A Declarative Language for the Coq Proof Assistant[*]

Pierre Corbineau

Institute for Computing and Information Science
Radboud University Nijmegen, Postbus 9010
6500GL Nijmegen, The Netherlands
`pierre.corbineau@cs.ru.nl`

Abstract. This paper presents a new proof language for the Coq proof assistant. This language uses the declarative style. It aims at providing a simple, natural and robust alternative to the existing $\mathcal{L}_{tac}$ tactic language. We give the syntax of our language, an informal description of its commands and its operational semantics. We explain how this language can be used to implement formal proof sketches. Finally, we present some extra features we wish to implement in the future.

1 Introduction

1.1 Motivations

An interactive proof assistant can be described as a state machine that is guided by the user from the 'statement ϕ to be proved' state to the 'QED' state. The system ensures that the state transitions (also known as *proof steps* in this context) are sound. The user's guidance is required because automated theorem proving in any reasonable logics is undecidable in theory and difficult in practice. This guidance can be provided either through some kind of text input or using a graphical interface and a pointing device. In this paper, we will focus on the former method.

The ML language developed for the LCF theorem prover [7] was a seminal work in this domain. The ML language was a fully-blown programming language with specific functions called *tactics* to modify the *proof state*. The tool itself consisted of an interpreter for the ML language. Thus a formal proof was merely a computer program. With this in mind, think of a person reading somebody else's formal proof, or even one of his/her own proof a couple of months after having written it. Similarly to what happens with source code, this person will have a lot of trouble understanding what is going on with the proof unless he/she has a very good memory or the proof is thoroughly documented. Of course, running the proof through the prover and looking at the output might help a little.

This illustrates a major inconvenience which still affects many popular proof languages used nowadays: they lack readability. Most proofs written are actually

---

write-only or rather write- and execute-only, since what the user is interested in when re-running the proof in is not really the input, but rather the output of the proof assistant, i.e. the sequence of proof states from the statement the theorem to 'QED'.

The idea behind declarative style proofs is to actually base the proof language on this sequence of proof states. This is indeed the very feature that makes the distinction between procedural proof languages (like ML tactics in LCF) and declarative proof languages. On the one hand, procedural languages emphasize proof methods (application of theorems, rewriting, proof by induction...), at the expense of a loss of precision on intermediate proof states: the intermediate states depend on the implementation on tactics instead of a formal semantics. On the other hand, declarative languages emphasize proof states but are less precise about the logical justification of the gap between one state and the next one.

1.2   Related Work

The first proof assistant to implement a declarative style proof language was the Mizar system, whose modern versions date back to the early 1990's. The Mizar system is a batch proof assistant: it compiles whole files and writes error messages in the body of the input text, so it is not exactly interactive, but its proof language has been an inspiration for all later designs [15].

Another important source in this subject is Lamport's *How to write a proof* [10] which takes the angle of the mathematician and provides a very simple system for proof notations, aimed at making proof verification as simple as possible.

John Harrison has been the first to develop a declarative proof language for an interactive proof assistant: the HOL88 theorem prover [8].

Donald Syme has developed the DECLARE [12,13] batch proof assistant for higher-order logic with declarative proofs in mind from the start. He also describes CAPW (Computer Aided Proof Writing) as a means of overcoming the verbosity of declarative proofs.

The first interactive proof assistant for which a declarative language has been widely adopted is Isabelle [11], with the Isar (Intelligible Semi-Automated Reasoning) language [14], designed by Markus Wenzel.

Freek Wiedijk also developed a light declarative solution [16] for John Harrison's own prover HOL Light [9].

For the Coq proof assistant [2], Mariusz Giero and Freek Wiedijk have built a set of tactics called the MMode [6] to provide an experimental mathematical mode which give a declarative flavor to Coq.

Recently, Claudio Sacerdoti-Coen added a declarative language to the Matita proof assistant [1].

1.3   A New Language for the Coq Proof Assistant

The Coq proof assistant is a Type Theory-based interactive proof assistant developed at INRIA. It has a strong user base both in the field of software and

hardware verification and in the field of formalized mathematics. It also has the reputation of being a tool whose procedural proof language $\mathcal{L}_{tac}$ has a very steep learning curve both at the beginner and advanced level.

Coq has been evolving quite extensively during the last decade, and the evolution has made it necessary to regularly update existing proofs to maintain compatibility with most recent versions of the tool.

Coq also has a documentation generation tool to do hyper-linked rendering of files containing proofs, but most proofs are written in a style that makes them hard to understand (even with syntax highlighting) unless you can actually run them, as was said earlier for other procedural proof languages.

Building on previous experience from Mariusz Giero, we have built a stable mainstream declarative proof language for Coq. This language was built to have the following characteristics:

readable. The designed language should use clear English words to make proof reading a feasible exercise.

natural. We want the language to use a structure similar to the ones used in textbook mathematics (e.g. for case analysis), not a bare sequence of meaningless commands.

maintainable. The new language should make it easy to upgrade the prover itself: behavior changes should only affect the proof locally.

stand-alone. The proof script should contain enough explicit information to be able to retrace the proof path without running Coq.

The Mizar language has been an important source of inspiration in this work but additional considerations had to be taken into account because the Calculus of Inductive Constructions (CIC) is much richer than Mizar's (essentially) first-order Set Theory.

One of the main issue is that of proof genericity: Coq proofs use a lot of inductive objects for lots of different applications (logical connectives, records, natural numbers, algebraic data-types, inductive relations...). Rather than enforcing the use of the most common inductive definitions, we want to be as generic as possible in the support we give for reasoning with these objects.

Finally, we want to give an extended support for proofs by induction by allowing multi-level induction proofs, using a very natural syntax to specify the different cases in the proof. The implementation was part of the official release 8.1 version of Coq.

1.4   Outline

We first describe some core features of our language, such as forward and backward steps, justifications, and partial conclusions. Then we give a formal syntax and a quick reference of the commands of our language, as well as an operational semantics. We go on by explaining how our language is indeed an implementation of the *formal proof sketches* [17] concept, and we define the notion of well-formed proof. We finally give some perspective for future work.

2   Informal Description

2.1   An Introductory Example

To give a sample of the declarative language, we provide here the proof of a simple lemma about Peano numbers: the `double` function is defined by $\mathtt{double}\,x = x + x$ and the $\mathtt{div}_2$ functions by:

$$\begin{cases} \mathtt{div}_2\,0 = 0 \\ \mathtt{div}_2\,1 = 0 \\ \mathtt{div}_2\,(S\,(S\,x)) = S\,(\mathtt{div}_2\,x) \end{cases}$$

The natural numbers are defined by means of an inductive type `nat` with two constructors `0` and `S` (successor function). The lemma states that $\mathtt{div}_2$ is the left inverse of `double`. We first give a proof of the lemma using the usual tactic language:

```
Lemma double_div2: forall n, div2 (double n) = n.
intro n.
induction n.
reflexivity.
unfold double in *|-*.
simpl.
rewrite <- plus_n_Sm.
rewrite IHn.
reflexivity.
Qed.
```

Now, we give the same proof using the new declarative language:

```
Lemma double_div2: forall n, div2 (double n) = n.
proof.
  let n:nat.
  per induction on n.
    suppose it is 0.
     reconsider thesis as (0=0).
     thus thesis.
    suppose it is (S m) and Hrec:thesis for m.
      have (div2 (double (S m))
            = div2 (S (S (double m)))).
          ~= (S (div2 (double m))).
      thus ~= (S m) by Hrec.
  end induction.
end proof.
Qed.
```

The proof consists of a simple induction on the natural number `n`. The first case is done by conversion (computation) to $0 = 0$ and the second case by

computation and by rewriting the induction hypothesis. Of course, you could have guessed that by simply reading the declarative proof.

2.2   Forward and Backward Proofs

The notions of declarative and procedural proofs are often confused with the notions of forward and backward proof. We believe those two notions are mostly orthogonal: the distinction between declarative and procedural style is that declarative proofs mention explicitly the intermediate proof steps, while procedural proofs explain what method is used to go to the next state without mentioning it, whereas the distinction between forward and backward proof is merely a notion of whether the proof is build bottom-up by putting together smaller proofs, or top-down by cutting a big proof obligation into smaller ones.

A possible reason for the confusion between declarative style and backwards proofs is that most declarative languages rely on the core command

$$\texttt{have } h : \phi \; justification$$

which introduces a new hypothesis h that asserts ϕ while using *justification* to check that it holds. This kind of command is the essence of forward proofs: it builds a new object — a proof of ϕ — from objects that already exists and are somehow mentioned in the justification.

In order to show that you can also use backwards steps in a declarative script, our language contains a command

$$\texttt{suffices } H_1 \texttt{ and } \ldots \texttt{ and } H_n \texttt{ to show } G \; justification$$

which acts as the dual of the `have` construction: it allows to replace the current statement to prove by sufficient conditions with stronger statements (as explained by the justification). For example, you can use this command to generalize your thesis before stating a proof by induction.

2.3   Justifications

When using a command to deduce a new statement from existing ones, or to prove that some statements suffice to prove a part of the conclusion, you need the proof assistant to fill in the gap. The justification is a hint to tell the proof assistant which proof objects are to be used in the process, and how they shall be used.

In our language, justifications are of the form

$$\texttt{by } \pi_1, \ldots, \pi_n \texttt{ using t}$$

the π_k are proof objects which can be hypotheses names (variables) but also more complex terms like application of a general result $H : (\forall x : \texttt{nat}, P\,x)$ to a specific object $n : \texttt{nat}$ to get a proof $(\texttt{H n}) : P\,n$. The expression $\texttt{t}$ is an optional tactic expression that will be used to prove the validity of the step.

The meaning of the π_k objects is that only they and their dependencies — if H is specified and H has type $P\,x$ then x is also implicitly added — will be in the local context in which the new statement is to be proved. If they are omitted then the new statement should be either a tautology or provable without any local hypothesis by the means of the provided tactic. If the user types `by *`, all local hypotheses can be used. The use of `by *` should be transient because it makes the proof more brittle: the justification depends too much on what happened before.

If specified, the tactic `t` will then be applied to the modified context and if necessary, the remaining subgoals will be treated with the `assumption` tactic, which looks for a hypothesis convertible with the conclusion. If `using t` is omitted then a default automation tactic is used. This tactic is a combination of `auto`, `congruence`[4] and `firstorder` [3].

When writing a sequence of deduction steps, it often happens that a statement is only used in the next step. In that case the statement might be anonymous. By using the `then` keyword instead of `have`, this anonymous statement will be added to the list of proof objects to be used in that particular justification.

If a justification fails, the proof assistant issues a warning

<div align="center">

`Warning: insufficient justification.`

</div>

rather than an error. This allows the user to write proofs from the outside in by filling the gaps, rather than linearly from start to end. In the CoqIDE interface, this warning is emphasized by coloring the corresponding command with an orange background. This way, a user reading a proof script will immediately identify where work still needs to be done.

2.4   Partial Conclusion

A (partial) conclusion step is a deduction step of which the statement is a structural subformula of the current thesis and hence can be counted as proven and removed from the thesis. For example, A is a partial conclusion for $A \wedge B$, it is also a partial conclusion for $A \vee B$. In the latter case, choosing to prove A implies a choice in the proof we want to make, by proving A rather than B. In our language, the user can type the command

<div align="center">

`thus` G $justification$

</div>

to provide a partial conclusion G proved using *justification*. If G is not proved, the usual warning is issued, but if G is not a sub-formula of the conclusion then an error occurs: the user is trying to prove the wrong formula.

More precisely, the notion of partial conclusion is a consequence of the encoding of logical connectives as inductive types in Coq. We will look for partial conclusions in the possible sub-terms of proof terms based on inductive type constructors.

In Figure 1, we have exposed the rules for the refinement of a conclusion G using a proof u of V possibly assuming $\bar{U}$ as sufficient conditions. The result is

$$\frac{T\sigma \equiv V \quad \sigma \text{ most general}}{x : T \Rightarrow_u \bar{y} : \bar{U}; \sigma \cup \{x \mapsto (u\,\bar{y})\}} \textbf{ match} \qquad \frac{x : T \Rightarrow_u \Gamma; \sigma}{\Gamma_l, x : T, \Gamma_r \Rightarrow_u \Gamma_l\sigma, \Gamma, \Gamma_r\sigma; \sigma} \textbf{ list}$$

$$\frac{\bar{x} \text{ fresh} \quad C_i \text{ constructor} \quad \bar{x} : \bar{T} \Rightarrow_u \Gamma; \sigma \quad \bar{x} : \bar{T} \vdash C_i\,\bar{p}\,\bar{x} : I\,\bar{p}}{z : (I\,\bar{p}) \Rightarrow_u \Gamma; \sigma \cup \{z \mapsto (C_i\,\bar{p}\,\bar{x})\sigma\}} \textbf{ constructor}$$

$$\frac{z \text{ fresh} \quad z : G \Rightarrow_u \bar{x} : \bar{T}, \sigma}{G \setminus (u : \bar{U} \vdash V) = \Sigma\bar{x} : \bar{T}; \lambda\,(y : \Sigma\bar{x} : \bar{T})\,.\texttt{match}\,y\,\texttt{with}\,\bar{x} \Rightarrow \sigma z\,\texttt{end}} \textbf{ gather}$$

Fig. 1. the partial conclusion mechanism

a pair of a new conclusion G' and a proof of $G \to G'$. If Γ is a list of typed variables and σ a substitution, then $\Gamma\sigma$ is formed by applying σ to all elements of Γ the following way:

- if $\sigma x \neq x$ then $x : T$ is discarded
- is $\sigma x = x$ then $x : T\sigma$ is added to $\Gamma\sigma$

The notation $\Sigma\bar{x} : \bar{T}$ stands for a (possibly dependent) tuple built using the standard (`and`, `prod`, `ex`, `sig`, `sigT`) binary type constructors.

Since a given problem may have several solutions, a simple search strategy has been adopted, and the **constructor** rule has been restricted to non-recursive types in order to keep the search space finite.

thesis	instruction	remaining thesis
$A \wedge B$	`thus` B	A
$(A \wedge B) \wedge C$	`thus` B	$A \wedge C$
$A \vee B$	`thus` B	$\square$
$(A \wedge B) \vee (C \wedge D)$	`thus` C	D
$\exists x : \texttt{nat}, P\,x$	`take` $(2 : \texttt{nat})$	$P\,2$
$\exists x : \texttt{nat}, P\,x$	`thus` $P\,2$	$\square$
$\exists x : \texttt{nat}, \exists y : \texttt{nat}, (P\,y \wedge R\,x\,y)$	`thus` $P\,2$	$\exists x : \texttt{nat}, R\,x\,2$
$A \wedge B$	`suffices to have` $x : \texttt{nat}$ `such that` $P\,x$ `to show` B	$A \wedge \exists x : \texttt{nat}, P\,x$

Fig. 2. examples of partial conclusions

On Fig. 2, we give various examples of uses for the partial conclusion mechanism. The special constant $\square$ stands for a conclusion where everything is proved. When using $P2$ as a partial conclusion for $\exists x : \texttt{nat}, P\,x$, even though a placeholder for `nat` should remain, this placeholder is filled by 2 because of typing constraints. Please note that instantiating an existential quantifier with a specific witness is an instance of this operation.

In a former version of this language, we kept so-called split thesis (i.e several conclusions) instead of assembling them again. We decided to abandon this feature because it added much confusion without extending significantly the

functionalities. In the case of the `suffices` construction, the part of the conclusion is removed and replaced by the sufficient conditions.

Using this mechanism helps the user to build the proof piece by piece using partial conclusions to simplify the thesis and thus keeping track of his/her progress.

3   Syntax and Semantics

3.1   Syntax

Figure 3 gives the complete formal syntax of the declarative language. the unbound non-terminals are *id* for identifiers, *num* for natural numbers, *term* and *type* for terms and types of the Calculus of Inductive Constructions, *pattern* refers to a pattern for matching against inductive objects. Patterns may be nested and contain the `as` keyword and the `_` wildcard, but no disjunctive pattern is allowed.

$$
\begin{aligned}
instruction \ ::= \ &\texttt{proof} \\
&| \ \texttt{assume} \ statement \ [\texttt{and} \ statement]^* \ [\texttt{and} \ (\texttt{we have})\text{-}clause]^? \\
&| \ (\texttt{let}, \texttt{be})\text{-}clause \\
&| \ (\texttt{given})\text{-}clause \\
&| \ (\texttt{consider})\text{-}clause \ \texttt{from} \ term \\
&| \ [\texttt{have}|\texttt{then}|\texttt{thus}|\texttt{hence}] \ statement \ justification \\
&| \ \texttt{thus}^? \ [\sim \ =|= \ \sim] \ [id\text{:}]^? term \ justification \\
&| \ \texttt{suffices} \ [(\texttt{to have}) - clause | statement \ [\texttt{and} \ statement]^*] \\
&\quad \texttt{to show} \ statement \ justification \\
&| \ [\texttt{claim}|\texttt{focus on}] \ statement \\
&| \ \texttt{take} \ term \\
&| \ \texttt{define} \ id \ [var[,var]^*]^? \ \texttt{as} \ term \\
&| \ \texttt{reconsider} \ [id|\texttt{thesis}[[num]]^?] \ \texttt{as} \ type \\
&| \ \texttt{per} \ [\texttt{cases}|\texttt{induction}] \ \texttt{on} \ term \\
&| \ \texttt{per cases of} \ type \ justification \\
&| \ \texttt{suppose} \ [id[,id]^* \ \texttt{and}]^? \ \texttt{is is} \ pattern \\
&\quad [\texttt{such that} \ statement \ [\texttt{and} \ statement]^*[\texttt{and} \ (\texttt{we have})\text{-}clause]^?]^? \\
&| \ \texttt{end} \ [\texttt{proof}|\texttt{claim}|\texttt{focus}|\texttt{cases}|\texttt{induction}] \\
&| \ \texttt{escape} \\
&| \ \texttt{return} \\[6pt]
\alpha,\beta\text{-}clause \ ::= \ &\alpha \ var[,var]^* \ [\beta \ \texttt{such that} \ statement \ [\texttt{and} \ statement]^* \\
&[\texttt{and} \ \alpha,\beta\text{-}clause]^?]^? \\[6pt]
statement \ ::= \ &[id\text{:}]^? type \\
&| \ \texttt{thesis} \\
&| \ \texttt{thesis for} \ id \\[6pt]
var \ ::= \ &id[\text{:}type]^? \\[6pt]
justification \ ::= \ &[\texttt{by} \ [*|term[,term]^*]]^? \ [\texttt{using} \ tactic]^?
\end{aligned}
$$

Fig. 3. Syntax for the declarative language

3.2   Commands Description

```
proof.
  ...
end proof.
```
This is the outermost block of any declarative proof. If several sub-goals existed when the `proof` command occurred, only the first one is proved in the declarative proof. If the proof is not complete when encountering `end proof`, then the proof is closed all the same, but with a warning, and `Qed` or `Defined` to save the proof will fail.

`have ` $h{:}\phi$ ` justification.`
`then ` $h{:}\phi$ ` justification.` This command adds a new hypothesis h of type ϕ in the context. If the justification fails, a warning is issued but the hypothesis is still added to the context. The `then` variant adds the previous fact to the list of objects used in the justification.

`thus ` $h{:}\phi$ ` justification.`
`hence ` $h{:}\phi$ ` justification.` These commands behave respectively like `have` and `then` but the proof of ϕ is used as a partial conclusion. This can end the proof or remove part of the proof obligations. These commands fail if ϕ is not a sub-formula of the thesis.

```
claim h : φ.
  ...
end claim.
```
This block contains a proof of ϕ which will be named h after `end claim`. If the subproof is not complete when encountering `end claim`, then the subproof is still closed, but with a warning, and `Qed` or `Defined` to save the proof later will fail.

```
focus on φ.
  ...
end focus.
```
This block is similar to the `claim` block, except that it leads to a partial conclusion. In a way, `focus` is to `claim` what `thus` is to `have`. This comes handy when the thesis is a conjunction and one of the conjuncts is an implication or a universal quantification: the `focus` block will allow to use local hypotheses.

`(thus) ~= ` t ` justification.`
`(thus) =~ ` t ` justification.` These commands can only be used if the last step was an equality $l = r$. t should be a term of the same type as l and r. If `~=` is used then the *justification* will be used to prove $l = t$ and the new statement will be $l = t$. Otherwise, the *justification* will be used to prove $t = r$ and the new statement will be $t = r$. When present, the `thus` keyword will trigger a conclusion step.

`suffices ` $H : \Phi$ ` to show ` Ψ ` justification.` This command allows to replace a part of the thesis by a sufficient condition, e.g. to strengthen it before

	simple	with previous step	opens sub-proof	iterated equality
intermediate step	have	then	claim	~=/=~
conclusive step	thus	hence	focus on	thus ~=/thus =~

Fig. 4. Synthetic classification of forward steps

starting an proof by induction. In the thesis, Ψ is then replaced by Φ. The justification should prove Ψ using Φ.

`assume G:`Ψ `...and we have x such that H:`Φ`.`

`let x be such that H:`Φ`.`   These commands are two different flavors for the introduction of hypotheses. They expect the thesis to be a product (implication or universal quantifier) of the shape $\Pi_i x_i : Ti.Gi$. It expects the T_i to be convertible with the provided hypotheses statements. This command is well-formed only if the missing types can be inferred.

`given` x `such that` H`:`Φ`.`

`consider` x `such that` H`:`Φ `from` G`.` `given` is similar to `let`, except that this command works up to elimination of tuples and dependent tuples such as conjunctions and existential quantifiers. Here the thesis could be $\exists x.\Phi' \to \Psi$ with Φ' convertible to Φ. The `consider` command takes an explicit object `G` to destruct instead of using an introduction rule.

`define` f $(x:T)$ `as` t`.` This command allows to defines objects locally. If parameters are given, a function (λ-abstraction) is defined.

`reconsider thesis as` T`.`

`reconsider` H `as` T`.`   These commands allows to replace the statement the thesis or a hypothesis with a convertible one, and fails if the provided statement is not convertible.

`take` t`.` This command allows to do a partial conclusion using an explicit proof object. This is especially useful when proving an existential statement: it allows to specify the existential witness.

`per cases on` t`.`
`— of` F $justification$`.`
`suppose` $x:H$`.`
`...`
`suppose` $x':H'$`.`
`...`
`end cases.`   This introduces a proof per cases on a disjunctive proof object t or a proof of the statement F derived from $justification$. The `per cases` command must immediately be followed by a `suppose` command which will introduce the first case. Further `suppose` commands or `end cases` can be typed even if the previous case is not complete. In that case a warning is issued. If t occurs in the thesis, you should use `suppose it is` instead of `suppose`.

`per induction on` t`.`
`— cases —`
`suppose it is` $patt$ `and` $x:H$`.`
`...`
`suppose it is` $patt'$ `and` $x':H'$`.`
`...`
`end cases.`   This introduces a proof per dependent cases or by induction. When doing the proof, t is substituted with $patt$ in the thesis. $patt$ must be a pattern for a value of the same type as t. It may contain arbitrary sub-patterns and `as` statements to bind names to sub-patterns.

Those name aliases are necessary to apply the induction hypothesis at multiple levels. If you are doing a proof by induction, you may use the `thesis for` construction in the `suppose it is` command to refer to an induction hypothesis. You may also write induction hypotheses explicitly.

`escape.`

`...`

`return.` This block allows to escape the declarative mode back to the tactic mode. When encountering the `return` instruction, this subproof should be closed, or else a warning is issued.

3.3   Operational Semantics

The purpose of this section is to give precise details about what happens to the proof state when you type a proof command. The proof state consists of a stack $\mathcal{S}$ that contains open proofs and markers to count open sub-proofs, and each sub-proof is either a judgement $\Gamma \vdash G$, where Γ is a list of types (or propositions) indexed by names and G is a type (or proposition) or the symbol $\square$ which stands for a closed subproof.

A proof script $\mathfrak{S}$ consists of the concatenation of instructions. The rules are given as big-step semantics $\mathfrak{S} \Rightarrow_{\mathfrak{T}} \mathcal{S}$ means that when proving theorem $\mathfrak{T}$, we reach state $\mathcal{S}$ when executing the script $\mathfrak{S}$. This means that any script allowing to reach the empty stack $[]$ is a complete proof of the theorem $\mathfrak{T}$.

For clarity, we only give here the rules for some essential rules, the remaining rules can be found in appendix A.

$$\frac{\mathfrak{T} = \{\Gamma \vdash G\}}{\texttt{proof.} \ \Rightarrow_{\mathfrak{T}} \ (\Gamma \vdash G); []} \qquad \frac{\mathfrak{S} \ \Rightarrow_{\mathfrak{T}} \ (\Gamma \vdash \square); []}{\mathfrak{S} \ \texttt{end proof.} \ \Rightarrow_{\mathfrak{T}} \ []}$$

$$\frac{\mathfrak{S} \ \Rightarrow_{\mathfrak{T}} \ (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash T \quad G \neq \square}{\mathfrak{S} \ \texttt{have} \ (x:T) \ j. \ \Rightarrow_{\mathfrak{T}} \ (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \ \Rightarrow_{\mathfrak{T}} \ (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash T \quad G \neq \square}{\mathfrak{S} \ \texttt{thus} \ (x:T) \ j. \ \Rightarrow_{\mathfrak{T}} \ (\Gamma; x:T \vdash G \setminus (x :\vdash T)); \mathcal{S}}$$

The $j \triangleright \Gamma \vdash G$ expression means that the justification j is sufficient to solve the problem $\Gamma \vdash G$. If it is not, the command issues a warning. The $\equiv$ relation is the conversion ($\beta\delta\iota\zeta$-equivalence) relation of the Calculus of Inductive Constructions (see [2]).

We write $L \sqsubseteq R$ whenever the L context can be obtained by decomposing tuples in the R context. The $\setminus$ operator is defined in Fig. 1, rule **gather**. We use the traditional λ notation for abstractions and Π for dependent products (either implication or universal quantification, depending on the context).

The distinction between $\textsf{cases}_d$ and $\textsf{cases}_{nd}$ is used to prevent the mixing of `suppose` with `suppose it is`. For simplicity, the appendix A omits the coverage condition for case analysis as well as the semantics for `escape` and `return`.

4   Proof Editing

4.1   Well-Formedness

If we drop the justification $(j \triangleright \dots)$ and completeness $(G = \Box)$ conditions in our formal semantics, we get a notion of well-formed proofs. Those proofs when run in Coq, are accepted with warnings but cannot be saved since the proof tree contains gaps.

This does not prevent the user from going further with the proof since the user is still able to use the result from the previous step. The smallest well formed proof is: `proof. end proof.`

Introduction steps such as `assume` have additional well-formedness requirements: the introduced hypotheses must match the available ones. The `given` construction allows a looser correspondence. The `reconsider` statements have to give a convertible type.

For proofs by induction, well-formedness requires the patterns to be of the correct type, and induction hypotheses to be build from the correct sub-objects in the pattern.

4.2   Formal Proof Sketches

We claim that well-formed but incomplete proofs in our language play the role of *formal proof sketches*: they ensure that hypotheses correspond to the current statement and that objects referred to exist and have a correct type. When avoiding the `by *` construction, justifications are preserved when adding extra commands inside the proof. In this sense our language supports incremental proof development.

The only thing that the user might have trouble doing when turning a textbook proof into a proof sketch in our language is ensuring that first-order objects are introduced before a statement refers to them, since a textbook proofs might not be topologically organized. The user will then be able to add new lines within blocks (mostly forward steps).

5   Conclusion

5.1   Further Work

Arbitrary relation composition. The first extension that is needed for our language is the support for iterated relations other than equality. This is possible as soon as a generalized transitivity lemma of the form $\forall xyz, x \, R_1 \, y \to y \, R_2 \, z \to x \, R_3 \, z$ is available.

Better automation. There is a need for a more precise and powerful automation for the default justification method, to be able to give better predictions of when a deduction step will be accepted. A specific need would be an extension of equality reasoning to arbitrary equivalence relations (setoids, PERs . . .).

Multiple induction. The support for induction is already quite powerful (support for deep patterns with multiple `as` bindings), but more can be done if we start considering multiple induction. It might be feasible to detect the induction scheme used (double induction, lexicographic induction ...) to build the corresponding proof on-the-fly.

Translation of procedural proofs. The declarative language offers a stable format for the preservation of old proofs over time. Since many Coq proofs in procedural style already exist, it will be necessary to translate them to this new format. The translation can be done in two ways: by generating a declarative proof either from the proof tree, or from the proof term. The latter will be more fine grain but might miss some aspects of the original procedural proof. The former looks more difficult to implement.

5.2   Conclusion

The new declarative language is now widely distributed, though not yet widely used, and we hope that this paper will help new users to discover our language. As a beginning, Jean-Marc Notin (INRIA Futurs) has translated a part of the Coq Standard Library to the declarative language. The implementation is quite stable and the automation, although not very predictable, offers a reasonable compromise between speed and power.

We really hope that this language will be a useful medium to make proof assistants more popular, especially in the mathematicians community and among undergraduate students. We believe that our language provides a helpful implementation of the formal proof sketch concept; this means it could be a language of choice for turning textbook proofs into formal proofs. It could also become a tool of choice for education.

In the context of collaborative proof repositories such as in [5], our language, together with other declarative languages, will fill the gap between the narrow proof assistant community and the general public: we aim at presenting big formal proofs to the public.

References

1. Coen, C.S.: Automatic generation of declarative scripts. CHAT: Connecting Humans and Type Checkers (December 2006)
2. The Coq Development Team. The Coq Proof Assistant Reference Manual – Version V8.1 (February 2007)
3. Corbineau, P.: First-order reasoning in the calculus of inductive constructions. In: Berardi, S., Coppo, M., Damiani, F. (eds.) TYPES 2003. LNCS, vol. 3085, pp. 162–177. Springer, Heidelberg (2004)
4. Corbineau, P.: Deciding equality in the constructor theory. In: Altenkirch, T., McBride, C. (eds.) TYPES 2006. LNCS, vol. 4502, pp. 78–92. Springer, Heidelberg (2007)

5. Corbineau, P., Kaliszyk, C.: Cooperative repositories for formal proofs. In: Kauers, M., Kerber, M., Miner, R., Windsteiger, W. (eds.) MKM/CALCULEMUS 2007. LNCS (LNAI), vol. 4573, pp. 221–234. Springer, Heidelberg (2007)
6. Giero, M., Wiedijk, F.: MMode, a mizar mode for the proof assistant coq. Technical report, ICIS, Radboud Universiteit Nijmegen (2004)
7. Gordon, M., Milner, R., Wadsworth, C.: Edinburgh LCF. LNCS, vol. 78. Springer, Heidelberg (1979)
8. Harrison, J.: A mizar mode for HOL. In: von Wright, J., Harrison, J., Grundy, J. (eds.) TPHOLs 1996. LNCS, vol. 1125, pp. 203–220. Springer, Heidelberg (1996)
9. Harrison, J.: The HOL Light manual, Version 2.20 (2006)
10. Lamport, L.: How to write a proof. American Mathematics Monthly 102(7), 600–608 (1995)
11. Paulson, L.: Isabelle. LNCS, vol. 828. Springer, Heidelberg (1994)
12. Syme, D.: DECLARE: A prototype declarative proof system for higher order logic. Technical report, University of Cambridge (1997)
13. Syme, D.: Three tactic theorem proving. In: Bertot, Y., Dowek, G., Hirschowitz, A., Paulin, C., Théry, L. (eds.) TPHOLs 1999. LNCS, vol. 1690, pp. 203–220. Springer, Heidelberg (1999)
14. Wenzel, M.: Isar - A generic interpretative approach to readable formal proof documents. In: Bertot, Y., Dowek, G., Hirschowitz, A., Paulin, C., Théry, L. (eds.) TPHOLs 1999. LNCS, vol. 1690, pp. 167–184. Springer, Heidelberg (1999)
15. Wenzel, M., Wiedijk, F.: A comparison of Mizar and Isar. Journal of Automated Reasoning 29(3-4), 389–411 (2002)
16. Wiedijk, F.: Mizar light for HOL light. In: Boulton, R.J., Jackson, P.B. (eds.) TPHOLs 2001. LNCS, vol. 2152, pp. 378–394. Springer, Heidelberg (2001)
17. Wiedijk, F.: Formal proof sketches. In: Berardi, S., Coppo, M., Damiani, F. (eds.) TYPES 2003. LNCS, vol. 3085, pp. 378–393. Springer, Heidelberg (2004)

A   Operational Semantic (Continued)

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma; l : T' \vdash G); \mathcal{S} \quad j, l \triangleright \Gamma \vdash T \quad G \neq \square}{\mathfrak{S} \; \texttt{then} \; (x : T) \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; l : T'; x : T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma; l : T' \vdash G); \mathcal{S} \quad j, l \triangleright \Gamma \vdash T \quad G \neq \square}{\mathfrak{S} \; \texttt{hence} \; (x : T) \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; l : T'; x : T \vdash G \setminus (x :\vdash T)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash r = u \quad G \neq \square}{\mathfrak{S} \; \texttt{\textasciitilde=} \; u \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; e : l = u \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash u = l \quad G \neq \square}{\mathfrak{S} \; \texttt{=\textasciitilde} \; u \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; e : u = r \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash r = u \quad G \neq \square}{\mathfrak{S} \; \texttt{thus \textasciitilde=} \; u \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; e : l = u \vdash G \setminus (\vdash l = u)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash u = l \quad G \neq \square}{\mathfrak{S} \; \texttt{thus =\textasciitilde} \; u \; j. \; \Rightarrow_{\mathfrak{T}} (\Gamma; e : u = r \vdash G \setminus (\vdash u = r)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{claim} \; (x:T).} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash T); \mathsf{claim}; (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{focus on} \; (x:T).} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash T); \mathsf{focus}; (\Gamma; x:T \vdash G \setminus (x :\vdash T)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma' \vdash \Box); \mathsf{claim}; (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S} \; \boxed{\texttt{end claim.}} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S}} \qquad \frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma' \vdash \Box); \mathsf{focus}; (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S} \; \boxed{\texttt{end focus.}} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad \Gamma \vdash t:T \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{take} \; t.} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G \setminus (t :\vdash T)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad \Gamma; \; x_1:T_1, \ldots, x_n:T_n \vdash t:T \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{define} \; f \; (x_1:T_1) \ldots (x_n:T_n) \; \texttt{as} \; t.} \; \Rightarrow_{\mathfrak{T}}}{(\Gamma; f := \lambda x_1:T_1 \ldots \lambda x_n:T_n.t \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash \Pi x_1:T_1' \ldots \Pi x_n:T_n'.G); \mathcal{S} \quad (T_1 \ldots T_n) \equiv (T_1' \ldots T_n')}{\mathfrak{S} \; \boxed{\texttt{assume/let} \; (x_1:T_1) \ldots (x_n:T_n).} \; \Rightarrow_{\mathfrak{T}} (\Gamma; \; x_1:T_1; \ldots; x_n:T_n \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash \Pi x_1:T_1' \ldots \Pi x_n:T_n'.G); \mathcal{S} \quad (T_1 \ldots T_m) \sqsubseteq (T_1' \ldots T_n')}{\mathfrak{S} \; \boxed{\texttt{given} \; (x_1:T_1) \ldots (x_m:T_m).} \; \Rightarrow_{\mathfrak{T}} (\Gamma; \; x_1:T_1; \ldots; x_m:T_m \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad \Gamma \vdash t:T \quad (T_1 \ldots T_n) \sqsubseteq (T) \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{consider} \; (x_1:T_1) \ldots (x_n:T_n) \; \texttt{from} \; t.} \; \Rightarrow_{\mathfrak{T}}}{(\Gamma; \; x_1:T_1; \ldots; x_n:T_n \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma; x:T' \vdash G); \mathcal{S} \quad T \equiv T' \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{reconsider} \; x \; \texttt{as} \; T.} \; \Rightarrow_{\mathfrak{T}} (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash T'); \mathcal{S} \quad T \equiv T' \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{reconsider thesis as} \; T.} \; \Rightarrow_{\mathfrak{T}} (\Gamma \vdash T); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma; x_1:T_1; \ldots; x_n:T_n \vdash T \quad G \neq \Box}{\mathfrak{S} \; \boxed{\texttt{suffices} \; (x_1:T_1) \ldots (x_n:T_n) \; \texttt{to show} \; T \; j.} \; \Rightarrow_{\mathfrak{T}}}{(\Gamma \vdash G \setminus (T_1; \ldots; T_n \vdash T)); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad j \triangleright \Gamma \vdash t:T}{\mathfrak{S} \; \boxed{\texttt{per cases of} \; T \; j.} \; \Rightarrow_{\mathfrak{T}} \mathsf{cases_{nd}}(t:T); (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} (\Gamma \vdash G); \mathcal{S} \quad \Gamma \vdash t:T}{\mathfrak{S} \; \boxed{\texttt{per cases on} \; t.} \; \Rightarrow_{\mathfrak{T}} \mathsf{cases}(t:T); (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_{\mathfrak{T}} \mathsf{cases}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S} \; \boxed{\texttt{suppose} \; (x_1:T_1) \ldots (x_n:T_n).} \; \Rightarrow_{\mathfrak{T}}}{(\Gamma; x_1:T_1; \ldots; x_n:T_n \vdash G); \mathsf{cases_{nd}}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma' \vdash \square); \mathsf{cases_{nd}}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{suppose}\ (x_1:T_1)\ldots(x_n:T_n).\ \Rightarrow_\mathfrak{T} (\Gamma; x_1:T_1;\ldots;x_n:T_n \vdash G); \mathsf{cases_{nd}}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} \mathsf{cases}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{suppose it is}\ p\ \texttt{and}\ (x_1:T_1)\ldots(x_n:T_n).\ \Rightarrow_\mathfrak{T} (\Gamma; x_1:T_1;\ldots;x_n:T_n \vdash G[p/t]); \mathsf{cases_d}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma' \vdash \square); \mathsf{cases_d}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{suppose it is}\ p\ \texttt{and}\ (x_1:T_1)\ldots(x_n:T_n).\ \Rightarrow_\mathfrak{T} (\Gamma; x_1:T_1;\ldots;x_n:T_n \vdash G[p/t]); \mathsf{cases_d}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma' \vdash \square); \mathsf{cases_{d/nd}}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{end cases.}\ \Rightarrow_\mathfrak{T} (\Gamma \vdash \square); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma \vdash G); \mathcal{S} \quad \Gamma \vdash t:T}{\mathfrak{S}\ \texttt{per induction on}\ t.\ \Rightarrow_\mathfrak{T} \mathsf{induction}(t:T); (\Gamma; x:T \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} \mathsf{induction}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{suppose it is}\ p\ \texttt{and}\ (x_1:T_1)\ldots(x_n:T_n).\ \Rightarrow_\mathfrak{T} (\Gamma; x_1:T_1;\ldots;x_n:T_n \vdash G[p/t]); \mathsf{induction}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma' \vdash \square); \mathsf{induction}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{suppose it is}\ p\ \texttt{and}\ (x_1:T_1)\ldots(x_n:T_n).\ \Rightarrow_\mathfrak{T} (\Gamma; x_1:T_1;\ldots;x_n:T_n \vdash G[p/t]); \mathsf{induction}(t:T); (\Gamma \vdash G); \mathcal{S}}$$

$$\frac{\mathfrak{S} \Rightarrow_\mathfrak{T} (\Gamma' \vdash \square); \mathsf{induction}(t:T); (\Gamma \vdash G); \mathcal{S}}{\mathfrak{S}\ \texttt{end induction.}\ \Rightarrow_\mathfrak{T} (\Gamma \vdash \square); \mathcal{S}}$$