



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

A Small Scale Reflection Extension for the COQ system

Georges Gonthier — Assia Mahboubi

MICROSOFT RESEARCH
INRIA



JOINT
CENTRE

N° ????

December 2007

Thème SYM

*R*apport
de recherche



A Small Scale Reflection Extension for the COQ system

Georges Gonthier* , Assia Mahboubi†

Thème SYM — Systèmes symboliques
Projets Composants Mathématiques, Centre Commun INRIA Microsoft Research

Rapport de recherche n° ???? — December 2007 — 49 pages

Abstract: This document describes a set of extensions to the proof scripting language of the COQ proof assistant. While these extensions were developed to support a particular proof methodology - small-scale reflection - most of them actually are of a quite general nature, improving the functionality of COQ in basic areas such as script layout and structuring, proof context management, and rewriting. Consequently, and in spite of the title of this document, most of the extensions described here should be of interest for all COQ users, whether they embrace small-scale reflection or not.

Key-words: proof assistants, formal proofs, Coq, small scale reflection, tactics

* Microsoft Research, Cambridge, R-U, Centre commun INRIA Microsoft Research

† INRIA Futurs, Centre commun INRIA Microsoft Research

A Small Scale Reflection Extension for the COQ system

Résumé : Ce rapport présente une extension de l'assistant à la preuve COQ, développé par G. Gonthier. Cette extension a été conçue pour améliorer le support d'une méthodologie de preuve formelle, appelée réflexion à petite échelle. Néanmoins, la majeure partie de ses apports sont des améliorations d'ordre général des fonctionnalités du système COQ comme la structuration des scripts, la gestion des contextes de preuve, et la réécriture. C'est pourquoi, en dépit du titre de ce document, la plupart des fonctionnalités décrites ici sont susceptibles d'intéresser tout utilisateur de COQ, qu'il utilise ou non les techniques de réflexion à petite échelle.

Mots-clés : assistants à la preuve, preuve formelle, Coq, réflexion à petite échelle, tactiques

Small-scale reflection is a formal proof methodology based on the pervasive use of computation with symbolic representations. Symbolic representations are usually hidden in traditional computational reflection (e.g., as used in the COQ `ring`, `romega` or `setoid` tactics): they are generated on-the-fly by some heuristic algorithm and directly fed to some decision or simplification procedure whose output is translated back to "logical" form before being displayed to the user. By contrast, in small-scale reflection symbolic representations are ubiquitous; the statements of many top-level lemmas, and of most proof subgoals, explicitly contain symbolic representations; translation between logical and symbolic representations is performed under the explicit, fine-grained control of the proof script.

The efficiency of small-scale reflection hinges on the fact that fixing a particular symbolic representation strongly directs the behaviour of a theorem-prover:

- Logical case analysis is done by enumerating the symbols according to their inductive type: the representation describes which cases should be considered.
- Many logical functions and predicates are represented by concrete functions on the symbolic representation, which can be computed once (part of) the symbolic representation of objects is known: the representation describes what should be done in each case.

Thus by controlling the representation we also control the automated behaviour of the theorem prover, which can be quite complex, for example if a predicate is represented by a sophisticated decision procedure. The real strength of small-scale reflection, however, is that even very simple representations provide useful procedures. For example, the truth-table representation of connectives, evaluated left-to-right on the Boolean representation of propositions, provides sufficient automation for most propositional reasoning.

Small-scale reflection defines a basis for dividing the proof workload between the user and the prover: the prover engine provides computation and database functions (via partial evaluation, and definition and type lookup, respectively), and the user script guides the execution of these functions, step by step. User scripts comprise three kinds of steps:

- Deduction steps directly specify part of the construction of the proof, either top down (so-called forward steps), or bottom-up (backward steps). A reflection step that switches between logical and symbolic representation is just a special kind of deductive step.
- Bookkeeping steps manage the proof context, introducing, renaming, discharging, or splitting constants and assumptions. Case-splitting on symbolic representations is an efficient way to drive the prover engine, because most of the data required for the splitting can be retrieved from the representation type, and because specialising a single representation often triggers the evaluation of several representation functions.
- Rewriting steps use equations to locally change parts of the goal or assumptions. Rewriting is often used to complement partial evaluation, bypassing unknown parameters (e.g., simplifying `b && false` to `false`). Obviously, it's also used to implement equational reasoning at the logical level, for instance, switching to different representation.

It is a characteristic of the small-scale reflection style that the three kinds of steps are roughly equinumerous, and interleaved; there are strong reasons for this, chief among them the fact that goals and contexts tend to grow rapidly through the partial evaluation of representations. This makes it impractical to embed most intermediate goals in the proof script - the so-called declarative style of proof, which hinges on the exclusive use of forward steps. This also means that subterm selection, especially in rewriting, is often an issue.

The basic COQ tactic language is not well adapted to small-scale reflection proofs. It is heavily biased towards backward steps, with little support for forward steps, or even script layout (these are deferred to the "vernacular", i.e., `Section/Module` layer of the input language). The support for rewriting is primitive, requiring a separate tactic for each basic step, and considerable user expertise (relying on undocumented system behaviour) to perform even simple subterm selection.

Many of the basic tactics, such as `intros`, `induction` and `inversion`, implement fragile context manipulation heuristics which hinder precise bookkeeping; on the other hand the under-utilised "intro patterns" provide excellent support for case splitting.

The extensions presented here were designed to improve the functionality of COQ in all those areas, providing:

- support for better script layout
- better support for forward steps
- common support for bookkeeping in all tactics
- common support for subterm selection in all tactics
- a unified interface for rewriting, definition expansion, and partial evaluation
- improved robustness with respect to evaluation and conversion
- support for reflection steps.

We should point out that only the last functionality is specific to small-scale reflection. All the others are of general use. Moreover most of these features are introduced not by adding new tactics, but by extending the functionality of existing ones: indeed we introduce only three new tactics, rename three others, but all subsume more than a dozen of the basic COQ tactics.

How to read this documentation

The syntax of the tactics is presented as follows:

- `terminals` are in typewriter font and $\langle non\ terminals \rangle$ are between angle brackets.
- Optional parts of the grammar are surrounded by `[]` brackets. These should not be confused with verbatim brackets `[|]`, which are mandatory delimiters.
- A vertical rule `|` indicates an alternative in the syntax, and should not be confused with a verbatim vertical rule between verbatim brackets `[|]`.
- A non empty list of non terminals (at least one item should be present) is represented by $\langle non\ terminals \rangle^+$. A possibly empty one is represented by $\langle non\ terminals \rangle^*$.
- In a list of non terminals, items are separated by blanks.

We follow the color scheme of the SSREFLECT mode for ProofGeneral (see section 1.1):

`tactic` or `Command` or `keyword` or `tactical`

Closing tactics/tacticals like `exact` or `by` (see section 5.2) are in red.

1 Usage

The implementation of the small-scale reflection extensions comprises a Caml extension module (`ssreflect.ml4`), which provides the tactic language, and a COQ vernacular file (`ssreflect.v`) which provides the basic theory of reflection (essentially, the `reflect` predicate and related lemmas), along with a few additional tactics and notations. The `ssrbool.v` file contains an extended toolkit for boolean reflection.

1.1 Installing SSREFLECT

We propose hereafter a way of compiling and installing (locally) SSREFLECT from standard COQ sources. Experimented COQ users may of course prefer and adopt different architecture choices.

1.1.1 Linux, Mac OS

You need Objective Caml version 3.06 or later installed on your computer. Make sure that you do not have any environment variable named `COQLIB` or `COQBIN`. Download and unpack the sources of the current version 8.1p12 of COQ here. The sample directory and file names given above refer to the `coq-8.1p12.tar.gz` archive.

The following instructions lead to a local installation of COQ and SSREFLECT.

- Unpack the archive. This results in creating a directory called `coq-8.1p12`.
- In this directory do :


```
./configure -opt -local
```

 (this builds the Makefile for the local installation)


```
make world
```

 (this compiles the whole stuff)
- If you get an error message complaining about `coqide`, skip it, you still have succeeded in compiling all what is necessary for SSREFLECT.
- In the `.../coq-8.1p12/bin` directory, you should have a `coqtop` file, which is the COQ toplevel.
- Set an environment variable `COQTOP` and a variable `COQLIB` both to `.../coq-8.1p12/`, and an environment variable `COQBIN` to `.../coq-8.1p12/bin`. Add


```
.../coq-8.1p12/bin
```

 to your `PATH` variable.
- Copy the `ssreflect.m14` file in the `.../coq-8.1p12` directory.
- In `.../coq-8.1p12`, do :


```
make ssreflect.cmx
```

 Note for Mac OS users on a PowerPC: If you did not include the `-opt` argument to `./configure`, then replace the above command by `make OCAMLOPT=ocamlopt.opt ssreflect.cmx`. Otherwise you may encounter a "stack overflow" error message.
- Now you can build the SSREFLECT toplevel, by doing in `.../coq-8.1p12`:


```
coqmktop -opt -o ssrcoq ssreflect.cmx
```
- If the system complains about a missing `gramlib.cmxa`, find the place where this `camlp4` library is installed (typing for example `locate gramlib.cmxa`), and try again the command:


```
coqmktop -opt -R /the/address/of/the/rep/containing/gramlib/ -o ssrcoq
ssreflect.cmx
```
- Move `ssrcoq` to the `bin` directory.

1.1.2 Windows + Cygwin

You need Objective Caml version 3.06 or later installed under cygwin. You need Emacs for Windows installed on your computer. Make sure that you do not have any environment variable, neither in cygwin nor in Windows named `COQLIB` or `COQBIN` or `COQLIB`. Download the sources of the current version of COQ here. The following instructions lead to a local installation of COQ and SSREFLECT. The sample directory and file names given above refer to the `coq-8.1p12.tar.gz` archive.

- Open a Cygwin console. Using the console:
- Unpack the archive. This results in creating a directory called `coq-8.1p12`.
- WARNING : The place where you unpack the sources should not be under a directory containing a space in its name (like "Documents and Settings" ...) otherwise the next step will fail.
- In this directory do:

```
./configure -opt -local
```

(this builds the Makefile for the local installation). Then do:

```
make world
```

(this compiles the whole stuff).

If you get an error message complaining about `coqide`, skip it, you still have succeeded in compiling all what is necessary for SSREFLECT.

- In the `.../coq-8.1p12/bin` directory, you should have a `coqtop` file, which is the COQ toplevel.
- Set an environment variable `COQTOP` and a variable `COQLIB` both to `.../coq-8.1p12/`
- Set an environment variable `COQBIN` to `.../coq-8.1p12/bin`
- Add the latter to your `PATH` variable. You may also add these variables as Windows environment variables and update your Windows path. Also make sure that the `PATH` variable of the system contains your directory `.../cygwin/bin`
- Copy the `ssreflect.ml4` file in the `.../coq-8.1p12` directory.
- In `.../coq-8.1p12`, do :

```
make ssreflect.cmx
```

- Now you can build the SSREFLECT toplevel, by doing in `.../coq-8.1p12`:

```
coqmktop -opt -o ssrcoq ssreflect.cmx
```

- If the system complains about a missing `gramlib.cmxa`, find the place where this `camlp4` library is installed (typing for example `locate gramlib.cmxa`), and try again the tactic :

```
coqmktop -opt -R /the/address/of/the/rep/containing/gramlib/ -o ssrcoq
ssreflect.cmx
```

- Move `ssrcoq` to the `bin` directory.

1.1.3 Proof General

SSREFLECT comes with a small configuration file `pg-ssr.el` to extend ProofGeneral (PG), a generic interface for proof assistants based on the customisable text editor Emacs, to the syntax of SSREFLECT. Following the installation instructions, unpack the sources of PG in a directory, for instance `.../ProofGeneral-3.7` and add to your `.emacs` file the following line:

```
(load-file ".../ProofGeneral-3.7/generic/proof-site.el" )
```

or

```
(load-file "C:\\...\\ProofGeneral-3.7\\generic\\proof-site.el")
```

if you are running Windows+Cygwin, where `...` is the location of your own ProofGeneral directory.

Immediately after the previous line, add this one:

```
(load-file ".../pg-ssr.el")
```

respectively

```
(load-file "...\\pg-ssr.el")
```

for Windows+Cygwin users where `...` is the location of your `ssr.el` file.

COQ sources have a `.v` extension. Opening any `.v` file should automatically launch ProofGeneral. Try this on a `foo.v` file. In the menu `ProofGeneral`, choose `Advanced...`, then `Customize`, then `COQ`, then `Coq Prog Name`. Change the value of the variable to `.../bin/ssrcoq` (resp. `...\\bin\\ssrcoq` for Windows+Cygwin users), where `...` is the location of your `coq-8.1pl2` directory.

1.2 Using SSREFLECT

You can use this `ssrcoq` executable to compile vernacular (`.v`) files with the `-compile` flag (this also avoids the synchronisation problems for `coqc` under the Windows OS):

```
ssrcoq -compile ssreflect.v
```

Every COQ vernacular file processed by `ssrcoq` should import the SSREFLECT library with the line:

```
Require Import ssreflect.
```

The tactics described below will not work properly without this library. (Obviously, the `ssreflect.vo` file created by the tactic line above also needs to be in the library path.)

The sample library `ssrbool` provides reflection and algebraic lemmas for the `bool` type.

You should also configure your proof IDE to use SSREFLECT (through Emacs customisation for Proof General, and by linking with `ssrcoq` for CoqIDE), and to recognise and highlight new keywords, tactical and tactic names (`is`, `last`, `by`, `done`, `have`, `suffices`, `without loss`, `congr`, `unlock`), and vernacular tactics (`Prenex Implicits`, `Import Prenex Implicits`). You may also find it helpful to use stronger highlighting for tactic and tacticals that terminate proof branches (`assumption`, `by`, `contradiction`, `discriminate`, `done`, `exact`, `omega`, `reflexivity`, `solve`, `tauto`), as this makes it easier to spot layout errors.

1.3 Compatibility issues

Every effort has been made to make the small-scale reflection extensions upward compatible with the basic COQ, but a few discrepancies were unavoidable:

- New keywords (`is`, `by`) might clash with variable, constant, tactic or tactical names, or with quasi-keywords in tactic or vernacular notation.
- New tactic(al)s names (`last`, `done`, `have`, `suffices`, `without loss`, `congr`, `unlock`) might clash with user tactic names.

- The extensions to the `rewrite .. in ...` tactic are partly incompatible with those introduced by version 8.1pl2. of COQ; specifically: `rewrite .. in (type of k)` or `rewrite .. in *` will not work. Moreover SSREFLECT looks for redexes more aggressively than `ltac`. Workaround: use an explicit rewrite direction (`<-` or `->`) after `rewrite` to disable the SSREFLECT interpretation of `rewrite`.
- New symbols (`//`, `/=`, `//=`) might clash with adjacent 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:

```
Require ssreflect.
Import ssreflect.SsrSyntax.
```

- Some user notations (in particular, defining an infix `';`) might interfere with "open term" syntax of tactics such as `have`, `set` and `pose`.
- The generalisation of `if` statements to non-Boolean conditions is turned off by SSREFLECT, because it is mostly subsumed by the `if ... is` construct (see 2.2), and because this avoids spurious type annotations. To use the generalised form, turn off the SSREFLECT Boolean `if` notation:

```
Close Scope boolean_if_scope.
```

2 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.

2.1 Pattern assignment

The SSREFLECT extension provides the following construct for irrefutable pattern matching, that is, destructuring assignment:

```
let: <pattern> := <term>1 in <term>2
```

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 standard `let`, requires an extra type annotation.

The `let:` construct is just (more legible) notation for the primitive Gallina expression

```
match <term>1 with <pattern> => <term>2 end
```

Due to limitations of the COQ v8 display API, a `let:` expression will always be displayed as the expansion of this form in terms of primitive `match` expressions.

The SSREFLECT destructuring assignment supports all the dependent match annotations; the full syntax is

```
let:<pattern>1 as <ident> in <pattern>2 := <term>1 return <term>2 in <term>3
```

where $\langle pattern \rangle_2$ is a type pattern and $\langle term \rangle_1$ and $\langle term \rangle_2$ are types.

When the `as` and `return` are both present, then $\langle ident \rangle$ is bound in both the type $\langle term \rangle_2$ and the expression $\langle term \rangle_3$; variables in the optional type pattern $\langle pattern \rangle_2$ are bound only in the type $\langle term \rangle_2$, and other variables in $\langle pattern \rangle_1$ are bound only in the expression $\langle term \rangle_3$, however.

2.2 Pattern conditional

The following construct can be used for a refutable pattern matching, that is, pattern testing:

```
if <term>1 is <pattern>1 then <term>2 else <term>3
```

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 analysing 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 <term>1 with <pattern> => <term>2 | _ => <term>3 end
```

Similarly, it will always be displayed as the expansion of this form in terms of primitive `match` expressions (where the default expression $\langle term \rangle_3$ may be replicated).

Explicit pattern testing also largely subsumes the generalisation of the `if` construct to all binary datatypes; compare

```
if <term> is inl _ then <term>l else <term>r
```

and

```
if <term> then <term>l else <term>r
```

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 §1.2, this restriction can be removed with the command:

```
Close Scope boolean_if_scope.
```

Like `let`: above, the `if...is` construct supports the dependent `match` annotations:

```
if <term>1 is <pattern>1 as <ident> in <pattern>2 return <term>2
  then <term>3
  else <term>4
```

As in `let`: the variable $\langle ident \rangle$ (and those in the type pattern $\langle pattern \rangle_2$) are bound in $\langle term \rangle_2$; $\langle ident \rangle$ is also bound in $\langle term \rangle_3$ (but not in $\langle term \rangle_4$), while the variables in $\langle pattern \rangle_1$ are bound only in $\langle term \rangle_3$.

2.3 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 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 d) := all (@null d) s.
```

Unfortunately, such higher-order expressions are quite frequent in representation functions, especially those which use COQ's `Structure` 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`.

This is accomplished by extending COQ's `Contextual Implicits` declaration to function types (in COQ, it is limited to (co)inductives): SSREFLECT implements a new vernacular command:

```
Prenex Implicits <ident>+.
```

Let us denote $const_1 \dots const_n$ the list of identifiers given to a `Prenex Implicits` command. The command checks that each $const_i$ is the name of a functional constant, whose implicit arguments are prenex, i.e., the first $n_i > 0$ arguments of $const_i$ are implicit. It then modifies COQ's syntax data so that the first n_i arguments of $const_i$ are inferred for every use of $const_i$ (except in `@consti`), and never displayed (except after a `Set Printing All` command).

Due to technicalities of the COQ v8 system architecture, the SSREFLECT `Prenex Implicits` command is strictly local to the COQ module in which it occurs: its scope ends with that module, and is not reopened when this module is `Imported`. However, SSREFLECT provides a command:

```
Import Prenex Implicits.
```

that reestablishes the `Prenex Implicits` status of all the directly visible constants. For example, if file `mod1.v` contains definitions for constants `k1` and `k2` together with

```
Prenex Implicits k1 k2.
```

and file `mod2.v` contains (other) definitions for `k2` and `k3` along with

```
Prenex Implicits k3.
```

then after

```
Require Import mod1.
Require Import mod2.
Import Prenex Implicits.
```

both `k1` and `k3` will have `Prenex Implicits`, but `k2`, which refers to the definition in `mod2.v`, will not.

Finally, due to another technicality, the `unfold k` tactic will not work after `Prenex Implicits k`; however, the equivalent tactic

```
rewrite /k.
```

described in section 6 will work.

2.4 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 $(_ : \langle term \rangle)$ syntax for inner arguments. In SSREFLECT, this can be replaced by the open syntax ‘`of <term>`’ or (equivalently) ‘`& <term>`’, which are both syntactically equivalent to the standard COQ $(_ : \langle term \rangle)$ expression.

Hence the following declaration:

```
Inductive list (A : Type) : Type := nil | cons of A & (list A).
```

defines a type which is syntactically equal to the type `list` of the COQ standard `List` library.

2.5 Wildcards

As in standard Gallina, the terms passed as arguments to SSREFLECT tactics can contain *holes*, materialised 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, specially in local definitions or type expressions.

Wildcards may be interpreted as abstractions (see for example sections 3.1, 5.6), or their content can be inferred from the whole context of the goal (see for example section 3.2).

3 Definitions

3.1 Definitions

The standard `pose` tactic allows to add a defined constant to a proof context. SSREFLECT generalises 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 standard `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 := ...` constructions. 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 of 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 dependant 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 generalisation of wildcards makes the use of the `pose` tactic resemble ML-like definitions of polymorphic functions.

3.2 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 standard COQ `set` tactic by supplying:

- an open syntax, similarly to the `pose` tactic;
- a more aggressive matching algorithm (which may cause some script incompatibilities with standard COQ);
- 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

$$\text{set } \langle \text{ident} \rangle [: \langle \text{term} \rangle_1] := [\langle \text{occ switch} \rangle] \langle \text{term} \rangle_2$$

$$\langle \text{occ switch} \rangle \equiv \{ [+|-] \langle \text{num} \rangle^* \}$$

where:

- $\langle ident \rangle$ is a fresh identifier chosen by the user.
- $\langle term \rangle_1$ an optional type annotation. The type annotation $\langle term \rangle_1$ can be given in open syntax (no surrounding parentheses). If no $\langle occ\ switch \rangle$ (described hereafter) is present, it is also the case for $\langle term \rangle_2$. On the other hand, in presence of $\langle occ\ switch \rangle$, parentheses surrounding $\langle term \rangle_2$ are mandatory.
- In the occurrence switch $\langle occ\ switch \rangle$, if the first element of the list is a $\langle num \rangle$, this element should be a number, and not an ltac variable. The empty list $\{\}$ is not interpreted as a valid occurrence switch (but as an empty clear switch, see section 4.3).

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 $\langle type \rangle_1$ annotation 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 $\langle term \rangle$ (and its $\langle type \rangle$), and stops at the first subterm it finds. Then the occurrences of this subterm selected by the optional $\langle occ\ switch \rangle$ are replaced by $\langle ident \rangle$ and a definition $\langle ident \rangle := \langle term \rangle$ is added to the context. If the $\langle occ\ switch \rangle$ is empty, then all the occurrences are abstracted.

3.2.1 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 $term$.

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 := ...in...)` $t_1 \dots t_n$, 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 `(let f := ... in ...)` of $term$. For instance the following tactic:

```
set t := (let f := (fun x y z => x + y + z) in f 1) 2.
```

transforms the goal

```
(let f := (fun x y z => x + y + z) in f) 1 2 3 = 6.
```

into `t 3 = 6` and adds the local definition of `f` 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 : nat` in the context.

- The type of the subterm match 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.

3.2.2 Occurrence selection

SSREFLECT provides a generic syntax for the selection of occurrences by their indexes. These *occurrence switches* are shared by all SSREFLECT tactic which require control on subterm selection like rewriting, generalisation, ...

An *occurrence switch* can be:

- A list `{+⟨num⟩*}` 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, like in standard COQ, 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, both in SSREFLECT and in standard COQ, 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 <= y -> t < S y` since `x < y` is now a notation for `S x <= y`.

- A list `{⟨num⟩+}`. This is equivalent to `{+⟨num⟩+}` but the list should start with a number, and not with an ltac variable.
- A list `{-⟨num⟩*}` 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).
```

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.

on the goal $f\ 2 + f\ 8 = f\ 2 + f\ 2$

- In particular, the switch `{+}` selects *all* the occurrences. This switch is useful to turn off the default behaviour of a tactic which automatically clears some assumptions (see section 4.3 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 pattern (_(*1*) + _(*2*))
```

3.3 Localisation

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 $fact_1 \dots fact_n$.

A tactic of the form:

```
set x := term in fact1 ... factn *.
```

also folds the definition of x in the goal.

The body of x is the first subterm matching *term* in $fact_1 \dots fact_n$.

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 added 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 scope of the `in` tactical is not limited to the localisation of definitions: for a complete description of the `in` tactical, see section 4.1.

4 Bookkeeping

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. Indeed, SSREFLECT scripts appear to divide evenly between bookkeeping, formal algebra (rewriting), and actual deduction. Although they are logically trivial, bookkeeping steps are extremely important because they define the structure of the dataflow 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 of COQ convenient, this 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.

4.1 The proof stack

During the course of a proof COQ always present the user with a *sequent* whose general form is

$$\begin{array}{l}
 c_i : T_i \\
 \dots \\
 d_j := e_j : T_j \\
 \dots \\
 F_k : P_k \\
 \dots \\
 \hline
 \text{forall } (x_\ell : T_\ell) \dots, \\
 \text{let } y_m := b_m \text{ in } \dots \text{ in} \\
 P_n \rightarrow \dots \rightarrow C
 \end{array}$$

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 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; about one out of every six tacticals is a **move**.

²Thus scripts that depend on bound variable names, e.g., via **intros** or **with**, are inherently fragile.

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_m_n` 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`. This is exactly what the specialized COQ tactic `intros m n le_m_n` would do, but `'=>'` is much more general (see 4.4).

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_m_n` 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`. The specialized COQ tactic `revert` does exactly this, but `':'` is much more general (see 4.3).

Because they are tacticals, `':'` and `'=>'` can be combined, as in

`move: m le_n_m => p le_n_p.`

simultaneously renames `m` and `le_m_n` into `p` and `le_p_n`, 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 `'=>'`. The COQ tactics lack uniformity in that they require an argument from the context but operate on the goal. Their `SSREFLECT` counterparts use the first variable or constant of the goal instead, so they are “purely deductive”: they do not use or change the proof context. There is no loss since `':'` can readily be used to supply the required variable; 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 4.4), 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*.

³The name `subnK` means “right cancellation rule for `nat` subtraction”.

- `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 localisation 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 5.5 for the general syntax and presentation of the `in` tactical.

4.2 Basic tactics

In this section we briefly present the three basic tactics performing context manipulations and the main backward chaining tool.

4.2.1 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 `~ False` into `False -> 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 4.4 and 4.3). These combinations mostly subsume the `intros`, `generalize`, `rename`, `clear` and `pattern` tactics.

4.2.2 The `case` tactic.

The `case` tactic, like in standard COQ, 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 4.6).

The SSREFLECT `case` tactic has a special behaviour on equalities.⁴ If the top assumption of the goal is an equality, the `case` tactic “deconstructs” it as a set of equalities on the constructor arguments of its left and right hand sides, as per the standard COQ 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 `False -> G` will succeed and prove the goal.

4.2.3 The `elim` tactic.

The `elim` tactic, like in standard COQ 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.

4.2.4 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.
```

corresponds to the following standard COQ tactic:

```
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.

This use of the `refine` tactic makes the SSREFLECT `apply` tactic considerably more robust than its standard COQ namesake, since it tries to match the goal up to expansion of constants and evaluation of subterms.

⁴The primitive COQ behaviour, rewriting right to left, is somewhat counterintuitive.

4.3 Discharge

The general syntax of the discharging tactical ‘:’ is:

$$\langle \text{tactic} \rangle [\langle \text{ident} \rangle]: \langle d\text{-item} \rangle_1 \dots \langle d\text{-item} \rangle_n [\langle \text{clear switch} \rangle]$$

where $n > 0$, and $\langle d\text{-item} \rangle$ and $\langle \text{clear switch} \rangle$ are defined as

$$\begin{aligned} \langle d\text{-item} \rangle &\equiv [\langle \text{occ switch} \rangle \mid \langle \text{clear switch} \rangle] \langle \text{term} \rangle \\ \langle \text{clear switch} \rangle &\equiv \{ \langle \text{ident} \rangle_1 \dots \langle \text{ident} \rangle_m \} \end{aligned}$$

with the following requirements:

- $\langle \text{tactic} \rangle$ must be one of the four basic tactics described in 4.2, i.e., **move**, **case**, **elim** or **apply**, or the **exact** tactic (section 5.2), or the application of the *view* tactical ‘/’ (section 7.2) to one of **move**, **case**, or **elim**.
- The optional $\langle \text{ident} \rangle$ specifies *equation generation* (section 4.5), and is only allowed if $\langle \text{tactic} \rangle$ is **move**, **case** or **elim**, or the application of the view tactical ‘/’ (section 7.2) to **case** or **elim**.
- An $\langle \text{occ switch} \rangle$ selects occurrences of $\langle \text{term} \rangle$, as in 3.2.2; $\langle \text{occ switch} \rangle$ is not allowed if $\langle \text{tactic} \rangle$ is **apply** or **exact**.
- A clear item $\langle \text{clear switch} \rangle$ specifies facts and constants to be deleted from the proof context (as per the **clear** tactic).

The ‘:’ tactical first *discharges* all the $\langle d\text{-item} \rangle$ s, right to left, and then performs $\langle \text{tactic} \rangle$, i.e., for each $\langle d\text{-item} \rangle$, starting with $\langle d\text{-item} \rangle_n$:

1. The SSREFLECT matching algorithm described in section 3.2 is used to find occurrences of $\langle \text{term} \rangle$ in the goal, after filling any holes ‘_’ in $\langle \text{term} \rangle$; however if $\langle \text{tactic} \rangle$ is **apply** or **exact** a different matching algorithm, described below, is used.⁵
2. These occurrences are replaced by a new variable, as per the standard COQ **generalize** tactic; in particular, if $\langle \text{term} \rangle$ is a fact, this adds an assumption to the goal.
3. If $\langle \text{term} \rangle$ is *exactly* the name of a constant or fact in the proof context, it is deleted from the context as per the COQ **clear** tactic, unless there is an $\langle \text{occ switch} \rangle$.

Finally, $\langle \text{tactic} \rangle$ is performed just after $\langle d\text{-item} \rangle_1$ has been generalized — that is, between steps 2 and 3 for $\langle d\text{-item} \rangle_1$. The names listed in the final $\langle \text{clear switch} \rangle$ (if it is present) are cleared first, before $\langle d\text{-item} \rangle_n$ is discharged.

Switches affect the discharging of a $\langle d\text{-item} \rangle$ as follows:

- An $\langle \text{occ switch} \rangle$ restricts generalization (step 2) to a specific subset of the occurrences of $\langle \text{term} \rangle$, as per 3.2.2, and prevents clearing (step 3).
- All the names specified by a $\langle \text{clear switch} \rangle$ are deleted from the context in step 3, possibly in addition to $\langle \text{term} \rangle$.

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**.

⁵Also, a slightly different variant is used for the first $\langle d\text{-item} \rangle$ of **case** and **elim**; see section 4.6.

- 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 generalized defined constants as variables, and not as local definitions. Thus, the definition of the constant is always erased, in contrast with the behaviour of the `in` tactical (see section 4.1).

4.3.1 Clear rules

The clear step will fail if $\langle term \rangle$ is a proof constant that appears in other facts; in that case either the facts should be cleared explicitly with a $\langle clear switch \rangle$, or the clear step should be disabled. The latter can be done by adding an $\langle occ switch \rangle$ or simply by putting parentheses around $\langle term \rangle$: both

```
move: (n).
```

and

```
move: {+}n.
```

generalize n without clearing n from the proof context.

The clear step will also fail if the $\langle clear switch \rangle$ contains a $\langle ident \rangle$ that is not in the *proof* context; in particular, SSREFLECT never clears a section constant.

If $\langle tactic \rangle$ is `move` or `case` and an equation $\langle ident \rangle$ is given, then clear (step 3) for $\langle d-item \rangle_1$ is suppressed (see section 4.5).

4.3.2 Matching for `apply` and `exact`

The matching algorithm for $\langle d-item \rangle$ s of `apply` and `exact` exploits the type of $\langle d-item \rangle_1$ to interpret wildcards in the other $\langle d-item \rangle$ and to determine which occurrences of these should be generalized. Therefore, $\langle occur switch \rangle$ es are not needed for `apply` and `exact`.

Indeed, the SSREFLECT tactic `apply`: $H \ x$ is equivalent to the standard COQ tactic

```
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 accomodate the ‘_’ $\langle d-item \rangle$, 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 $(Hfg : \text{forall } x, f \ x = g \ x)$, into $g \ a = g \ b$. This tactic is equivalent (see section 4.1) to:

```
refine (trans_equal (Hfg _) _).
```

and this is a common idiom for applying transitivity on the left hand side of an equation.

4.4 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 conatants and fact. Furthermore it is very common to immediately

destructure 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

$$\langle \text{tactic} \rangle \Rightarrow \langle i\text{-item} \rangle_1 \dots \langle i\text{-item} \rangle_n$$

where $\langle \text{tactic} \rangle$ can be any tactic, $n > 0$ and

$$\begin{aligned} \langle i\text{-item} \rangle &\equiv \langle i\text{-pattern} \rangle \mid [\langle \text{clear switch} \rangle] \langle s\text{-item} \rangle \mid \langle \text{clear switch} \rangle \\ \langle s\text{-item} \rangle &\equiv /= \mid // \mid // = \\ \langle i\text{-pattern} \rangle &\equiv \langle \text{ident} \rangle \mid _ \mid ? \mid * \mid -> \mid <- \mid [\langle i\text{-item} \rangle_1^* \mid \dots \mid \langle i\text{-item} \rangle_n^*] \end{aligned}$$

subject to the condition that at least every other $\langle i\text{-item} \rangle$ should be an $\langle i\text{-pattern} \rangle$ (for any two consecutive $\langle i\text{-item} \rangle_i \langle i\text{-item} \rangle_{i+1}$, either $\langle i\text{-item} \rangle_i$ or $\langle i\text{-item} \rangle_{i+1}$ is an $\langle i\text{-pattern} \rangle$).

The ‘=>’ tactical first executes $\langle \text{tactic} \rangle$, then the $\langle i\text{-item} \rangle$ s, left to right, i.e., starting from $\langle i\text{-item} \rangle_1$. An $\langle s\text{-item} \rangle$ specifies a simplification operation; a $\langle \text{clear switch} \rangle$ specifies context pruning as in 4.3. The $\langle i\text{-pattern} \rangle$ s are quite similar to COQ’s *intro patterns*; each performs an introduction operation, i.e., pops some variables or assumptions from the goal.

An $\langle s\text{-item} \rangle$ can simplify the set of subgoals or the subgoal themselves:

- // removes all the “trivial” subgoals that can be resolved by the SSREFLECT tactic `done` described in 5.2, i.e., it executes `try done`.
- /= simplifies the goal by performing partial evaluation, as per the COQ tactic `simpl`.⁶
- // = combines both kinds of simplification; it is equivalent to /= //, i.e., `simpl`; `try done`.

When an $\langle s\text{-item} \rangle$ bears a $\langle \text{clear switch} \rangle$, then the $\langle \text{clear switch} \rangle$ is executed *after* the $\langle s\text{-item} \rangle$, e.g., `{IHn} //` will solve some subgoals, possibly using the fact `IHn`, and will erase `IHn` from the context of the remaining subgoals.

SSREFLECT supports the following $\langle i\text{-pattern} \rangle$ s:

- $\langle \text{ident} \rangle$ pops the top variable, assumption, or local definition into a new constant, fact, or defined constant $\langle \text{ident} \rangle$, respectively. As in COQ, 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. Unlike COQ, SSREFLECT only generates names that cannot appear later in the user script.⁷
- $_$ 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 $\langle i\text{-item} \rangle$ has been executed ($\langle s\text{-item} \rangle$ s can erase goals or terms where the constant appears).
- * pops all the remaining apparent variables/assumptions as anonymous constants/facts. Unlike ? and `move` the * $\langle i\text{-item} \rangle$ 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 `anon a : bool` and `anon b : bool` to the context; it takes a second `move=> *` to add `anon hyp : anon a = anon b`.

⁶Except /= does not expand the local definitions created by the SSREFLECT `in` tactical.

⁷Generated names are of the form “anon x” and contain internal whitespace.

- \rightarrow (resp. \leftarrow) 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 6), and finally deletes the anonymous fact from the context.
- $[\langle i\text{-item} \rangle_1^* | \dots | \langle i\text{-item} \rangle_m^*]$, when it is the very *first* $\langle i\text{-pattern} \rangle$ after $\langle tactic \rangle \Rightarrow$ tactical and $\langle tactic \rangle$ is not a `move`, is a *branching* $\langle i\text{-pattern} \rangle$. It executes the sequence $\langle i\text{-item} \rangle_i^*$ on the i^{th} subgoal produced by $\langle tactic \rangle$. The execution of $\langle tactic \rangle$ should thus generate exactly m subgoals, unless the [...] $\langle i\text{-pattern} \rangle$ comes after an initial `//` or `//=` $\langle s\text{-item} \rangle$ that closes some of the goals produced by $\langle tactic \rangle$, in which case exactly m subgoals should remain after the $\langle s\text{-item} \rangle$, or we have the trivial branching $\langle i\text{-pattern} \rangle []$, which always does nothing, regardless of the number of remaining subgoals.
- $[\langle i\text{-item} \rangle_1^* | \dots | \langle i\text{-item} \rangle_m^*]$, when it is *not* the first $\langle i\text{-pattern} \rangle$ or when $\langle tactic \rangle$ is a `move`, is a *destructing* $\langle i\text{-pattern} \rangle$. It starts by destructing the top variable, using the SSREFLECT `case` tactic described in 4.2. It then behaves as the corresponding branching $\langle i\text{-pattern} \rangle$, executing the sequence $\langle i\text{-item} \rangle_i^*$ in the i^{th} subgoal generated by the case analysis; unless we have the trivial destructing $\langle i\text{-pattern} \rangle []$, the latter should generate exactly m subgoals, i.e., the top variable should have an inductive type with exactly m constructors.⁸ While it is good style to use the $\langle i\text{-item} \rangle_i^*$ to pop the variables and assumptions corresponding to each constructor, this is not enforced by SSREFLECT.

Note that SSREFLECT does not support the alternative COQ syntax $(\langle ipat \rangle, \dots, \langle ipat \rangle)$ for destructing intro-patterns.

The rules for interpreting branching and destructing $\langle i\text{-pattern} \rangle$ are motivated by the fact that it would be pointless to have a branching pattern if $\langle tactic \rangle$ is a `move`, and in most of the remaining cases $\langle tactic \rangle$ 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 decompositions, such as splitting a pair, and `case` should be used for actual destructions, in particular for type families (see 4.6) and proof by contradiction.

The trivial branching $\langle i\text{-pattern} \rangle$ can be used to force the branching interpretation, e.g.,

```
case=> [] [a b] c.
move=> [[a b] c].
case; case=> a b c.
```

are all equivalent.

4.5 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 generalise or rewrite such a definition; on the other hand, it is automatically expanded during computation, whereas expanding the equation `En` requires explicit rewriting.

⁸More precisely, it should have a quantified inductive type with a assumptions and $m - a$ constructors.

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 anon n` and `H : S anon 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.

4.6 Type families

When the top assumption of a goal is 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 the parameters of the type to be cased (resp. eliminated).

A specific `/ switch` indicates the type family parameters of the type of a *d-item* immediately following this `/ switch`, using the syntax:

$$[\text{case|elim}] : \langle \text{hypothesis} \rangle^+ / \langle \text{d-item} \rangle^*$$

The $\langle \text{d-item} \rangle$ s on the right side of the `/ switch` are discharged as described in section 4.3. The case analysis or elimination will be done on the type of the top assumption after these discharge operations. Every $\langle \text{hypothesis} \rangle$ preceding the `/ switch` is interpreted as arguments of this type, which should be an instance of an inductive type family. These $\langle \text{hypothesis} \rangle$ are actually not generalised, but rather selected for substitution.

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 : forall (s : list A) (x : A), 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 `length l * 2 = length (nil ++ l)` and `forall (s : list A)(x : A), length l * 2 = length (add_last x s++1)` 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, e.g., again starting with the above goal, the command:

```
move=> l; case E: {1 3}l / (lastP l)=>[|s x].
```

adds `E : l = nil` and `E : l = 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.

5 Control flow

5.1 Indentation and bullets

The linear development of COQ scripts barely gives information on the structure of the proof. Replaying a proof after some changes in the statement to be proved is quite not informative without the help of display information to distinguish the various branches of case analysis for instance.

To help the user in this organisation 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 highlights four nested levels of branching, while the deepest case we have ever encountered is three. Indeed, the use of “`simpl` and closing” switches, of terminators (see above section 5.2) and selectors (see section 5.3) is powerful enough to avoid most of the time more than two levels of indentation.

Here is a slice of such a structured script:

```
case E1: (abezoutn _ _) => [[| k1] [| k2]].
- rewrite !muln0 !gexpn0 mulg1 => H1.
  move/eqP: (sym_equal F0); rewrite -H1 orderg1 eqn_mul1.
  by case/andP; move/eqP.
- rewrite muln0 gexpn0 mulg1 => H1.
  have F1: dvdn t (t * S k2 - 1).
    apply: (@dvdn_trans (orderg x)); first by rewrite F0; exact: dvdn_mull.
```

```

    rewrite orderg_dvd; apply/eqP; apply: (mulg_inj1 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 muln.
  by apply/eqP; rewrite -dvdn1.
+ by move: P1; case: (t) => [| [| s1]].
- rewrite muln0 gexpn0 mul1g => H1.
...

```

5.2 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 practise that the proof of any subgoal should end with a tactic which *fails if it does not solve the current goal*. Standard COQ already provides some tactics of this kind, like **discriminate**, **contradiction** or **assumption**.

SSREFLECT provides a generic tactical which turns any tactic into a closing one. Its general syntax is:

```
by <tactic>.
```

The ltac expression:

```
by [<tactic>1 | [<tactic>2 | ...].
```

is equivalent to:

```
[by <tactic>1 | by <tactic>2 | ...].
```

and this form should be preferred to the former.

In the script provided as example in section 5.1, the paragraph corresponding to each subcase 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.
```

Conversly, 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 6.3). The iterator tactical `do` is presented in section 5.4. This tactic can be customised 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. In its applied form, the tactic:

```
exact MyLemma.
```

is equivalent to

```
by apply MyLemma.
```

This applied form is supported by the `: discharge` tactical, and the tactic:

```
exact: MyLemma.
```

is equivalent to:

```
by apply: MyLemma.
```

(see section 4.3.2 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 parenthesised 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).
```

5.3 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 linearisation 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:

```
<tactic>1; last by <tactic>2.
```

tries to solve the last subgoal generated by $\langle \text{tactic} \rangle_1$ using the $\langle \text{tactic} \rangle_2$, and fails if it does not succeeds. Its analogous

```
<tactic>1; first by <tactic>2.
```

tries to solve the first subgoal generated by $\langle \text{tactic} \rangle_1$ using the tactic $\langle \text{tactic} \rangle_2$, 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 $\langle tactic \rangle$. It is equivalent to:

$\langle tactic \rangle$; **first last**.

More generally, the tactic:

$\langle tactic \rangle$; **last** $\langle strict\ num \rangle$ **first**.

where $\langle strict\ num \rangle$ is a natural number argument having the value k , rotates the n subgoals G_1, \dots, G_n generated by $\langle tactic \rangle$. The first subgoal becomes G_{n+1-k} and the circular order of subgoals remains unchanged.

Conversely, the tactic:

$\langle tactic \rangle$; **first** $\langle strict\ num \rangle$ **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** combined with the branching syntax of **ltac**: if the tactic $\langle tactic \rangle_1$ generates n subgoals on a given goal, then the tactic

$\langle tactic \rangle_1$; **last** k [$\langle tactic \rangle_1$ | ... | $\langle tactic \rangle_m$] || $\langle tactic \rangle_2$.

where k is a natural number applies $\langle tactic \rangle_1$ to the $*n-k+1$ -th goal, ..., $\langle tactic \rangle_m$ to the $n-k+2-m$ -th goal and $\langle tactic \rangle_2$ 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=> 1]; idtac.
```

creates a goal with four subgoals, the first and the last being **nat -> True**, the second and the third being **True** with respectively $k : \text{nat}$ and $1 : \text{nat}$ in their context.

5.4 Iteration

SSREFLECT offers an accurate control on the repetition of tactic tactics, thanks to the **do** tactical, whose general syntax is:

$$do[\langle mult \rangle][\langle tactic \rangle_1 | \dots | \langle tactic \rangle_n]$$

where $\langle mult \rangle$ 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 [$\langle tactic \rangle_1$ | ... | $\langle tactic \rangle_n$].

is equivalent to the standard **ltac** expression:

first [$\langle tactic \rangle_1$ | ... | $\langle tactic \rangle_n$].

The optional multiplier $\langle mult \rangle$ specifies how many times the action of $\langle tactic \rangle$ should be repeated on the current subgoal.

There are four kinds of multipliers:

- $n!$: the step $\langle tactic \rangle$ is repeated exactly n times (where n is a positive integer argument).
- $!$: the step $\langle tactic \rangle$ is repeated as many times as possible, and done at least once.
- $?$: the step $\langle tactic \rangle$ is repeated as many times as possible, optionally.
- $n?$: the step $\langle tactic \rangle$ 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 1! rewrite mult_comm.
```

rewrites exactly one time the lemma `mult_com` in all the subgoals generated by `<tactic>`, and fails if this rewrite is not possible in one 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 6.

5.5 Localisation

In sections 3.3 and 4.1, we have already presented the *localisation* tactical `in`, whose general syntax is:

$$\langle \text{tactic} \rangle \text{ in } \langle \text{hypothesis} \rangle^+ [*]$$

where $\langle \text{hypothesis} \rangle^+$ is a non empty list of fact names in the context. On the left side of `in`, $\langle \text{tactic} \rangle$ can be `move`, `case`, `elim`, `rewrite`, `set`, or any tactic formed with the general iteration tactical `do` (see section 5.4).

The operation described by $\langle \text{tactic} \rangle$ is performed in the facts listed in $\langle \text{hypothesis} \rangle^+$ and in the goal if a `*` ends the list.

The `in` tactical successively:

- generalises the selected hypotheses, possibly “protecting” the goal if `*` is not present,
- performs $\langle \text{tactic} \rangle$, on the obtained goal,
- reintroduces the generalised 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`.

5.6 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 5.2) and of indentation makes syntactically explicit the portion of the script building the proof of the intermediate statement.

5.6.1 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`: $\langle term \rangle$.

This tactic supports open syntax for $\langle term \rangle$. Apart from the open syntax, when $\langle term \rangle$ does not contain any wildcard, this tactic is almost⁹ equivalent to the standard COQ:

`assert` $\langle term \rangle$.

Applied to a goal G , it generates a first subgoal requiring a proof of $\langle term \rangle$ is the context of G . The difference with the standard COQ tactic is that the second subgoal generated is of the form $\langle term \rangle \rightarrow G$, where $\langle term \rangle$ becomes the new top assumption, instead of being introduced with a fresh name.

Like in the case of the `pose` tactic (see section 3.1), the types of the holes are abstracted in $\langle term \rangle$. 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 behaviour of the defective `have` tactic makes it possible to generalise it in the following general construction:

$$\text{have } [\langle clear\ switch \rangle] \langle i\text{-item} \rangle^* : \langle term \rangle_1 \text{ [by } \langle tactic \rangle \mid := \langle term \rangle_2]$$

Open syntax is supported for $\langle term \rangle_1$. For the description of *i-items* and clear switches see section 4.4. The first mode of the `have` tactic, which opens a subproof for an intermediate result, uses tactics of the form:

`have` *cl-switch* *i-item*₁ ... *i-item*_n : *term* **by** *tactic*.

which behave like:

`have`: *term*; **first** **by** *tactic*.
`move=>` *cl-switch* *i-item*₁ ... *i-item*_n.

Note that the *clear items* are performed *before* 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.

Hence the standard COQ:

`assert` *term*.

is in fact equivalent¹⁰ up to the open syntax to:

`have` `?: term`.

The **by** feature is specially 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 addnA`.

The possibility of using *i-items* supplies a very concise syntax for the further use of the intermediate step. For instance,

⁹The `assert` tactic creates a ζ redex, whereas the `have` tactic creates a β redex, and it introduces the lemma under an automatically chosen fresh name.

¹⁰again, except that the kind of redex created is different

```
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.

An other frequent use of the *i-items* 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 like for instance:

```
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 $H : \text{Type} \rightarrow \text{Prop}$. 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.

5.6.2 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 inverted
- 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.

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 $(\text{wlog_statement} \rightarrow G) \rightarrow G$ and `wlog_statement-> G`.

SSREFLECT implements this kind of reasoning step through the `without loss` tactic, whose short name is `wlog`. The general syntax of `without loss` is:

`wlog` [*cl-item*] [*i-item*]* : [*ident*₁ ... *ident*_n] / *term*.

where *ident*₁ ... *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, respectively *term*-> *G* and (*term*-> *G*)-> *G*.

If the optional list *ident*₁ ... *ident*_n is present on the left side of /, these constants are generalised on top of the first *term*-> *G* subgoal. In the second subgoal, the tactic:

`move=>` *cl-item* *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.
```

Note The list of generalised constants on the left side of the / switch can contain clear items between constants. These clear operations are intertwined with the generalisation ones, which helps in particular avoiding dependency issues while generalising some facts.

6 Rewriting

The generalised 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. We have therefore defined an extended `rewrite` tactic that unifies and combines most of the rewriting functionalities.

6.1 An extended `rewrite` tactic

The main improvements brought to the standard COQ `rewrite` tactic are:

- Whereas the primitive `rewrite` tactic can only perform a single rewriting operation in the goal or in the context, the extended `rewrite` can perform an entire series of such operations in any subset of the goal and/or context;
- The SSREFLECT `rewrite` tactic 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 3.3) performs rewriting in both the context and the goal.

A `rewrite` step *rstep* has the general form:

$$[(r\text{-prefix})]\langle r\text{-item} \rangle$$

where:

$$\begin{aligned} \langle r\text{-prefix} \rangle &\equiv [-] [\langle mult \rangle] [\langle occ\text{-switch} \rangle | \langle clear\ switch \rangle] [\langle term \rangle] \\ \langle r\text{-item} \rangle &\equiv [/] \langle term \rangle | \langle s\text{-item} \rangle \end{aligned}$$

An $\langle r\text{-prefix} \rangle$ contains annotations to qualify where and how the rewrite operation should be performed:

- The optional initial $-$ indicates the direction of the rewriting of $\langle r\text{-item} \rangle$.
- The multiplier $\langle mult \rangle$ (see section 5.4) 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 $\langle r\text{-item} \rangle$. The optional *redex switch* $[\langle term \rangle]$, which should always be surrounded by brackets, gives explicitly this rewrite pattern. It can be any term. If no explicit redex switch is present the rewrite pattern to be matched is inferred from the $\langle r\text{-item} \rangle$.
- This optional $\langle term \rangle$, or the $\langle r\text{-item} \rangle$, may be preceded by an occurrence switch (see section 5.3) or a clear item (see section 4.3), these two possibilities being exclusive. An occurrence switch selects the occurrences of the rewrite pattern which should be affected by the rewrite operation.

An $\langle r\text{-item} \rangle$ can be:

- A *simplification r-item*, represented by a $\langle s\text{-item} \rangle$ (see section 4.4). In that case, $\langle r\text{-prefix} \rangle$ es are not supported. 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 in principle¹¹ equivalent to:

```
unfold my_def.
```

Conversely:

```
rewrite -/my_def.
```

is equivalent to:

```
fold my_def.
```

Warning The combination of redex switch with `unfold` $\langle r\text{-item} \rangle$ is not yet implemented.

When an `unfold` r-item is combined with a redex pattern, a conversion operation is performed. A tactic of the form:

```
rewrite -[term1] /term2.
```

¹¹The implementation of these `fold/unfold` tactics do not call standard COQ `fold` and `unfold`.

is equivalent to:

```
change term1 with term2.
```

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 + _(*1*)) 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 of the form:

$$(x_1 : A_1) \dots (x_n : A_n) eq \text{term}_1 \text{term}_2$$

where eq is the Leibniz equality or a registered setoid equality. In the case of setoid relations, the only supported r-prefix is the directional $-$.

- A list of terms (t_1, \dots, t_n) , each t_i having the form:

$$(x_1 : A_1) \dots (x_n : A_n) eq \text{term}_1 \text{term}_2$$

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 $(_ : \text{term})$, where term has again the form:

$$(x_1 : A_1) \dots (x_n : A_n) eq \text{term}_1 \text{term}_2$$

The tactic:

```
rewrite (\_ : term)
```

is in fact equivalent to the standard COQ:

```
cutrewrite (term).
```

6.2 Remarks and examples

6.2.1 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 SSREFLECT and standard COQ.

In a rewrite tactic of the form:

```
rewrite occ switch[term1] term2.
```

$term_1$ is the explicit rewrite redex and $term_2$ is the rewrite rule. This execution of this tactic unfolds as follows:

- First $\langle term \rangle_1$ and $\langle term \rangle_2$ are $\beta\iota$ normalised. Then $\langle term \rangle_2$ 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 3.2) determines the first subgoal matching the rewrite pattern, which is given by $\langle term \rangle_1$, if an explicit redex pattern `switch` is provided, or by $\langle term \rangle_2$ in the opposite case. All the occurrences of this subterm in the goal are candidates for rewriting.
- Then only the occurrences coded by $\langle occ\ switch \rangle$ (see again section 3.2) are finally selected for rewriting.
- The left hand side of $\langle term \rangle_2$ 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 $\langle term \rangle_2$.
- Finally the goal is $\beta\iota$ normalised.

6.2.2 Chained rewritings

The possibility to chain rewrite operations in a single tactic makes scripts more compact and gathers in a single command line a bunch of chirurgical 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 /= my_eq //=. 
```

unfolds `my_def` in the goal, simplifies it, rewrites `my_eq`, simplifies again and closed 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 + S n = S (m + n).
```

```
Proof. by move=> m n; elim: m. Qed.
```

```
Lemma addSnnS : forall m n, S m + n = m + S n.
```

```
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](#).

6.2.3 Explicit redex switches are matched first

If an $\langle r\text{-prefix} \rangle$ involves a *redex switch*, the first step is to find a subterm matching this redex pattern, independently from the left hand side t_1 of the equality the user wants to rewrite.

For instance, if $H : \text{forall } t \ u, \ t + u = u + t$ is in the context of a goal $x + y = y + x$, the tactic:

```
rewrite [y + _]H.
```

transforms the goal into $x + y = y + x$.

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 $H : \text{forall } t \ u, \ t + u * 0 = t$ 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 $x + y * 4 = x + 2 * 0$.

6.2.4 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 $_ + y + 0$.

6.2.5 Occurrence selection and repetition

Occurrence selection has priority over repetition switches. This means the repetition of a rewrite tactic thanks to 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).
```

Moreover, this last tactic generates *three* other 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.

6.2.6 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`.

6.2.7 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. The two main cases are:

- SSREFLECT never accepts to rewrite empty patterns like:

```
Lemma foo : forall x : unit, x = 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 + g3 = 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.

6.3 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 `/= simpl` 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 occurrences switches (see section 3.2.2) 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`, which is *not convertible* to `t`. 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 = 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.
```

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 `cube` coded by the `<occ switch>` with `(Hypermap cube_mon3)` 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.

6.4 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 [strict num] 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 optional `<strict num>` forces the number of arguments for which the tactic should generate equality proof obligations.

This tactic supports equalities between applications with dependent arguments. Anyway as in standard COQ, 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`.

Being an arbitrary SSREFLECT term, $\langle term \rangle$ shall 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 $y + x - x = y$.

7 Views and reflection

The bookkeeping facilities presented in section 4 are crafted to ease simultaneous introduction/-generalisation of facts and casing, naming ... operations. It also a common practise to make a stack operation immediately follow an *interpretation* of the fact being pushed or popped, that is to say 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.

7.1 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.
```

corresponds to the standard COQ tactic:

```
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 4), the SSREFLECT tactic:

```
elim/V: x=> y.
```

corresponds to the standard COQ tactic:

```
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.

The elimination view mechanism is compatible with the equation name generation (see section 4.5).

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 `nil = nil` in a context featuring `E : l = nil` and the second to prove `add_last u v = add_last u v`, in a context containing `E : l = add_last u v`.

7.2 Interpreting assumptions

Interpreting an assumption in the context of a proof is applying it a correspondence lemma before generalising, and/or decomposing it. For instance, with the extensive use of boolean reflection (see section 7.4), 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 standard COQ 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 `HPn : P n` which has been introduced from the initial statement into `HQn : 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 4.3) 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 4.5):

```

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 behaviour 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.
```

7.2.1 Specialising assumptions

The special case when the *head symbol* of the view lemma is a wildcard is used to interpret an assumption by *specialising* 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
```

7.3 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 7.2, the following script, where $\sim\sim$ denotes the boolean negation:

```
Goal forall a, P (( $\sim\sim$  a) || a).
move=> a; apply/PQequiv.
```

transforms the goal into $Q (\sim\sim a)$, and is equivalent to:

```
Goal forall a, P (( $\sim\sim$  a) || a).
move=> a; apply: (iffR (PQequiv _ _)).
```

where `iffLR` 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.

Warning Goal interpretation view tactics are *not* compatible with the bookkeeping tactical `=>` since this would be redundant with the `apply: V=> _` construction.

7.4 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 ∨ B` is either a proof of `A` or a proof of `B`. The tactics `left` and `right` change the goal `A ∨ B` to `A` and `B`, respectively; dually, the tactic `case` reduces the goal `A ∨ B => G` to two subgoals `A => G` and `B => 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.

7.5 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 standard equivalence predicate is defined in Coq as:

```
Definition iff (A B:Prop) := (A -> B) /\ (B -> A).
```

where \wedge 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 /\ b2 -> b1` and `~ (b1 /\ b2) -> ~~ (b1 || 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, a -> b -> a /\ b.
move=> a b Ha Hb; apply/andP.
```

changes the goal `a /\ b` to `a && b`.

The same tactic `apply/andP` 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.

7.6 General mechanism for interpreting goals and assumptions

7.6.1 Specialising assumptions

The SSREFLECT tactic:

```
move/(_ <term>_1 ... <term>_n)
```

is equivalent to the tactic:

```
intro top; generalise (top <term>_1 ... <term>_n); clear top.
```

where `top` is a fresh name for introducing the top assumption of the current goal.

7.6.2 Interpreting assumptions

The general form of an assumption view tactic is:

$$[\text{move|case}]/\langle term \rangle_0.$$

The term $\langle term \rangle_0$, 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 7.5).

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 $\langle term \rangle_0$ is neither a double implication nor an instance of the **reflect** predicate, then the tactic automatically generalises a term of the form:

$$\langle term \rangle_0 \langle term \rangle_1 \dots \langle term \rangle_n$$

where the terms $\langle term \rangle_1 \dots \langle term \rangle_n$ instantiate the possible quantified variables of $\langle term \rangle_0$, in order for $(\langle term \rangle_0 \langle term \rangle_1 \dots \langle term \rangle_n \text{ top})$ to be well typed.

- If the type of $\langle term \rangle_0$ is an equivalence, or an instance of the **reflect** predicate, it generalises a term of the form:

$$\langle term \rangle_{vh} (\langle term \rangle_0 \langle term \rangle_1 \dots \langle term \rangle_n)$$

where the term $\langle term \rangle_{vh}$ inserted is called an *assumption interpretation view hint*.

- It finally clears **top**.

For a **case**/ $\langle term \rangle_0$ tactic, the generalisation step is replaced by a case analysis step.

View hints are declared by the user (see section 7.8) 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 $\langle term \rangle_0$ provided to the tactic the appropriate view hint in the database to be inserted.

If $\langle term \rangle_0$ 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 $\langle term \rangle_0$ 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 $\sim (a \wedge b)$. In fact¹² 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 `introNF` using `c := true`.

7.6.3 Interpreting goals

A goal interpretation view tactic of the form:

```
apply/<term>0.
```

applied to a goal `top` is interpreted in the following way:

- If the type of $\langle term \rangle_0$ is not an instance of the `reflect` predicate, nor an equivalence, then the term $\langle term \rangle_0$ is applied to the current goal `top`, possibly inserting implicit arguments.
- If the type of $\langle term \rangle_0$ 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 $(\langle term \rangle_{vh} (\langle term \rangle_0 _ \dots _))$ to the current goal, possibly inserting implicit arguments.

Like assumption interpretation view hints, goal interpretation ones are user defined lemmas stored (see section 7.8) in the `Hint View` database bridging the possible gap between the type of $\langle term \rangle_0$ and the type of the goal.

7.7 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/<term>l/r.
```

The term $\langle term \rangle_l$ is the view lemma applied to the left hand side of the equality, $\langle term \rangle_r$ 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.
```

transforms the goal $\sim\sim (b1 \parallel b2) = b3$ into two subgoals, respectively $\sim\sim (b1 \parallel b2) \rightarrow b3$ and $b3 \rightarrow \sim\sim (b1 \parallel b2)$.

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 $\sim\sim (b1 \parallel b2) = b3$ generates the subgoals $\sim\sim b1 \wedge \sim\sim b2 \rightarrow b3$ and $b3 \rightarrow \sim\sim b1 \wedge \sim\sim b2$.

¹²The current state of the proof shall be displayed by the `Show Proof` command of Coq proof mode.

7.8 Declaring new Hint Views

The user can declare his own hints for the view mechanism, following the syntax used in `ssrbool.v`:

```
Hint View for <tactic>/ <name> [|<num>].
```

where $\langle tactic \rangle \in \{\text{move}, \text{apply}\}$, $\langle name \rangle$ is the name of the lemma to be declared as a hint, and $\langle num \rangle$ a natural number. If `move` is used as $\langle tactic \rangle$, 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 $\langle num \rangle$ 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// <name> [|<num>].
```

with a double slash `//`, declares hint views for right hand sides of double views.

See the files `ssreflect.v` and `ssrbool.v` for examples.

8 SSREFLECT searching tool

SSREFLECT proposes an extension of the `Search` command of standard COQ. Its syntax is:

```
Search [[~]<string>* [<pattern>] [[<pattern>]1 ... <pattern>]n] [inside/outside M1 ... Mn].
```

This tactic returns the list of defined constants matching the given criteria:

- $[[\sim]\langle string \rangle]^*$ is a sequence of strings, which should all appear in the name of the returned constants. The optional \sim prefixes strings that are required *not* to appear.
- $\langle pattern \rangle$ should be a subterm of the *conclusion* of the lemmas found by the command. If a lemma features an occurrence of this pattern only in one or several of its assumptions, it will not be selected by the searching tool.
- $[\langle pattern \rangle_1 \dots \langle pattern \rangle_n]$ is a list of SSREFLECT terms between brackets, which may include types, that are required to appear in the returned constants.
- `inside` $M_1 \dots M_n$ limits the search to the signature of modules $M_1 \dots M_n$. `outside` $M_1 \dots M_n$ limits the search to all the open modules but ones in the given list. The command:

```
Search inside M.
```

is hence a way of obtaining the complete signature of the module M .

The command:

```
Search ~"add" leq [muln] inside ssrnat.
```

looks for all the lemmas defined in the module `ssrnat`, whose conclusion contain an inequality, whose statement contain an occurrence of natural number multiplication, and whose name does not contain the string “add”.

9 Conclusion

libs and stats.

Synopsis

Parameters

$\langle term \rangle$	Gallina term, possibly containing wildcards
$\langle pattern \rangle$	cf $\langle term \rangle$
$\langle ident \rangle$	standard COQ identifier
$\langle string \rangle$	standard COQ string
$\langle num \rangle$	standard COQ numeral or Ltac variable denoting a standard COQ $\langle num \rangle$ ¹³
$\langle strict\ num \rangle$	standard COQ numeral
$\langle view\ hint \rangle$	global constant stored in the view hints database through a Hint View command
$\langle tactic \rangle$	standard COQ tactic or SSREFLECT tactic
$\langle hypothesis \rangle$	named assumption or constant in the local context of a goal
$\langle fix_body \rangle$	standard COQ fix_body

Items and switches

$\langle occ\ switch \rangle$	$\{[+ -]\langle num \rangle^*\}$	occurrence switch
$\langle clear\ switch \rangle$	$\{\langle name \rangle^+\}$	clear item
$\langle d_item \rangle$	$[\langle occ\ switch \rangle \langle clear\ switch \rangle] \langle term \rangle$	discharge item
$\langle s_item \rangle$	$/= // // =$	simplification switch
$\langle i_item \rangle$	$\langle clear\ switch \rangle \langle s_item \rangle \langle i_pattern \rangle$	introduction item
$\langle i_pattern \rangle$	$\langle name \rangle _ ? * \rightarrow \leftarrow [\langle i_item \rangle_1^* \dots \langle i_item \rangle_m^*]$	intro pattern
$\langle mult_mark \rangle$	$? !$	multiplier mark
$\langle mult \rangle$	$[\langle strict\ num \rangle] \langle mult_mark \rangle$	multiplier
$\langle r_prefix \rangle$	$[-] [\langle mult \rangle] [\langle occ_switch \rangle \langle clear\ switch \rangle] [\langle term \rangle]$	rewrite prefix
$\langle r_item \rangle$	$[/] \langle term \rangle \langle s_item \rangle$	rewrite item
$\langle r_step \rangle$	$[\langle r_prefix \rangle] \langle r_item \rangle$	rewrite step

Tactics

<code>pose</code> $\langle name \rangle := \langle term \rangle$	local definition
<code>pose</code> $\langle name \rangle \langle ident \rangle^+ := \langle term \rangle$	local functional definition
<code>pose fix</code> $\langle fix-body \rangle$	local fixpoint definition
<code>pose cofix</code> $\langle fix-body \rangle$	local cofixpoint definition
<code>set</code> $\langle name \rangle [:\langle term \rangle_1] := [\langle occ switch \rangle] \langle term \rangle_2$	abbreviation
<code>move</code>	context \leftrightarrow goal moves
<code>done</code>	closing
<code>have</code> $[\langle clear switch \rangle][\langle i-item \rangle] : \langle term \rangle$ <code>by</code> $\langle tactic \rangle$	forward chaining
<code>suff</code> $[\langle clear switch \rangle][\langle i-item \rangle] : \langle term \rangle$ <code>by</code> $\langle tactic \rangle$	forward chaining
<code>wlog</code> $[\langle clear switch \rangle] [\langle i-item \rangle^*] : [\langle ident \rangle_1 \dots \langle ident \rangle_n] / \langle term \rangle$	forward chaining
<code>rewrite</code> $\langle rstep \rangle^+$	rewrite
<code>unlock</code>	unlock
<code>congr</code> $\langle term \rangle$	congruence

Tacticals

$\langle tactic \rangle [\langle name \rangle] : \langle d-item \rangle_1 \dots \langle d-item \rangle_n [\langle clear switch \rangle]$	discharge tactical
$\langle tactic \rangle \Rightarrow \langle i-item \rangle^+$	introduction tactical
$\langle tactic \rangle$ <code>in</code> $\langle hypothesis \rangle^+ [*]$	localisation tactical
<code>do</code> $[\langle mult \rangle][\langle tactic \rangle_1 \mid \dots \mid \langle tactic \rangle_n]$	iteration tactical
$\langle tactic \rangle_1$; <code>first</code> $\langle tactic \rangle_2$	selector
$\langle tactic \rangle_1$; <code>last</code> $\langle tactic \rangle_2$	selector
$\langle tactic \rangle$; <code>last</code> $[\langle strict num \rangle]$ <code>first</code>	rotation of subgoals
$\langle tactic \rangle$; <code>first</code> $[\langle strict num \rangle]$ <code>last</code>	rotation of subgoals
<code>by</code> $\langle tactic \rangle$	closing tactical
<code>exact</code> $\langle name \rangle$	closing application

Commands

<code>Prenex Implicits</code> $[\langle name \rangle]^+$	prenex implicits declaration
<code>Import Prenex Implicits</code>	enable prenex implicits
<code>Hint View for</code> $[\text{move apply}] / \langle name \rangle [\langle num \rangle]$	view hint declaration
<code>Hint View for</code> <code>apply</code> $// \langle name \rangle [\langle num \rangle]$	right hand side double view hint declaration



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399