eauto := [debug] [(dfs) | (bfs)] [depth]`

This command allows more global customization of the type class resolution tactic. The semantics of the options are:

- `debug` In debug mode, the trace of successfully applied tactics is printed.
- `dfs`, `bfs` This sets the search strategy to depth-first search (the default) or breadth-first search.
- `depth` This sets the depth limit of the search.

20.6.19 `Set Typeclasses Debug [Verbosity num]`

These options allow to see the resolution steps of typeclasses that are performed during search. The `Debug` option is synonymous to `Debug Verbosity 1`, and `Debug Verbosity 2` provides more information (tried tactics, shelving of goals, etc. ...).

20.6.20 Set Refine Instance Mode

The option `Refine Instance Mode` allows to switch the behavior of instance declarations made through the `Instance` command.

- When it is on (the default), instances that have unsolved holes in their proof-term silently open the proof mode with the remaining obligations to prove.
- When it is off, they fail with an error instead.

Chapter 21

Omega: a solver of quantifier-free problems in Presburger Arithmetic

Pierre Crégut

21.1 Description of `omega`

`omega` solves a goal in Presburger arithmetic, i.e. a universally quantified formula made of equations and inequations. Equations may be specified either on the type `nat` of natural numbers or on the type `Z` of binary-encoded integer numbers. Formulas on `nat` are automatically injected into `Z`. The procedure may use any hypothesis of the current proof session to solve the goal.

Multiplication is handled by `omega` but only goals where at least one of the two multiplicands of products is a constant are solvable. This is the restriction meant by “Presburger arithmetic”.

If the tactic cannot solve the goal, it fails with an error message. In any case, the computation eventually stops.

21.1.1 Arithmetical goals recognized by `omega`

`omega` applied only to quantifier-free formulas built from the connectors

`/\, \/ , ~ , ->`

on atomic formulas. Atomic formulas are built from the predicates

`=, le, lt, gt, ge`

on `nat` or from the predicates

`=, <, <=, >, >=`

on `Z`. In expressions of type `nat`, `omega` recognizes

`plus, minus, mult, pred, S, O`

and in expressions of type `Z`, omega recognizes

`+`, `-`, `*`, `Z.succ`, and constants.

All expressions of type `nat` or `Z` not built on these operators are considered abstractly as if they were arbitrary variables of type `nat` or `Z`.

21.1.2 Messages from omega

When omega does not solve the goal, one of the following errors is generated:

Error messages:

1. omega can't solve this system
This may happen if your goal is not quantifier-free (if it is universally quantified, try `intros` first; if it contains existentials quantifiers too, omega is not strong enough to solve your goal). This may happen also if your goal contains arithmetical operators unknown from omega. Finally, your goal may be really wrong!
2. omega: Not a quantifier-free goal
If your goal is universally quantified, you should first apply `intro` as many time as needed.
3. omega: Unrecognized predicate or connective: *ident*
4. omega: Unrecognized atomic proposition: *prop*
5. omega: Can't solve a goal with proposition variables
6. omega: Unrecognized proposition
7. omega: Can't solve a goal with non-linear products
8. omega: Can't solve a goal with equality on *type*

21.2 Using omega

The omega tactic does not belong to the core system. It should be loaded by

```
Coq < Require Import Omega.
```

```
Coq < Open Scope Z_scope.
```

Example 3:

```
Coq < Goal forall m n : Z, 1 + 2 * m <> 2 * n.
1 subgoal
```

```
=====
forall m n : Z, 1 + 2 * m <> 2 * n
```

```
Coq < intros; omega.
No more subgoals.
```

Example 4:

```
Coq < Goal forall z:Z, z > 0 -> 2 * z + 1 > z.
1 subgoal
```

```
=====
forall z : Z, z > 0 -> 2 * z + 1 > z
```

```
Coq < intro; omega.
No more subgoals.
```

21.3 Options

```
Unset Stable Omega
```

This deprecated option (on by default) is for compatibility with Coq pre 8.5. It resets internal name counters to make executions of `omega` independent.

```
Unset Omega UseLocalDefs
```

This option (on by default) allows `omega` to use the bodies of local variables.

```
Set Omega System Set Omega Action
```

These two options (off by default) activate the printing of debug information.

21.4 Technical data

21.4.1 Overview of the tactic

- The goal is negated twice and the first negation is introduced as an hypothesis.
- Hypothesis are decomposed in simple equations or inequations. Multiple goals may result from this phase.
- Equations and inequations over `nat` are translated over \mathbb{Z} , multiple goals may result from the translation of substraction.
- Equations and inequations are normalized.
- Goals are solved by the *OMEGA* decision procedure.
- The script of the solution is replayed.

21.4.2 Overview of the *OMEGA* decision procedure

The *OMEGA* decision procedure involved in the `omega` tactic uses a small subset of the decision procedure presented in

"The Omega Test: a fast and practical integer programming algorithm for dependence analysis", William Pugh, Communication of the ACM, 1992, p 102-114.

Here is an overview, look at the original paper for more information.

- Equations and inequations are normalized by division by the GCD of their coefficients.
- Equations are eliminated, using the Banerjee test to get a coefficient equal to one.
- Note that each inequation defines a half space in the space of real value of the variables.
- Inequations are solved by projecting on the hyperspace defined by cancelling one of the variable. They are partitioned according to the sign of the coefficient of the eliminated variable. Pairs of inequations from different classes define a new edge in the projection.
- Redundant inequations are eliminated or merged in new equations that can be eliminated by the Banerjee test.
- The last two steps are iterated until a contradiction is reached (success) or there is no more variable to eliminate (failure).

It may happen that there is a real solution and no integer one. The last steps of the Omega procedure (dark shadow) are not implemented, so the decision procedure is only partial.

21.5 Bugs

- The simplification procedure is very dumb and this results in many redundant cases to explore.
- Much too slow.
- Certainly other bugs! You can report them to <https://coq.inria.fr/bugs/>.

Chapter 22

Micromega: tactics for solving arithmetic goals over ordered rings

Frédéric Besson and Evgeny Makarov

22.1 Short description of the tactics

The `Psatz` module (`Require Import Psatz.`) gives access to several tactics for solving arithmetic goals over \mathbb{Z} , \mathbb{Q} , and \mathbb{R} :¹. It is also possible to get the tactics for integers by a `Require Import Lia`, rationals `Require Import Lqa` and reals `Require Import Lra`.

- `lia` is a decision procedure for linear integer arithmetic (see Section 22.4);
- `nia` is an incomplete proof procedure for integer non-linear arithmetic (see Section 22.6);
- `lra` is a decision procedure for linear (real or rational) arithmetic (see Section 22.3);
- `nra` is an incomplete proof procedure for non-linear (real or rational) arithmetic (see Section 22.5);
- `psatz D n` where D is \mathbb{Z} or \mathbb{Q} or \mathbb{R} , and n is an optional integer limiting the proof search depth is an incomplete proof procedure for non-linear arithmetic. It is based on John Harrison's HOL Light driver to the external prover `csdp`². Note that the `csdp` driver is generating a *proof cache* which makes it possible to rerun scripts even without `csdp` (see Section 22.7).

The tactics solve propositional formulas parameterized by atomic arithmetic expressions interpreted over a domain $D \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$. The syntax of the formulas is the following:

$$\begin{aligned} F &::= A \mid P \mid \text{True} \mid \text{False} \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid F_1 \leftrightarrow F_2 \mid F_1 \rightarrow F_2 \mid \neg F \\ A &::= p_1 = p_2 \mid p_1 > p_2 \mid p_1 < p_2 \mid p_1 \geq p_2 \mid p_1 \leq p_2 \\ p &::= c \mid x \mid \neg p \mid p_1 - p_2 \mid p_1 + p_2 \mid p_1 \times p_2 \mid p^n \end{aligned}$$

¹Support for `nat` and `N` is obtained by pre-processing the goal with the `zify` tactic.

²Sources and binaries can be found at <https://projects.coin-or.org/Csdp>

where c is a numeric constant, $x \in D$ is a numeric variable, the operators $-$, $+$, \times are respectively subtraction, addition, product, p^n is exponentiation by a constant n , P is an arbitrary proposition. For \mathbb{Q} , equality is not Leibniz equality $=$ but the equality of rationals $==$.

For \mathbb{Z} (resp. \mathbb{Q}), c ranges over integer constants (resp. rational constants). For \mathbb{R} , the tactic recognizes as real constants the following expressions:

```
c ::= R0 | R1 | Rmul(c, c) | Rplus(c, c) | Rminus(c, c) | IZR z | IQR q
    | Rdiv(c, c) | Rinv c
```

where z is a constant in \mathbb{Z} and q is a constant in \mathbb{Q} . This includes integer constants written using the decimal notation *i.e.*, $c\%R$.

22.2 Positivstellensatz refutations

The name `psatz` is an abbreviation for *positivstellensatz* – literally positivity theorem – which generalizes Hilbert’s *nullstellensatz*. It relies on the notion of *Cone*. Given a (finite) set of polynomials S , $Cone(S)$ is inductively defined as the smallest set of polynomials closed under the following rules:

$$\frac{p \in S}{p \in Cone(S)} \quad \frac{p^2 \in Cone(S)}{p \in Cone(S)} \quad \frac{p_1 \in Cone(S) \quad p_2 \in Cone(S) \quad \bowtie \in \{+, *\}}{p_1 \bowtie p_2 \in Cone(S)}$$

The following theorem provides a proof principle for checking that a set of polynomial inequalities does not have solutions.³

Theorem 1 *Let S be a set of polynomials.*

If -1 belongs to $Cone(S)$ then the conjunction $\bigwedge_{p \in S} p \geq 0$ is unsatisfiable.

A proof based on this theorem is called a *positivstellensatz* refutation. The tactics work as follows. Formulas are normalized into conjunctive normal form $\bigwedge_i C_i$ where C_i has the general form $(\bigwedge_{j \in S_i} p_j \bowtie 0) \rightarrow False$ and $\bowtie \in \{>, \geq, =\}$ for $D \in \{\mathbb{Q}, \mathbb{R}\}$ and $\bowtie \in \{\geq, =\}$ for \mathbb{Z} . For each conjunct C_i , the tactic calls an oracle which searches for -1 within the cone. Upon success, the oracle returns a *cone expression* that is normalized by the `ring` tactic (see chapter 25) and checked to be -1 .

22.3 lra: a decision procedure for linear real and rational arithmetic

The `lra` tactic is searching for *linear* refutations using Fourier elimination.⁴ As a result, this tactic explores a subset of the *Cone* defined as

$$LinCone(S) = \left\{ \sum_{p \in S} \alpha_p \times p \mid \alpha_p \text{ are positive constants} \right\}.$$

The deductive power of `lra` is the combined deductive power of `ring_simplify` and `fourier`. There is also an overlap with the `field` tactic *e.g.*, $x = 10 * x / 10$ is solved by `lra`.

³Variants deal with equalities and strict inequalities.

⁴More efficient linear programming techniques could equally be employed.

22.4 lia: a tactic for linear integer arithmetic

The tactic `lia` offers an alternative to the `omega` and `romega` tactic (see Chapter 21). Roughly speaking, the deductive power of `lia` is the combined deductive power of `ring_simplify` and `omega`. However, it solves linear goals that `omega` and `romega` do not solve, such as the following so-called *omega nightmare* [130].

```
Coq < Goal forall x y,
      27 <= 11 * x + 13 * y <= 45 ->
      -10 <= 7 * x - 9 * y <= 4 -> False.
```

The estimation of the relative efficiency of `lia` vs `omega` and `romega` is under evaluation.

High level view of `lia`. Over \mathbb{R} , *positivstellensatz* refutations are a complete proof principle.⁵ However, this is not the case over \mathbb{Z} . Actually, *positivstellensatz* refutations are not even sufficient to decide linear *integer* arithmetic. The canonical example is $2 * x = 1 \rightarrow \text{False}$ which is a theorem of \mathbb{Z} but not a theorem of \mathbb{R} . To remedy this weakness, the `lia` tactic is using recursively a combination of:

- linear *positivstellensatz* refutations;
- cutting plane proofs;
- case split.

Cutting plane proofs are a way to take into account the discreteness of \mathbb{Z} by rounding up (rational) constants up-to the closest integer.

Theorem 2 *Let p be an integer and c a rational constant.*

$$p \geq c \Rightarrow p \geq \lceil c \rceil$$

For instance, from $2x = 1$ we can deduce

- $x \geq 1/2$ which cut plane is $x \geq \lceil 1/2 \rceil = 1$;
- $x \leq 1/2$ which cut plane is $x \leq \lfloor 1/2 \rfloor = 0$.

By combining these two facts (in normal form) $x - 1 \geq 0$ and $-x \geq 0$, we conclude by exhibiting a *positivstellensatz* refutation: $-1 \equiv x - 1 + -x \in \text{Cone}(\{x - 1, x\})$.

Cutting plane proofs and linear *positivstellensatz* refutations are a complete proof principle for integer linear arithmetic.

Case split enumerates over the possible values of an expression.

Theorem 3 *Let p be an integer and c_1 and c_2 integer constants.*

$$c_1 \leq p \leq c_2 \Rightarrow \bigvee_{x \in [c_1, c_2]} p = x$$

Our current oracle tries to find an expression e with a small range $[c_1, c_2]$. We generate $c_2 - c_1$ subgoals which contexts are enriched with an equation $e = i$ for $i \in [c_1, c_2]$ and recursively search for a proof.

⁵In practice, the oracle might fail to produce such a refutation.

22.5 nra: a proof procedure for non-linear arithmetic

The `nra` tactic is an experimental proof procedure for non-linear arithmetic. The tactic performs a limited amount of non-linear reasoning before running the linear prover of `lra`. This pre-processing does the following:

- If the context contains an arithmetic expression of the form $e[x^2]$ where x is a monomial, the context is enriched with $x^2 \geq 0$;
- For all pairs of hypotheses $e_1 \geq 0, e_2 \geq 0$, the context is enriched with $e_1 \times e_2 \geq 0$.

After this pre-processing, the linear prover of `lra` searches for a proof by abstracting monomials by variables.

22.6 nia: a proof procedure for non-linear integer arithmetic

The `nia` tactic is a proof procedure for non-linear integer arithmetic. It performs a pre-processing similar to `nra`. The obtained goal is solved using the linear integer prover `lia`.

22.7 psatz: a proof procedure for non-linear arithmetic

The `psatz` tactic explores the *Cone* by increasing degrees – hence the depth parameter n . In theory, such a proof search is complete – if the goal is provable the search eventually stops. Unfortunately, the external oracle is using numeric (approximate) optimization techniques that might miss a refutation.

To illustrate the working of the tactic, consider we wish to prove the following Coq goal.

```
Coq < Goal forall x, -x^2 >= 0 -> x - 1 >= 0 -> False.
```

Such a goal is solved by `intro x; psatz Z 2`. The oracle returns the cone expression $2 \times (x - 1) + (x - 1) \times (x - 1) + \mathbf{-x^2}$ (polynomial hypotheses are printed in bold). By construction, this expression belongs to $\text{Cone}(\{-x^2, x - 1\})$. Moreover, by running `ring` we obtain -1 . By Theorem 1, the goal is valid.

Chapter 23

Extraction of programs in Objective Caml and Haskell

Jean-Christophe Filliâtre and Pierre Letouzey

We present here the COQ extraction commands, used to build certified and relatively efficient functional programs, extracting them from either COQ functions or COQ proofs of specifications. The functional languages available as output are currently OCAML, HASKELL and SCHEME. In the following, “ML” will be used (abusively) to refer to any of the three.

Before using any of the commands or options described in this chapter, the extraction framework should first be loaded explicitly via `Require Extraction`, or via the more robust `From Coq Require Extraction`. Note that in earlier versions of Coq, these commands and options were directly available without any preliminary `Require`.

```
Coq < Require Extraction.  
[Loading ML file extraction_plugin.cmxs ... done]
```

23.1 Generating ML code

The next two commands are meant to be used for rapid preview of extraction. They both display extracted term(s) inside COQ.

```
Extraction qualid .
```

Extraction of a constant or module in the COQ toplevel.

```
Recursive Extraction qualid1 ... qualidn.
```

Recursive extraction of all the globals (or modules) *qualid*₁ ... *qualid*_{*n*} and all their dependencies in the COQ toplevel.

All the following commands produce real ML files. User can choose to produce one monolithic file or one file per COQ library.

`Extraction "file" qualid1 ... qualidn.`

Recursive extraction of all the globals (or modules) *qualid*₁ ... *qualid*_{*n*} and all their dependencies in one monolithic file *file*. Global and local identifiers are renamed according to the chosen ML language to fulfill its syntactic conventions, keeping original names as much as possible.

`Extraction Library ident.`

Extraction of the whole COQ library *ident.v* to an ML module *ident.ml*. In case of name clash, identifiers are here renamed using prefixes `coq_` or `Coq_` to ensure a session-independent renaming.

`Recursive Extraction Library ident.`

Extraction of the COQ library *ident.v* and all other modules *ident.v* depends on.

`Separate Extraction qualid1 ... qualidn.`

Recursive extraction of all the globals (or modules) *qualid*₁ ... *qualid*_{*n*} and all their dependencies, just as `Extraction "file"`, but instead of producing one monolithic file, this command splits the produced code in separate ML files, one per corresponding Coq .v file. This command is hence quite similar to `Recursive Extraction Library`, except that only the needed parts of Coq libraries are extracted instead of the whole. The naming convention in case of name clash is the same one as `Extraction Library`: identifiers are here renamed using prefixes `coq_` or `Coq_`.

The following command is meant to help automatic testing of the extraction, see for instance the `test-suite` directory in the COQ sources.

`Extraction TestCompile qualid1 ... qualidn.`

All the globals (or modules) *qualid*₁ ... *qualid*_{*n*} and all their dependencies are extracted to a temporary Ocaml file, just as in `Extraction "file"`. Then this temporary file and its signature are compiled with the same Ocaml compiler used to built COQ. This command succeeds only if the extraction and the Ocaml compilation succeed (and it fails if the current target language of the extraction is not Ocaml).

23.2 Extraction options

23.2.1 Setting the target language

The ability to fix target language is the first and more important of the extraction options. Default is Ocaml.

`Extraction Language Ocaml.`

`Extraction Language Haskell.`

`Extraction Language Scheme.`

23.2.2 Inlining and optimizations

Since Objective Caml is a strict language, the extracted code has to be optimized in order to be efficient (for instance, when using induction principles we do not want to compute all the recursive calls but only the needed ones). So the extraction mechanism provides an automatic optimization routine that will be called each time the user want to generate Ocaml programs. The optimizations can be split in two groups: the type-preserving ones – essentially constant inlining and reductions – and the non type-preserving ones – some function abstractions of dummy types are removed when it is deemed safe in order to have more elegant types. Therefore some constants may not appear in the resulting monolithic Ocaml program. In the case of modular extraction, even if some inlining is done, the inlined constant are nevertheless printed, to ensure session-independent programs.

Concerning Haskell, type-preserving optimizations are less useful because of laziness. We still make some optimizations, for example in order to produce more readable code.

The type-preserving optimizations are controlled by the following COQ options:

`Unset Extraction Optimize.`

Default is `Set`. This controls all type-preserving optimizations made on the ML terms (mostly reduction of dummy beta/iota redexes, but also simplifications on Cases, etc). Put this option to `Unset` if you want a ML term as close as possible to the Coq term.

`Set Extraction Conservative Types.`

Default is `Unset`. This controls the non type-preserving optimizations made on ML terms (which try to avoid function abstraction of dummy types). Turn this option to `Set` to make sure that $e : \tau$ implies that $e' : \tau'$ where e' and τ' are the extracted code of e and τ respectively.

`Set Extraction KeepSingleton.`

Default is `Unset`. Normally, when the extraction of an inductive type produces a singleton type (i.e. a type with only one constructor, and only one argument to this constructor), the inductive structure is removed and this type is seen as an alias to the inner type. The typical example is `sig`. This option allows disabling this optimization when one wishes to preserve the inductive structure of types.

`Unset Extraction AutoInline.`

Default is `Set`. The extraction mechanism inlines the bodies of some defined constants, according to some heuristics like size of bodies, uselessness of some arguments, etc. Those heuristics are not always perfect; if you want to disable this feature, do it by `Unset`.

`Extraction [Inline|NoInline] qualid1 ... qualidn.`

In addition to the automatic inline feature, you can tell to inline some more constants by the `Extraction Inline` command. Conversely, you can forbid the automatic inlining of some specific constants by the `Extraction NoInline` command. Those two commands enable a precise control of what is inlined and what is not.

`Print Extraction Inline.`

Prints the current state of the table recording the custom inlinings declared by the two previous commands.

`Reset Extraction Inline.`

Puts the table recording the custom inlinings back to empty.

Inlining and printing of a constant declaration. A user can explicitly ask for a constant to be extracted by two means:

- by mentioning it on the extraction command line
- by extracting the whole COQ module of this constant.

In both cases, the declaration of this constant will be present in the produced file. But this same constant may or may not be inlined in the following terms, depending on the automatic/custom inlining mechanism.

For the constants non-explicitly required but needed for dependency reasons, there are two cases:

- If an inlining decision is taken, whether automatically or not, all occurrences of this constant are replaced by its extracted body, and this constant is not declared in the generated file.
- If no inlining decision is taken, the constant is normally declared in the produced file.

23.2.3 Extra elimination of useless arguments

The following command provides some extra manual control on the code elimination performed during extraction, in a way which is independent but complementary to the main elimination principles of extraction (logical parts and types).

```
Extraction Implicit qualid [ ident1 ... identn ].
```

This experimental command allows declaring some arguments of *qualid* as implicit, i.e. useless in extracted code and hence to be removed by extraction. Here *qualid* can be any function or inductive constructor, and *ident*_{*i*} are the names of the concerned arguments. In fact, an argument can also be referred by a number indicating its position, starting from 1.

When an actual extraction takes place, an error is normally raised if the `Extraction Implicit` declarations cannot be honored, that is if any of the implicated variables still occurs in the final code. This behavior can be relaxed via the following option:

```
Unset Extraction SafeImplicits.
```

Default is Set. When this option is Unset, a warning is emitted instead of an error if some implicated variables still occur in the final code of an extraction. This way, the extracted code may be obtained nonetheless and reviewed manually to locate the source of the issue (in the code, some comments mark the location of these remaining implicated variables). Note that this extracted code might not compile or run properly, depending of the use of these remaining implicated variables.

23.2.4 Realizing axioms

Extraction will fail if it encounters an informative axiom not realized (see Section 23.2.4). A warning will be issued if it encounters a logical axiom, to remind the user that inconsistent logical axioms may lead to incorrect or non-terminating extracted terms.

It is possible to assume some axioms while developing a proof. Since these axioms can be any kind of proposition or object or type, they may perfectly well have some computational content. But a program must be a closed term, and of course the system cannot guess the program which realizes an axiom. Therefore, it is possible to tell the system what ML term corresponds to a given axiom.

`Extract Constant qualid => string .`

Give an ML extraction for the given constant. The *string* may be an identifier or a quoted string.

`Extract Inlined Constant qualid => string .`

Same as the previous one, except that the given ML terms will be inlined everywhere instead of being declared via a `let`.

Note that the `Extract Inlined Constant` command is sugar for an `Extract Constant` followed by a `Extraction Inline`. Hence a `Reset Extraction Inline` will have an effect on the realized and inlined axiom.

Of course, it is the responsibility of the user to ensure that the ML terms given to realize the axioms do have the expected types. In fact, the strings containing realizing code are just copied to the extracted files. The extraction recognizes whether the realized axiom should become a ML type constant or a ML object declaration.

Example:

```
Coq < Axiom X:Set.
Coq < Axiom x:X.
Coq < Extract Constant X => "int".
Coq < Extract Constant x => "0".
```

Notice that in the case of type scheme axiom (i.e. whose type is an arity, that is a sequence of product finished by a sort), then some type variables have to be given. The syntax is then:

`Extract Constant qualid string1 ... stringn => string .`

The number of type variables is checked by the system.

Example:

```
Coq < Axiom Y : Set -> Set -> Set.
Coq < Extract Constant Y "'a" "'b" => " 'a*'b " .
```

Realizing an axiom via `Extract Constant` is only useful in the case of an informative axiom (of sort `Type` or `Set`). A logical axiom have no computational content and hence will not appears in extracted terms. But a warning is nonetheless issued if extraction encounters a logical axiom. This warning reminds user that inconsistent logical axioms may lead to incorrect or non-terminating extracted terms.

If an informative axiom has not been realized before an extraction, a warning is also issued and the definition of the axiom is filled with an exception labeled `AXIOM TO BE REALIZED`. The user must then search these exceptions inside the extracted file and replace them by real code.

The system also provides a mechanism to specify ML terms for inductive types and constructors. For instance, the user may want to use the ML native boolean type instead of COQ one. The syntax is the following:

`Extract Inductive qualid => string [string ... string] optstring .`

Give an ML extraction for the given inductive type. You must specify extractions for the type itself (first *string*) and all its constructors (between square brackets). If given, the final optional string should contain a function emulating pattern-matching over this inductive type. If this optional string is not given, the ML extraction must be an ML inductive datatype, and the native pattern-matching of the language will be used.

For an inductive type with k constructor, the function used to emulate the match should expect $(k + 1)$ arguments, first the k branches in functional form, and then the inductive element to destruct. For instance, the match branch `| S n => foo` gives the functional form `(fun n -> foo)`. Note that a constructor with no argument is considered to have one unit argument, in order to block early evaluation of the branch: `| O => bar` leads to the functional form `(fun () -> bar)`. For instance, when extracting `nat` into `int`, the code to provide has type: `(unit->'a)->(int->'a)->int->'a`.

As for `Extract Inductive`, this command should be used with care:

- The ML code provided by the user is currently *not* checked at all by extraction, even for syntax errors.
- Extracting an inductive type to a pre-existing ML inductive type is quite sound. But extracting to a general type (by providing an ad-hoc pattern-matching) will often *not* be fully rigorously correct. For instance, when extracting `nat` to Ocaml's `int`, it is theoretically possible to build `nat` values that are larger than Ocaml's `max_int`. It is the user's responsibility to be sure that no overflow or other bad events occur in practice.
- Translating an inductive type to an ML type does *not* magically improve the asymptotic complexity of functions, even if the ML type is an efficient representation. For instance, when extracting `nat` to Ocaml's `int`, the function `mult` stays quadratic. It might be interesting to associate this translation with some specific `Extract Constant` when primitive counterparts exist.

Example: Typical examples are the following:

```
Coq < Extract Inductive unit => "unit" [ "()" ].
Coq < Extract Inductive bool => "bool" [ "true" "false" ].
Coq < Extract Inductive sumbool => "bool" [ "true" "false" ].
```

When extracting to OCAML, if an inductive constructor or type has arity 2 and the corresponding string is enclosed by parentheses, and the string meets OCAML's lexical criteria for an infix symbol, then the rest of the string is used as infix constructor or type.

```
Coq < Extract Inductive list => "list" [ "[]" "(::)" ].
Coq < Extract Inductive prod => "(*)" [ "(,)" ].
```

As an example of translation to a non-inductive datatype, let's turn `nat` into Ocaml's `int` (see caveat above):

```
Coq < Extract Inductive nat => int [ "0" "succ" ]
      "(fun fO fS n -> if n=0 then fO () else fS (n-1))".
```

23.2.5 Avoiding conflicts with existing filenames

When using `Extraction Library`, the names of the extracted files directly depends from the names of the COQ files. It may happen that these filenames are in conflict with already existing files, either in the standard library of the target language or in other code that is meant to be linked with the extracted code. For instance the module `List` exists both in COQ and in Ocaml. It is possible to instruct the extraction not to use particular filenames.

```
Extraction Blacklist ident ... ident.
```

Instruct the extraction to avoid using these names as filenames for extracted code.

```
Print Extraction Blacklist.
```

Show the current list of filenames the extraction should avoid.

```
Reset Extraction Blacklist.
```

Allow the extraction to use any filename.

For Ocaml, a typical use of these commands is `Extraction Blacklist String List`.

23.3 Differences between COQ and ML type systems

Due to differences between COQ and ML type systems, some extracted programs are not directly typable in ML. We now solve this problem (at least in Ocaml) by adding when needed some unsafe casting `Obj.magic`, which give a generic type `'a` to any term.

For example, here are two kinds of problem that can occur:

- If some part of the program is *very* polymorphic, there may be no ML type for it. In that case the extraction to ML works alright but the generated code may be refused by the ML type-checker. A very well known example is the *distr-pair* function:

```
Definition dp :=
  fun (A B:Set) (x:A) (y:B) (f:forall C:Set, C->C) => (f A x, f B y).
```

In Ocaml, for instance, the direct extracted term would be

```
let dp x y f = Pair((f () x), (f () y))
```

and would have type

```
dp : 'a -> 'a -> (unit -> 'a -> 'b) -> ('b,'b) prod
```

which is not its original type, but a restriction.

We now produce the following correct version:

```
let dp x y f = Pair ((Obj.magic f () x), (Obj.magic f () y))
```

- Some definitions of COQ may have no counterpart in ML. This happens when there is a quantification over types inside the type of a constructor; for example:

```
Inductive anything : Type := dummy : forall A:Set, A -> anything.
```

which corresponds to the definition of an ML dynamic type. In Ocaml, we must cast any argument of the constructor `dummy`.

Even with those unsafe castings, you should never get error like “segmentation fault”. In fact even if your program may seem ill-typed to the Ocaml type-checker, it can’t go wrong: it comes from a Coq well-typed terms, so for example inductives will always have the correct number of arguments, etc.

More details about the correctness of the extracted programs can be found in [99].

We have to say, though, that in most “realistic” programs, these problems do not occur. For example all the programs of Coq library are accepted by Caml type-checker without any `Obj.magic` (see examples below).

23.4 Some examples

We present here two examples of extractions, taken from the COQ Standard Library. We choose OCAML as target language, but all can be done in the other dialects with slight modifications. We then indicate where to find other examples and tests of Extraction.

23.4.1 A detailed example: Euclidean division

The file `Euclid` contains the proof of Euclidean division (theorem `eucl_dev`). The natural numbers defined in the example files are unary integers defined by two constructors *O* and *S*:

```
Coq < Inductive nat : Set :=
  | O : nat
  | S : nat -> nat.
```

This module contains a theorem `eucl_dev`, whose type is

```
forall b:nat, b > 0 -> forall a:nat, diveucl a b
```

where `diveucl` is a type for the pair of the quotient and the modulo, plus some logical assertions that disappear during extraction. We can now extract this program to OCAML:

```
Coq < Require Extraction.
Coq < Require Import Euclid Wf_nat.
Coq < Extraction Inline gt_wf_rec lt_wf_rec induction_ltof2.
Coq < Recursive Extraction eucl_dev.
type nat =
| O
| S of nat
type sumbool =
| Left
| Right
(** val sub : nat -> nat -> nat **)
let rec sub n m =
  match n with
  | O -> n
  | S k -> (match m with
            | O -> n
            | S l -> sub k l)
(** val le_lt_dec : nat -> nat -> sumbool **)
let rec le_lt_dec n m =
  match n with
  | O -> Left
  | S n0 -> (match m with
             | O -> Right
             | S m0 -> le_lt_dec n0 m0)
(** val le_gt_dec : nat -> nat -> sumbool **)
let le_gt_dec =
  le_lt_dec
type diveucl =
| Divex of nat * nat
(** val eucl_dev : nat -> nat -> diveucl **)
```

```

let rec eucl_dev n m =
  let s = le_gt_dec n m in
  (match s with
  | Left ->
    let d = let y = sub m n in eucl_dev n y in
    let Divex (q, r) = d in Divex ((S q), r)
  | Right -> Divex (O, m))

```

The inlining of `gt_wf_rec` and others is not mandatory. It only enhances readability of extracted code. You can then copy-paste the output to a file `euclid.ml` or let COQ do it for you with the following command:

```
Extraction "euclid" eucl_dev.
```

Let us play the resulting program:

```

# #use "euclid.ml";;
type nat = O | S of nat
type sumbool = Left | Right
val minus : nat -> nat -> nat = <fun>
val le_lt_dec : nat -> nat -> sumbool = <fun>
val le_gt_dec : nat -> nat -> sumbool = <fun>
type diveucl = Divex of nat * nat
val eucl_dev : nat -> nat -> diveucl = <fun>
# eucl_dev (S (S O)) (S (S (S (S (S O)))));;
- : diveucl = Divex (S (S O), S O)

```

It is easier to test on OCAML integers:

```

# let rec nat_of_int = function 0 -> O | n -> S (nat_of_int (n-1));;
val nat_of_int : int -> nat = <fun>
# let rec int_of_nat = function 0 -> 0 | S p -> 1+(int_of_nat p);;
val int_of_nat : nat -> int = <fun>
# let div a b =
  let Divex (q,r) = eucl_dev (nat_of_int b) (nat_of_int a)
  in (int_of_nat q, int_of_nat r);;
val div : int -> int -> int * int = <fun>
# div 173 15;;
- : int * int = (11, 8)

```

Note that these `nat_of_int` and `int_of_nat` are now available via a mere `Require Import ExtrOcamlIntConv` and then adding these functions to the list of functions to extract. This file `ExtrOcamlIntConv.v` and some others in `plugins/extraction/` are meant to help building concrete program via extraction.

23.4.2 Extraction's horror museum

Some pathological examples of extraction are grouped in the file `test-suite/success/extraction.v` of the sources of COQ.

23.4.3 Users' Contributions

Several of the COQ Users' Contributions use extraction to produce certified programs. In particular the following ones have an automatic extraction test:

- `additions`
- `bdds`
- `canon-bdds`
- `chinese`
- `continuations`
- `coq-in-coq`
- `exceptions`
- `firing-squad`
- `founify`
- `graphs`
- `higman-cf`
- `higman-nw`
- `hardware`
- `multiplier`
- `search-trees`
- `stalmarck`

`continuations` and `multiplier` are a bit particular. They are examples of developments where `Obj.magic` are needed. This is probably due to an heavy use of impredicativity. After compilation, those two examples run nonetheless, thanks to the correction of the extraction [99].

Chapter 24

PROGRAM

Matthieu Sozeau

We present here the `PROGRAM` tactic commands, used to build certified COQ programs, elaborating them from their algorithmic skeleton and a rich specification [135]. It can be thought of as a dual of extraction (see Chapter 23). The goal of `PROGRAM` is to program as in a regular functional programming language whilst using as rich a specification as desired and proving that the code meets the specification using the whole COQ proof apparatus. This is done using a technique originating from the “Predicate subtyping” mechanism of PVS[132], which generates type-checking conditions while typing a term constrained to a particular type. Here we insert existential variables in the term, which must be filled with proofs to get a complete COQ term. `PROGRAM` replaces the `PROGRAM` tactic by Catherine Parent [121] which had a similar goal but is no longer maintained.

The languages available as input are currently restricted to COQ’s term language, but may be extended to OCAML, HASKELL and others in the future. We use the same syntax as COQ and permit to use implicit arguments and the existing coercion mechanism. Input terms and types are typed in an extended system (RUSSELL) and interpreted into COQ terms. The interpretation process may produce some proof obligations which need to be resolved to create the final term.

24.1 Elaborating programs

The main difference from COQ is that an object in a type $T : \mathbf{Set}$ can be considered as an object of type $\{x : T \mid P\}$ for any wellformed $P : \mathbf{Prop}$. If we go from T to the subset of T verifying property P , we must prove that the object under consideration verifies it. RUSSELL will generate an obligation for every such coercion. In the other direction, RUSSELL will automatically insert a projection.

Another distinction is the treatment of pattern-matching. Apart from the following differences, it is equivalent to the standard `match` operation (see Section 4.5.3).

- Generation of equalities. A `match` expression is always generalized by the corresponding equality. As an example, the expression:

```
match x with
| 0 => t
| S n => u
end.
```

will be first rewritten to:

```
(match x as y return (x = y -> _) with
| 0 => fun H : x = 0 -> t
| S n => fun H : x = S n -> u
end) (eq_refl n).
```

This permits to get the proper equalities in the context of proof obligations inside clauses, without which reasoning is very limited.

- **Generation of inequalities.** If a pattern intersects with a previous one, an inequality is added in the context of the second branch. See for example the definition of `div2` below, where the second branch is typed in a context where $\forall p, _ <> S(Sp)$.
- **Coercion.** If the object being matched is coercible to an inductive type, the corresponding coercion will be automatically inserted. This also works with the previous mechanism.

There are options to control the generation of equalities and coercions.

- **Unset Program Cases** This deactivates the special treatment of pattern-matching generating equalities and inequalities when using `PROGRAM` (it is on by default). All pattern-matchings and let-patterns are handled using the standard algorithm of Coq (see Section 17) when this option is deactivated.
- **Unset Program Generalized Coercion** This deactivates the coercion of general inductive types when using `PROGRAM` (the option is on by default). Coercion of subset types and pairs is still active in this case.

24.1.1 Syntactic control over equalities

To give more control over the generation of equalities, the typechecker will fall back directly to COQ's usual typing of dependent pattern-matching if a `return` or `in` clause is specified. Likewise, the `if` construct is not treated specially by `PROGRAM` so boolean tests in the code are not automatically reflected in the obligations. One can use the `dec` combinator to get the correct hypotheses as in:

```
Coq < Program Definition id (n : nat) : { x : nat | x = n } :=
  if dec (leb n 0) then 0
  else S (pred n).
id has type-checked, generating 2 obligations
Solving obligations automatically...
2 obligations remaining
Obligation 1 of id:
(forall n : nat, (n <=? 0) = true -> (fun x : nat => x = n) 0).

Obligation 2 of id:
(forall n : nat,
  (n <=? 0) = false -> (fun x : nat => x = n) (S (Init.Nat.pred n))).
```

The let tupling construct `let (x1, ..., xn) := t in b` does not produce an equality, contrary to the let pattern construct `let ' (x1, ..., xn) := t in b`. Also, `term :>` explicitly asks the system to coerce *term* to its support type. It can be useful in notations, for example:


```
Coq < Notation " x `= y " := (@eq _ (x :>) (y :>)) (only parsing).
```

This notation denotes equality on subset types using equality on their support types, avoiding uses of proof-irrelevance that would come up when reasoning with equality on the subset types themselves.

The next two commands are similar to their standard counterparts `Definition` (see Section 1.3.2) and `Fixpoint` (see Section 1.3.4) in that they define constants. However, they may require the user to prove some goals to construct the final definitions.

24.1.2 Program Definition `ident := term`.

This command types the value `term` in `RUSSELL` and generates proof obligations. Once solved using the commands shown below, it binds the final COQ term to the name `ident` in the environment.

Error messages:

1. `ident` already exists

Variants:

1. Program Definition `ident :term1 := term2`.
It interprets the type `term1`, potentially generating proof obligations to be resolved. Once done with them, we have a COQ type `term'1`. It then checks that the type of the interpretation of `term2` is coercible to `term'1`, and registers `ident` as being of type `term'1` once the set of obligations generated during the interpretation of `term2` and the aforementioned coercion derivation are solved.
2. Program Definition `ident binder1...bindern :term1 := term2`.
This is equivalent to
Program Definition `ident : forall binder1...bindern, term1 := fun binder1...bindern => term2`.

Error messages:

1. In environment ... the term: `term2` does not have type `term1`.
Actually, it has type `term3`.

See also: Sections 6.10.1, 6.10.2, 8.7.5

24.1.3 Program Fixpoint `ident params {order} : type := term`

The structural fixpoint operator behaves just like the one of Coq (see Section 1.3.4), except it may also generate obligations. It works with mutually recursive definitions too.

```
Coq < Program Fixpoint div2 (n : nat) : { x : nat | n = 2 * x \/ n = 2 * x + 1 } :=
  match n with
  | S (S p) => S (div2 p)
  | _      => O
  end.
```

Solving obligations automatically...
4 obligations remaining

Here we have one obligation for each branch (branches for 0 and (S 0) are automatically generated by the pattern-matching compilation algorithm).

```
Coq < Obligation 1.
1 subgoal

p, x : nat
o : p = x + (x + 0) \ / p = x + (x + 0) + 1
=====
S (S p) = S (x + S (x + 0)) \ / S (S p) = S (x + S (x + 0) + 1)
```

One can use a well-founded order or a measure as termination orders using the syntax:

```
Coq < Program Fixpoint div2 (n : nat) {measure n} :
  { x : nat | n = 2 * x \ / n = 2 * x + 1 } :=
  match n with
  | S (S p) => S (div2 p)
  | _ => 0
  end.
```

The order annotation can be either:

- `measure f (R) ?` where `f` is a value of type `X` computed on any subset of the arguments and the optional (parenthesised) term `(R)` is a relation on `X`. By default `X` defaults to `nat` and `R` to `lt`.
- `wf R x` which is equivalent to `measure x (R)`.

Caution When defining structurally recursive functions, the generated obligations should have the prototype of the currently defined functional in their context. In this case, the obligations should be transparent (e.g. defined using `Defined`) so that the guardedness condition on recursive calls can be checked by the kernel's type-checker. There is an optimization in the generation of obligations which gets rid of the hypothesis corresponding to the functional when it is not necessary, so that the obligation can be declared opaque (e.g. using `Qed`). However, as soon as it appears in the context, the proof of the obligation is *required* to be declared transparent.

No such problems arise when using measures or well-founded recursion.

24.1.4 Program Lemma *ident* : type.

The RUSSELL language can also be used to type statements of logical properties. It will generate obligations, try to solve them automatically and fail if some unsolved obligations remain. In this case, one can first define the lemma's statement using `Program Definition` and use it as the goal afterwards. Otherwise the proof will be started with the elaborated version as a goal. The `Program` prefix can similarly be used as a prefix for `Variable`, `Hypothesis`, `Axiom` etc...

24.2 Solving obligations

The following commands are available to manipulate obligations. The optional identifier is used when multiple functions have unsolved obligations (e.g. when defining mutually recursive blocks). The optional tactic is replaced by the default one if not specified.

- `[Local|Global] Obligation Tactic := expr` Sets the default obligation solving tactic applied to all obligations automatically, whether to solve them or when starting to prove one, e.g. using `Next`. `Local` makes the setting last only for the current module. Inside sections, `local` is the default.
- `Show Obligation Tactic` Displays the current default tactic.
- `Obligations [of ident]` Displays all remaining obligations.
- `Obligation num [of ident]` Start the proof of obligation `num`.
- `Next Obligation [of ident]` Start the proof of the next unsolved obligation.
- `Solve Obligations [of ident] [with expr]` Tries to solve each obligation of `ident` using the given tactic or the default one.
- `Solve All Obligations [with expr]` Tries to solve each obligation of every program using the given tactic or the default one (useful for mutually recursive definitions).
- `Admit Obligations [of ident]` Admits all obligations (does not work with structurally recursive programs).
- `Preterm [of ident]` Shows the term that will be fed to the kernel once the obligations are solved. Useful for debugging.
- `Set Transparent Obligations` Control whether all obligations should be declared as transparent (the default), or if the system should infer which obligations can be declared opaque.
- `Set Hide Obligations` Control whether obligations appearing in the term should be hidden as implicit arguments of the special constant `Program.Tactics.obligation`.
- `Set Shrink Obligations` *Deprecated since 8.7* This option (on by default) controls whether obligations should have their context minimized to the set of variables used in the proof of the obligation, to avoid unnecessary dependencies.

The module `Coq.Program.Tactics` defines the default tactic for solving obligations called `program_simpl`. Importing `Coq.Program.Program` also adds some useful notations, as documented in the file itself.

24.3 Frequently Asked Questions

- **Ill-formed recursive definitions** This error can happen when one tries to define a function by structural recursion on a subset object, which means the Coq function looks like:

```
Program Fixpoint f (x : A | P) := match x with A b => f b end.
```

Supposing $b : A$, the argument at the recursive call to `f` is not a direct subterm of `x` as `b` is wrapped inside an `exist` constructor to build an object of type $\{x : A \mid P\}$. Hence the definition is rejected by the guardedness condition checker. However one can use wellfounded recursion on subset objects like this:

```
Program Fixpoint f (x : A | P) { measure (size x) } :=
  match x with A b => f b end.
```

One will then just have to prove that the measure decreases at each recursive call. There are three drawbacks though:

1. A measure function has to be defined;
2. The reduction is a little more involved, although it works well using lazy evaluation;
3. Mutual recursion on the underlying inductive type isn't possible anymore, but nested mutual recursion is always possible.

Chapter 25

The `ring` and `field` tactic families

Bruno Barras, Benjamin Grégoire, Assia Mahboubi, Laurent Théry¹

This chapter presents the tactics dedicated to deal with ring and field equations.

25.1 What does this tactic do?

`ring` does associative-commutative rewriting in ring and semi-ring structures. Assume you have two binary functions \oplus and \otimes that are associative and commutative, with \oplus distributive on \otimes , and two constants 0 and 1 that are unities for \oplus and \otimes . A *polynomial* is an expression built on variables V_0, V_1, \dots and constants by application of \oplus and \otimes .

Let an *ordered product* be a product of variables $V_{i_1} \otimes \dots \otimes V_{i_n}$ verifying $i_1 \leq i_2 \leq \dots \leq i_n$. Let a *monomial* be the product of a constant and an ordered product. We can order the monomials by the lexicographic order on products of variables. Let a *canonical sum* be an ordered sum of monomials that are all different, i.e. each monomial in the sum is strictly less than the following monomial according to the lexicographic order. It is an easy theorem to show that every polynomial is equivalent (modulo the ring properties) to exactly one canonical sum. This canonical sum is called the *normal form* of the polynomial. In fact, the actual representation shares monomials with same prefixes. So what does `ring`? It normalizes polynomials over any ring or semi-ring structure. The basic use of `ring` is to simplify ring expressions, so that the user does not have to deal manually with the theorems of associativity and commutativity.

Examples:

1. In the ring of integers, the normal form of $x(3 + yx + 25(1 - z)) + zx$ is $28x + (-24)xz + xxy$.

`ring` is also able to compute a normal form modulo monomial equalities. For example, under the hypothesis that $2x^2 = yz + 1$, the normal form of $2(x + 1)x - x - zy$ is $x + 1$.

¹based on previous work from Patrick Loiseleur and Samuel Boutin

25.2 The variables map

It is frequent to have an expression built with $+$ and \times , but rarely on variables only. Let us associate a number to each subterm of a ring expression in the GALLINA language. For example in the ring `nat`, consider the expression:

```
(plus (mult (plus (f (5)) x) x)
      (mult (if b then (4) else (f (3))) (2)))
```

As a ring expression, it has 3 subterms. Give each subterm a number in an arbitrary order:

```
0  $\mapsto$  if b then (4) else (f (3))
1  $\mapsto$  (f (5))
2  $\mapsto$  x
```

Then normalize the “abstract” polynomial

$$((V_1 \otimes V_2) \oplus V_2) \oplus (V_0 \otimes 2)$$

In our example the normal form is:

$$(2 \otimes V_0) \oplus (V_1 \otimes V_2) \oplus (V_2 \otimes V_2)$$

Then substitute the variables by their values in the variables map to get the concrete normal polynomial:

```
(plus (mult (2) (if b then (4) else (f (3))))
      (plus (mult (f (5)) x) (mult x x)))
```

25.3 Is it automatic?

Yes, building the variables map and doing the substitution after normalizing is automatically done by the tactic. So you can just forget this paragraph and use the tactic according to your intuition.

25.4 Concrete usage in Coq

The `ring` tactic solves equations upon polynomial expressions of a ring (or semi-ring) structure. It proceeds by normalizing both hand sides of the equation (w.r.t. associativity, commutativity and distributivity, constant propagation, rewriting of monomials) and comparing syntactically the results.

`ring_simplify` applies the normalization procedure described above to the terms given. The tactic then replaces all occurrences of the terms given in the conclusion of the goal by their normal forms. If no term is given, then the conclusion should be an equation and both hand sides are normalized. The tactic can also be applied in a hypothesis.

The tactic must be loaded by `Require Import Ring`. The ring structures must be declared with the `Add Ring` command (see below). The ring of booleans is predefined; if one wants to use the tactic on `nat` one must first require the module `ArithRing` (exported by `Arith`); for `Z`, do `Require Import ZArithRing` or simply `Require Import ZArith`; for `N`, do `Require Import NArithRing` or `Require Import NArith`.

Example:

```

Coq < Require Import ZArith.
Coq < Open Scope Z_scope.
Coq < Goal forall a b c:Z,
    (a + b + c)^2 =
    a * a + b^2 + c * c + 2 * a * b + 2 * a * c + 2 * b * c.
1 subgoal

=====
forall a b c : Z,
  (a + b + c) ^ 2 = a * a + b ^ 2 + c * c + 2 * a * b + 2 * a * c + 2 * b * c
Coq < intros; ring.
No more subgoals.

Coq < Goal forall a b:Z, 2*a*b = 30 ->
    (a+b)^2 = a^2 + b^2 + 30.
1 subgoal

=====
forall a b : Z, 2 * a * b = 30 -> (a + b) ^ 2 = a ^ 2 + b ^ 2 + 30
Coq < intros a b H; ring [H].
No more subgoals.

```

Variants:

1. `ring [term1 ... termn]` decides the equality of two terms modulo ring operations and rewriting of the equalities defined by `term1 ... termn`. Each of `term1 ... termn` has to be a proof of some equality $m = p$, where m is a monomial (after “abstraction”), p a polynomial and $=$ the corresponding equality of the ring structure.
2. `ring_simplify [term1 ... termn] t1...tm in ident` performs the simplification in the hypothesis named `ident`.

Warning: `ring_simplify term1; ring_simplify term2` is not equivalent to `ring_simplify term1 term2`. In the latter case the variables map is shared between the two terms, and common subterm t of `term1` and `term2` will have the same associated variable number. So the first alternative should be avoided for terms belonging to the same ring theory.

Error messages:

1. not a valid ring equation The conclusion of the goal is not provable in the corresponding ring theory.
2. arguments of `ring_simplify` do not have all the same type
`ring_simplify` cannot simplify terms of several rings at the same time. Invoke the tactic once per ring structure.
3. cannot find a declared ring structure over term No ring has been declared for the type of the terms to be simplified. Use `Add Ring` first.
4. cannot find a declared ring structure for equality term Same as above is the case of the `ring` tactic.

25.5 Adding a ring structure

Declaring a new ring consists in proving that a ring signature (a carrier set, an equality, and ring operations: `Ring_theory.ring_theory` and `Ring_theory.semi_ring_theory`) satisfies the ring axioms. Semi-rings (rings without $+$ inverse) are also supported. The equality can be either Leibniz equality, or any relation declared as a setoid (see 27.2.2). The definition of ring and semi-rings (see module `Ring_theory`) is:

```
Record ring_theory : Prop := mk_rt {
  Radd_0_l      : forall x, 0 + x == x;
  Radd_sym      : forall x y, x + y == y + x;
  Radd_assoc    : forall x y z, x + (y + z) == (x + y) + z;
  Rmul_1_l      : forall x, 1 * x == x;
  Rmul_sym      : forall x y, x * y == y * x;
  Rmul_assoc    : forall x y z, x * (y * z) == (x * y) * z;
  Rdistr_l      : forall x y z, (x + y) * z == (x * z) + (y * z);
  Rsub_def      : forall x y, x - y == x + -y;
  Ropp_def      : forall x, x + (- x) == 0
}.
```

```
Record semi_ring_theory : Prop := mk_srt {
  SRadd_0_l     : forall n, 0 + n == n;
  SRadd_sym     : forall n m, n + m == m + n ;
  SRadd_assoc   : forall n m p, n + (m + p) == (n + m) + p;
  SRmul_1_l     : forall n, 1*n == n;
  SRmul_0_l     : forall n, 0*n == 0;
  SRmul_sym     : forall n m, n*m == m*n;
  SRmul_assoc   : forall n m p, n*(m*p) == (n*m)*p;
  SRdistr_l     : forall n m p, (n + m)*p == n*p + m*p
}.
```

This implementation of `ring` also features a notion of constant that can be parameterized. This can be used to improve the handling of closed expressions when operations are effective. It consists in introducing a type of *coefficients* and an implementation of the ring operations, and a morphism from the coefficient type to the ring carrier type. The morphism needs not be injective, nor surjective.

As an example, one can consider the real numbers. The set of coefficients could be the rational numbers, upon which the ring operations can be implemented. The fact that there exists a morphism is defined by the following properties:

```
Record ring_morph : Prop := mkmorph {
  morph0       : [c0] == 0;
  morph1       : [c1] == 1;
  morph_add    : forall x y, [x +! y] == [x]+[y];
  morph_sub    : forall x y, [x -! y] == [x]-[y];
  morph_mul    : forall x y, [x *! y] == [x]*[y];
  morph_opp    : forall x, [-!x] == -[x];
  morph_eq     : forall x y, x=?!y = true -> [x] == [y]
}.
```



```

Record semi_morph : Prop := mkRmorph {
  Smorph0 : [c0] == 0;
  Smorph1 : [cI] == 1;
  Smorph_add : forall x y, [x +! y] == [x]+[y];
  Smorph_mul : forall x y, [x *! y] == [x]*[y];
  Smorph_eq : forall x y, x?!=y = true -> [x] == [y]
}.

```

where `c0` and `cI` denote the 0 and 1 of the coefficient set, `+`!, `*`!, `-`! are the implementations of the ring operations, `==` is the equality of the coefficients, `?+!` is an implementation of this equality, and `[x]` is a notation for the image of `x` by the ring morphism.

Since \mathbb{Z} is an initial ring (and \mathbb{N} is an initial semi-ring), it can always be considered as a set of coefficients. There are basically three kinds of (semi-)rings:

abstract rings to be used when operations are not effective. The set of coefficients is \mathbb{Z} (or \mathbb{N} for semi-rings).

computational rings to be used when operations are effective. The set of coefficients is the ring itself. The user only has to provide an implementation for the equality.

customized ring for other cases. The user has to provide the coefficient set and the morphism.

This implementation of ring can also recognize simple power expressions as ring expressions. A power function is specified by the following property:

```

Section POWER.
  Variable Cpow : Set.
  Variable Cp_phi : N -> Cpow.
  Variable rpow : R -> Cpow -> R.

  Record power_theory : Prop := mkpow_th {
    rpow_pow_N : forall r n, req (rpow r (Cp_phi n)) (pow_N rI rmul r n)
  }.

End POWER.

```

The syntax for adding a new ring is `Add Ring name : ring (mod1, ..., mod2)`. The name is not relevant. It is just used for error messages. The term *ring* is a proof that the ring signature satisfies the (semi-)ring axioms. The optional list of modifiers is used to tailor the behavior of the tactic. The following list describes their syntax and effects:

abstract declares the ring as abstract. This is the default.

decidable term declares the ring as computational. The expression *term* is the correctness proof of an equality test `?=!` (which should be evaluable). Its type should be of the form `forall x y, x?!=y = true -> x == y`.

morphism term declares the ring as a customized one. The expression *term* is a proof that there exists a morphism between a set of coefficient and the ring carrier (see `Ring_theory.ring_morph` and `Ring_theory.semi_morph`).

setoid $term_1$ $term_2$ forces the use of given setoid. The expression $term_1$ is a proof that the equality is indeed a setoid (see `Setoid.Setoid_Theory`), and $term_2$ a proof that the ring operations are morphisms (see `Ring_theory.ring_eq_ext` and `Ring_theory.sring_eq_ext`). This modifier needs not be used if the setoid and morphisms have been declared.

constants [\mathcal{L}_{tac}] specifies a tactic expression that, given a term, returns either an object of the coefficient set that is mapped to the expression via the morphism, or returns `InitialRing.NotConstant`. The default behavior is to map only 0 and 1 to their counterpart in the coefficient set. This is generally not desirable for non trivial computational rings.

preprocess [\mathcal{L}_{tac}] specifies a tactic that is applied as a preliminary step for `ring` and `ring_simplify`. It can be used to transform a goal so that it is better recognized. For instance, `S n` can be changed to `plus 1 n`.

postprocess [\mathcal{L}_{tac}] specifies a tactic that is applied as a final step for `ring_simplify`. For instance, it can be used to undo modifications of the preprocessor.

power_tac $term$ [\mathcal{L}_{tac}] allows `ring` and `ring_simplify` to recognize power expressions with a constant positive integer exponent (example: x^2). The term $term$ is a proof that a given power function satisfies the specification of a power function ($term$ has to be a proof of `Ring_theory.power_theory`) and \mathcal{L}_{tac} specifies a tactic expression that, given a term, “abstracts” it into an object of type `N` whose interpretation via `Cp_phi` (the evaluation function of power coefficient) is the original term, or returns `InitialRing.NotConstant` if not a constant coefficient (i.e. \mathcal{L}_{tac} is the inverse function of `Cp_phi`). See files `plugins/setoid_ring/ZArithRing.v` and `plugins/setoid_ring/RealField.v` for examples. By default the tactic does not recognize power expressions as ring expressions.

sign $term$ allows `ring_simplify` to use a minus operation when outputting its normal form, i.e. writing $x - y$ instead of $x + (-y)$. The term $term$ is a proof that a given sign function indicates expressions that are signed ($term$ has to be a proof of `Ring_theory.get_sign`). See `plugins/setoid_ring/InitialRing.v` for examples of sign function.

div $term$ allows `ring` and `ring_simplify` to use monomials with coefficient other than 1 in the rewriting. The term $term$ is a proof that a given division function satisfies the specification of an euclidean division function ($term$ has to be a proof of `Ring_theory.div_theory`). For example, this function is called when trying to rewrite $7x$ by $2x = z$ to tell that $7 = 3 * 2 + 1$. See `plugins/setoid_ring/InitialRing.v` for examples of div function.

Error messages:

1. `bad ring structure` The proof of the ring structure provided is not of the expected type.
2. `bad lemma for decidability of equality` The equality function provided in the case of a computational ring has not the expected type.
3. `ring operation should be declared as a morphism` A setoid associated to the carrier of the ring structure as been found, but the ring operation should be declared as morphism. See 27.2.2.

25.6 How does it work?

The code of `ring` is a good example of tactic written using *reflection*. What is reflection? Basically, it is writing COQ tactics in COQ, rather than in OCAML. From the philosophical point of view, it is using the ability of the Calculus of Constructions to speak and reason about itself. For the `ring` tactic we used COQ as a programming language and also as a proof environment to build a tactic and to prove it correctness.

The interested reader is strongly advised to have a look at the file `Ring_polynom.v`. Here a type for polynomials is defined:

```
Inductive PExpr : Type :=
| PEc : C -> PExpr
| PEX : positive -> PExpr
| PEadd : PExpr -> PExpr -> PExpr
| PEsab : PExpr -> PExpr -> PExpr
| PEMul : PExpr -> PExpr -> PExpr
| PEopp : PExpr -> PExpr
| PEPow : PExpr -> N -> PExpr.
```

Polynomials in normal form are defined as:

```
Inductive Pol : Type :=
| Pc : C -> Pol
| Pinj : positive -> Pol -> Pol
| PX : Pol -> positive -> Pol -> Pol.
```

where `Pinj n P` denotes P in which V_i is replaced by V_{i+n} , and `PX P n Q` denotes $P \otimes V_1^n \oplus Q'$, Q' being Q where V_i is replaced by V_{i+1} .

Variables maps are represented by list of ring elements, and two interpretation functions, one that maps a variables map and a polynomial to an element of the concrete ring, and the second one that does the same for normal forms:

```
Definition PEval : list R -> PExpr -> R := [...].
Definition Pphi_dev : list R -> Pol -> R := [...].
```

A function to normalize polynomials is defined, and the big theorem is its correctness w.r.t interpretation, that is:

```
Definition norm : PExpr -> Pol := [...].
Lemma Pphi_dev_ok :
  forall l pe npe, norm pe = npe -> PEval l pe == Pphi_dev l npe.
```

So now, what is the scheme for a normalization proof? Let p be the polynomial expression that the user wants to normalize. First a little piece of ML code guesses the type of p , the ring theory T to use, an abstract polynomial ap and a variables map v such that p is $\beta\delta\iota$ -equivalent to $(PEval\ v\ ap)$. Then we replace it by $(Pphi_dev\ v\ (norm\ ap))$, using the main correctness theorem and we reduce it to a concrete expression p' , which is the concrete normal form of p . This is summarized in this diagram:

$$\begin{array}{lcl}
 p & \rightarrow_{\beta\delta\iota} & (PEval\ v\ ap) \\
 & & \text{= (by the main correctness theorem)} \\
 p' & \leftarrow_{\beta\delta\iota} & (Pphi_dev\ v\ (norm\ ap))
 \end{array}$$

The user do not see the right part of the diagram. From outside, the tactic behaves like a $\beta\delta\iota$ simplification extended with AC rewriting rules. Basically, the proof is only the application of the main correctness theorem to well-chosen arguments.

25.7 Dealing with fields

The `field` tactic is an extension of the `ring` to deal with rational expression. Given a rational expression $F = 0$. It first reduces the expression F to a common denominator $N/D = 0$ where N and D are two ring expressions. For example, if we take $F = (1 - 1/x)x - x + 1$, this gives $N = (x - 1)x - x^2 + x$ and $D = x$. It then calls `ring` to solve $N = 0$. Note that `field` also generates non-zero conditions for all the denominators it encounters in the reduction. In our example, it generates the condition $x \neq 0$. These conditions appear as one subgoal which is a conjunction if there are several denominators. Non-zero conditions are *always* polynomial expressions. For example when reducing the expression $1/(1 + 1/x)$, two side conditions are generated: $x \neq 0$ and $x + 1 \neq 0$. Factorized expressions are broken since a field is an integral domain, and when the equality test on coefficients is complete w.r.t. the equality of the target field, constants can be proven different from zero automatically.

The tactic must be loaded by `Require Import Field`. New field structures can be declared to the system with the `Add Field` command (see below). The field of real numbers is defined in module `RealField` (in `textttplugins/setoid_ring`). It is exported by module `Rbase`, so that requiring `Rbase` or `Reals` is enough to use the field tactics on real numbers. Rational numbers in canonical form are also declared as a field in module `Qcanon`.

Example:

```
Coq < Require Import Reals.
Coq < Open Scope R_scope.
Coq < Goal forall x, x <> 0 ->
      (1 - 1/x) * x - x + 1 = 0.
1 subgoal

=====
forall x : R, x <> 0 -> (1 - 1 / x) * x - x + 1 = 0
Coq < intros; field; auto.
No more subgoals.

Coq < Goal forall x y, y <> 0 -> y = x -> x/y = 1.
1 subgoal

=====
forall x y : R, y <> 0 -> y = x -> x / y = 1
Coq < intros x y H H1; field [H1]; auto.
No more subgoals.
```

Variants:

1. `field [term1 ... termn]` decides the equality of two terms modulo field operations and rewriting of the equalities defined by `term1 ... termn`. Each of `term1 ... termn` has to be a proof of some equality $m = p$, where m is a monomial (after “abstraction”), p a polynomial and $=$ the corresponding equality of the field structure. Beware that rewriting works with the equality $m = p$ only if p is a polynomial since rewriting is handled by the underlying `ring` tactic.
2. `field_simplify` performs the simplification in the conclusion of the goal, $F_1 = F_2$ becomes $N_1/D_1 = N_2/D_2$. A normalization step (the same as the one for rings) is then applied to N_1, D_1 ,

N_2 and D_2 . This way, polynomials remain in factorized form during the fraction simplifications. This yields smaller expressions when reducing to the same denominator since common factors can be canceled.

3. `field_simplify [term1 ... termn]` performs the simplification in the conclusion of the goal using the equalities defined by `term1 ... termn`.
4. `field_simplify [term1 ... termn] t1 ... tm` performs the simplification in the terms `t1 ... tm` of the conclusion of the goal using the equalities defined by `term1 ... termn`.
5. `field_simplify in H` performs the simplification in the assumption `H`.
6. `field_simplify [term1 ... termn] in H` performs the simplification in the assumption `H` using the equalities defined by `term1 ... termn`.
7. `field_simplify [term1 ... termn] t1 ... tm in H` performs the simplification in the terms `t1 ... tm` of the assumption `H` using the equalities defined by `term1 ... termn`.
8. `field_simplify_eq` performs the simplification in the conclusion of the goal removing the denominator. $F_1 = F_2$ becomes $N_1 D_2 = N_2 D_1$.
9. `field_simplify_eq [term1 ... termn]` performs the simplification in the conclusion of the goal using the equalities defined by `term1 ... termn`.
10. `field_simplify_eq in H` performs the simplification in the assumption `H`.
11. `field_simplify_eq [term1 ... termn] in H` performs the simplification in the assumption `H` using the equalities defined by `term1 ... termn`.

25.8 Adding a new field structure

Declaring a new field consists in proving that a field signature (a carrier set, an equality, and field operations: `Field_theory.field_theory` and `Field_theory.semi_field_theory`) satisfies the field axioms. Semi-fields (fields without $+$ inverse) are also supported. The equality can be either Leibniz equality, or any relation declared as a setoid (see 27.2.2). The definition of fields and semi-fields is:

```
Record field_theory : Prop := mk_field {
  F_R : ring_theory rO rI radd rmul rsub ropp req;
  F_1_neq_0 : ~ 1 == 0;
  Fdiv_def : forall p q, p / q == p * / q;
  Finv_1 : forall p, ~ p == 0 -> / p * p == 1
}.
```

```
Record semi_field_theory : Prop := mk_sfield {
  SF_SR : semi_ring_theory rO rI radd rmul req;
  SF_1_neq_0 : ~ 1 == 0;
  SFdiv_def : forall p q, p / q == p * / q;
  SFinv_1 : forall p, ~ p == 0 -> / p * p == 1
}.
```

The result of the normalization process is a fraction represented by the following type:

```
Record linear : Type := mk_linear {
  num : PExpr C;
  denum : PExpr C;
  condition : list (PExpr C)
}.
```

where `num` and `denum` are the numerator and denominator; `condition` is a list of expressions that have appeared as a denominator during the normalization process. These expressions must be proven different from zero for the correctness of the algorithm.

The syntax for adding a new field is `Add Field name : field (mod1, ..., mod2)`. The `name` is not relevant. It is just used for error messages. `field` is a proof that the field signature satisfies the (semi-)field axioms. The optional list of modifiers is used to tailor the behavior of the tactic. Since field tactics are built upon ring tactics, all modifiers of the `Add Ring` apply. There is only one specific modifier:

completeness term allows the field tactic to prove automatically that the image of non-zero coefficients are mapped to non-zero elements of the field. *term* is a proof of `forall x y, [x] == [y] -> x != y = true`, which is the completeness of equality on coefficients w.r.t. the field equality.

25.9 History of ring

First Samuel Boutin designed the tactic `ACDSimpl`. This tactic did lot of rewriting. But the proofs terms generated by rewriting were too big for COQ's type-checker. Let us see why:

```
Coq < Goal forall x y z:Z, x + 3 + y + y * z = x + 3 + y + z * y.
1 subgoal

=====
forall x y z : Z, x + 3 + y + y * z = x + 3 + y + z * y
Coq < intros; rewrite (Z.mul_comm y z); reflexivity.
Coq < Save toto.
Coq < Print toto.
toto =
fun x y z : Z =>
eq_ind_r (fun z0 : Z => x + 3 + y + z0 = x + 3 + y + z * y) eq_refl
  (Z.mul_comm y z)
  : forall x y z : Z, x + 3 + y + y * z = x + 3 + y + z * y
Argument scopes are [Z_scope Z_scope Z_scope]
```

At each step of rewriting, the whole context is duplicated in the proof term. Then, a tactic that does hundreds of rewriting generates huge proof terms. Since `ACDSimpl` was too slow, Samuel Boutin rewrote it using reflection (see his article in TACS'97 [19]). Later, the stuff was rewritten by Patrick Loiseleur: the new tactic does not any more require `ACDSimpl` to compile and it makes use of $\beta\delta\iota$ -reduction not only to replace the rewriting steps, but also to achieve the interleaving of computation and reasoning (see 25.10). He also wrote a few ML code for the `Add Ring` command, that allow to register new rings dynamically.

Proofs terms generated by `ring` are quite small, they are linear in the number of \oplus and \otimes operations in the normalized terms. Type-checking those terms requires some time because it makes a large use of the conversion rule, but memory requirements are much smaller.

25.10 Discussion

Efficiency is not the only motivation to use reflection here. `ring` also deals with constants, it rewrites for example the expression $34+2*x-x+12$ to the expected result $x+46$. For the tactic `ACDSimpl`, the only constants were 0 and 1. So the expression $34+2*(x-1)+12$ is interpreted as $V_0 \oplus V_1 \otimes (V_2 \ominus 1) \oplus V_3$, with the variables mapping $\{V_0 \mapsto 34; V_1 \mapsto 2; V_2 \mapsto x; V_3 \mapsto 12\}$. Then it is rewritten to $34-x+2*x+12$, very far from the expected result. Here rewriting is not sufficient: you have to do some kind of reduction (some kind of *computation*) to achieve the normalization.

The tactic `ring` is not only faster than a classical one: using reflection, we get for free integration of computation and reasoning that would be very complex to implement in the classic fashion.

Is it the ultimate way to write tactics? The answer is: yes and no. The `ring` tactic uses intensively the conversion rule of CIC, that is replaces proof by computation the most as it is possible. It can be useful in all situations where a classical tactic generates huge proof terms. Symbolic Processing and Tautologies are in that case. But there are also tactics like `auto` or `linear` that do many complex computations, using side-effects and backtracking, and generate a small proof term. Clearly, it would be significantly less efficient to replace them by tactics using reflection.

Another idea suggested by Benjamin Werner: reflection could be used to couple an external tool (a rewriting program or a model checker) with COQ. We define (in COQ) a type of terms, a type of *traces*, and prove a correction theorem that states that *replaying traces* is safe w.r.t some interpretation. Then we let the external tool do every computation (using side-effects, backtracking, exception, or others features that are not available in pure lambda calculus) to produce the trace: now we can check in Coq that the trace has the expected semantic by applying the correction lemma.

Chapter 26

Nsatz: tactics for proving equalities in integral domains

Loïc Pottier

The tactic `nsatz` proves goals of the form

$$\begin{aligned} &\forall X_1, \dots, X_n \in A, \\ &P_1(X_1, \dots, X_n) = Q_1(X_1, \dots, X_n), \dots, P_s(X_1, \dots, X_n) = Q_s(X_1, \dots, X_n) \\ &\vdash P(X_1, \dots, X_n) = Q(X_1, \dots, X_n) \end{aligned}$$

where $P, Q, P_1, Q_1, \dots, P_s, Q_s$ are polynomials and A is an integral domain, i.e. a commutative ring with no zero divisor. For example, A can be \mathbb{R}, \mathbb{Z} , or \mathbb{Q} . Note that the equality $=$ used in these goals can be any setoid equality (see [27.2.2](#)), not only Leibnitz equality.

It also proves formulas

$$\begin{aligned} &\forall X_1, \dots, X_n \in A, \\ &P_1(X_1, \dots, X_n) = Q_1(X_1, \dots, X_n) \wedge \dots \wedge P_s(X_1, \dots, X_n) = Q_s(X_1, \dots, X_n) \\ &\rightarrow P(X_1, \dots, X_n) = Q(X_1, \dots, X_n) \end{aligned}$$

doing automatic introductions.

26.1 Using the basic tactic `nsatz`

Load the `Nsatz` module: `Require Import Nsatz.`
and use the tactic `nsatz`.

26.2 More about `nsatz`

Hilbert's Nullstellensatz theorem shows how to reduce proofs of equalities on polynomials on a commutative ring A with no zero divisor to algebraic computations: it is easy to see that if a polynomial P in $A[X_1, \dots, X_n]$ verifies $cP^r = \sum_{i=1}^s S_i P_i$, with $c \in A$, $c \neq 0$, r a positive integer, and the S_i s in $A[X_1, \dots, X_n]$, then P is zero whenever polynomials P_1, \dots, P_s are zero (the converse is also true when A is an algebraic closed field: the method is complete).

So, proving our initial problem can reduce into finding S_1, \dots, S_s, c and r such that $c(P - Q)^r = \sum_i S_i(P_i - Q_i)$, which will be proved by the tactic `ring`.

This is achieved by the computation of a Groebner basis of the ideal generated by $P_1 - Q_1, \dots, P_s - Q_s$, with an adapted version of the Buchberger algorithm.

This computation is done after a step of *reification*, which is performed using *Type Classes* (see 20).

The `Nsatz` module defines the tactic `nsatz`, which can be used without arguments:

```
nsatz
```

or with the syntax:

```
nsatz with radicalmax:=number%N strategy:=number%Z parameters:=list
of variables variables:=list of variables
where:
```

- `radicalmax` is a bound when for searching r s.t. $c(P - Q)^r = \sum_{i=1..s} S_i(P_i - Q_i)$
- `strategy` gives the order on variables X_1, \dots, X_n and the strategy used in Buchberger algorithm (see [72] for details):
 - `strategy = 0`: reverse lexicographic order and newest s-polynomial.
 - `strategy = 1`: reverse lexicographic order and sugar strategy.
 - `strategy = 2`: pure lexicographic order and newest s-polynomial.
 - `strategy = 3`: pure lexicographic order and sugar strategy.
- `parameters` is the list of variables X_{i_1}, \dots, X_{i_k} among X_1, \dots, X_n which are considered as parameters: computation will be performed with rational fractions in these variables, i.e. polynomials are considered with coefficients in $R(X_{i_1}, \dots, X_{i_k})$. In this case, the coefficient c can be a non constant polynomial in X_{i_1}, \dots, X_{i_k} , and the tactic produces a goal which states that c is not zero.
- `variables` is the list of the variables in the decreasing order in which they will be used in Buchberger algorithm. If `variables = (@nil R)`, then `lvar` is replaced by all the variables which are not in `parameters`.

See file `Nsatz.v` for many examples, specially in geometry.

Chapter 27

Generalized rewriting

Matthieu Sozeau

This chapter presents the extension of several equality related tactics to work over user-defined structures (called setoids) that are equipped with ad-hoc equivalence relations meant to behave as equalities. Actually, the tactics have also been generalized to relations weaker than equivalences (e.g. rewriting systems). The toolbox also extends the automatic rewriting capabilities of the system, allowing the specification of custom strategies for rewriting.

This documentation is adapted from the previous setoid documentation by Claudio Sacerdoti Coen (based on previous work by Clément Renard). The new implementation is a drop-in replacement for the old one,¹ hence most of the documentation still applies.

The work is a complete rewrite of the previous implementation, based on the type class infrastructure. It also improves on and generalizes the previous implementation in several ways:

- **User-extensible algorithm.** The algorithm is separated in two parts: generations of the rewriting constraints (done in ML) and solving of these constraints using type class resolution. As type class resolution is extensible using tactics, this allows users to define general ways to solve morphism constraints.
- **Sub-relations.** An example extension to the base algorithm is the ability to define one relation as a subrelation of another so that morphism declarations on one relation can be used automatically for the other. This is done purely using tactics and type class search.
- **Rewriting under binders.** It is possible to rewrite under binders in the new implementation, if one provides the proper morphisms. Again, most of the work is handled in the tactics.
- **First-class morphisms and signatures.** Signatures and morphisms are ordinary Coq terms, hence they can be manipulated inside Coq, put inside structures and lemmas about them can be proved inside the system. Higher-order morphisms are also allowed.
- **Performance.** The implementation is based on a depth-first search for the first solution to a set of constraints which can be as fast as linear in the size of the term, and the size of the proof term is linear in the size of the original term. Besides, the extensibility allows the user to customize the proof search if necessary.

¹Nicolas Tabareau helped with the gluing.

27.1 Introduction to generalized rewriting

27.1.1 Relations and morphisms

A parametric *relation* R is any term of type $\text{forall } (x_1:T_1) \dots (x_n:T_n), \text{ relation } A$. The expression A , which depends on $x_1 \dots x_n$, is called the *carrier* of the relation and R is said to be a relation over A ; the list x_1, \dots, x_n is the (possibly empty) list of parameters of the relation.

Example 1 (Parametric relation) *It is possible to implement finite sets of elements of type A as unordered list of elements of type A . The function $\text{set_eq} : \text{forall } (A : \text{Type}), \text{ relation } (\text{list } A)$ satisfied by two lists with the same elements is a parametric relation over $(\text{list } A)$ with one parameter A . The type of set_eq is convertible with $\text{forall } (A : \text{Type}), \text{ list } A \rightarrow \text{list } A \rightarrow \text{Prop}$.*

An instance of a parametric relation R with n parameters is any term $(R \ t_1 \dots t_n)$.

Let R be a relation over A with n parameters. A term is a parametric proof of reflexivity for R if it has type $\text{forall } (x_1:T_1) \dots (x_n:T_n), \text{ reflexive } (R \ x_1 \dots x_n)$. Similar definitions are given for parametric proofs of symmetry and transitivity.

Example 2 (Parametric relation (cont.)) *The set_eq relation of the previous example can be proved to be reflexive, symmetric and transitive.*

A parametric unary function f of type $\text{forall } (x_1:T_1) \dots (x_n:T_n), A_1 \rightarrow A_2$ covariantly respects two parametric relation instances R_1 and R_2 if, whenever x, y satisfy $R_1 \ x \ y$, their images $(f \ x)$ and $(f \ y)$ satisfy $R_2 \ (f \ x) \ (f \ y)$. An f that respects its input and output relations will be called a unary covariant *morphism*. We can also say that f is a monotone function with respect to R_1 and R_2 . The sequence x_1, \dots, x_n represents the parameters of the morphism.

Let R_1 and R_2 be two parametric relations. The *signature* of a parametric morphism of type $\text{forall } (x_1:T_1) \dots (x_n:T_n), A_1 \rightarrow A_2$ that covariantly respects two instances I_{R_1} and I_{R_2} of R_1 and R_2 is written $I_{R_1} ++ I_{R_2}$. Notice that the special arrow $++$, which reminds the reader of covariance, is placed between the two relation instances, not between the two carriers. The signature relation instances and morphism will be typed in a context introducing variables for the parameters.

The previous definitions are extended straightforwardly to n -ary morphisms, that are required to be simultaneously monotone on every argument.

Morphisms can also be contravariant in one or more of their arguments. A morphism is contravariant on an argument associated to the relation instance R if it is covariant on the same argument when the inverse relation R^{-1} ($\text{inverse } R$ in Coq) is considered. The special arrow $-->$ is used in signatures for contravariant morphisms.

Functions having arguments related by symmetric relations instances are both covariant and contravariant in those arguments. The special arrow $==>$ is used in signatures for morphisms that are both covariant and contravariant.

An instance of a parametric morphism f with n parameters is any term $f \ t_1 \dots t_n$.

Example 3 (Morphisms) *Continuing the previous example, let $\text{union} : \text{forall } (A : \text{Type}), \text{ list } A \rightarrow \text{list } A \rightarrow \text{list } A$ perform the union of two sets by appending one list to the other. union is a binary morphism parametric over A that respects the relation instance $(\text{set_eq } A)$. The latter condition is proved by showing $\text{forall } (A : \text{Type}) \ (S1 \ S1' \ S2 \ S2' : \text{list } A), \text{ set_eq } A \ S1 \ S1' \rightarrow \text{set_eq } A \ S2 \ S2' \rightarrow \text{set_eq } A \ (\text{union } A \ S1 \ S2) \ (\text{union } A \ S1' \ S2')$.*

The signature of the function `union` is `set_eq A ==> set_eq A ==> set_eq A` for all `A`.

Example 4 (Contravariant morphism) The division function `Rdiv: R -> R -> R` is a morphism of signature `le ++> le --> le` where `le` is the usual order relation over real numbers. Notice that division is covariant in its first argument and contravariant in its second argument.

Leibniz equality is a relation and every function is a morphism that respects Leibniz equality. Unfortunately, Leibniz equality is not always the intended equality for a given structure.

In the next section we will describe the commands to register terms as parametric relations and morphisms. Several tactics that deal with equality in COQ can also work with the registered relations. The exact list of tactic will be given in Sect. 27.2.2. For instance, the tactic `reflexivity` can be used to close a goal `R n n` whenever `R` is an instance of a registered reflexive relation. However, the tactics that replace in a context `C[]` one term with another one related by `R` must verify that `C[]` is a morphism that respects the intended relation. Currently the verification consists in checking whether `C[]` is a syntactic composition of morphism instances that respects some obvious compatibility constraints.

Example 5 (Rewriting) Continuing the previous examples, suppose that the user must prove `set_eq int (union int (union int S1 S2) S2) (f S1 S2)` under the hypothesis `H: set_eq int S2 (@nil int)`. It is possible to use the `rewrite` tactic to replace the first two occurrences of `S2` with `@nil int` in the goal since the context `set_eq int (union int (union int S1 nil) nil) (f S1 S2)`, being a composition of morphisms instances, is a morphism. However the tactic will fail replacing the third occurrence of `S2` unless `f` has also been declared as a morphism.

27.1.2 Adding new relations and morphisms

A parametric relation `Aeq: forall (y1:β1 ...ym:βm), relation (A t1 ...tn) over (A:αi->...αn->Type)` can be declared with the following command:

```
Add Parametric Relation (x1:T1)...(xn:Tn): (A t1 ...tn) (Aeq t'1 ...t'm)
[reflexivity proved by refl]
[symmetry proved by sym]
[transitivity proved by trans]
as id.
```

after having required the `Setoid` module with the `Require Setoid` command.

The identifier `id` gives a unique name to the morphism and it is used by the command to generate fresh names for automatically provided lemmas used internally.

Notice that the carrier and relation parameters may refer to the context of variables introduced at the beginning of the declaration, but the instances need not be made only of variables. Also notice that `A` is *not* required to be a term having the same parameters as `Aeq`, although that is often the case in practice (this departs from the previous implementation).

In case the carrier and relations are not parametric, one can use the command `Add Relation` instead, whose syntax is the same except there is no local context.

The proofs of reflexivity, symmetry and transitivity can be omitted if the relation is not an equivalence relation. The proofs must be instances of the corresponding relation definitions: e.g. the proof of reflexivity must have a type convertible to `reflexive (A t1 ...tn) (Aeq t'1 ...t'n)`. Each proof may refer to the introduced variables as well.

Example 6 (Parametric relation) For *Leibniz equality*, we may declare: *Add Parametric Relation* ($A : \text{Type}$) : A ($@eq\ A$)
[reflexivity proved by @refl_equal A]
 ...

Some tactics (*reflexivity*, *symmetry*, *transitivity*) work only on relations that respect the expected properties. The remaining tactics (*replace*, *rewrite* and derived tactics such as *autorewrite*) do not require any properties over the relation. However, they are able to replace terms with related ones only in contexts that are syntactic compositions of parametric morphism instances declared with the following command.

```
Add Parametric Morphism ( $x_1 : T_1$ ) ... ( $x_k : T_k$ ) : ( $f\ t_1 \dots t_n$ )
with signature sig
as id.
Proof
...
Qed
```

The command declares f as a parametric morphism of signature *sig*. The identifier *id* gives a unique name to the morphism and it is used as the base name of the type class instance definition and as the name of the lemma that proves the well-definedness of the morphism. The parameters of the morphism as well as the signature may refer to the context of variables. The command asks the user to prove interactively that f respects the relations identified from the signature.

Example 7 We start the example by assuming a small theory over homogeneous sets and we declare set equality as a parametric equivalence relation and union of two sets as a parametric morphism.

```
Coq < Require Export Setoid.
Coq < Require Export Relation_Definitions.
Coq < Set Implicit Arguments.
Coq < Parameter set : Type -> Type.
Coq < Parameter empty : forall A, set A.
Coq < Parameter eq_set : forall A, set A -> set A -> Prop.
Coq < Parameter union : forall A, set A -> set A -> set A.
Coq < Axiom eq_set_refl : forall A, reflexive _ (eq_set (A:=A)).
Coq < Axiom eq_set_sym : forall A, symmetric _ (eq_set (A:=A)).
Coq < Axiom eq_set_trans : forall A, transitive _ (eq_set (A:=A)).
Coq < Axiom empty_neutral : forall A (S : set A), eq_set (union S (empty A)) S.
Coq < Axiom union_compat :
  forall (A : Type),
    forall x x' : set A, eq_set x x' ->
      forall y y' : set A, eq_set y y' ->
        eq_set (union x y) (union x' y').
Coq < Add Parametric Relation A : (set A) (@eq_set A)
  reflexivity proved by (eq_set_refl (A:=A))
```

```

    symmetry proved by (eq_set_sym (A:=A))
    transitivity proved by (eq_set_trans (A:=A))
    as eq_set_rel.
Coq < Add Parametric Morphism A : (@union A) with
    signature (@eq_set A) ==> (@eq_set A) ==> (@eq_set A) as union_mor.
Coq < Proof. exact (@union_compat A). Qed.

```

It is possible to reduce the burden of specifying parameters using (maximally inserted) implicit arguments. If A is always set as maximally implicit in the previous example, one can write:

```

Coq < Add Parametric Relation A : (set A) eq_set
    reflexivity proved by eq_set_refl
    symmetry proved by eq_set_sym
    transitivity proved by eq_set_trans
    as eq_set_rel.
Coq < Add Parametric Morphism A : (@union A) with
    signature eq_set ==> eq_set ==> eq_set as union_mor.
Coq < Proof. exact (@union_compat A). Qed.

```

We proceed now by proving a simple lemma performing a rewrite step and then applying reflexivity, as we would do working with Leibniz equality. Both tactic applications are accepted since the required properties over `eq_set` and `union` can be established from the two declarations above.

```

Coq < Goal forall (S: set nat),
    eq_set (union (union S empty) S) (union S S).
Coq < Proof. intros. rewrite empty_neutral. reflexivity. Qed.

```

The tables of relations and morphisms are managed by the type class instance mechanism. The behavior on section close is to generalize the instances by the variables of the section (and possibly hypotheses used in the proofs of instance declarations) but not to export them in the rest of the development for proof search. One can use the `Existing Instance` command to do so outside the section, using the name of the declared morphism suffixed by `_Morphism`, or use the `Global` modifier for the corresponding class instance declaration (see §27.2.1) at definition time. When loading a compiled file or importing a module, all the declarations of this module will be loaded.

27.1.3 Rewriting and non reflexive relations

To replace only one argument of an n -ary morphism it is necessary to prove that all the other arguments are related to themselves by the respective relation instances.

Example 8 *To replace `(union S empty)` with `S` in `(union (union S empty) S) (union S S)` the rewrite tactic must exploit the monotony of `union` (axiom `union_compat` in the previous example). Applying `union_compat` by hand we are left with the goal `eq_set (union S S) (union S S)`.*

When the relations associated to some arguments are not reflexive, the tactic cannot automatically prove the reflexivity goals, that are left to the user.

Setoids whose relation are partial equivalence relations (PER) are useful to deal with partial functions. Let R be a PER. We say that an element x is defined if $R \ x \ x$. A partial function whose domain comprises all the defined elements only is declared as a morphism that respects R . Every time a rewriting step is performed the user must prove that the argument of the morphism is defined.

Example 9 Let $eq0$ be $\text{fun } x \ y \Rightarrow x = y \wedge x \neq 0$ (the smaller PER over non zero elements). Division can be declared as a morphism of signature $eq \Rightarrow eq0 \Rightarrow eq$. Replace x with y in $\text{div } x \ n = \text{div } y \ n$ opens the additional goal $eq0 \ n \ n$ that is equivalent to $n=n \wedge n \neq 0$.

27.1.4 Rewriting and non symmetric relations

When the user works up to relations that are not symmetric, it is no longer the case that any covariant morphism argument is also contravariant. As a result it is no longer possible to replace a term with a related one in every context, since the obtained goal implies the previous one if and only if the replacement has been performed in a contravariant position. In a similar way, replacement in an hypothesis can be performed only if the replaced term occurs in a covariant position.

Example 10 (Covariance and contravariance) Suppose that division over real numbers has been defined as a morphism of signature $Z.\text{div}: Z.\text{lt} \text{ ++> } Z.\text{lt} \text{ --> } Z.\text{lt}$ (i.e. $Z.\text{div}$ is increasing in its first argument, but decreasing on the second one). Let $<$ denotes $Z.\text{lt}$. Under the hypothesis $H: x < y$ we have $k < x / y \rightarrow k < x / x$, but not $k < y / x \rightarrow k < x / x$. Dually, under the same hypothesis $k < x / y \rightarrow k < y / y$ holds, but $k < y / x \rightarrow k < y / y$ does not. Thus, if the current goal is $k < x / x$, it is possible to replace only the second occurrence of x (in contravariant position) with y since the obtained goal must imply the current one. On the contrary, if $k < x / x$ is an hypothesis, it is possible to replace only the first occurrence of x (in covariant position) with y since the current hypothesis must imply the obtained one.

Contrary to the previous implementation, no specific error message will be raised when trying to replace a term that occurs in the wrong position. It will only fail because the rewriting constraints are not satisfiable. However it is possible to use the `at` modifier to specify which occurrences should be rewritten.

As expected, composing morphisms together propagates the variance annotations by switching the variance every time a contravariant position is traversed.

Example 11 Let us continue the previous example and let us consider the goal $x / (x / x) < k$. The first and third occurrences of x are in a contravariant position, while the second one is in covariant position. More in detail, the second occurrence of x occurs covariantly in (x / x) (since division is covariant in its first argument), and thus contravariantly in $x / (x / x)$ (since division is contravariant in its second argument), and finally covariantly in $x / (x / x) < k$ (since $<$, as every transitive relation, is contravariant in its first argument with respect to the relation itself).

27.1.5 Rewriting in ambiguous setoid contexts

One function can respect several different relations and thus it can be declared as a morphism having multiple signatures.

Example 12 Union over homogeneous lists can be given all the following signatures: $eq \Rightarrow eq \Rightarrow eq$ (eq being the equality over ordered lists) $\text{set_eq} \Rightarrow \text{set_eq} \Rightarrow \text{set_eq}$ (set_eq being the equality over unordered lists up to duplicates), $\text{multiset_eq} \Rightarrow \text{multiset_eq} \Rightarrow \text{multiset_eq}$ (multiset_eq being the equality over unordered lists).

To declare multiple signatures for a morphism, repeat the `Add Morphism` command.

When morphisms have multiple signatures it can be the case that a rewrite request is ambiguous, since it is unclear what relations should be used to perform the rewriting. Contrary to the previous

implementation, the tactic will always choose the first possible solution to the set of constraints generated by a rewrite and will not try to find *all* possible solutions to warn the user about.

27.2 Commands and tactics

27.2.1 First class setoids and morphisms

The implementation is based on a first-class representation of properties of relations and morphisms as type classes. That is, the various combinations of properties on relations and morphisms are represented as records and instances of these classes are put in a hint database. For example, the declaration:

```
Add Parametric Relation (x1 : T1) ... (xn : Tk) : (A t1 ... tn) (Aeq t'1 ... t'm)
[reflexivity proved by refl]
[symmetry proved by sym]
[transitivity proved by trans]
as id.
```

is equivalent to an instance declaration:

```
Instance (x1 : T1) ... (xn : Tk) => id : @Equivalence (A t1 ... tn) (Aeq t'1 ... t'm) :=
[Equivalence_Reflexive := refl]
[Equivalence_Symmetric := sym]
[Equivalence_Transitive := trans].
```

The declaration itself amounts to the definition of an object of the record type `Coq.Classes.RelationClasses.Equivalence` and a hint added to the `typeclass_instances` hint database. Morphism declarations are also instances of a type class defined in `Classes.Morphisms`. See the documentation on type classes [20](#) and the theories files in `Classes` for further explanations.

One can inform the rewrite tactic about morphisms and relations just by using the typeclass mechanism to declare them using `Instance` and `Context` vernacular commands. Any object of type `Proper` (the type of morphism declarations) in the local context will also be automatically used by the rewriting tactic to solve constraints.

Other representations of first class setoids and morphisms can also be handled by encoding them as records. In the following example, the projections of the setoid relation and of the morphism function can be registered as parametric relations and morphisms.

Example 13 (First class setoids) `Coq < Require Import Relation_Definitions Setoid.`

```
Coq < Record Setoid: Type :=
{ car:Type;
  eq:car->car->Prop;
  refl: reflexive _ eq;
  sym: symmetric _ eq;
  trans: transitive _ eq
}.

Coq < Add Parametric Relation (s : Setoid) : (@car s) (@eq s)
  reflexivity proved by (refl s)
  symmetry proved by (sym s)
  transitivity proved by (trans s) as eq_rel.
```

```

Coq < Record Morphism (S1 S2:Setoid): Type :=
  { f:car S1 ->car S2;
    compat: forall (x1 x2: car S1), eq S1 x1 x2 -> eq S2 (f x1) (f x2) }.

Coq < Add Parametric Morphism (S1 S2 : Setoid) (M : Morphism S1 S2) :
  (@f S1 S2 M) with signature (@eq S1 ==> @eq S2) as apply_mor.

Coq < Proof. apply (compat S1 S2 M). Qed.

Coq < Lemma test: forall (S1 S2:Setoid) (m: Morphism S1 S2)
  (x y: car S1), eq S1 x y -> eq S2 (f _ _ m x) (f _ _ m y).

Coq < Proof. intros. rewrite H. reflexivity. Qed.

```

27.2.2 Tactics enabled on user provided relations

The following tactics, all prefixed by `setoid_`, deal with arbitrary registered relations and morphisms. Moreover, all the corresponding unprefixd tactics (i.e. `reflexivity`, `symmetry`, `transitivity`, `replace`, `rewrite`) have been extended to fall back to their prefixed counterparts when the relation involved is not Leibniz equality. Notice, however, that using the prefixed tactics it is possible to pass additional arguments such as `using relation`.

```

setoid_reflexivity
setoid_symmetry [in ident]
setoid_transitivity
setoid_rewrite [orientation] term [at occs] [in ident]
setoid_replace term with term [in ident] [using relation term] [by tactic]

```

The `using relation` arguments cannot be passed to the unprefixd form. The latter argument tells the tactic what parametric relation should be used to replace the first tactic argument with the second one. If omitted, it defaults to the `DefaultRelation` instance on the type of the objects. By default, it means the most recent `Equivalence` instance in the environment, but it can be customized by declaring new `DefaultRelation` instances. As Leibniz equality is a declared equivalence, it will fall back to it if no other relation is declared on a given type.

Every derived tactic that is based on the unprefixd forms of the tactics considered above will also work up to user defined relations. For instance, it is possible to register hints for `autorewrite` that are not proof of Leibniz equalities. In particular it is possible to exploit `autorewrite` to simulate normalization in a term rewriting system up to user defined equalities.

27.2.3 Printing relations and morphisms

The `Print Instances` command can be used to show the list of currently registered `Reflexive` (using `Print Instances Reflexive`), `Symmetric` or `Transitive` relations, `Equivalences`, `PreOrders`, `PERs`, and `Morphisms` (implemented as `Proper` instances). When the rewriting tactics refuse to replace a term in a context because the latter is not a composition of morphisms, the `Print Instances` commands can be useful to understand what additional morphisms should be registered.

27.2.4 Deprecated syntax and backward incompatibilities

Due to backward compatibility reasons, the following syntax for the declaration of setoids and morphisms is also accepted.

```
Add Setoid A Aeq ST as ident
```

where *Aeq* is a congruence relation without parameters, *A* is its carrier and *ST* is an object of type `(Setoid_Theory A Aeq)` (i.e. a record packing together the reflexivity, symmetry and transitivity lemmas). Notice that the syntax is not completely backward compatible since the identifier was not required.

```
Add Morphism f:ident.
```

```
Proof.
```

```
...
```

```
Qed.
```

The latter command also is restricted to the declaration of morphisms without parameters. It is not fully backward compatible since the property the user is asked to prove is slightly different: for *n*-ary morphisms the hypotheses of the property are permuted; moreover, when the morphism returns a proposition, the property is now stated using a bi-implication in place of a simple implication. In practice, porting an old development to the new semantics is usually quite simple.

Notice that several limitations of the old implementation have been lifted. In particular, it is now possible to declare several relations with the same carrier and several signatures for the same morphism. Moreover, it is now also possible to declare several morphisms having the same signature. Finally, the replace and rewrite tactics can be used to replace terms in contexts that were refused by the old implementation. As discussed in the next section, the semantics of the new `setoid_rewrite` command differs slightly from the old one and `rewrite`.

27.3 Extensions

27.3.1 Rewriting under binders

Warning: Due to compatibility issues, this feature is enabled only when calling the `setoid_rewrite` tactics directly and not `rewrite`.

To be able to rewrite under binding constructs, one must declare morphisms with respect to pointwise (setoid) equivalence of functions. Example of such morphisms are the standard `all` and `ex` combinators for universal and existential quantification respectively. They are declared as morphisms in the `Classes.Morphisms_Prop` module. For example, to declare that universal quantification is a morphism for logical equivalence:

```
Coq < Instance all_iff_morphism (A : Type) :
    Proper (pointwise_relation A iff ==> iff) (@all A).

Coq < Proof. simpl_relation.
1 subgoal

  A : Type
  =====
  Proper (pointwise_relation A iff ==> iff) (all (A:=A))
1 subgoal

  A : Type
  x, y : A -> Prop
  H : pointwise_relation A iff x y
```

```
=====
all x <-> all y
```

One then has to show that if two predicates are equivalent at every point, their universal quantifications are equivalent. Once we have declared such a morphism, it will be used by the `setoid_rewrite` tactic each time we try to rewrite under an `all` application (products in `Prop` are implicitly translated to such applications).

Indeed, when rewriting under a lambda, binding variable x , say from $P\ x$ to $Q\ x$ using the relation `iff`, the tactic will generate a proof of `pointwise_relation A iff (fun x => P x) (fun x => Q x)` from the proof of `iff (P x) (Q x)` and a constraint of the form `Proper (pointwise_relation A iff ==> ?)` `m` will be generated for the surrounding morphism `m`.

Hence, one can add higher-order combinators as morphisms by providing signatures using pointwise extension for the relations on the functional arguments (or whatever subrelation of the pointwise extension). For example, one could declare the `map` combinator on lists as a morphism:

```
Coq < Instance map_morphism `{Equivalence A eqA, Equivalence B eqB} :
  Proper ((eqA ==> eqB) ==> list_equiv eqA ==> list_equiv eqB) (@map A B).
```

where `list_equiv` implements an equivalence on lists parameterized by an equivalence on the elements.

Note that when one does rewriting with a lemma under a binder using `setoid_rewrite`, the application of the lemma may capture the bound variable, as the semantics are different from `rewrite` where the lemma is first matched on the whole term. With the new `setoid_rewrite`, matching is done on each subterm separately and in its local environment, and all matches are rewritten *simultaneously* by default. The semantics of the previous `setoid_rewrite` implementation can almost be recovered using the `at 1` modifier.

27.3.2 Sub-relations

Sub-relations can be used to specify that one relation is included in another, so that morphisms signatures for one can be used for the other. If a signature mentions a relation R on the left of an arrow `==>`, then the signature also applies for any relation S that is smaller than R , and the inverse applies on the right of an arrow. One can then declare only a few morphisms instances that generate the complete set of signatures for a particular constant. By default, the only declared subrelation is `iff`, which is a subrelation of `impl` and `inverse impl` (the dual of implication). That's why we can declare only two morphisms for conjunction: `Proper (impl ==> impl ==> impl)` and `Proper (iff ==> iff ==> iff)` and. This is sufficient to satisfy any rewriting constraints arising from a rewrite using `iff`, `impl` or `inverse impl` through `and`.

Sub-relations are implemented in `Classes.Morphisms` and are a prime example of a mostly user-space extension of the algorithm.

27.3.3 Constant unfolding

The resolution tactic is based on type classes and hence regards user-defined constants as transparent by default. This may slow down the resolution due to a lot of unifications (all the declared `Proper` instances are tried at each node of the search tree). To speed it up, declare your constant as rigid for proof search using the command `Typeclasses Opaque` (see §20.6.7).

27.4 Strategies for rewriting

27.4.1 Definitions

The generalized rewriting tactic is based on a set of strategies that can be combined to obtain custom rewriting procedures. Its set of strategies is based on Elan's rewriting strategies [102]. Rewriting strategies are applied using the tactic `rewrite_strat s` where s is a strategy expression. Strategies are defined inductively as described by the following grammar:

$s, t, u ::=$	(s)	strategy
	c	lemma
	$<- c$	lemma, right-to-left
	<code>fail</code>	failure
	<code>id</code>	identity
	<code>refl</code>	reflexivity
	<code>progress s</code>	progress
	<code>try s</code>	failure catch
	$s ; u$	composition
	<code>choice s t</code>	left-biased choice
	<code>repeat s</code>	iteration (+)
	<code>any s</code>	iteration (*)
	<code>subterm s</code>	one subterm
	<code>subterms s</code>	all subterms
	<code>innermost s</code>	innermost first
	<code>outermost s</code>	outermost first
	<code>bottomup s</code>	bottom-up
	<code>topdown s</code>	top-down
	<code>hints hintdb</code>	apply hint
	<code>terms c...c</code>	any of the terms
	<code>eval redexpr</code>	apply reduction
	<code>fold c</code>	fold expression

Actually a few of these are defined in term of the others using a primitive fixpoint operator:

```

try s      = choice s id
any s      = fix u. try (s ; u)
repeat s   = s ; any s
bottomup s = fix bu. (choice (progress (subterms bu)) s) ; try bu
topdown s  = fix td. (choice s (progress (subterms td))) ; try td
innermost s = fix i. (choice (subterm i) s)
outermost s = fix o. (choice s (subterm o))

```

The basic control strategy semantics are straightforward: strategies are applied to subterms of the term to rewrite, starting from the root of the term. The lemma strategies unify the left-hand-side of the lemma with the current subterm and on success rewrite it to the right-hand-side. Composition can be used to continue rewriting on the current subterm. The fail strategy always fails while the identity strategy succeeds without making progress. The reflexivity strategy succeeds, making progress using a reflexivity proof of rewriting. Progress tests progress of the argument strategy and fails if no progress

was made, while `try` always succeeds, catching failures. Choice is left-biased: it will launch the first strategy and fall back on the second one in case of failure. One can iterate a strategy at least 1 time using `repeat` and at least 0 times using `any`.

The `subterm` and `subterms` strategies apply their argument strategy *s* to respectively one or all subterms of the current term under consideration, left-to-right. `subterm` stops at the first subterm for which *s* made progress. The composite strategies `innermost` and `outermost` perform a single innermost or outermost rewrite using their argument strategy. Their counterparts `bottomup` and `topdown` perform as many rewritings as possible, starting from the bottom or the top of the term.

Hint databases created for `autorewrite` can also be used by `rewrite_strat` using the `hints` strategy that applies any of the lemmas at the current subterm. The `terms` strategy takes the lemma names directly as arguments. The `eval` strategy expects a reduction expression (see §8.7) and succeeds if it reduces the subterm under consideration. The `fold` strategy takes a term *c* and tries to *unify* it to the current subterm, converting it to *c* on success, it is stronger than the tactic `fold`.

27.4.2 Usage

`rewrite_strat s [in ident]:`

Rewrite using the strategy *s* in hypothesis *ident* or the conclusion.

Error messages:

1. Nothing to rewrite. If the strategy failed.
2. No progress made. If the strategy succeeded but made no progress.
3. Unable to satisfy the rewriting constraints. If the strategy succeeded and made progress but the corresponding rewriting constraints are not satisfied.

The `setoid_rewrite c` tactic is basically equivalent to `rewrite_strat (outermost c)`.

Chapter 28

Asynchronous and Parallel Proof Processing

Enrico Tassi

This chapter explains how proofs can be asynchronously processed by Coq. This feature improves the reactivity of the system when used in interactive mode via CoqIDE. In addition, it allows Coq to take advantage of parallel hardware when used as a batch compiler by decoupling the checking of statements and definitions from the construction and checking of proofs objects.

This feature is designed to help dealing with huge libraries of theorems characterized by long proofs. In the current state, it may not be beneficial on small sets of short files.

This feature has some technical limitations that may make it unsuitable for some use cases.

For example, in interactive mode, some errors coming from the kernel of Coq are signaled late. The type of errors belonging to this category are universe inconsistencies.

At the time of writing, only opaque proofs (ending with `Qed` or `Admitted`) can be processed asynchronously.

Finally, asynchronous processing is disabled when running CoqIDE in Windows. The current implementation of the feature is not stable on Windows. It can be enabled, as described below at [28.3](#), though doing so is not recommended.

28.1 Proof annotations

To process a proof asynchronously Coq needs to know the precise statement of the theorem without looking at the proof. This requires some annotations if the theorem is proved inside a `Section` (see [Section 2.4](#)).

When a section ends, Coq looks at the proof object to decide which section variables are actually used and hence have to be quantified in the statement of the theorem. To avoid making the construction of proofs mandatory when ending a section, one can start each proof with the `Proof using` command ([Section 7.1.5](#)) that declares which section variables the theorem uses.

The presence of `Proof using` is needed to process proofs asynchronously in interactive mode.

It is not strictly mandatory in batch mode if it is not the first time the file is compiled and if the file itself did not change. When the proof does not begin with `Proof using`, the system records in an auxiliary file, produced along with the `.vo` file, the list of section variables used.

Automatic suggestion of proof annotations

The command `Set Suggest Proof Using` makes Coq suggest, when a `Qed` command is processed, a correct proof annotation. It is up to the user to modify the proof script accordingly.

28.2 Proof blocks and error resilience

Coq 8.6 introduces a mechanism for error resiliency: in interactive mode Coq is able to completely check a document containing errors instead of bailing out at the first failure.

Two kind of errors are supported: errors occurring in vernacular commands and errors occurring in proofs.

To properly recover from a failing tactic, Coq needs to recognize the structure of the proof in order to confine the error to a sub proof. Proof block detection is performed by looking at the syntax of the proof script (i.e. also looking at indentation). Coq comes with four kind of proof blocks, and an ML API to add new ones.

curly blocks are delimited by `{` and `}`, see [7](#)

par blocks are atomic, i.e. just one tactic introduced by the `par :` goal selector

indent blocks end with a tactic indented less than the previous one

bullet blocks are delimited by two equal bullet signs at the same indentation level

28.2.1 Caveats

When a vernacular command fails the subsequent error messages may be bogus, i.e. caused by the first error. Error resiliency for vernacular commands can be switched off passing `-async-proofs-command-error-resilience off` to CoqIDE.

An incorrect proof block detection can result into an incorrect error recovery and hence in bogus errors. Proof block detection cannot be precise for bullets or any other non well parenthesized proof structure. Error resiliency can be turned off or selectively activated for any set of block kind passing to CoqIDE one of the following options: `-async-proofs-tactic-error-resilience off`, `-async-proofs-tactic-error-resilience all`, `-async-proofs-tactic-error-resilience blocktype1, ..., blocktypen`. Valid proof block types are: “curly”, “par”, “indent”, “bullet”.

28.3 Interactive mode

At the time of writing the only user interface supporting asynchronous proof processing is CoqIDE.

When CoqIDE is started, two Coq processes are created. The master one follows the user, giving feedback as soon as possible by skipping proofs, which are delegated to the worker process. The worker process, whose state can be seen by clicking on the button in the lower right corner of the main CoqIDE window, asynchronously processes the proofs. If a proof contains an error, it is reported in red in the label of the very same button, that can also be used to see the list of errors and jump to the corresponding line.

If a proof is processed asynchronously the corresponding `Qed` command is colored using a lighter color than usual. This signals that the proof has been delegated to a worker process (or will be processed lazily if the `-async-proofs lazy` option is used). Once finished, the worker process will provide

the proof object, but this will not be automatically checked by the kernel of the main process. To force the kernel to check all the proof objects, one has to click the button with the gears. Only then are all the universe constraints checked.

Caveats

The number of worker processes can be increased by passing CoqIDE the `-async-proofs-j n` flag. Note that the memory consumption increases too, since each worker requires the same amount of memory as the master process. Also note that increasing the number of workers may reduce the reactivity of the master process to user commands.

To disable this feature, one can pass the `-async-proofs off` flag to CoqIDE. Conversely, on Windows, where the feature is disabled by default, pass the `-async-proofs on` flag to enable it.

Proofs that are known to take little time to process are not delegated to a worker process. The threshold can be configured with `-async-proofs-delegation-threshold`. Default is 0.03 seconds.

28.4 Batch mode

When Coq is used as a batch compiler by running `coqc` or `coqtop -compile`, it produces a `.vo` file for each `.v` file. A `.vo` file contains, among other things, theorems statements and proofs. Hence to produce a `.vo` Coq need to process all the proofs of the `.v` file.

The asynchronous processing of proofs can decouple the generation of a compiled file (like the `.vo` one) that can be loaded by `Require` from the generation and checking of the proof objects. The `-quick` flag can be passed to `coqc` or `coqtop` to produce, quickly, `.vio` files. Alternatively, when using a `Makefile` produced by `coq_makefile`, the `quick` target can be used to compile all files using the `-quick` flag.

A `.vio` file can be loaded using `Require` exactly as a `.vo` file but proofs will not be available (the `Print` command produces an error). Moreover, some universe constraints might be missing, so universes inconsistencies might go unnoticed. A `.vio` file does not contain proof objects, but proof tasks, i.e. what a worker process can transform into a proof object.

Compiling a set of files with the `-quick` flag allows one to work, interactively, on any file without waiting for all the proofs to be checked.

When working interactively, one can fully check all the `.v` files by running `coqc` as usual.

Alternatively one can turn each `.vio` into the corresponding `.vo`. All `.vio` files can be processed in parallel, hence this alternative might be faster. The command `coqtop -schedule-vio2vo 2 a b c` can be used to obtain a good scheduling for 2 workers to produce `a.vo`, `b.vo`, and `c.vo`. When using a `Makefile` produced by `coq_makefile`, the `vio2vo` target can be used for that purpose. Variable `J` should be set to the number of workers, e.g. `make vio2vo J=2`. The only caveat is that, while the `.vo` files obtained from `.vio` files are complete (they contain all proof terms and universe constraints), the satisfiability of all universe constraints has not been checked globally (they are checked to be consistent for every single proof). Constraints will be checked when these `.vo` files are (recursively) loaded with `Require`.

There is an extra, possibly even faster, alternative: just check the proof tasks stored in `.vio` files without producing the `.vo` files. This is possibly faster because all the proof tasks are independent, hence one can further partition the job to be done between workers. The `coqtop -schedule-vio-checking 6 a b c` command can be used to obtain a good scheduling for 6 workers to check all the proof tasks of `a.vio`, `b.vio`, and `c.vio`. Auxiliary files are used to predict how long a proof task will take, assuming it will take the same amount of time it took last time. When

using a Makefile produced by `coq_makefile`, the `checkproofs` target can be used to check all `.vio` files. Variable `J` should be set to the number of workers, e.g. `make checkproofs J=6`. As when converting `.vio` files to `.vo` files, universe constraints are not checked to be globally consistent. Hence this compilation mode is only useful for quick regression testing and on developments not making heavy use of the *Type* hierarchy.

28.5 Limiting the number of parallel workers

Many Coq processes may run on the same computer, and each of them may start many additional worker processes. The `coqworkmgr` utility lets one limit the number of workers, globally.

The utility accepts the `-j` argument to specify the maximum number of workers (defaults to 2). `coqworkmgr` automatically starts in the background and prints an environment variable assignment like `COQWORKMGR_SOCKET=localhost:45634`. The user must set this variable in all the shells from which Coq processes will be started. If one uses just one terminal running the bash shell, then `export `coqworkmgr -j 4`` will do the job.

After that, all Coq processes, e.g. `coqide` and `coqc`, will honor the limit, globally.

Chapter 29

Polymorphic Universes

Matthieu Sozeau

29.1 General Presentation

The status of Universe Polymorphism is experimental.

This section describes the universe polymorphic extension of Coq. Universe polymorphism makes it possible to write generic definitions making use of universes and reuse them at different and sometimes incompatible universe levels.

A standard example of the difference between universe *polymorphic* and *monomorphic* definitions is given by the identity function:

```
Coq < Definition identity {A : Type} (a : A) := a.
```

By default, constant declarations are monomorphic, hence the identity function declares a global universe (say `Top.1`) for its domain. Subsequently, if we try to self-apply the identity, we will get an error:

```
Coq < Fail Definition selfid := identity (@identity).  
The command has indeed failed with message:  
The term "@identity" has type "forall A : Type@{Top.1}, A -> A"  
while it is expected to have type "?A"  
(unable to find a well-typed instantiation for  
"?A": cannot ensure that "Type@{Top.1+1}" is a subtype of  
"Type@{Top.1}").
```

Indeed, the global level `Top.1` would have to be strictly smaller than itself for this self-application to typecheck, as the type of `(@identity)` is `forall (A : Type@Top.1), A -> A` whose type is itself `Type@Top.1+1`.

A universe polymorphic identity function binds its domain universe level at the definition level instead of making it global.

```

Coq < Polymorphic Definition pidentity {A : Type} (a : A) := a.
pidentity is defined

Coq < About pidentity.
pidentity@{Top.2} : forall A : Type@{Top.2}, A -> A
(* Top.2 != *)
pidentity is universe polymorphic
Argument A is implicit and maximally inserted
Argument scopes are [type_scope _]
pidentity is transparent
Expands to: Constant Top.pidentity

```

It is then possible to reuse the constant at different levels, like so:

```

Coq < Definition selfpid := pidentity (@pidentity).
selfpid is defined

```

Of course, the two instances of `pidentity` in this definition are different. This can be seen when `Set Printing Universes` is on:

```

Coq < Print selfpid.
selfpid =
pidentity@{Top.3} (@pidentity@{Top.4})
      : forall A : Type@{Top.4}, A -> A
(* Top.3 Top.4 != Top.4 < Top.3
      *)
Argument scopes are [type_scope _]

```

Now `pidentity` is used at two different levels: at the head of the application it is instantiated at `Top.3` while in the argument position it is instantiated at `Top.4`. This definition is only valid as long as `Top.4` is strictly smaller than `Top.3`, as show by the constraints. Note that this definition is monomorphic (not universe polymorphic), so the two universes (in this case `Top.3` and `Top.4`) are actually global levels.

Inductive types can also be declared universes polymorphic on universes appearing in their parameters or fields. A typical example is given by monoids:

```

Coq < Polymorphic Record Monoid := { mon_car :> Type; mon_unit : mon_car;
      mon_op : mon_car -> mon_car -> mon_car }.
Monoid is defined
mon_car is defined
mon_unit is defined
mon_op is defined

Coq < Print Monoid.
Polymorphic NonCumulative Record Monoid : Type@{Top.6+1}
:= Build_Monoid
{ mon_car : Type@{Top.6};
  mon_unit : mon_car;
  mon_op : mon_car -> mon_car -> mon_car }
For Build_Monoid: Argument scopes are [type_scope _ function_scope]

```

The `Monoid`'s carrier universe is polymorphic, hence it is possible to instantiate it for example with `Monoid` itself. First we build the trivial unit monoid in `Set`:

```
Coq < Definition unit_monoid : Monoid :=
  {| mon_car := unit; mon_unit := tt; mon_op x y := tt |}.
unit_monoid is defined
```

From this we can build a definition for the monoid of Set-monoids (where multiplication would be given by the product of monoids).

```
Coq < Polymorphic Definition monoid_monoid : Monoid.
Coq <   refine (@Build_Monoid Monoid unit_monoid (fun x y => x)).
Coq < Defined.
Coq < Print monoid_monoid.
Polymorphic monoid_monoid@{Top.10} =
{|
mon_car := Monoid@{Set};
mon_unit := unit_monoid;
mon_op := fun x _ : Monoid@{Set} => x |}
      : Monoid@{Top.10}
(* Top.10 |= Set < Top.10
   *)
monoid_monoid is universe polymorphic
```

As one can see from the constraints, this monoid is “large”, it lives in a universe strictly higher than Set.

29.2 Polymorphic, Monomorphic

As shown in the examples, polymorphic definitions and inductives can be declared using the Polymorphic prefix. There also exists an option `Set Universe Polymorphism` which will implicitly prepend it to any definition of the user. In that case, to make a definition producing global universe constraints, one can use the Monomorphic prefix. Many other commands support the Polymorphic flag, including:

- Lemma, Axiom, and all the other “definition” keywords support polymorphism.
- Variables, Context, Universe and Constraint in a section support polymorphism. This means that the universe variables (and associated constraints) are discharged polymorphically over definitions that use them. In other words, two definitions in the section sharing a common variable will both get parameterized by the universes produced by the variable declaration. This is in contrast to a “monomorphic” variable which introduces global universes and constraints, making the two definitions depend on the *same* global universes associated to the variable.
- Hint {Resolve, Rewrite} will use the auto/rewrite hint polymorphically, not at a single instance.

29.3 Cumulative, NonCumulative

Polymorphic inductive types, coinductive types, variants and records can be declared cumulative using the Cumulative. Alternatively, there is an option `Set Polymorphic Inductive`

Cumulativity which when set, makes all subsequent *polymorphic* inductive definitions cumulative. When set, inductive types and the like can be enforced to be *non-cumulative* using the `NonCumulative` prefix. Consider the examples below.

```
Coq < Polymorphic Cumulative Inductive list {A : Type} :=
  | nil : list
  | cons : A -> list -> list.

Coq < Print list.
Polymorphic Cumulative Inductive
list@{Top.13} (A : Type@{Top.13}) : Type@{max(Top.13, Set)} :=
  nil : list@{Top.13} | cons : A -> list@{Top.13} -> list@{Top.13}
(* Top.13 |=

  ~@{Top.13} <= ~@{Top.15} iff
  *)
For list: Argument A is implicit and maximally inserted
For nil: Argument A is implicit and maximally inserted
For cons: Argument A is implicit and maximally inserted
For list: Argument scope is [type_scope]
For nil: Argument scope is [type_scope]
For cons: Argument scopes are [type_scope _ _]
```

When printing `list`, the part of the output of the form $\sim@{i} <= \sim@{j}$ iff indicates the universe constraints in order to have the subtyping $E[\Gamma] \vdash \text{list}@{i} A \leq_{\beta\delta\iota\zeta\eta} \text{list}@{j} B$ (for fully applied instances of `list`) whenever $E[\Gamma] \vdash A =_{\beta\delta\iota\zeta\eta} B$. In the case of `list` there is no constraint! This also means that any two instances of `list` are convertible: $E[\Gamma] \vdash \text{list}@{i} A =_{\beta\delta\iota\zeta\eta} \text{list}@{j} B$ whenever $E[\Gamma] \vdash A =_{\beta\delta\iota\zeta\eta} B$ and furthermore their corresponding (when fully applied to convertible arguments) constructors. See Chapter 4 for more details on convertibility and subtyping. The following is an example of a record with non-trivial subtyping relation:

```
Coq < Polymorphic Cumulative Record packType := {pk : Type}.

Coq < Print packType.
Polymorphic Cumulative Record packType : Type@{Top.20+1}
:= Build_packType
{ pk : Type@{Top.20} }
(* Top.20 |=

  ~@{Top.20} <= ~@{Top.21} iff
  Top.20 <= Top.21
  *)
For Build_packType: Argument scope is [type_scope]
```

Notice that as expected, `packType@{i}` and `packType@{j}` are convertible if and only if $i = j$.

Cumulative inductive types, coninductive types, variants and records only make sense when they are universe polymorphic. Therefore, an error is issued whenever the user uses the `Cumulative` or `NonCumulative` prefix in a monomorphic context. Notice that this is not the case for the option `Set Polymorphic Inductive Cumulativity`. That is, this option, when set, makes all subsequent *polymorphic* inductive declarations cumulative (unless, of course the `NonCumulative` prefix is used) but has no effect on *monomorphic* inductive declarations. Consider the following examples.

```
Coq < Monomorphic Cumulative Inductive Unit := unit.
Toplevel input, characters 0-46:
> Monomorphic Cumulative Inductive Unit := unit.
> ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
```

An example of a proof using cumulativity

Coq Reference Manual, V8.7.1, December 16, 2017

```
Coq < Defined.
funext_down is defined
```

29.4 Global and local universes

Each universe is declared in a global or local environment before it can be used. To ensure compatibility, every *global* universe is set to be strictly greater than **Set** when it is introduced, while every *local* (i.e. polymorphically quantified) universe is introduced as greater or equal to **Set**.

29.5 Conversion and unification

The semantics of conversion and unification have to be modified a little to account for the new universe instance arguments to polymorphic references. The semantics respect the fact that definitions are transparent, so indistinguishable from their bodies during conversion.

This is accomplished by changing one rule of unification, the first-order approximation rule, which applies when two applicative terms with the same head are compared. It tries to short-cut unfolding by comparing the arguments directly. In case the constant is universe polymorphic, we allow this rule to fire only when unifying the universes results in instantiating a so-called flexible universe variables (not given by the user). Similarly for conversion, if such an equation of applicative terms fail due to a universe comparison not being satisfied, the terms are unfolded. This change implies that conversion and unification can have different unfolding behaviors on the same development with universe polymorphism switched on or off.

29.6 Minimization

Universe polymorphism with cumulativity tends to generate many useless inclusion constraints in general. Typically at each application of a polymorphic constant f , if an argument has expected type $\text{Type}@{i}$ and is given a term of type $\text{Type}@{j}$, a $j \leq i$ constraint will be generated. It is however often the case that an equation $j = i$ would be more appropriate, when f 's universes are fresh for example. Consider the following example:

```
Coq < Definition id0 := @pidentity nat 0.
id0 is defined

Coq < Print id0.
id0@{} = pidentity@{Set} 0
      : nat
id0 is universe polymorphic
```

This definition is elaborated by minimizing the universe of `id` to level **Set** while the more general definition would keep the fresh level i generated at the application of `id` and a constraint that $\text{Set} \leq i$. This minimization process is applied only to fresh universe variables. It simply adds an equation between the variable and its lower bound if it is an atomic universe (i.e. not an algebraic `max()` universe).

The option `Unset Universe Minimization ToSet` disallows minimization to the sort **Set** and only collapses floating universes between themselves.

29.7 Explicit Universes

The syntax has been extended to allow users to explicitly bind names to universes and explicitly instantiate polymorphic definitions.

29.7.1 Universe *ident*.

In the monomorphic case, this command declares a new global universe named *ident*. It supports the polymorphic flag only in sections, meaning the universe quantification will be discharged on each section definition independently. One cannot mix polymorphic and monomorphic declarations in the same section.

29.7.2 Constraint *ident ord ident*.

This command declares a new constraint between named universes. The order relation can be one of $<$, \leq or $=$. If consistent, the constraint is then enforced in the global environment. Like `Universe`, it can be used with the `Polymorphic` prefix in sections only to declare constraints discharged at section closing time. One cannot declare a global constraint on polymorphic universes.

Error messages:

1. Undeclared universe *ident*.
2. Universe inconsistency

29.7.3 Polymorphic definitions

For polymorphic definitions, the declaration of (all) universe levels introduced by a definition uses the following syntax:

```
Coq < Polymorphic Definition le@{i j} (A : Type@{i}) : Type@{j} := A.
Coq < Print le.
le@{i j} =
fun A : Type@{i} => A
  : Type@{i} -> Type@{j}
(* i j | = i <= j
   *)
le is universe polymorphic
Argument scope is [type_scope]
```

During refinement we find that j must be larger or equal than i , as we are using $A : \text{Type}@i \leq \text{Type}@j$, hence the generated constraint. At the end of a definition or proof, we check that the only remaining universes are the ones declared. In the term and in general in proof mode, introduced universe names can be referred to in terms. Note that local universe names shadow global universe names. During a proof, one can use `Show Universes` to display the current context of universes.

Definitions can also be instantiated explicitly, giving their full instance:

```
Coq < Check (pidentity@{Set}).
pidentity@{Set}
  : ?A -> ?A
where
```

```
?A : [ |- Set]
Coq < Universes k l.
Coq < Check (le@{k l}).
le@{k l}
  : Type@{k} -> Type@{l}
(*   |= k <= l
    *)
```

User-named universes and the anonymous universe implicitly attached to an explicit *Type* are considered rigid for unification and are never minimized. Flexible anonymous universes can be produced with an underscore or by omitting the annotation to a polymorphic definition.

```
Coq <   Check (fun x => x) : Type -> Type.
(fun x : Type@{Top.48} => x) : Type@{Top.48} -> Type@{Top.49}
  : Type@{Top.48} -> Type@{Top.49}
(* Top.48 Top.49 |= Top.48 <= Top.49
   *)

Coq <   Check (fun x => x) : Type -> Type@{__}.
(fun x : Type@{Top.50} => x) : Type@{Top.50} -> Type@{Top.50}
  : Type@{Top.50} -> Type@{Top.50}
(* Top.50 |=   *)

Coq <   Check le@{k __}.
le@{k k}
  : Type@{k} -> Type@{k}

Coq <   Check le.
le@{Top.53 Top.53}
  : Type@{Top.53} -> Type@{Top.53}
(* Top.53 |=   *)
```

29.7.4 Unset Strict Universe Declaration.

The command `Unset Strict Universe Declaration` allows one to freely use identifiers for universes without declaring them first, with the semantics that the first use declares it. In this mode, the universe names are not associated with the definition or proof once it has been defined. This is meant mainly for debugging purposes.

Chapter 30

Miscellaneous extensions

30.1 Program derivation

Coq comes with an extension called `Derive`, which supports program derivation. Typically in the style of Bird and Meertens or derivations of program refinements. To use the `Derive` extension it must first be required with `Require Coq.Derive.Derive`. When the extension is loaded, it provides the following command.

30.1.1 `Derive ident1 SuchThat term As ident2`

The name `ident1` can appear in `term`. This command opens a new proof presenting the user with a goal for `term` in which the name `ident1` is bound to a existential variables `?x` (formally, there are other goals standing for the existential variables but they are shelved, as described in Section 8.17.4).

When the proof ends two constants are defined:

- The first one is name `ident1` and is defined as the proof of the shelved goal (which is also the value of `?x`). It is always transparent.
- The second one is name `ident2`. It has type `term`, and its body is the proof of the initially visible goal. It is opaque if the proof ends with `Qed`, and transparent if the proof ends with `Defined`.

Example:

```
Coq < Require Coq.derive.Derive.
Coq < Require Import Coq.Numbers.Natural.Peano.NPeano.
Coq < Section P.
Coq < Variables (n m k:nat).
Coq < Derive p SuchThat ((k*n)+(k*m) = p) As h.
1 focused subgoal
(shelved: 1)

  n, m, k : nat
  p := ?Goal : nat
  =====
  k * n + k * m = p
Coq < Proof.
```

```

1 focused subgoal
(shelved: 1)

n, m, k : nat
p := ?Goal : nat
=====
k * n + k * m = p
Coq < rewrite <- Nat.mul_add_distr_l.
1 focused subgoal
(shelved: 1)

n, m, k : nat
p := ?Goal : nat
=====
k * (n + m) = p
Coq < subst p.
1 focused subgoal
(shelved: 1)

n, m, k : nat
=====
k * (n + m) = ?Goal
Coq < reflexivity.
No more subgoals.
Coq < Qed.
Coq < End P.
Coq < Print p.
p = fun n m k : nat => k * (n + m)
    : nat -> nat -> nat -> nat
Argument scopes are [nat_scope nat_scope nat_scope]
Coq < Check h.
h
    : forall n m k : nat, k * n + k * m = p n m k

```

Any property can be used as *term*, not only an equation. In particular, it could be an order relation specifying some form of program refinement or a non-executable property from which deriving a program is convenient.

Bibliography

- [1] David Aspinall. Proof general. <https://proofgeneral.github.io/>.
- [2] Ph. Audebaud. Partial Objects in the Calculus of Constructions. In *Proceedings of the sixth Conf. on Logic in Computer Science*. IEEE, 1991.
- [3] Ph. Audebaud. CC+ : an extension of the Calculus of Constructions with fixpoints. In B. Nordström and K. Petersson and G. Plotkin, editor, *Proceedings of the 1992 Workshop on Types for Proofs and Programs*, pages 21–34, 1992. Also Research Report LIP-ENS-Lyon.
- [4] Ph. Audebaud. *Extension du Calcul des Constructions par Points fixes*. PhD thesis, Université Bordeaux I, 1992.
- [5] L. Augustsson. Compiling Pattern Matching. In *Conference Functional Programming and Computer Architecture*, 1985.
- [6] H. Barendregt. Lambda Calculi with Types. Technical Report 91-19, Catholic University Nijmegen, 1991. In *Handbook of Logic in Computer Science*, Vol II.
- [7] H. Barendregt and T. Nipkow, editors. *Types for Proofs and Programs*, volume 806 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994.
- [8] H.P. Barendregt. *The Lambda Calculus its Syntax and Semantics*. North-Holland, 1981.
- [9] B. Barras. *Auto-validation d'un système de preuves avec familles inductives*. Thèse de doctorat, Université Paris 7, 1999.
- [10] J.L. Bates and R.L. Constable. Proofs as Programs. *ACM transactions on Programming Languages and Systems*, 7, 1985.
- [11] M.J. Beeson. *Foundations of Constructive Mathematics, Metamathematical Studies*. Springer-Verlag, 1985.
- [12] G. Bellin and J. Ketonen. A decision procedure revisited : Notes on direct logic, linear logic and its implementation. *Theoretical Computer Science*, 95:115–142, 1992.
- [13] Stefano Berardi and Mario Coppo, editors. *Types for Proofs and Programs, International Workshop TYPES'95, Torino, Italy, June 5-8, 1995, Selected Papers*, volume 1158 of *Lecture Notes in Computer Science*. Springer, 1996.
- [14] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. An EATCS series. Springer Verlag, 2004.

- [15] E. Bishop. *Foundations of Constructive Analysis*. McGraw-Hill, 1967.
- [16] Mathieu Boespflug, Maxime Dénès, and Benjamin Grégoire. Full reduction at full throttle. In Jean-Pierre Jouannaud and Zhong Shao, editors, *Certified Programs and Proofs - First International Conference, CPP 2011, Kenting, Taiwan, December 7-9, 2011. Proceedings*, volume 7086 of *Lecture Notes in Computer Science*, pages 362–377. Springer, 2011.
- [17] S. Boutin. Certification d’un compilateur ML en Coq. Master’s thesis, Université Paris 7, September 1992.
- [18] S. Boutin. *Réflexions sur les quotients*. thèse d’université, Paris 7, April 1997.
- [19] S. Boutin. Using reflection to build efficient and certified decision procedures. In Martin Abadi and Takahashi Ito, editors, *TACS’97*, volume 1281 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [20] R.S. Boyer and J.S. Moore. *A computational logic*. ACM Monograph. Academic Press, 1979.
- [21] Paul Callaghan, Zhaohui Luo, James McKinna, and Robert Pollack, editors. *Types for Proofs and Programs, International Workshop, TYPES 2000, Durham, UK, December 8-12, 2000, Selected Papers*, volume 2277 of *Lecture Notes in Computer Science*. Springer, 2002.
- [22] Laurent Chicli, Loïc Pottier, and Carlos Simpson. Mathematical quotients and quotient types in coq. In Geuvers and Wiedijk [66].
- [23] R.L. Constable et al. *Implementing Mathematics with the Nuprl Proof Development System*. Prentice-Hall, 1986.
- [24] Th. Coquand. *Une Théorie des Constructions*. PhD thesis, Université Paris 7, January 1985.
- [25] Th. Coquand. An Analysis of Girard’s Paradox. In *Symposium on Logic in Computer Science*, Cambridge, MA, 1986. IEEE Computer Society Press.
- [26] Th. Coquand. Metamathematical Investigations of a Calculus of Constructions. In P. Oddifredi, editor, *Logic and Computer Science*. Academic Press, 1990. INRIA Research Report 1088, also in [64].
- [27] Th. Coquand. A New Paradox in Type Theory. In *Proceedings 9th Int. Congress of Logic, Methodology and Philosophy of Science*, August 1991.
- [28] Th. Coquand. Pattern Matching with Dependent Types. In Nordström et al. [116].
- [29] Th. Coquand. Infinite objects in Type Theory. In H. Barendregt and T. Nipkow, editors, *Types for Proofs and Programs*, volume 806 of *Lecture Notes in Computer Science*, pages 62–78. Springer-Verlag, 1993.
- [30] Th. Coquand and G. Huet. Constructions : A Higher Order Proof System for Mechanizing Mathematics. In *EUROCAL’85*, volume 203 of *Lecture Notes in Computer Science*, Linz, 1985. Springer-Verlag.
- [31] Th. Coquand and G. Huet. Concepts Mathématiques et Informatiques formalisés dans le Calcul des Constructions. In The Paris Logic Group, editor, *Logic Colloquium’85*. North-Holland, 1987.

- [32] Th. Coquand and G. Huet. The Calculus of Constructions. *Information and Computation*, 76(2/3), 1988.
- [33] Th. Coquand and C. Paulin-Mohring. Inductively defined types. In P. Martin-Löf and G. Mints, editors, *Proceedings of Colog'88*, volume 417 of *Lecture Notes in Computer Science*. Springer-Verlag, 1990.
- [34] P. Corbineau. A declarative language for the coq proof assistant. In M. Miculan, I. Scagnetto, and F. Honsell, editors, *TYPES '07, Cividale del Friuli, Revised Selected Papers*, volume 4941 of *Lecture Notes in Computer Science*, pages 69–84. Springer, 2007.
- [35] C. Cornes. *Conception d'un langage de haut niveau de représentation de preuves*. Thèse de doctorat, Université Paris 7, November 1997.
- [36] Cristina Cornes and Delphine Terasse. Automating inversion of inductive predicates in coq. In Berardi and Coppo [13], pages 85–104.
- [37] J. Courant. Explicitation de preuves par récurrence implicite. Master's thesis, DEA d'Informatique, ENS Lyon, September 1994.
- [38] Haskell B. Curry, Robert Feys, and William Craig. *Combinatory Logic*, volume 1. North-Holland, 1958. §9E.
- [39] N.J. de Bruijn. Lambda-Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem. *Indag. Math.*, 34, 1972.
- [40] N.J. de Bruijn. A survey of the project Automath. In J.P. Seldin and J.R. Hindley, editors, *to H.B. Curry : Essays on Combinatory Logic, Lambda Calculus and Formalism*. Academic Press, 1980.
- [41] D. de Rauglaudre. Camlp4 version 1.07.2. In Camlp4 distribution, 1998.
- [42] D. Delahaye. Information retrieval in a coq proof library using type isomorphisms. In *Proceedings of TYPES '99, Lökeberg*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
- [43] D. Delahaye. A Tactic Language for the System Coq. In *Proceedings of Logic for Programming and Automated Reasoning (LPAR), Reunion Island*, volume 1955 of *Lecture Notes in Computer Science*, pages 85–95. Springer-Verlag, November 2000.
- [44] D. Delahaye and M. Mayero. *Field*: une procédure de décision pour les nombres réels en COQ. In *Journées Francophones des Langages Applicatifs, Pontarlier*. INRIA, Janvier 2001.
- [45] R. di Cosmo. *Isomorphisms of Types: from λ -calculus to information retrieval and language design*. Progress in Theoretical Computer Science. Birkhauser, 1995. ISBN-0-8176-3763-X.
- [46] G. Dowek. Naming and scoping in a mathematical vernacular. Research Report 1283, INRIA, 1990.
- [47] G. Dowek. *Démonstration automatique dans le Calcul des Constructions*. PhD thesis, Université Paris 7, December 1991.
- [48] G. Dowek. L'indécidabilité du filtrage du troisième ordre dans les calculs avec types dépendants ou constructeurs de types. *Compte-Rendus de l'Académie des Sciences*, I, 312(12):951–956, 1991. The undecidability of Third Order Pattern Matching in Calculi with Dependent Types or Type Constructors.

- [49] G. Dowek. A second order pattern matching algorithm in the cube of typed λ -calculi. In *Proceedings of Mathematical Foundation of Computer Science*, volume 520 of *Lecture Notes in Computer Science*, pages 151–160. Springer-Verlag, 1991. Also INRIA Research Report.
- [50] G. Dowek. A Complete Proof Synthesis Method for the Cube of Type Systems. *Journal Logic Computation*, 3(3):287–315, June 1993.
- [51] G. Dowek. The undecidability of pattern matching in calculi where primitive recursive functions are representable. *Theoretical Computer Science*, 107(2):349–356, 1993.
- [52] G. Dowek. Third order matching is decidable. *Annals of Pure and Applied Logic*, 69:135–155, 1994.
- [53] G. Dowek. Lambda-calculus, combinators and the comprehension schema. In *Proceedings of the second international conference on typed lambda calculus and applications*, 1995.
- [54] G. Dowek, A. Felty, H. Herbelin, G. Huet, C. Murthy, C. Parent, C. Paulin-Mohring, and B. Werner. The Coq Proof Assistant User’s Guide Version 5.8. Technical Report 154, INRIA, May 1993.
- [55] P. Dybjer. Inductive sets and families in Martin-Löf’s type theory and their set-theoretic semantics: An inversion principle for Martin-Löf’s type theory. In G. Huet and G. Plotkin, editors, *Logical Frameworks*, volume 14, pages 59–79. Cambridge University Press, 1991.
- [56] Roy Dyckhoff. Contraction-free sequent calculi for intuitionistic logic. *The Journal of Symbolic Logic*, 57(3), September 1992.
- [57] J.-C. Filliâtre. Une procédure de décision pour le calcul des prédicats direct. Étude et implémentation dans le système COQ. Master’s thesis, DEA d’Informatique, ENS Lyon, September 1994.
- [58] J.-C. Filliâtre. A decision procedure for direct predicate calculus. Research report 96–25, LIP-ENS-Lyon, 1995.
- [59] J.-C. Filliâtre. *Preuve de programmes impératifs en théorie des types*. Thèse de doctorat, Université Paris-Sud, July 1999.
- [60] J.-C. Filliâtre. Formal Proof of a Program: Find. Submitted to *Science of Computer Programming*, January 2000.
- [61] J.-C. Filliâtre and N. Magaud. Certification of sorting algorithms in the system COQ. In *Theorem Proving in Higher Order Logics: Emerging Trends*, 1999.
- [62] J.-C. Filliâtre. Verification of non-functional programs using interpretations in type theory. *Journal of Functional Programming*, 13(4):709–745, July 2003. [English translation of [59]].
- [63] E. Fleury. Implantation des algorithmes de Floyd et de Dijkstra dans le Calcul des Constructions. Rapport de Stage, July 1990.
- [64] Projet Formel. The Calculus of Constructions. Documentation and user’s guide, Version 4.10. Technical Report 110, INRIA, 1989.

- [65] Jean-Baptiste-Joseph Fourier. *Fourier's method to solve linear inequations/equations systems*. Gauthier-Villars, 1890.
- [66] H. Geuvers and F. Wiedijk, editors. *Types for Proofs and Programs, Second International Workshop, TYPES 2002, Berg en Dal, The Netherlands, April 24-28, 2002, Selected Papers*, volume 2646 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
- [67] E. Giménez. Codifying guarded definitions with recursive schemes. In *Types'94 : Types for Proofs and Programs*, volume 996 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994. Extended version in LIP research report 95-07, ENS Lyon.
- [68] E. Giménez. An application of co-inductive types in coq: verification of the alternating bit protocol. In *Workshop on Types for Proofs and Programs*, number 1158 in *Lecture Notes in Computer Science*, pages 135–152. Springer-Verlag, 1995.
- [69] E. Giménez. *Un calcul des constructions infinies et son application à la vérification de systèmes communicants*. PhD thesis, École Normale Supérieure de Lyon, 1996.
- [70] E. Giménez. A tutorial on recursive types in coq. Technical report, INRIA, March 1998.
- [71] E. Giménez and P. Castéran. A tutorial on [co-]inductive types in coq. available at <http://coq.inria.fr/doc>, January 2005.
- [72] Alessandro Giovini, Teo Mora, Gianfranco Niesi, Lorenzo Robbiano, and Carlo Traverso. "one sugar cube, please" or selection strategies in the buchberger algorithm. In *Proceedings of the ISSAC'91, ACM Press*, pages 5–4, 1991.
- [73] J.-Y. Girard. Une extension de l'interprétation de Gödel à l'analyse, et son application à l'élimination des coupures dans l'analyse et la théorie des types. In *Proceedings of the 2nd Scandinavian Logic Symposium*. North-Holland, 1970.
- [74] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris 7, 1972.
- [75] J.-Y. Girard, Y. Lafont, and P. Taylor. *Proofs and Types*. Cambridge Tracts in Theoretical Computer Science 7. Cambridge University Press, 1989.
- [76] Georges Gonthier, Beta Ziliani, Aleksandar Nanevski, and Derek Dreyer. How to make ad hoc proof automation less ad hoc. *SIGPLAN Not.*, 46(9):163–175, September 2011.
- [77] Benjamin Grégoire and Xavier Leroy. A compiled implementation of strong reduction. In Mitchell Wand and Simon L. Peyton Jones, editors, *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02), Pittsburgh, Pennsylvania, USA, October 4-6, 2002.*, pages 235–246. ACM, 2002.
- [78] John Harrison. Metatheory and reflection in theorem proving: A survey and critique. Technical Report CRC-053, SRI International Cambridge Computer Science Research Centre., 1995.
- [79] D. Hirschhoff. Écriture d'une tactique arithmétique pour le système COQ. Master's thesis, DEA IARFA, Ecole des Ponts et Chaussées, Paris, September 1994.

- [80] Martin Hofmann and Thomas Streicher. The groupoid interpretation of type theory. In *Proceedings of the meeting Twenty-five years of constructive type theory*. Oxford University Press, 1998.
- [81] W.A. Howard. The formulae-as-types notion of constructions. In J.P. Seldin and J.R. Hindley, editors, *to H.B. Curry : Essays on Combinatory Logic, Lambda Calculus and Formalism*. Academic Press, 1980. Unpublished 1969 Manuscript.
- [82] G. Huet. Programming of future generation computers. In *Proceedings of TAPSOFT87*, volume 249 of *Lecture Notes in Computer Science*, pages 276–286. Springer-Verlag, 1987.
- [83] G. Huet. Extending the calculus of constructions with type:type. Unpublished, 1988.
- [84] G. Huet. Induction principles formalized in the Calculus of Constructions. In K. Fuchi and M. Nivat, editors, *Programming of Future Generation Computers*. Elsevier Science, 1988. Also in [82].
- [85] G. Huet, editor. *Logical Foundations of Functional Programming*. The UT Year of Programming Series. Addison-Wesley, 1989.
- [86] G. Huet. The Constructive Engine. In R. Narasimhan, editor, *A perspective in Theoretical Computer Science. Commemorative Volume for Gift Siromoney*. World Scientific Publishing, 1989. Also in [64].
- [87] G. Huet. The gallina specification language : A case study. In *Proceedings of 12th FST/TCS Conference, New Delhi*, volume 652 of *Lecture Notes in Computer Science*, pages 229–240. Springer-Verlag, 1992.
- [88] G. Huet. Residual theory in λ -calculus: a formal development. *J. Functional Programming*, 4,3:371–394, 1994.
- [89] G. Huet and J.-J. Lévy. Call by need computations in non-ambiguous linear term rewriting systems. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic, Essays in Honor of Alan Robinson*. The MIT press, 1991. Also research report 359, INRIA, 1979.
- [90] G. Huet and G. Plotkin, editors. *Logical Frameworks*. Cambridge University Press, 1991.
- [91] G. Huet and G. Plotkin, editors. *Logical Environments*. Cambridge University Press, 1992.
- [92] J. Ketonen and R. Weyhrauch. A decidable fragment of Predicate Calculus. *Theoretical Computer Science*, 32:297–307, 1984.
- [93] S.C. Kleene. *Introduction to Metamathematics*. Bibliotheca Mathematica. North-Holland, 1952.
- [94] J.-L. Krivine. *Lambda-calcul types et modèles*. Etudes et recherche en informatique. Masson, 1990.
- [95] A. Laville. Comparison of priority rules in pattern matching and term rewriting. *Journal of Symbolic Computation*, 11:321–347, 1991.
- [96] F. Leclerc and C. Paulin-Mohring. Programming with Streams in Coq. A case study : The Sieve of Eratosthenes. In H. Barendregt and T. Nipkow, editors, *Types for Proofs and Programs, Types’93*, volume 806 of *LNCS*. Springer-Verlag, 1994.

- [97] Gyesik Lee and Benjamin Werner. Proof-irrelevant model of CC with predicative induction and judgmental equality. *Logical Methods in Computer Science*, 7(4), 2011.
- [98] X. Leroy. The ZINC experiment: an economical implementation of the ML language. Technical Report 117, INRIA, 1990.
- [99] P. Letouzey. A new extraction for coq. In Geuvers and Wiedijk [66].
- [100] L. Puel and A. Suárez. Compiling Pattern Matching by Term Decomposition. In *Conference Lisp and Functional Programming*, ACM. Springer-Verlag, 1990.
- [101] Z. Luo. *An Extended Calculus of Constructions*. PhD thesis, University of Edinburgh, 1990.
- [102] Sebastiaan P. Luttik and Eelco Visser. Specification of rewriting strategies. In *2nd International Workshop on the Theory and Practice of Algebraic Specifications (ASF+SDF'97), Electronic Workshops in Computing*. Springer-Verlag, 1997.
- [103] Assia Mahboubi and Enrico Tassi. Canonical Structures for the working Coq user. In Sandrine Blazy, Christine Paulin, and David Pichardie, editors, *ITP 2013, 4th Conference on Interactive Theorem Proving*, volume 7998 of *LNCS*, pages 19–34, Rennes, France, 2013. Springer.
- [104] P. Manoury. A User's Friendly Syntax to Define Recursive Functions as Typed λ -Terms. In *Types for Proofs and Programs, TYPES'94*, volume 996 of *LNCS*, June 1994.
- [105] P. Manoury and M. Simonot. Automatizing termination proofs of recursively defined functions. *TCS*, 135(2):319–343, 1994.
- [106] L. Maranget. Two Techniques for Compiling Lazy Pattern Matching. Technical Report 2385, INRIA, 1994.
- [107] Conor McBride. Elimination with a motive. In Callaghan et al. [21], pages 197–216.
- [108] A. Miquel. A model for impredicative type systems with universes, intersection types and subtyping. In *Proceedings of the 15th Annual IEEE Symposium on Logic in Computer Science (LICS'00)*. IEEE Computer Society Press, 2000.
- [109] A. Miquel. The implicit calculus of constructions: Extending pure type systems with an intersection type binder and subtyping. In *Proceedings of the fifth International Conference on Typed Lambda Calculi and Applications (TLCA01), Krakow, Poland*, number 2044 in *LNCS*. Springer-Verlag, 2001.
- [110] A. Miquel. *Le Calcul des Constructions implicite: syntaxe et sémantique*. PhD thesis, Université Paris 7, dec 2001.
- [111] A. Miquel and B. Werner. The not so simple proof-irrelevant model of cc. In Geuvers and Wiedijk [66], pages 240–258.
- [112] C. Muñoz. *Un calcul de substitutions pour la représentation de preuves partielles en théorie de types*. Thèse de doctorat, Université Paris 7, 1997. Version en anglais disponible comme rapport de recherche INRIA RR-3309.
- [113] C. Muñoz. Démonstration automatique dans la logique propositionnelle intuitionniste. Master's thesis, DEA d'Informatique Fondamentale, Université Paris 7, September 1994.

- [114] B. Nordström. Terminating general recursion. *BIT*, 28, 1988.
- [115] B. Nordström, K. Peterson, and J. Smith. *Programming in Martin-Löf's Type Theory*. International Series of Monographs on Computer Science. Oxford Science Publications, 1990.
- [116] B. Nordström, K. Petersson, and G. Plotkin, editors. *Proceedings of the 1992 Workshop on Types for Proofs and Programs*. Available by ftp at site ftp.inria.fr, 1992.
- [117] P. Odifreddi, editor. *Logic and Computer Science*. Academic Press, 1990.
- [118] P. Martin-Löf. *Intuitionistic Type Theory*. Studies in Proof Theory. Bibliopolis, 1984.
- [119] C. Parent. Developing certified programs in the system Coq- The Program tactic. Technical Report 93-29, Ecole Normale Supérieure de Lyon, October 1993. Also in [7].
- [120] C. Parent. *Synthèse de preuves de programmes dans le Calcul des Constructions Inductives*. PhD thesis, Ecole Normale Supérieure de Lyon, 1995.
- [121] C. Parent. Synthesizing proofs from programs in the Calculus of Inductive Constructions. In *Mathematics of Program Construction '95*, volume 947 of LNCS. Springer-Verlag, 1995.
- [122] M. Parigot. Recursive Programming with Proofs. *Theoretical Computer Science*, 94(2):335–356, 1992.
- [123] M. Parigot, P. Manoury, and M. Simonot. ProPre : A Programming language with proofs. In A. Voronkov, editor, *Logic Programming and automated reasoning*, number 624 in LNCS, St. Petersburg, Russia, July 1992. Springer-Verlag.
- [124] C. Paulin-Mohring. Extracting F_ω 's programs from proofs in the Calculus of Constructions. In *Sixteenth Annual ACM Symposium on Principles of Programming Languages*, Austin, January 1989. ACM.
- [125] C. Paulin-Mohring. *Extraction de programmes dans le Calcul des Constructions*. PhD thesis, Université Paris 7, January 1989.
- [126] C. Paulin-Mohring. Inductive Definitions in the System Coq - Rules and Properties. In M. Bezem and J.-F. Groote, editors, *Proceedings of the conference Typed Lambda Calculi and Applications*, number 664 in LNCS. Springer-Verlag, 1993. Also LIP research report 92-49, ENS Lyon.
- [127] C. Paulin-Mohring. *Le système Coq. Thèse d'habilitation*. ENS Lyon, January 1997.
- [128] C. Paulin-Mohring and B. Werner. Synthesis of ML programs in the system Coq. *Journal of Symbolic Computation*, 15:607–640, 1993.
- [129] K.V. Prasad. Programming with broadcasts. In *Proceedings of CONCUR'93*, volume 715 of LNCS. Springer-Verlag, 1993.
- [130] W. Pugh. The omega test: a fast and practical integer programming algorithm for dependence analysis. *Communication of the ACM*, pages 102–114, 1992.
- [131] J. Rouyer. Développement de l'Algorithme d'Unification dans le Calcul des Constructions. Technical Report 1795, INRIA, November 1992.

- [132] John Rushby, Sam Owre, and N. Shankar. Subtypes for specifications: Predicate subtyping in PVS. *IEEE Transactions on Software Engineering*, 24(9):709–720, September 1998.
- [133] A. Saïbi. Axiomatization of a lambda-calculus with explicit-substitutions in the Coq System. Technical Report 2345, INRIA, December 1994.
- [134] H. Saidi. Résolution d'équations dans le système λ de gödel. Master's thesis, DEA d'Informatique Fondamentale, Université Paris 7, September 1994.
- [135] Matthieu Sozeau. Subset coercions in Coq. In *TYPES'06*, volume 4502 of *LNCS*, pages 237–252. Springer, 2007.
- [136] Matthieu Sozeau and Nicolas Oury. First-Class Type Classes. In *TPHOLs'08*, 2008.
- [137] T. Streicher. Semantical investigations into intensional type theory, 1993. Habilitationsschrift, LMU Munchen.
- [138] Lemme Team. Pcoq a graphical user-interface for Coq. <http://www-sop.inria.fr/lemme/pcoq/>.
- [139] The Coq Development Team. The Coq Proof Assistant Reference Manual Version 7.2. Technical Report 255, INRIA, February 2002.
- [140] D. Terrasse. Traduction de TYPOL en COQ. Application à Mini ML. Master's thesis, IARFA, September 1992.
- [141] L. Théry, Y. Bertot, and G. Kahn. Real theorem provers deserve real user-interfaces. Research Report 1684, INRIA Sophia, May 1992.
- [142] A.S. Troelstra and D. van Dalen. *Constructivism in Mathematics, an introduction*. Studies in Logic and the foundations of Mathematics, volumes 121 and 123. North-Holland, 1988.
- [143] P. Wadler. Efficient compilation of pattern matching. In S.L. Peyton Jones, editor, *The Implementation of Functional Programming Languages*. Prentice-Hall, 1987.
- [144] P. Weis and X. Leroy. *Le langage Caml*. InterEditions, 1993.
- [145] B. Werner. *Une théorie des constructions inductives*. Thèse de doctorat, Université Paris 7, 1994.

Global Index

$+$, 267
 $||$, 267
 $*$, 109, 115
 $*$ (command), 177
 $+$, 109, 115
 $+$ (command), 177
 $-$, 115
 $-$ (command), 177
 $/$, 115
 $;$, 262
 $;$ [...|...|...], 262
 $<$, 115
 \leq , 115
 $>$, 115
 \geq , 115
 $?$, 185
 $?=$, 115
 $[>...|...|...]$, 262
 $\%$, 375
 $\&$, 110
 $_$, 46
 $_CoqProject$, 400
 $\dots : \dots$ (ssreflect), 314
 $\dots \Rightarrow \dots$ (ssreflect), 317
 \dots in \dots (ssreflect), 310, 326
 $@$, 95
 $:$, 265
 $\{$, 177
 $\{A\} + \{B\}$, 110
 $\{x:A \mid (P \ x)\}$, 110
 $\{x:A \ \& \ (P \ x)\}$, 110
 $\}$, 177
 $\langle \dots \rangle$, 99
 $\{\dots\}$, 99
2-level approach, 254

 $A*B$, 109
 $A+\{B\}$, 110

$A+B$, 109
Abbreviations, 380
Abort, 176
About, 153
Absolute names, 86
abstract, 276
abstract: ... (ssreflect), 317
abstractions, 45
absurd, 108, 206
absurd_set, 111
Acc, 113
Acc_inv, 113
Acc_rect, 113
Add Field, 255, 501
Add LoadPath, 164
Add ML Path, 165
Add Morphism, 515
Add Parametric Morphism, 510
Add Parametric Relation, 509
Add Printing Coercion, 443
Add Printing If *ident*, 74
Add Printing Let *ident*, 73
Add Rec LoadPath, 164
Add Rec ML Path, 165
Add Relation, 509
Add Ring, 255, 496
Add Setoid, 514
admit, 206
Admit Obligations, 491
Admitted, 63, 174
algebraic universe, 122
all, 107
and, 107
and_rect, 111
app, 118
appcontext
 in pattern, 272
applications, 46

- apply, 187
- apply ... in, 190
- apply ... with, 187
- apply/... (ssreflect), 357
- apply/.../... (ssreflect), 357
- apply: ... (ssreflect), 314
- Arguments, 90–92, 95, 232, 375
- Arithmetical notations, 115
- Arity, 129
- arity of inductive type, 128
- assert, 202
- assert as, 203
- assert by, 202
- Associativity, 366
- assumption
 - global, 123
 - local, 123
- assumption, 185
- Asymmetric Patterns, 431
- Asynchronous and Parallel Proof Processing
 - presentation, 519
- auto, 236
- autoapply, 465
- Automatic Coercions Import, 444
- Automatic Introduction, 181
- autorewrite, 238
- autounfold, 238
- Axiom, 50
- Axiom (and coercions), 442

- Back, 166
- BackTo, 166
- Backtrack, 166
- Bad Magic Number, 163
- β -reduction, 125, 126
- Bind Scope, 376
- binders, 43
- Binding list, 184
- BNF metasyntax, 41
- bool, 109
- bool_choice, 110
- Boolean Equality Schemes, 385
- Bracketing Last Introduction
 - Pattern, 196
- btauto, 254
- Bullet Behavior, 178
- Bullets, 177

- by ... (ssreflect), 323
- byte-code, 393

- Calculus of Inductive Constructions, 121
- Canonical Structure, 97
- Canonical Structures
 - presentation, 449
- case, 209
- Case Analysis Schemes, 385
- case: ... (ssreflect), 313
- case: ... / ... (ssreflect), 320
- case_eq, 209
- Cases, 427
- Cast, 46
 - (...: ...), 46
- cbn, 232
- cbv, 230
- Cd, 164
- change, 229
- change ... in, 229
- Check, 155
- Choice, 110
- Choice2, 110
- CIC, 121
- Class, 463
- classical_left, 254
- classical_right, 254
- Clauses, 229
- clear, 197
- clear dependent, 197
- clearbody, 197
- Close Scope, 374
- Coercion, 100, 441, 442
- Coercions, 100
 - and modules, 444
 - and records, 444
 - and sections, 444
 - classes, 439
 - Funclass, 440
 - identity, 440
 - inheritance graph, 441
 - presentation, 439
 - Sortclass, 440
- cofix, 225
- CoFixpoint, 60, 62
- CoFixpoint ... where ..., 370

- CoInductive, [57](#), [65](#)
- CoInductive (and coercions), [442](#)
- Collection, [175](#)
- Combined Scheme, [385](#)
- Comments, [41](#)
- compare, [252](#)
- Compiled files, [162](#)
- Compute, [156](#)
- compute, [230](#), [231](#)
- congruence, [249](#)
- Congruence Verbose, [250](#)
- conj, [107](#)
- Conjecture, [50](#)
- Connectives, [107](#)
- Constant, [51](#)
- constr_eq, [251](#)
- Constraint, [529](#)
- constructor, [191](#)
- Context, [464](#)
- context
 - in expression, [273](#)
 - in pattern, [272](#)
- context of parameters, [128](#)
- Contextual Implicit, [94](#)
- contradict, [207](#)
- contradiction, [207](#)
- Contributions, [119](#)
- Conversion rules, [125](#)
- Conversion tactics, [229](#)
- coq-tex, [413](#)
- coq_Makefile, [400](#)
- coqc, [393](#)
- coqchk, [393](#)
- coqdep, [405](#)
- coqdoc, [405](#)
- coqide, [415](#)
- coqmktop, [399](#)
- coqtop, [393](#)
- Corollary, [61](#)
- CreateHintDb, [239](#)
- Cumulative, [525](#)
- cut, [204](#)
- cutrewrite, [227](#)
- cycle, [256](#)
-
- Datatypes, [109](#)
- Debug Auto, [237](#)
- Debug Cbv, [231](#)
- Debug Eauto, [237](#)
- Debug RAKAM, [234](#)
- Debug Trivial, [237](#)
- Debugger, [399](#)
- Decidable Equality Schemes, [385](#)
- decide equality, [252](#)
- declaration
 - global, [123](#)
 - local, [123](#)
- Declarations, [49](#)
- Declare Implicit Tactic, [246](#)
- Declare Instance, [464](#)
- Declare Left Step, [228](#)
- Declare ML Module, [163](#)
- Declare Module, [81](#)
- Declare Right Step, [229](#)
- decompose, [202](#)
- decompose record, [202](#)
- decompose sum, [202](#)
- Default Goal Selector, [183](#)
- Default Proof Using, [175](#)
- Default Timeout, [167](#), [168](#)
- Defined, [63](#), [174](#)
- Definition, [51](#)
- definition
 - global, [123](#)
 - inductive, [128](#)
 - local, [123](#)
- Definitions, [51](#)
- Delimit Scope, [375](#)
- δ -reduction, [51](#), [126](#)
- Dependencies, [405](#)
- dependent destruction, [214](#)
- dependent induction, [213](#)
- dependent induction ...
 - generalizing, [214](#)
- dependent inversion, [220](#)
- dependent inversion ... as, [220](#)
- dependent inversion ... as ...
 - with, [221](#)
- dependent inversion ... with, [221](#)
- dependent inversion_clear, [221](#)
- dependent inversion_clear ...
 - as, [221](#)
- dependent inversion_clear ...
 - as ... with, [221](#)

dependent inversion_clear ...
 with, 221
 dependent rewrite \rightarrow , 253
 dependent rewrite \leftarrow , 253
 Derive, 531
 Derive Dependent Inversion, 389
 Derive Dependent
 Inversion_clear, 389
 Derive Inversion, 389
 Derive Inversion_clear, 389
 destruct, 207
 dintuition, 247
 discriminate, 216
 discrR, 117
 Disjunctive/conjunctive introduction patterns,
 194
 do, 266
 do ... [...] (ssreflect),
 325
 double induction, 213
 Drop, 167
 dtauto, 246

 eapply, 187
 eapply ... in, 191
 eassert, 203
 eassert as, 203
 eassert by, 203
 eassumption, 185
 easy, 239
 eauto, 237
 ecase, 209
 econstructor, 192
 edestruct, 208
 ediscriminate, 216
 eelim, 212
 eenough, 204
 eenough as, 204
 eenough by, 204
 eexact, 185
 eexists, 192
 einduction, 211
 einjection, 218
 eleft, 192
 elim ... using, 212
 elim/... (ssreflect), 349
 elim: ... (ssreflect), 313

Elimination
 Empty elimination, 136
 Singleton elimination, 136
 Elimination Schemes, 385
 Elimination sorts, 135
 elimtype, 213
 Emacs, 413
 End, 78, 80, 81
 enough, 203
 enough as, 203
 enough by, 204
 Environment, 49, 51
 Environment variables, 394
 epose, 201
 epose proof, 203
 eq, 108
 eq_add_S, 111
 eq_ind_r, 108
 eq_rec_r, 108
 eq_rect, 108, 111
 eq_rect_r, 108
 eq_refl, 108
 eq_S, 111
 eq_sym, 108
 eq_trans, 108
 Equality, 108
 Equality introduction patterns, 194
 eremember, 201
 erewrite, 227
 eright, 192
 error, 111
 eset, 201
 esimplify_eq, 252
 esplit, 192
 η -expansion, 126
 Eval, 155
 eval
 in Ltac, 274
 evar, 205
 ex, 107
 ex2, 107
 ex_intro, 107
 ex_intro2, 107
 exact, 185
 exactly_once, 268
 Example, 51
 Exc, 111

- exfalse, 207
- exist, 110
- exist2, 110
- Existential, 176
- Existing Class, 463
- Existing Instance, 464
- Existing Instances, 464
- exists, 107, 192
- exists2, 107
- existT, 110
- existT2, 110
- Explicitly given implicit arguments, 95
- Export, 85
- Extract Constant, 480
- Extract Inductive, 481
- Extraction, 477
- Extraction, 156, 477
- Extraction AutoInline, 479
- Extraction Blacklist, 482
- Extraction Conservative Types, 479
- Extraction Implicit, 480
- Extraction Inline, 479
- Extraction KeepSingleton, 479
- Extraction Language, 478
- Extraction Library, 477
- Extraction NoInline, 479
- Extraction Optimize, 479
- Extraction SafeImplicits, 480
- f_equal, 108, 251
- f_equal i , 108
- Fact, 61
- fail, 269
- False, 107
- false, 109
- False_rec, 111
- False_rect, 111
- field, 255, 500
- field_simplify, 255, 500
- field_simplify_eq, 255, 500
- first, 267
- first ... last (ssreflect), 324
- firstorder, 249
- Firstorder Depth, 249
- Firstorder Solver, 249
- firstorder using, 249
- firstorder with, 249
- firstorder *tactic*, 249
- Fix, 138
- fix, 224
- fix *ident $_i$* {...}, 48
- fix_eq, 113
- Fix_F, 113
- Fix_F_eq, 113
- Fix_F_inv, 113
- Fixpoint, 57, 62
- Fixpoint ... where ..., 370
- flat_map, 118
- Focus, 177
- fold, 235
- fold_left, 118
- fold_right, 118
- forall ..., ..., 46
- form*, 43
- fourier, 256
- fresh
 - in Ltac, 274
- From Require, 163
- fst, 109
- fun
 - in Ltac, 270
- fun ... => ..., 45
- Function, 75
- function_scope, 378
- functional induction, 215, 387
- functional inversion, 253
- Functional Scheme, 386
- Gallina, 41, 65
- gallina, 414
- ge, 112
- Generalizable Variables, 99
- generalize, 204
- generalize dependent, 205
- gfail, 269
- give_up, 206, 259
- Global, 172
- Global Arguments, 91, 92
- Global environment, 123
- Global Set, 154
- Global Unset, 154, 155
- Goal, 173
- goal, 183

- Goal clauses, 229
- Grab Existential Variables, 176
- gt, 112
- guard
 - in Ltac, 275
- Guarded, 180
- has_evar, 251
- have: ... (ssreflect), 327
- have: ... := ... (ssreflect), 327
- head, 118
- Head normal form, 128
- Hide Obligations, 491
- Hint, 239
- Hint Constructors, 241
- Hint Cut, 242
- Hint Extern, 242
- Hint Immediate, 240
- Hint Mode, 243
- Hint Opaque, 241
- Hint Resolve, 240
- Hint Rewrite, 244
- Hint Transparent, 241
- Hint Unfold, 241
- Hint View (ssreflect), 358
- Hints databases, 239
- hnf, 232
- Hypotheses, 50
- Hypothesis, 50
- Hypothesis (and coercions), 442
- Hyps Limit, 180
- I, 107
- ident*, 41
- identity, 109
- Identity Coercion, 442
- idtac, 269
- if ... then ... else, 70
- IF_then_else, 107
- iff, 107
- Implicit Arguments, 94
- Implicit arguments, 88
- Implicit Types, 98
- Import, 84
- Include, 79, 81
- induction, 209
- Inductive, 52, 55, 65
- Inductive (and coercions), 442
- Inductive definitions, 52
- Inductive ... where ..., 370
- Infix, 369
- Info, 277
- Info Auto, 237
- Info Eauto, 237
- Info Level, 277
- Info Trivial, 237
- injection, 217
- injection ... as, 218
- inl, 109
- inleft, 110
- Inline, 81
- inr, 109
- inright, 110
- Inspect, 153
- Instance, 464
- instantiate, 205
- integer*, 42
- Interpretation scopes, 374
- intro, 193
- intro after, 194
- intro at bottom, 194
- intro at top, 194
- intro before, 194
- Introduction patterns, 194
- intros, 193
- intros *intro_pattern*, 194
- intros until, 193, 194
- intuition, 247
- Intuition Iff Unfolding, 248
- Intuition Negation Unfolding, 248
- inversion, 219
- inversion ... as, 219
- inversion ... as ... in, 220
- inversion ... in, 220
- inversion ... using, 221, 389
- inversion ... using ... in, 221
- inversion_clear, 219
- inversion_clear ... as, 220
- inversion_clear ... as ... in, 220
- inversion_clear ... in, 220
- inversion_sigma, 221
- ι -reduction, 126, 138, 141
- is_evar, 251

- is_var, 251
- IsSucc, 111
- Keep Proof Equalities, 218, 219
- λ -calculus, 122
- lapply, 188
- last ... first (ssreflect), 324
- L^AT_EX, 413
- lazy, 230
- lazymatch
 - in Ltac, 272
- lazymatch goal
 - in Ltac, 273
- lazymatch reverse goal
 - in Ltac, 273
- le, 112
- le_n, 112
- le_S, 112
- left, 110, 192
- Lemma, 61
- length, 118
- Let, 52
- let
 - in Ltac, 270
- let ... := ... in ..., 46
- let '... in, 72
- let ... in, 71
- let rec
 - in Ltac, 270
- let-in, 46
- Let-in definitions, 46
- Lexical conventions, 41
- lia, 473, 475
- Libraries, 86
- Load, 161
- Load Verbose, 162
- Loadpath, 87
- Local, 172
- Local Arguments, 91, 93
- Local Axiom, 50
- Local Coercion, 441, 442
- Local context, 123
- local context, 173
- Local Definition, 51
- Local Set, 154
- Local Strategy, 171
- Local Unset, 154, 155
- Locate, 161, 370
- Locate File, 165
- Locate Library, 165
- Locate Ltac, 161
- Locate Module, 86
- Locate Term, 161
- Logical paths, 87
- Loose Hint Behavior, 245
- lra, 473
- lt, 112
- Ltac
 - eval, 274
 - fresh, 274
 - fun, 270
 - guard, 275
 - lazymatch, 272
 - lazymatch goal, 273
 - lazymatch reverse goal, 273
 - let, 270
 - let rec, 270
 - match, 271
 - match goal, 273
 - match reverse goal, 273
 - multimatch, 271
 - multimatch goal, 273
 - multimatch reverse goal, 273
 - numgoals, 275
 - type of, 274
 - type_term, 274
 - uconstr, 274
- Ltac, 276
- Ltac Batch Debug, 278
- Ltac Debug, 278
- Ltac Profiling, 279
- ltac:(...), 103
- Makefile, 400
- Man pages, 414
- map, 118
- match
 - in Ltac, 271
- match...with...end, 46, 70, 135
- match goal
 - in Ltac, 273
- match reverse goal
 - in Ltac, 273
- Maximal Implicit Insertion, 95

- ML-like patterns, 70, 427
- mod, 115
- Module, 79, 80
- Module Type, 80
- Modules, 78
- Monomorphic, 525
- move, 198
- move (ssreflect), 313
- move eq : ... (ssreflect), 320
- move/... (ssreflect), 351, 353, 355
- move: ... (ssreflect), 310
- move: ... => ... (ssreflect), 310
- move=> ... (ssreflect), 310
- mult, 111
- mult_n_O, 111
- mult_n_Sm, 111
- multimatch
 - in Ltac, 271
- multimatch goal
 - in Ltac, 273
- multimatch reverse goal
 - in Ltac, 273
- n_Sn, 111
- Naming introduction patterns, 194
- nat, 109
- nat_case, 113
- nat_double_ind, 113
- nat_scope, 115
- native code, 393
- native_compute, 230, 231
- Next Obligation, 491
- nia, 473, 476
- NonCumulative, 525
- None, 109
- Nonrecursive Elimination
 - Schemes, 385
- Normal form, 128
- not, 107
- not_eq_S, 111
- Notation, 365, 380
- Notations for lists, 118
- Notations for real numbers, 116
- notT, 114
- notypeclasses refine, 186
- now, 239
- nra, 473, 476
- nsatz, 505
- nth, 118
- num, 42
- numgoals
 - in Ltac, 275
- O, 109
- O_S, 111
- Obligation, 491
- Obligation Tactic, 491
- Obligations, 491
- Occurrences clauses, 184
- omega, 255, 469
- Omega Action, 471
- Omega System, 471
- Omega UseLocalDefs, 471
- once, 268
- Opaque, 170
- Open Scope, 374
- Optimize Heap, 181
- Optimize Proof, 181
- option, 109
- Options of the command line, 394
- or, 107
- or_introl, 107
- or_intror, 107
- pair, 109
- pairT, 114
- Parameter, 50
- Parameter (and coercions), 442
- Parameters, 50
- Parsing Explicit, 97
- pattern, 235
- Peano's arithmetic, 115
- Physical paths, 87
- plus, 111
- plus_n_O, 111
- plus_n_Sm, 111
- Polymorphic, 525
- Polymorphic Inductive
 - Cumulativity, 525
- pose, 201
- pose ... := ... (ssreflect), 305
- pose cofix ... := ... (ssreflect), 305

- pose fix ... := ...
 (ssreflect), 305
- pose proof, 203
- Positivity, 130
- Precedences, 366
- pred, 111
- pred_Sn, 111
- Preterm, 491
- Primitive Projections, 69
- Primitive projections, 69
- Print, 153
- Print All, 153
- Print All Dependencies, 156
- Print Assumptions, 156
- Print Canonical Projections, 98
- Print Classes, 443
- Print Coercion Paths, 443
- Print Coercions, 443
- Print Extraction Inline, 479
- Print Fields, 255
- Print Grammar constr, 367
- Print Grammar pattern, 367
- Print Graph, 443
- Print Hint, 244
- Print HintDb, 244
- Print Implicit, 96
- Print Libraries, 163
- Print LoadPath, 165
- Print Ltac, 277
- Print Ltac Signatures, 277
- Print ML Modules, 164
- Print ML Path, 165
- Print Module, 86
- Print Module Type, 86
- Print Opaque Dependencies, 156
- Print Options, 155
- Print Rings, 255
- Print Scope, 380
- Print Scopes, 380
- Print Section, 153
- Print Sorted Universes, 101
- Print Strategies, 171
- Print Strategy, 171
- Print Table Printing If, 74
- Print Table Printing Let, 73
- Print Tables, 155
- Print Term, 153
- Print Transparent Dependencies, 156
- Print Universes, 101
- Print Visibility, 379
- Printing All, 100
- Printing Coercions, 443
- Printing Compact Contexts, 169
- Printing Dependent Evars Line, 169
- Printing Depth, 168, 169
- Printing Existential Instances, 102
- Printing Implicit, 96
- Printing Implicit Defensive, 96
- Printing Matching, 72
- Printing Notations, 370
- Printing Primitive Projection
 Compatibility, 69
- Printing Primitive Projection
 Parameters, 69
- Printing Projections, 67
- Printing Records, 67
- Printing Synth, 73
- Printing Unfocused, 169
- Printing Universes, 101
- Printing Width, 168
- Printing Wildcard, 73
- prod, 109
- prodT, 114
- products, 46, 122
- Program, 487
- Program Cases, 488
- Program Definition, 489
- Program Fixpoint, 489
- Program Generalized Coercion, 488
- Program Instance, 460, 464
- Program Lemma, 490
- Programming, 109
- progress, 266
- proj1, 107
- proj2, 107
- projT1, 110
- projT2, 110
- Prompt, 173
- Proof, 62, 174
- Proof editing, 173
- PROOF GENERAL, 413

- Proof term, 173
- Proof using, 174
- Proof with, 245
- Prop**, 43, 121
- Proposition, 61
- psatz, 473
- Pwd, 164
- Qed, 62, 174
- Qed exporting, 276
- qualid*, 95
- Qualified identifiers, 86
- Quantifiers, 107
- Quit, 167
- quote, 254, 288
- Record, 65
- Record Elimination Schemes, 385
- Recursion**, 113
- Recursive arguments, 139
- Recursive Extraction, 477
- Recursive Extraction Library, 477
- red, 231
- Redirect, 167
- refine, 185
- Refine Instance Mode, 468
- refl_identity, 109
- reflexivity, 252
- Refolding Reduction, 234
- Regular Subst Tactic, 228
- Remark, 61
- remember, 201
- Remove Hints, 244
- Remove LoadPath, 165
- Remove Printing Coercion, 443
- Remove Printing If *ident*, 74
- Remove Printing Let *ident*, 73
- rename, 200
- repeat, 266
- replace ... with, 227
- Require, 162
- Require Export, 162
- Require Import, 162
- Reserved Notation, 369
- Reset, 165
- Reset Extraction Inline, 479
- Reset Initial, 166
- Reset Ltac Profile, 279
- Resource file, 394
- Restart, 176
- rev, 118
- Reversible Pattern Implicit, 94
- revert, 197
- revert dependent, 197
- revgoals, 258
- rewrite, 225
- rewrite →, 226
- rewrite <-, 226
- rewrite ... at, 226
- rewrite ... by, 226
- rewrite ... in, 226
- rewrite ... (ssreflect), 333
- rewrite_strat, 518
- Rewriting Schemes, 385
- right, 110, 192
- ring, 255, 493, 494
- ring_simplify, 255, 494
- rtauto, 248
- S, 109
- Scheme, 383
- Scheme Equality, 383
- Schemes**, 383
- Script file, 161
- Search, 156
- Search ... (ssreflect), 359
- Search Output Name Only, 168
- Search Output Name Only mode**, 168
- SearchAbout, 158
- SearchHead, 158
- SearchPattern, 159
- SearchRewrite, 160
- Section, 78
- Sections**, 77
- Separate Extraction, 477
- Set, 154
- Set**, 43, 121
- set, 200
- set ... := ... (ssreflect), 306
- setoid_reflexivity, 514
- setoid_replace, 514
- setoid_rewrite, 514
- setoid_symmetry, 514
- setoid_transitivity, 514
- shelve, 258

- shelve_unifiable, 259
- Short Module Printing, 86
- Show, 179
- Show Conjectures, 179
- Show Existentials, 180
- Show Intro, 179
- Show Intros, 180
- Show Ltac Profile, 279
- Show Match, 180
- Show Obligation Tactic, 491
- Show Proof, 179
- Show Script, 179
- Show Universes, 180
- Shrink Abstract, 276
- Shrink Obligations, 491
- sig, 110
- sig2, 110
- sigT, 110
- sigT2, 110
- Silent, 168
- Silent mode**, 168
- simpl, 232
- simpl ... in, 233
- simple apply, 188
- simple apply ... in, 191
- simple destruct, 209
- simple eapply ... in, 191
- simple induction, 213
- simple inversion, 221
- simple inversion ... as, 221
- simple notypeclasses refine, 186
- simple refine, 186
- simplify_eq, 252
- snd, 109
- solve, 268
- Solve Obligations, 491
- Some, 109
- sort, 44
- Sorts**, 43, 121
- specialize, 204
- specif*, 43
- split, 192
- split_Rabs, 117
- split_Rmult, 118
- Stable Omega, 471
- start ltac profiling, 280
- stepl, 228
- stepr, 229
- stop ltac profiling, 280
- Strategy, 171
- Strict Implicit, 94
- Strict Universe Declaration, 530
- string*, 42
- Strongly Strict Implicit, 94
- Structural Injection, 218
- Structure, 444
- SubClass, 443
- subgoal, 183
- subst, 228
- Substitution**, 123
- Subtyping rules, 127
- suff: ... (ssreflect), 330
- suffices: ... (ssreflect), 330
- Suggest Proof Using, 175
- sum, 109
- sumbool, 110
- sumor, 110
- swap, 257
- sym_not_eq, 108
- symmetry, 252
- symmetry in, 252
- tactic*, 183
- Tactic Compat Context, 272
- Tactic macros**, 259
- Tactic Notation, 381
- Tacticals**, 262
 - : , 265
 - tactic*₁; *tactic*₂, 262
 - tactic*₀; [*tactic*₁ | ... | *tactic*_{*n*}], 262
 - tactic*₀; [*tactic*₁ | ... | *tactic*_{*n*}], 262
 - abstract, 276
 - do, 266
 - exactly_once, 268
 - fail, 269
 - first, 267
 - gfail, 269
 - idtac, 269
 - once, 268
 - +, 267
 - ||, 267
 - repeat, 266
 - solve, 268
 - time, 270

- timeout, 269
- transparent_abstract, 276
- try, 266
- tryif, 268
- Tactics, 183**
- tail, 118
- tauto, 246
- Template polymorphism, 132**
- term*, 44
- Terms, 43**
- Test, 154, 155
- Test Printing If for *ident*, 74
- Test Printing Let for *ident*, 73
- Theorem, 61, 173
- Theories, 105**
- Time, 167
- time, 270
- Timeout, 167
- timeout, 269
- transitivity, 252
- Transparent, 170
- Transparent Obligations, 491
- transparent_abstract, 276
- trivial, 237
- True, 107
- true, 109
- try, 266
- tryif, 268
- tt, 109
- Type, 43, 121**
- type*, 43
- type of
 - in Ltac, 274
- type of constructor, 130
- type_scope, 377
- type_term
 - in Ltac, 274
- Typeclass Resolution After
 - Apply, 467
- Typeclass Resolution For
 - Conversion, 467
- Typeclasses Debug, 467
- Typeclasses Debug Verbosity, 467
- Typeclasses Dependency Order, 466
- Typeclasses eauto, 467
- typeclasses eauto, 465
- Typeclasses Filtered
 - Unification, 466
- Typeclasses Legacy Resolution, 466
- Typeclasses Limit Intros, 466
- Typeclasses Modulo Eta, 466
- Typeclasses Opaque, 465
- Typeclasses Strict Resolution, 467
- Typeclasses Transparent, 465
- Typeclasses Unique Instances, 467
- Typeclasses Unique Solutions, 467
- Typing rules, 124**
 - App, 125, 204
 - Ax-Prop, 124
 - Ax-Set, 124
 - Ax-Type, 124
 - Const, 124
 - Constr, 129
 - Conv, 127, 193, 229
 - Fix, 139
 - Ind, 129
 - Lam, 125, 193
 - Let, 125, 193
 - match, 138
 - Prod (impredicative Set), 143
 - Prod-Prop, 125
 - Prod-Set, 125
 - Prod-Type, 125
 - Var, 124
- uconstr
 - in Ltac, 274
- Undelimit Scope, 375
- Undo, 176
- Unfocus, 177
- Unfocused, 177
- unfold, 234
- unfold ...in, 234
- unify, 251
- unit, 109
- Universal Lemma Under
 - Conjunction, 190
- Universe, 529
- Universe Minimization ToSet, 528
- Universe Polymorphism, 525
- Universes**

presentation, 523
Unset, 154, 155
Unshelve, 259

value, 111
Variable, 50
variable, 123
Variable (and coercions), 442
Variables, 50
Variant, 52
vm_compute, 230, 231

Warnings, 168
Well founded induction, 113
Well foundedness, 113
well_founded, 113
without loss: ... / ...
 (ssreflect), 330
wlog: ... / ... (ssreflect),
 330

 ζ -reduction, 126

Tactics Index

`+`, 267
`||`, 267
`;`, 262
`:[...|...|...]`, 262
`[>...|...|...]`, 262
`... : ... (ssreflect)`, 314
`... => ... (ssreflect)`, 317
`... in ... (ssreflect)`, 310, 326
`;`, 265

`abstract`, 276
`abstract: ... (ssreflect)`, 317
`absurd`, 206
`admit`, 206
`apply`, 187
`apply ... in`, 190
`apply ... with`, 187
`apply/... (ssreflect)`, 357
`apply/.../... (ssreflect)`, 357
`apply: ... (ssreflect)`, 314
`assert`, 202
`assert as`, 203
`assert by`, 202
`assumption`, 185
`auto`, 236
`autoapply`, 465
`autorewrite`, 238
`autounfold`, 238

`btauto`, 254
`by ... (ssreflect)`, 323

`case`, 209
`case: ... (ssreflect)`, 313
`case: ... / ... (ssreflect)`, 320
`case_eq`, 209
`cbn`, 232
`cbv`, 230

`change`, 229
`change ... in`, 229
`classical_left`, 254
`classical_right`, 254
`clear`, 197
`clear dependent`, 197
`clearbody`, 197
`cofix`, 225
`compare`, 252
`compute`, 230, 231
`congruence`, 249
`constr_eq`, 251
`constructor`, 191
`contradict`, 207
`contradiction`, 207
`cut`, 204
`cutrewrite`, 227
`cycle`, 256

`decide equality`, 252
`decompose`, 202
`decompose record`, 202
`decompose sum`, 202
`dependent destruction`, 214
`dependent induction`, 213
`dependent induction ...`
 `generalizing`, 214
`dependent inversion`, 220
`dependent inversion ... as`, 220
`dependent inversion ... as ...`
 `with`, 221
`dependent inversion ... with`, 221
`dependent inversion_clear`, 221
`dependent inversion_clear ...`
 `as`, 221
`dependent inversion_clear ...`
 `as ... with`, 221
`dependent inversion_clear ...`
 `with`, 221

- dependent rewrite \rightarrow , 253
- dependent rewrite \leftarrow , 253
- destruct, 207
- dintuition, 247
- discriminate, 216
- discrR, 117
- do, 266
- do ... [...] (ssreflect), 325
- double induction, 213
- dtauto, 246
- eapply, 187
- eapply ... in, 191
- eassert, 203
- eassert as, 203
- eassert by, 203
- eassumption, 185
- easy, 239
- eauto, 237
- ecase, 209
- econstructor, 192
- edestruct, 208
- ediscriminate, 216
- eelim, 212
- eenough, 204
- eenough as, 204
- eenough by, 204
- eexact, 185
- eexists, 192
- einduction, 211
- einjection, 218
- eleft, 192
- elim ... using, 212
- elim/... (ssreflect), 349
- elim: ... (ssreflect), 313
- elimtype, 213
- enough, 203
- enough as, 203
- enough by, 204
- epose, 201
- epose proof, 203
- eremember, 201
- erewrite, 227
- eright, 192
- eset, 201
- esimplify_eq, 252
- esplit, 192
- evar, 205
- exact, 185
- exactly_once, 268
- exfalse, 207
- exists, 192
- f_equal, 251
- fail, 269
- field, 255, 500
- field_simplify, 255, 500
- field_simplify_eq, 255, 500
- first, 267
- first ... last (ssreflect), 324
- firstorder, 249
- firstorder using, 249
- firstorder with, 249
- firstorder *tactic*, 249
- fix, 224
- fold, 235
- fourier, 256
- functional induction, 215, 387
- functional inversion, 253
- generalize, 204
- generalize dependent, 205
- gfail, 269
- give_up, 206, 259
- has_evar, 251
- have: ... (ssreflect), 327
- have: ... := ... (ssreflect), 327
- hnf, 232
- idtac, 269
- induction, 209
- injection, 217
- injection ... as, 218
- instantiate, 205
- intro, 193
- intro after, 194
- intro at bottom, 194
- intro at top, 194
- intro before, 194
- intros, 193
- intros *intro_pattern*, 194
- intros until, 193, 194

intuition, 247
 inversion, 219
 inversion ... as, 219
 inversion ... as ... in, 220
 inversion ... in, 220
 inversion ... using, 221, 389
 inversion ... using ... in, 221
 inversion_clear, 219
 inversion_clear ... as, 220
 inversion_clear ... as ... in, 220
 inversion_clear ... in, 220
 inversion_sigma, 221
 is_evar, 251
 is_var, 251

 lapply, 188
 last ... first (ssreflect), 324
 lazy, 230
 left, 192
 lia, 473, 475
 lra, 473

 move, 198
 move (ssreflect), 313
 move eq : ... (ssreflect), 320
 move/... (ssreflect), 351, 353, 355
 move: ... (ssreflect), 310
 move: ... => ... (ssreflect), 310
 move=> ... (ssreflect), 310

 native_compute, 230, 231
 nia, 473, 476
 notypeclasses refine, 186
 now, 239
 nra, 473, 476
 nsatz, 505

 omega, 255, 469
 once, 268

 pattern, 235
 pose, 201
 pose ... := ... (ssreflect), 305
 pose cofix ... := ... (ssreflect), 305

 pose fix ... := ... (ssreflect), 305
 pose proof, 203
 progress, 266
 psatz, 473

 quote, 254, 288

 red, 231
 refine, 185
 reflexivity, 252
 remember, 201
 rename, 200
 repeat, 266
 replace ... with, 227
 revert, 197
 revert dependent, 197
 revgoals, 258
 rewrite, 225
 rewrite ->, 226
 rewrite <-, 226
 rewrite ... at, 226
 rewrite ... by, 226
 rewrite ... in, 226
 rewrite ... (ssreflect), 333
 rewrite_strat, 518
 right, 192
 ring, 255, 493, 494
 ring_simplify, 255, 494
 rtauto, 248

 set, 200
 set ... := ... (ssreflect), 306
 setoid_reflexivity, 514
 setoid_replace, 514
 setoid_rewrite, 514
 setoid_symmetry, 514
 setoid_transitivity, 514
 shelve, 258
 shelve_unifiable, 259
 simpl, 232
 simpl ... in, 233
 simple apply, 188
 simple apply ... in, 191
 simple destruct, 209
 simple eapply ... in, 191
 simple induction, 213
 simple inversion, 221

simple inversion ... as, 221
simple notypeclasses refine, 186
simple refine, 186
simplify_eq, 252
solve, 268
specialize, 204
split, 192
split_Rabs, 117
split_Rmult, 118
start ltac profiling, 280
stepl, 228
stepr, 229
stop ltac profiling, 280
subst, 228
suff: ... (ssreflect), 330
suffices: ... (ssreflect), 330
swap, 257
symmetry, 252
symmetry in, 252

tauto, 246
time, 270
timeout, 269
transitivity, 252
transparent_abstract, 276
trivial, 237
try, 266
tryif, 268
typeclasses eauto, 465

unfold, 234
unfold ...in, 234
unify, 251

vm_compute, 230, 231

without loss: ... / ...
 (ssreflect), 330
wlog: ... / ... (ssreflect),
 330

Vernacular Commands Index

- `*` (command), 177
- `+` (command), 177
- `-` (command), 177
- `{`, 177
- `}`, 177
- Abort, 176
- About, 153
- Add Field, 255, 501
- Add LoadPath, 164
- Add ML Path, 165
- Add Morphism, 515
- Add Parametric Morphism, 510
- Add Parametric Relation, 509
- Add Printing Coercion, 443
- Add Printing If *ident*, 74
- Add Printing Let *ident*, 73
- Add Rec LoadPath, 164
- Add Rec ML Path, 165
- Add Relation, 509
- Add Ring, 255, 496
- Add Setoid, 514
- Admit Obligations, 491
- Admitted, 63, 174
- Arguments, 90–92, 95, 232, 375
- Axiom, 50
- Axiom (and coercions), 442
- Back, 166
- BackTo, 166
- Backtrack, 166
- Bind Scope, 376
- Canonical Structure, 97
- Cd, 164
- Check, 155
- Class, 463
- Close Scope, 374
- Coercion, 100, 441, 442
- CoFixpoint, 60, 62
- CoFixpoint ... where ..., 370
- CoInductive, 57, 65
- CoInductive (and coercions), 442
- Collection, 175
- Combined Scheme, 385
- Compute, 156
- Conjecture, 50
- Constraint, 529
- Context, 464
- Corollary, 61
- CreateHintDb, 239
- Cumulative, 525
- Declare Implicit Tactic, 246
- Declare Instance, 464
- Declare Left Step, 228
- Declare ML Module, 163
- Declare Module, 81
- Declare Right Step, 229
- Defined, 63, 174
- Definition, 51
- Delimit Scope, 375
- Derive, 531
- Derive Dependent Inversion, 389
- Derive Dependent
 Inversion_clear, 389
- Derive Inversion, 389
- Derive Inversion_clear, 389
- Drop, 167
- End, 78, 80, 81
- Eval, 155
- Example, 51
- Existential, 176
- Existing Class, 463
- Existing Instance, 464
- Existing Instances, 464
- Export, 85

- Extract Constant, 480
- Extract Inductive, 481
- Extraction, 156, 477
- Extraction Blacklist, 482
- Extraction Implicit, 480
- Extraction Inline, 479
- Extraction Language, 478
- Extraction Library, 477
- Extraction NoInline, 479

- Fact, 61
- Fixpoint, 57, 62
- Fixpoint ... where ..., 370
- Focus, 177
- From Require, 163
- Function, 75
- Functional Scheme, 386

- Generalizable Variables, 99
- Global, 172
- Global Arguments, 91, 92
- Global Set, 154
- Global Unset, 154, 155
- Goal, 173
- Grab Existential Variables, 176
- Guarded, 180

- Hint, 239
- Hint Constructors, 241
- Hint Cut, 242
- Hint Extern, 242
- Hint Immediate, 240
- Hint Mode, 243
- Hint Opaque, 241
- Hint Resolve, 240
- Hint Rewrite, 244
- Hint Transparent, 241
- Hint Unfold, 241
- Hint View (ssreflect), 358
- Hypotheses, 50
- Hypothesis, 50
- Hypothesis (and coercions), 442

- Identity Coercion, 442
- Implicit Types, 98
- Import, 84
- Include, 79, 81
- Inductive, 52, 55, 65
- Inductive (and coercions), 442
- Inductive ... where ..., 370
- Infix, 369
- Info, 277
- Inline, 81
- Inspect, 153
- Instance, 464

- Lemma, 61
- Let, 52
- Load, 161
- Load Verbose, 162
- Local, 172
- Local Arguments, 91, 93
- Local Axiom, 50
- Local Coercion, 441, 442
- Local Definition, 51
- Local Set, 154
- Local Strategy, 171
- Local Unset, 154, 155
- Locate, 161, 370
- Locate File, 165
- Locate Library, 165
- Locate Ltac, 161
- Locate Module, 86
- Locate Term, 161
- Ltac, 276

- Module, 79, 80
- Module Type, 80
- Monomorphic, 525

- Next Obligation, 491
- NonCumulative, 525
- Notation, 365, 380

- Obligation, 491
- Obligation Tactic, 491
- Obligations, 491
- Opaque, 170
- Open Scope, 374
- Optimize Heap, 181
- Optimize Proof, 181

- Parameter, 50
- Parameter (and coercions), 442
- Parameters, 50
- Polymorphic, 525

- Preterm, [491](#)
- Print, [153](#)
- Print All, [153](#)
- Print All Dependencies, [156](#)
- Print Assumptions, [156](#)
- Print Canonical Projections, [98](#)
- Print Classes, [443](#)
- Print Coercion Paths, [443](#)
- Print Coercions, [443](#)
- Print Extraction Inline, [479](#)
- Print Fields, [255](#)
- Print Grammar constr, [367](#)
- Print Grammar pattern, [367](#)
- Print Graph, [443](#)
- Print Hint, [244](#)
- Print HintDb, [244](#)
- Print Implicit, [96](#)
- Print Libraries, [163](#)
- Print LoadPath, [165](#)
- Print Ltac, [277](#)
- Print Ltac Signatures, [277](#)
- Print ML Modules, [164](#)
- Print ML Path, [165](#)
- Print Module, [86](#)
- Print Module Type, [86](#)
- Print Opaque Dependencies, [156](#)
- Print Options, [155](#)
- Print Rings, [255](#)
- Print Scope, [380](#)
- Print Scopes, [380](#)
- Print Section, [153](#)
- Print Sorted Universes, [101](#)
- Print Strategies, [171](#)
- Print Strategy, [171](#)
- Print Table Printing If, [74](#)
- Print Table Printing Let, [73](#)
- Print Tables, [155](#)
- Print Term, [153](#)
- Print Transparent Dependencies, [156](#)
- Print Universes, [101](#)
- Print Visibility, [379](#)
- Program Definition, [489](#)
- Program Fixpoint, [489](#)
- Program Instance, [460](#), [464](#)
- Program Lemma, [490](#)
- Proof, [62](#), [174](#)
- Proof using, [174](#)
- Proof with, [245](#)
- Proposition, [61](#)
- Pwd, [164](#)
- Qed, [62](#), [174](#)
- Qed exporting, [276](#)
- Quit, [167](#)
- Record, [65](#)
- Recursive Extraction, [477](#)
- Recursive Extraction Library, [477](#)
- Redirect, [167](#)
- Remark, [61](#)
- Remove Hints, [244](#)
- Remove LoadPath, [165](#)
- Remove Printing Coercion, [443](#)
- Remove Printing If *ident*, [74](#)
- Remove Printing Let *ident*, [73](#)
- Require, [162](#)
- Require Export, [162](#)
- Require Import, [162](#)
- Reserved Notation, [369](#)
- Reset, [165](#)
- Reset Extraction Inline, [479](#)
- Reset Initial, [166](#)
- Reset Ltac Profile, [279](#)
- Restart, [176](#)
- Scheme, [383](#)
- Scheme Equality, [383](#)
- Search, [156](#)
- Search ... (ssreflect), [359](#)
- SearchAbout, [158](#)
- SearchHead, [158](#)
- SearchPattern, [159](#)
- SearchRewrite, [160](#)
- Section, [78](#)
- Separate Extraction, [477](#)
- Set, [154](#)
- Show, [179](#)
- Show Conjectures, [179](#)
- Show Existentials, [180](#)
- Show Intro, [179](#)
- Show Intros, [180](#)
- Show Ltac Profile, [279](#)
- Show Match, [180](#)
- Show Obligation Tactic, [491](#)

Show Proof, [179](#)
Show Script, [179](#)
Show Universes, [180](#)
Solve Obligations, [491](#)
Strategy, [171](#)
Structure, [444](#)
SubClass, [443](#)

Tactic Notation, [381](#)
Test, [154](#), [155](#)
Test Printing If for *ident*, [74](#)
Test Printing Let for *ident*, [73](#)
Theorem, [61](#), [173](#)
Time, [167](#)
Timeout, [167](#)
Transparent, [170](#)
Typeclasses eauto, [467](#)
Typeclasses Opaque, [465](#)
Typeclasses Transparent, [465](#)

Undelimit Scope, [375](#)
Undo, [176](#)
Unfocus, [177](#)
Unfocused, [177](#)
Universe, [529](#)
Unset, [154](#), [155](#)
Unshelve, [259](#)

Variable, [50](#)
Variable (and coercions), [442](#)
Variables, [50](#)
Variant, [52](#)

Vernacular Options Index

Asymmetric Patterns, [431](#)
Automatic Coercions Import, [444](#)
Automatic Introduction, [181](#)

Boolean Equality Schemes, [385](#)
Bracketing Last Introduction
 Pattern, [196](#)
Bullet Behavior, [178](#)

Case Analysis Schemes, [385](#)
Congruence Verbose, [250](#)
Contextual Implicit, [94](#)

Debug Auto, [237](#)
Debug Cbv, [231](#)
Debug Eauto, [237](#)
Debug RAKAM, [234](#)
Debug Trivial, [237](#)
Decidable Equality Schemes, [385](#)
Default Goal Selector, [183](#)
Default Proof Using, [175](#)
Default Timeout, [167](#), [168](#)

Elimination Schemes, [385](#)
Extraction AutoInline, [479](#)
Extraction Conservative Types,
 [479](#)
Extraction KeepSingleton, [479](#)
Extraction Optimize, [479](#)
Extraction SafeImplicits, [480](#)

Firstorder Depth, [249](#)
Firstorder Solver, [249](#)

Hide Obligations, [491](#)
Hyps Limit, [180](#)

Implicit Arguments, [94](#)
Info Auto, [237](#)
Info Eauto, [237](#)
Info Level, [277](#)
Info Trivial, [237](#)
Intuition Iff Unfolding, [248](#)
Intuition Negation Unfolding, [248](#)

Keep Proof Equalities, [218](#), [219](#)

Loose Hint Behavior, [245](#)
Ltac Batch Debug, [278](#)
Ltac Debug, [278](#)
Ltac Profiling, [279](#)

Maximal Implicit Insertion, [95](#)

Nonrecursive Elimination
 Schemes, [385](#)

Omega Action, [471](#)
Omega System, [471](#)
Omega UseLocalDefs, [471](#)

Parsing Explicit, [97](#)
Polymorphic Inductive
 Cumulativity, [525](#)
Primitive Projections, [69](#)
Printing All, [100](#)
Printing Coercions, [443](#)
Printing Compact Contexts, [169](#)
Printing Dependent Evars Line,
 [169](#)
Printing Depth, [168](#), [169](#)
Printing Existential Instances,
 [102](#)
Printing Implicit, [96](#)
Printing Implicit Defensive, [96](#)
Printing Matching, [72](#)
Printing Notations, [370](#)
Printing Primitive Projection
 Compatibility, [69](#)

- Printing Primitive Projection
 - Parameters, [69](#)
- Printing Projections, [67](#)
- Printing Records, [67](#)
- Printing Synth, [73](#)
- Printing Unfocused, [169](#)
- Printing Universes, [101](#)
- Printing Width, [168](#)
- Printing Wildcard, [73](#)
- Program Cases, [488](#)
- Program Generalized Coercion, [488](#)
- Record Elimination Schemes, [385](#)
- Refine Instance Mode, [468](#)
- Refolding Reduction, [234](#)
- Regular Subst Tactic, [228](#)
- Reversible Pattern Implicit, [94](#)
- Rewriting Schemes, [385](#)
- Search Output Name Only, [168](#)
- Short Module Printing, [86](#)
- Shrink Abstract, [276](#)
- Shrink Obligations, [491](#)
- Silent, [168](#)
- Stable Omega, [471](#)
- Strict Implicit, [94](#)
- Strict Universe Declaration, [530](#)
- Strongly Strict Implicit, [94](#)
- Structural Injection, [218](#)
- Suggest Proof Using, [175](#)
- Tactic Compat Context, [272](#)
- Transparent Obligations, [491](#)
- Typeclass Resolution After
 - Apply, [467](#)
- Typeclass Resolution For
 - Conversion, [467](#)
- Typeclasses Debug, [467](#)
- Typeclasses Debug Verbosity, [467](#)
- Typeclasses Dependency Order, [466](#)
- Typeclasses Filtered
 - Unification, [466](#)
- Typeclasses Legacy Resolution, [466](#)
- Typeclasses Limit Intros, [466](#)
- Typeclasses Modulo Eta, [466](#)
- Typeclasses Strict Resolution, [467](#)
- Typeclasses Unique Instances, [467](#)
- Typeclasses Unique Solutions, [467](#)
- Universal Lemma Under
 - Conjunction, [190](#)
- Universe Minimization ToSet, [528](#)
- Universe Polymorphism, [525](#)
- Warnings, [168](#)

Index of Error Messages

- ident* already exists, 50–52, 63, 489
- already exists, 61
- Argument of match does not evaluate to a term, 271
- arguments of `ring_simplify` do not have all the same type, 495
- Attempt to save an incomplete proof, 174
- bad lemma for decidability of equality, 498
- Bad magic number, 163
- bad ring structure, 498
- Can't find file *ident* on loadpath, 162
- cannot be used as a hint, 240, 241
- Cannot build functional inversion principle, 76
- Cannot define graph for *ident*..., 76
- Cannot define principle(s) for *ident*..., 76
- cannot find a declared ring structure for equality term, 495
- cannot find a declared ring structure over term, 495
- Cannot find induction information on *qualid*, 216
- Cannot find inversion information for hypothesis *ident*, 253
- Cannot find library foo in loadpath, 163
- Cannot find the source class of *qualid*, 441
- Cannot handle mutually (co)inductive records., 69
- Cannot infer a term for this placeholder, 90, 186
- Cannot load *qualid*: no physical path bound to *dirpath*, 163
- Cannot move *ident*₁ after *ident*₂: it depends on *ident*₂, 198
- Cannot move *ident*₁ after *ident*₂: it occurs in the type of *ident*₂, 198
- Cannot recognize *class*₁ as a source class of *qualid*, 441
- Cannot solve the goal, 268
- Cannot use mutual definition with well-founded recursion or measure, 76
- Compiled library *ident.vo* makes inconsistent assumptions over library *qualid*, 163
- Condition not satisfied, 276
- does not denote an evaluable constant, 234
- does not respect the uniform inheritance condition, 441
- Failed to progress, 266
- File not found on loadpath : , 164
- Found target class *class* instead of *class*₂, 441
- Funclass cannot be a source class, 441
- goal does not satisfy the expected preconditions, 218
- Goal is solvable by congruence but some arguments are missing. Try congruence with ..., replacing metavariables by arbitrary terms., 250
- Hypothesis *ident* must contain at least one Function, 253
- I don't know how to handle dependent equality, 250
- In environment ... the term: *term*₂ does not have type *term*₁, 489
- invalid argument, 186
- Invalid backtrack, 166
- is already a coercion, 441
- is already used, 193, 200
- is not a function, 441
- is not a local definition, 197
- is not a module, 85
- is not an inductive type, 241
- is used in conclusion, 204

- is used in hypothesis, 204
- is used in the conclusion, 197
- is used in the hypothesis, 197
- Loading of ML object file forbidden in a native COQ, 164
- Module/section *module* not found, 158
- must be a transparent constant, 442
- name *ident* is already used, 193
- No applicable tactic, 267
- No argument name *ident*, 76
- No discriminable equalities, 217
- No evars, 251
- No focused proof, 173, 179
- No focused proof (No proof-editing in progress), 176
- No focused proof to restart, 176
- No matching clauses for match, 271
- No matching clauses for match goal, 273
- No primitive equality found, 216
- No product even after head-reduction, 193
- No progress made, 518
- No such assumption, 185, 207
- No such binder, 184
- no such entry, 166
- No such goal, 179, 266
- No such goal. Focus next goal with bullet *bullet.*, 178
- No such goal. Try unfocusing with *}.*, 178
- No such hypothesis, 194, 197, 198, 200, 236
- No such hypothesis in current goal, 193, 194
- No such label *ident*, 80
- Non exhaustive pattern-matching, 437
- Non strictly positive occurrence of *ident* in *type*, 54
- not a context variable, 274
- not a defined object, 153
- Not a discriminable equality, 216
- Not a primitive equality, 218
- Not a projectable equality but a discriminable one, 218
- Not a proposition or a type, 202
- not a valid ring equation, 495
- Not a variable or hypothesis, 251
- Not an evar, 251
- Not an exact proof, 185
- Not an inductive goal with 1 constructor, 192
- Not an inductive goal with 2 constructors, 192
- Not an inductive product, 191, 210
- Not convertible, 229
- not declared, 441
- Not enough constructors, 191
- Not equal, 251
- Not reducible, 231
- Not the right number of induction arguments, 216
- Not the right number of missing arguments, 184, 187
- Not unifiable, 251
- Nothing to do, it is an equality between convertible terms, 218
- Nothing to inject, 218
- Nothing to rewrite, 518
- omega can't solve this system, 470
- omega: Can't solve a goal with equality on *type*, 470
- omega: Can't solve a goal with non-linear products, 470
- omega: Can't solve a goal with proposition variables, 470
- omega: Not a quantifier-free goal, 470
- omega: Unrecognized atomic proposition: *prop*, 470
- omega: Unrecognized predicate or connective: *ident*, 470
- omega: Unrecognized proposition, 470
- Proof is not complete, 202, 276
- quote: not a simple fixpoint, 254, 289
- Records declared with the keyword *Record* or *Structure* cannot be recursive., 68
- Require is not allowed inside a module or a module type, 163
- ring *operation* should be declared as a morphism, 498
- Signature components for label *ident* do not match, 80
- Sortclass cannot be a source class, 441
- Statement without assumptions, 190
- Tactic Failure *message* (level *n*), 269

Tactic generated a subgoal identical to the original goal, 226

terms do not have convertible types, 227

The conclusion is not a substitutive equation, 252

The conclusion of *type* is not valid; it must be built from *ident*, 54

The file *ident.vo* contains library *dirpath* and not library *dirpath'*, 163

The recursive argument must be specified, 76

The reference *qualid* was not found in the current environment, 156, 170

The term provided does not end with an equation, 226

The term *form* has type ... which should be Set, Prop or Type, 61

The term *term* has type *type* while it is expected to have type *type*, 51

The variable *ident* is already defined, 200

The *numth* argument of *ident* must be *ident'* in *type*, 55

This is not the last opened module, 80

This is not the last opened module type, 81

This is not the last opened section, 78

This proof is focused, but cannot be unfocused this way, 177

This tactic has more than one success, 268

Unable to apply, 190

Unable to find an instance for the variables *ident* ... *ident*, 187

Unable to find an instance for the variables *ident* ... *ident*, 210

Unable to satisfy the rewriting constraints, 518

Unable to unify ... with ..., 187, 252

Undeclared universe *ident*, 529

Universe inconsistency, 122, 529

Unknown inductive type, 180

Variable *ident* is already declared, 203

Wrong bullet *bullet1* : Bullet *bullet2* is mandatory here., 178

Wrong bullet *bullet1* : Current bullet *bullet2* is not finished., 178

List of Figures

1.1	Syntax of terms	44
1.2	Syntax of terms (continued)	45
1.3	Syntax of sentences	49
2.1	Syntax for the definition of <code>Record</code>	65
2.2	Syntax for constructing elements of a <code>Record</code> using named fields	66
2.3	Syntax for <code>Record</code> projections	68
2.4	Syntax of modules	79
2.5	Syntax for explicitly giving implicit arguments	95
3.1	Notations in the initial state	106
3.2	Syntax of formulas	106
3.3	Syntax of data-types and specifications	109
3.4	Definition of the scope for integer arithmetics (<code>Z_scope</code>)	116
3.5	Definition of the scope for natural numbers (<code>nat_scope</code>)	116
3.6	Definition of the scope for real arithmetics (<code>R_scope</code>)	117
3.7	Definition of the scope for lists (<code>list_scope</code>)	118
9.1	Syntax of the tactic language	263
9.2	Syntax of the tactic language (continued)	264
9.3	Tactic toplevel definitions	265
10.1	Definition of the permutation predicate	292
10.2	Permutation tactic	293
10.3	Deciding intuitionistic propositions (1)	294
10.4	Deciding intuitionistic propositions (2)	295
10.5	Type isomorphism axioms	296
10.6	Type isomorphism tactic (1)	297
10.7	Type isomorphism tactic (2)	298
12.1	Syntax of the variants of <code>Notation</code>	371
16.1	COQIDE main screen	416
16.2	COQIDE: the query window	418
18.1	Syntax of classes	440