

# *Dependent Types Ensure Partial Correctness of Theorem Provers*

Andrew W. Appel

*Princeton University, 35 Olden Street, Princeton, NJ 08544, USA*  
(e-mail: `appel@princeton.edu`)

Amy P. Felty,

*University of Ottawa, 800 King Edward Ave., Ottawa, Ontario K1N 6N5, Canada*  
(e-mail: `afelty@site.uottawa.ca`)

---

## Abstract

Static type systems in programming languages allow many errors to be detected at compile time that wouldn't be detected until runtime otherwise. Dependent types are more expressive than the type systems in most programming languages, so languages that have them should allow programmers to detect more errors earlier. In this paper, using the Twelf system, we show that dependent types in the logic programming setting can be used to ensure partial correctness of programs which implement theorem provers, and thus avoid runtime errors in proof search and proof construction. We present two examples: a tactic-style interactive theorem prover and a union-find decision procedure.

---

## 1 Introduction

Many theorem proving systems implement tactics and tacticals, which provide flexible goal-directed proof search. Tactics reduce goals to subgoals, while tacticals are primitives for combining tactics into larger ones that can perform multiple proof steps. They also allow programming of proof search strategies. Some of the first provers using this style of proof search (e.g. LCF (Gordon *et al.*, 1979) and HOL (Gordon & Melham, 1993)) were written in ML, whose pattern matching, exception handling, and polymorphic type system are useful in writing tactics concisely. Felty (1993) showed that Lambda Prolog's (Nadathur & Miller, 1988) higher-order unification, backtracking, and polymorphic type system provided a more expressive notation for writing tactics and tacticals. Specifically, higher-order abstract syntax is more useful and expressive than ML pattern matching, backtracking is more concise than exception handling, but Lambda Prolog's prenex-polymorphic type system is essentially similar to ML's.

In this paper we will discuss the advantages of a dependent type system over ML-style polymorphism for writing theorem provers. Dependent types could be used in a functional language (such as ML) or a logic-programming language (such as Lambda Prolog); we use the logic-programming language Twelf (Pfenning & Schürmann, 1999). This means that the style of prover we illustrate is similar to those presented

by Felty (1993), but the issue of ML-style types vs. dependent types is orthogonal to the issue of ML-style or Prolog-style control and data structures. A decade of experience with tactical Felty’s prover has shown that this technique is expressive and powerful, and could be used as the core of a full interactive theorem prover similar in strength to many existing provers such as HOL and Isabelle that have been used in a variety of large-scale applications; we expect that the dependently typed variant of Felty’s approach would scale just as well.

A problem in the implementation of theorem provers (tactical and other) is that they may have bugs. That is, the “proof” constructed by the prover may not be valid. There are at least two ways that industrial-strength theorem provers defend against invalid proofs:

- Edinburgh LCF (and Isabelle (Paulson, 1994), HOL, etc.) have an unforgeable abstract data type *theorem*. An attempt by a prover to construct an invalid proof will be detected at run time when some (privileged) proof-constructor function detects mismatched arguments.
- Coq (Barras *et al.*, 1998) (and Elf, Twelf, etc.) require provers to construct proof witnesses that can be checked (in principle) by a small and reliable type-checker that’s independent of any (complex, unreliable) theorem prover. Provers in Coq and Twelf have usually been written in ML (Caml and Standard ML, respectively); although each of these systems contains a dependently typed language (functional and logic-programming, respectively), that language is meant for describing objects in the object logic, and not as a language for programming provers.

But in each case, bugs in the tactics (or other proof-search algorithm) will be detected only when the tactics are executed: either when they attempt to use a proof constructor with bad arguments, or when a proof witness fails the type-checker.

Static type systems (such as ML’s) have the advantage over dynamic type systems (such as Lisp’s) that many errors are detected much earlier, without needing to run the program on an adequate sample of test cases. The languages in which the above-described theorem provers are implemented – Standard ML, Caml, Lambda Prolog– all have static checking. But ML-style polymorphism is not strong enough to catch all programming errors.

We had experience in building a complicated tactical prover prototype (in Lambda Prolog) for a proof-carrying code application (Appel & Felty, 2000). We had a collection of complicated, ad-hoc tactics (as required by Gödel’s incompleteness theorem, any sufficiently powerful prover will be complicated and ad-hoc). As we maintained the prover, from time to time we found that it built invalid proofs; to debug, we had to do runtime tracing of appropriately stripped-down test cases to isolate the problem.

As we will show, using a dependently typed programming language can yield a partial correctness (i.e., soundness) guarantee for a theorem prover: if the implementation type-checks, then any proof (or subproof) that it builds will be valid. There is no total correctness (i.e., completeness) guarantee: that is, the prover might still

infinite-loop or be incomplete in some other way – i.e., fail with a run-time exception (in ML) or backtracking failure (Prolog).

The source code for all our examples can be found at [www.cs.princeton.edu/~appel/prover/](http://www.cs.princeton.edu/~appel/prover/).

## 2 LF and Twelf

The logical framework LF (Harper *et al.*, 1993) allows the specification of logics, and implementations of LF such as Twelf (Pfenning & Schürmann, 1999) allow checking of proofs in those logics. Another view of LF/Twelf is that it is a higher-order dependently typed logic-programming language: Prolog, with higher-order abstract syntax (as in Lambda Prolog), well-scoped dynamic clauses, and a dependent type system. We will make use of both views of LF/Twelf: we will specify an object logic (e.g., first-order logic or higher-order logic), and we will also do Prolog-like programming of the prover tactics.

The Twelf (Pfenning & Schürmann, 1999) implementation of LF has (partial) type inference, proof search (i.e., Prolog-style backtracking), and constraint domains (e.g., the theory of the rational numbers). Twelf has a distinguished type *type*, the type of all types (and the type of logic-programming goals). A *constructor declaration* declares an axiom or inference rule of a logic, or a logic-programming data-constructor, or a logic-programming clause. A *definition* can be used to make a theorem, a lemma, or a defined function or predicate.<sup>1</sup>

*An object logic.* We will use Twelf to write theorem provers. We begin by defining operators and axioms of an object logic: here we use first-order logic, which is encoded by the Twelf declarations in Figure 1. Everything we do in this paper also works for higher-order and other logics; but we wish to simplify the presentation.

The declaration  $i : \text{type}$  declares the type  $i$  of individuals (over which the quantifiers range), and  $o : \text{type}$  declares the type of logical formulas (booleans). The constant  $pf$  is a dependent type constructor: for any formula  $A$ ,  $pf(A)$  is a type; we interpret this type to mean, “proofs of the formula  $A$ .” That is, if  $p$  has type  $pf(\text{false imp true})$ , then  $p$  must be a proof of *false imp true*.

The `%use` declaration brings in the (built-in) theory of the rational numbers, with constants 0, 1, 2, 3/2, 248/83, and so on, and operators  $+$ ,  $\times$ ,  $>$ ,  $\geq$ . Though we don’t need the full power of the rationals – we use numbers only to index elements of our hypothesis list – this is Twelf’s preferred number system, so it’s simplest to just use it. We define a datatype constructor  $const : \text{rational} \rightarrow i$  to inject rational constants into our logic’s element type.

We define infix operators *imp*, *and*, and *or* to construct formulas. The proof-constructor *and<sub>i</sub>* (and-introduction) can be read as, “function taking a proof of  $A$  and returning (function taking a proof of  $B$  and returning a proof of  $A$  and  $B$ ).” Thus if  $p1 : pf(A)$  and  $p2 : pf(B)$ , then *and<sub>i</sub>*  $p1$   $p2 : pf(A$  and  $B)$ .

<sup>1</sup> Or even a logic-programming clause justified by a proof, though we won’t use that capability in this paper.

---

```

i      : type .
o      : type .
pf     : o → type .
%use inequality/rationals .
const  : rational → i .

imp    : o → o → o .      %infix right 10 imp .
imp_i  : (pf A → pf B) → pf (A imp B) .
imp_e  : pf (A imp B) → pf A → pf B .

and    : o → o → o .      %infix right 12 and .
and_i  : pf A → pf B → pf (A and B) .
and_e1 : pf (A and B) → pf A .
and_e2 : pf (A and B) → pf B .

or     : o → o → o .      %infix right 11 or .
or_i1  : pf A → pf (A or B) .
or_i2  : pf B → pf (A or B) .
or_e   : pf (A or B) → (pf A → pf C) → (pf B → pf C) → pf C .

forall : (i → o) → o .
forall_i : ({x : i} pf (A x)) → pf (forall A) .
forall_e : pf (forall A) → {x : i} pf (A x) .

exists : (i → o) → o .
exists_i : {x : i} pf (A x) → pf (exists A) .
exists_e : pf (exists A) → ({x : i} pf (A x) → pf B) → pf B .

false  : o .
false_e : pf false → pf A .

```

---

Fig. 1. First-order logic.

The proof-constructor *imp\_i* (implication-introduction) can be read as, “function taking (function from proof of *A* to proof of *B*) and returning proof of (*A imp B*).” Twelf’s function notation uses square brackets for lambda, thus  $([p] \text{ and\_i } p \ p)$  is a function with formal parameter *p* and result  $(\text{and\_i } p \ p)$ . Alternately, we can read  $\text{imp\_i}[p : \text{pf } A] \ Q(p)$  to mean, assuming *A* is true (with proof *p*), then the expression  $Q(p)$  is a proof of *B*; thus *A imp B*.

In the following lemma, represented as a Twelf definition, we apply *imp\_i* to this function to get the proof in the body of the definition.

$$\text{lemma1} : \text{pf } (A \text{ imp } (A \text{ and } A)) = \text{imp\_i } ([p : \text{pf } A] \text{ and\_i } p \ p).^2$$

As in most presentations of lambda-calculus, the lambda (square brackets) has a syntactic scope that extends as far as possible to the right; Twelf can reconstruct

<sup>2</sup> Unbound capitalized variables are implicitly universally quantified, so Twelf would internally reconstruct this definition to

$$\text{lemma1} : \{A : o\} \text{pf } (A \text{ imp } (A \text{ and } A)) = [A : o] \text{imp\_i } ([p : \text{pf } A] \text{ and\_i } p \ p).$$

where the curly braces construct dependent types: the type of *lemma1* is, in effect, “function from formula (call it *A*) to proofs of *A imp (A and A)*.”

the type of the function argument; and our *and* binds tighter than *imp*; so we could also write

$$\text{lemma1} : \text{pf} (A \text{ imp } A \text{ and } A) = \text{imp\_i} [p] \text{ and\_i } p \text{ p} .$$

Using this style of definition and proof, we introduce some useful definitions and lemmas:

$$\begin{aligned} \text{not} & : o \rightarrow o = [A] A \text{ imp } \text{false} . \\ \text{not\_i} & : (\text{pf } A \rightarrow \text{pf } \text{false}) \rightarrow \text{pf} (\text{not } A) = \text{imp\_i} . \\ \text{not\_e} & : \text{pf}(\text{not } A) \rightarrow \text{pf } A \rightarrow \text{pf } \text{false} = \text{imp\_e} . \\ \text{true} & : o = \text{not } \text{false} . \\ \text{true\_i} & : \text{pf}(\text{true}) = \text{not\_i} [p] p . \end{aligned}$$

Notational definitions in Twelf are like type abbreviations in ML: the type-checker can freely expand them when type-checking. Furthermore, the type-checker's unifier uses rules of beta-eta equivalence. Thus, the proof of the true-introduction rule *true\_i* must have type *pf(true)* which is equivalent (by definition) to *pf(not false)*; the right-hand-side of *true\_i* is *not\_i [p] p* whose type is indeed *pf(not false)*. Note that even though *not\_i* is defined to be *imp\_i*, it is really the special case where the *B* in *imp\_i* is instantiated with *false*.

These definitions – including the proofs of the lemmas *not\_i*, *not\_e*, *true\_i* – are type-checked by the system, so they can't be invalid. This means that we don't really need a prover at all; we could just write proofs by hand (as definitions) and check them in Twelf's type-checker; and in fact such a method can be quite effective and useful (Appel, 2000).

However, we wish to automate: we will write a program to produce proofs semi-automatically or automatically, guided by tactical hints. Since Twelf's support for interactive I/O is minimal, in the prototype we do “interactive” tactical proving by editing proof-scripts.

### 3 A theorem prover using tactics and tacticals

Our prover manipulates *goals*, which are data structures of the form  $h_1, \dots, h_n \vdash h$ , where each of the  $h_i$  is a hypothesis, represented as a formula with attached proof. For  $h_1, \dots, h_n$  we assume that the proof is already constructed. The conclusion  $h$  is also a formula with attached proof; typically we have not yet found the proof, so its “attached proof” is an uninstantiated logic variable.

The Twelf declarations for such data structures are as follows. *hyp* is the type of a single hypothesis, and *hyps* is a list of hypotheses:

$$\begin{aligned} \text{hyp} & : \text{type} . \\ \text{hyps} & : \text{type} . \end{aligned}$$

An individual hypothesis is a pair of some formula *A* and a proof of that formula; we declare the nonassociative infix constructor *by* to construct such formula/proof

pairs:

$$\text{by} : \{A : o\} \text{ pf}(A) \rightarrow \text{hyp}. \quad \%infix \text{ none } 5 \text{ by}.$$

This is a dependently typed constructor. Thus,  $(\text{true by true}_i)$  is well typed, but  $(\text{false by true}_i)$  is ill typed, even though  $\text{false}$  is a formula and  $\text{true}_i$  is a proof – it’s the wrong type of proof.

In order to write  $A \text{ by } P$  instead of  $\text{by } A \ P$ , we declare  $\text{by}$  as an infix operator (nonassociative, binding tightness 5) using the `%infix` declaration shown above.

To make hypothesis lists we declare two constructors for  $\text{hyps}$ , where our  $\text{cons}$  is an infix comma:

$$\begin{aligned} \text{nil} & : \text{hyps}. \\ , & : \text{hyp} \rightarrow \text{hyps} \rightarrow \text{hyps}. \quad \%infix \text{ right } 4 \text{ , } . \end{aligned}$$

Now we can declare the  $\text{goal}$  type with its infix constructor  $\vdash^3$ .

$$\begin{aligned} \text{goal} & : \text{type}. \\ \vdash & : \text{hyps} \rightarrow \text{hyp} \rightarrow \text{goal}. \quad \%infix \text{ none } 3 \vdash . \\ \& & : \text{goal} \rightarrow \text{goal} \rightarrow \text{goal}. \quad \%infix \text{ right } 2 \& . \\ \text{allp} & : (\text{pf } A \rightarrow \text{goal}) \rightarrow \text{goal}. \\ \text{alli} & : (i \rightarrow \text{goal}) \rightarrow \text{goal}. \\ \text{tt} & : \text{goal}. \end{aligned}$$

In addition to the basic goal  $h_1, \dots, h_n \vdash h$  we have compound goals  $G_1 \& G_2$  to represent the case where the use of a tactic results in several subgoals (remaining proof obligations). The empty goal  $\text{tt}$  is the identity for  $\&$  and indicates no remaining proof obligations. As we will explain later, we need separate constructors  $\text{allp}$  and  $\text{alli}$  because Twelf is not a polymorphic language. This implementation of goals can be viewed as the Twelf version of a similar implementation in Lambda Prolog (Felty, 1993). The programs which manipulate them, in particular the tacticals and the *maptac* program below, are similar also. They do not make any essential use of dependent types, and thus do not contribute to the partial correctness of our tacticals. It is mainly the type of the  $\text{by}$  constructor introduced above that is important for guaranteeing partial correctness of our tacticals.

*Tactics.* A tactic is a procedure which takes a goal as input and returns subgoals that remain to be proven. We first show some simple tactics that implement the application of inference rules and lemmas, and later show some more complex tactics which perform some proof search. We first need the type  $\text{tac}$  of *tactic names*, and then we define the names of some tactics:

<sup>3</sup> Identifiers in the real Twelf system must be written in ASCII, of course, so we use the symbol  $\vdash$  for  $\vdash$ .

```

tac          : type.
initial_tac : tac.
and_r_tac   : tac.
and_l_tac   : rational → tac.
imp_r_tac   : tac.
imp_l_tac   : rational → tac.

```

We define *tactic* as the interpreter relation for the logic program; that is, the expression *tactic T G<sub>1</sub> G<sub>2</sub>* is a logic-programming goal that applies the tactic named *T* to the proof obligation *G<sub>1</sub>*, resulting in remaining proof obligations *G<sub>2</sub>*.

```
tactic : tac → goal → goal → type.
```

Finally, we define clauses for the *tactic* relation. Generally, there are one or two clauses for each tactic-name. Examples are:

```

t1 : tactic initial_tac (Hs ⊢ A by P) tt ←
      nth_item N (A by P) Hs.

t2 : tactic and_r_tac (Hs ⊢ (A and B) by (and_i P1 P2))
      (Hs ⊢ A by P1 & Hs ⊢ B by P2).

t3 : tactic imp_r_tac (Hs ⊢ (A imp B) by (imp_i P1))
      (allp [p2 : pf A](A by p2 , Hs ⊢ B by (P1 p2))).

t9 : tactic (and_l_tac N) (Hs ⊢ C by P)
      ((A by (and_e1 Q)) , (B by (and_e2 Q)) , Hs ⊢ C by P) ←
      nth_item N ((A and B) by Q) Hs.

t11 : tactic (imp_l_tac N) (Hs ⊢ C by P)
      ((Hs ⊢ A by P2) & ((B by (imp_e P1 P2)) , Hs ⊢ C by P)) ←
      nth_item N ((A imp B) by P1) Hs.

```

The lines *t1*, *t2*, ... can be understood as logic-programming clauses, where ← is used instead of the Prolog or Lambda Prolog :- . Thus, the rule *t1* might be written in Lambda Prolog as

```

tactic initial_tac (Hs |- (A by P)) tt :-
  nth_item N (A by P) Hs.

```

where the data constructors |- and by are infix (of course, in Lambda Prolog the type-checker can't check soundness of the tactic).

The operational interpretation of a Prolog clause *H* :- *G<sub>1</sub>*; *G<sub>2</sub>*; *G<sub>3</sub>* or a Twelf clause *H* ← *G<sub>1</sub>* ← *G<sub>2</sub>* ← *G<sub>3</sub>* is, first match the head *H* against the current goal. If that matches, try and satisfy subgoal *G<sub>1</sub>*; if that matches, satisfy subgoal *G<sub>2</sub>*, and so on. Twelf, like Prolog, uses backtracking (so that if *G<sub>2</sub>* fails, then a different way of satisfying *G<sub>1</sub>* is tried, and so on).

The supporting clauses for *nth\_item N H Hs* are straightforward (typed) Prolog, and define the relation that the *N*th item of *Hs* is precisely *H*:

```

nth_item    : rational → hyp → hyps → type.
nth_item1   : nth_item 1 H1 (H1 , Hs).
nth_itemN   : nth_item N H1 (H2 , Hs) ← nth_item (N - 1) H1 Hs.

```

Thus, *initial\_tac* matches a goal  $Hs \vdash A$  by  $P$  if there exists an  $N$  such that the hypothesis  $A$  by  $P$  is the  $N$ th item of  $Hs$  (in Isabelle this is called *assume\_tac*).

We can let Prolog backtracking find the right  $N$  for *initial\_tac* because the subgoals are trivial, but for *andL\_tac* it would be unwise to rely on this, because *andL\_tac* has nontrivial subgoals. Therefore the user must supply a number when using *andL\_tac*, but has the option of supplying a Prolog unification variable, which causes *nth\_item* to do a backtracking search for an assumption of the form  $A$  and  $B$ .

The tactic implementation of most of the right introduction rules of our object logic is straightforward. The input goal contains the conclusion paired with its proof, and the output goal contains the hypotheses paired with their proofs. If there is more than one subgoal, they are connected by  $\&$ , as in *andR\_tac*. Rules which use nested implication or quantification in Twelf such as *imp\_i* and *forall\_i* in Figure 1 must use one of the *all* goal constructors in their tactic implementations. For example, the argument to *imp\_i* is a function from proofs of  $A$  to proofs of  $B$ . In the tactic version (*t3* above), the use of *allp* introduces a bound variable  $p2$  to represent an arbitrary proof of  $A$  which gets paired with  $A$  and added to the assumption list  $Hs$  of the subgoal.

The tactics for the left introduction rules are implemented so that they perform forward proof from hypotheses. An argument is given to indicate the position in the hypothesis list of the hypothesis to which the rule should be applied. The partial proofs are constructed and added to the hypothesis lists of the subgoals.

For each of the left introduction rules, we provide a second version of the tactic which removes the hypothesis to which the specified rule is applied when forming the subgoal. For example, for and-elimination, we have:

```

t10: tactic (andL_tacR N) (Hs1 ⊢ C by P)
      ((A by (and_e1 Q)) , (B by (and_e2 Q)) , Hs2 ⊢ C by P) ←
      nth_and_rest N ((A and B) by Q) Hs1 Hs2.

```

where *nth\_and\_rest* is a logic-programming predicate which finds the  $N$ th formula in  $Hs1$  and returns the set of hypotheses  $Hs2$  with the  $N$ th one removed. Such tactics are useful in writing automated proof search procedures so that they can avoid repeatedly applying the same rule to the same hypothesis.

*More tactics.* Using these general principles, it's easy to implement a large variety of tactics. Here we show three more:

```

forallR_tac : tac.
forallL_tac : rational → tac.
resolve2_tac : (pf A1 → pf A2 → pf B) → tac.

```



$t7$ : *tactic forall<sub>r</sub>\_tac* ( $\Gamma \vdash (\text{forall } A) \text{ by } (\text{forall}_i P)$ )  
 $(\text{alli } [t : i](\Gamma \vdash (A \ t) \text{ by } (P \ t)))$ .  
 $t17$ : *tactic forall<sub>l</sub>\_tac*  $N$  ( $\Gamma \vdash C \text{ by } P$ )  
 $((A \ X) \text{ by } (\text{forall}_e Q \ X)) ; \Gamma \vdash C \text{ by } P \leftarrow$   
 $\text{nth\_item } N ((\text{forall } A) \text{ by } Q) \Gamma$ .  
 $t25$ : *tactic resolve2\_tac* ( $\text{Thm} : \text{pf } A1 \rightarrow \text{pf } A2 \rightarrow \text{pf } B$ )  
 $(\Gamma \vdash B \text{ by } (\text{Thm } P1 \ P2))$   
 $(\Gamma \vdash A1 \text{ by } P1 \ \& \ \Gamma \vdash A2 \text{ by } P2)$ .

To prove a universally quantified formula  $\forall x:A(x)$ , *forall<sub>r</sub>\_tac* introduces an *alli* goal; then clause *m4* (shown below) will dynamically create an atom of type *i*, so that the subgoal, in effect, is to prove *A* with the new atom substituted for *x*. The substitution is handled entirely by the Twelf metalogic (the same would be true in Lambda Prolog).

To make use of a universally quantified hypothesis, *forall<sub>l</sub>\_tac* uses the argument *N* to select the *N*th hypothesis from the assumptions, which must be of the form *forall A* (equivalently, *forall [x] A(x)*). A logic variable *X* is introduced to instantiate the bound variable in *A*. It can later be unified with a term that is needed to complete the proof. Then *A X* is unified with the hypothesis in the goal formula; although this is higher-order unification (which is undecidable in general), extensive experience with the use of similar tactics in Lambda Prolog has found them to work fine in practice. We can also write a version of this tactic that allows the user to provide the instantiation term *X* at the time the tactic is applied. We do this by adding *X* to the first argument as follows:

*forall<sub>l</sub>\_tacx* : *rational*  $\rightarrow i \rightarrow \text{tac}$ .  
 $t17x$ : *tactic forall<sub>l</sub>\_tacx*  $N \ X$  ( $\Gamma \vdash C \text{ by } P$ )  
 $((A \ X) \text{ by } (\text{forall}_e Q \ X)) ; \Gamma \vdash C \text{ by } P \leftarrow$   
 $\text{nth\_item } N ((\text{forall } A) \text{ by } Q) \Gamma$ .

We have also shown an example of a resolution tactic. Given some theorem *T* of the form,  $\text{pf}(A_1) \rightarrow \text{pf}(A_2) \rightarrow \text{pf}(B)$ , the tactic *resolve2\_tacT* matches a goal *B* and produces subgoals *A*<sub>1</sub> and *A*<sub>2</sub>. A minor disadvantage of doing this in a well typed way is that we need a different tactic for 2-premise theorems than for 3-premise theorems, and so on. Note that the user need not type in a proof term for the *Thm* argument directly. Instead, the name of a previously defined Twelf declaration which expresses a lemma can be given, as long as it has the right type. By Twelf definition expansion, this name is equivalent to the term it abbreviates.

*Tacticals*. Tacticals implement basic control mechanisms which allow simple tactics to be combined into more complex ones, and can be used as a programming language to implement search procedures. Most tacticals assume the input goal is a basic goal (constructed using  $\vdash$  in our prover). In the logic programming setting, we first implement a *maptac* tactical which applies tactics to compound goals, reducing

them to basic goals before passing them on to other tacticals and tactics.

```

mptac : tac → goal → goal → type.
m1 :    mptac T tt tt.
m2 :    mptac T (InG1 & InG2) (OutG1 & OutG2) ←
        mptac T InG1 OutG1 ← mptac T InG2 OutG2.
m3 :    mptac T (allp InG) (allp OutG) ←
        {p} mptac T (InG p) (OutG p).
m4 :    mptac T (alli InG) (alli OutG) ←
        {t} mptac T (InG t) (OutG t).
m5 :    mptac T (Hs ⊢ A by P) OutG ←
        tactic T (Hs ⊢ A by P) OutG.

```

This tactical reduces the goal to subgoals in a manner consistent with the meaning of the top-level goal constructor. In the clauses for the *all* constructors, the quantification within goals is transferred to quantification in Twelf. For example, *allp* quantifies over proofs in the object logic; in the *m3* clause, *p* is introduced as an arbitrary proof to replace the bound variable in *InG*. After completion of the Twelf subgoal, *OutG* is also an abstraction over *p*.

Since *mptac* has the same type as *tactic*, we could have dispensed with *mptac* and written *m1*; ... ; *m4* as clauses for *tactic*; but this would allow the user less control of how and when the tactics are applied.

Some common tacticals found in most tactic-style theorem provers are implemented in Twelf with the following clauses.

```

idtac :    tac.
then :    tac → tac → tac.    %infix left 2 then.
orelse :  tac → tac → tac.    %infix left 2 orelse.
repeat :  tac → tac.
try :    tac → tac.
complete : tac → tac.

tactical1 : tactic idtac G G.
tactical2 : tactic (T1 then T2) InG OutG ←
            tactic T1 InG MidG ← mptac T2 MidG OutG.
tactical3 : tactic (T1 orelse T2) InG OutG ← tactic T1 InG OutG.
tactical4 : tactic (T1 orelse T2) InG OutG ← tactic T2 InG OutG.
tactical5 : tactic (repeat T) InG OutG ←
            tactic ((T then (repeat T)) orelse idtac) InG OutG.
tactical6 : tactic (try T) InG OutG ← tactic (T orelse idtac) InG OutG.
tactical7 : tactic (complete T) InG tt ←
            tactic T InG OutG ← goalreduce OutG tt.

```

The *idtac* tactical returns the goal unchanged and is used mainly in programming search strategies for ending a series of multiple proof steps. The *then* tactical performs the composition of tactics. The *orelse* tactical is also useful in programming search strategies and allows choice of tactics. The *repeat* tactical repeatedly applies

a tactic until it can no longer be applied and is defined in terms of the others. The *try* tactical prevents failure of the given argument tactic by using *idtac* when tactic *T* fails. Finally the *complete* tactical tries to completely solve the given goal. It uses *goalreduce* (not shown) to simplify compound goal expressions by removing occurrences of *tt* from them. For example, applying multiple tactics could result in goal structures such as (*allp* (*[x]tt & tt*)) whose only subgoals are *tt* and so should reduce to *tt*.

#### 4 A more intricate tactic

An important property of a tactical prover is that it is *extensible*, so that its users can write their own tactics. It is in the checking of user-defined tactics that the dependent type system is particularly useful. To illustrate, we will show a specialized tactic of the kind that some user might write.

Suppose we have a sum-of-products assertion,

$$\begin{aligned} C = & (A_{11} \wedge A_{21} \wedge A_{31} \wedge \top) \vee \\ & (A_{12} \wedge A_{22} \wedge \top) \vee \\ & (A_{13} \wedge A_{23} \wedge A_{33} \wedge A_{43} \wedge \top) \vee \\ & \perp \end{aligned}$$

and we want to prove *C* implies *D*, where we know  $A_{i1} \vdash D$ ,  $A_{i2} \vdash D$ ,  $A_{i3} \vdash D$ , for a particular *i*. To handle this we can write a tactic *sumprod(i)*.

This kind of situation comes up, for example, in proving properties of a program that fetches from an ML-style sum-of-products datatype. Suppose some value *x* belongs to an ML datatype that has three constructors (disjuncts), which take values that are all records (of 3 elements, 2 elements, and 4 elements, respectively). We would like to fetch and use the 2nd record field even before doing the case-analysis that tells us which disjunct applies. To do this “hoist” operation, we need to prove that the second field exists (in each disjunct) and has the right properties. The *sumprod* tactic will be useful in such proofs. But clearly it’s a very specialized situation – therefore this tactic will be user-defined, not provided by default.

We start with two preliminary lemmas. The specialized subtactics of *sumprod* will apply these specialized lemmas:

$$\begin{aligned} \text{or\_imp} : & \text{pf } (A \text{ imp } C) \rightarrow \text{pf } (B \text{ imp } C) \rightarrow \text{pf } ((A \text{ or } B) \text{ imp } C) = \\ & [p1 : \text{pf } (A \text{ imp } C)] [p2 : \text{pf } (B \text{ imp } C)] \\ & \text{imp\_i } [p3 : \text{pf } (A \text{ or } B)] \\ & \text{or\_e } p3 ([p4 : \text{pf } A] \text{ imp\_e } p1 p4) ([p5 : \text{pf } B] \text{ imp\_e } p2 p5) . \\ \text{and\_imp} : & \text{pf } (B \text{ imp } D) \rightarrow \text{pf } (A \text{ and } B \text{ imp } D) = \\ & [p1 : \text{pf } (B \text{ imp } D)] \text{ imp\_i } [p2 : \text{pf } (A \text{ and } B)] \text{ imp\_e } p1 (\text{and\_e2 } p2) . \end{aligned}$$

We start with an auxiliary tactic *prodn(j)* that converts the goal  $Hs \vdash (A_1 \wedge A_2 \wedge \dots \wedge A_n \wedge \top) \rightarrow D$  to the goal  $A_j, Hs \vdash D$ .

```

prodn: rational → tac .
t136 : tactic (prodn 1) (Hs ⊢ A and As imp D by (imp_i [p] P (and_e1 p)))
      (allp [p] (A by p , Hs ⊢ D by P p)) .
t137 : tactic (prodn N) (Hs ⊢ A and As imp D by and_imp P) G ←
      tactic (prodn (N - 1)) (Hs ⊢ As imp D by P) G .

```

Finally, the tactic *sumprod*(*i*) transforms the goal  $Hs \vdash (\bigvee_i \bigwedge_j A_{ij}) \rightarrow D$  to the goal  $(A_{i1}, Hs \vdash D) \& \dots \& (A_{in}, Hs \vdash D)$ :

```

sumprod: rational → tac .
t134 : tactic (sumprod N) (Hs ⊢ false imp D by (imp_i false_e)) tt .
t135 : tactic (sumprod N)
      (Hs ⊢ (A or As) imp D by (or_imp P1 P2)) (G1 & G2) ←
      tactic (prodn N) (Hs ⊢ A imp D by P1) G1 ←
      tactic (sumprod N) (Hs ⊢ As imp D by P2) G2 .

```

To see how the dependent type system ensures that we got this right, let's examine the typechecking of rule *t135*. As reconstructed by Twelf's typechecker, we have,

```

t135 :
  {N:rational} {Hs:hyps} {As:o} {D:o} {P2:pf (As imp D)} {G2:goal} {A:o}
  {P1:pf (A imp D)} {G1:goal}
  tactic (sumprod N) (Hs |- As imp D by P2) G2
    -> tactic (prodn N) (Hs |- A imp D by P1) G1
    -> tactic (sumprod N) (Hs |- A or As imp D by or_imp P1 P2) (G1 & G2) .

```

Here we have explicit metalevel quantification (using curly braces) of all the implicitly quantified logical variables  $N$ ,  $Hs$ ,  $As$ ,  $D$ , etc. The type of  $P1$  was inferred from the expression  $A \text{ imp } D \text{ by } P1$ : it must be  $pf(A \text{ imp } D)$ . Therefore the use of  $P1$  in the expression  $or\_imp P1 P2$  typechecks.

But suppose we had mistakenly written the rule *t135* with

$$A \text{ or } As \text{ imp } D \text{ by } and\_imp P2 .$$

Then this rule wouldn't type-check, and Twelf would report the error,

```

Type mismatch
Expected: pf ('A or 'As imp 'D)
Found:   pf (X1 and 'As imp 'D)

```

When writing tactics such as this (but quite a bit messier) in Lambda Prolog, we found that mismatches between tactics and the lemmas that they apply were one of the two common sources of errors in the prover; such errors do not impede us in Twelf. The other kind of error – incompleteness via infinite loops or backtracking failure – continues to be bothersome, of course: dependent types do not save us there.

## 5 Union-Find

Not only tactical provers, but also other decision procedures can be dependently typed to ensure partial correctness. For example, in decision procedures for equality, the standard efficient union-find algorithm with path compression (Aho *et al.*, 1974) is often used to represent equivalence classes.

For each equivalence class, the algorithm maintains a canonical representative. As new equalities are learned (from some other source), the algorithm is instructed (by a *union a b* command) to merge the two equivalence classes to which *a* and *b* belong. To query the data structure, the *find a B* command seeks the canonical representative of the class to which *a* belongs, and unifies it with *B*. In the context of our theorem prover, *find* must also produce a proof that  $a = b$ .

We have implemented a union-find prover in Twelf. Assuming that the logic-programming engine efficiently indexes atomic dynamic clauses<sup>4</sup>, it should run in  $O(N \alpha(N))$ , where  $\alpha(N)$  is the inverse Ackermann function.

In our example, we add an equality primitive  $==$  to the logic, along with some axioms. Union-find will maintain and query canonical representatives of equivalence classes:

```

==: i → i → o.  %infix none 20 == .
refl : pf(A == A).
symm : pf(A == B) → pf(B == A).
trans : pf(A == B) → pf(B == C) → pf(A == C).

```

Some of the important constructors and predicates used in this example are declared as follows.

```

⊢: hyps → pf A → goal.  %infix none 3 ⊢ .
union : pf(X == Y) → hyps.
find : {x}{y}pf(x == y) → hyps.
canonical: i → type.

```

Assume we have a function *f* and some primitive equality facts:

```

f : rational → i.
u35 : pf(f 3 == f 5).
u79 : pf(f 7 == f 9).
u75 : pf(f 7 == f 5).

find2 : pf(A == B) → pf(C == B) → pf(A == C) =
[pAB][pCB] trans pAB (symm pCB).

```

A typical query that our union-find can answer is,

```

union u35, union u79, union u75, find (f 9) X P9, find (f 3) X P3, nil
⊢ f 9 == f 3 by find2 P9 P3.

```

<sup>4</sup> Dynamic clauses will be explained in this section. Twelf does not index dynamic clauses, so a real test of our program's efficiency has not yet been performed.

In this prover, the “hypotheses” to the left of the turnstile  $\vdash$  are treated as commands to the union-find engine. Associated with each command is a proof: *union*  $P$  (where  $P$  is a proof of  $A == B$ ) is a command to union the sets to which  $A$  and  $B$  belong. *find*  $X Y P$  is a command to find the canonical representative of  $X$ , unify it with  $Y$ , and construct a proof that  $X == Y$ ; this proof is then unified with  $P$ . Thus, by the time *nil* is reached, the proof to the right of the turnstile in our example, *find2*  $P2 P3$ , must be a proof of  $f\ 9 == f\ 3$ .

How could such a query fail? In our example, the only possible point is where the command *find*  $(f\ 3)\ X\ P3$  is executed: here,  $X$  has already been instantiated to the canonical representative of  $f\ 9$ , so if that is not the same as the canonical representative of  $f\ 3$ , the *find* command will fail and backtrack. In this example, such failure does not occur.

Our program introduces dynamic clauses of the form *canon*  $X\ Z\ Pxz$  to indicate that  $Z$  is the canonical representative of  $X$ , with proof  $Pxz$ :

$$\text{canon} : \{x : i\}\{y : i\} \text{pf}(x == y) \rightarrow \text{type}.$$

That is, these clauses of the Prolog program will be created at runtime by the execution of other clauses. Standard Prolog has **assert** and **retract** to add and delete clauses to/from the fact database; both LambdaProlog and Twelf have a dynamically scoped version of this feature, in which dynamically added clauses are automatically removed when the goals containing them complete successfully, or when backtracking occurs. A Twelf clause such as

$$c : \text{expr1} \leftarrow \{d : \text{expr2}\} \text{expr3}.$$

would operate as follows: if the top-level goal matches *expr1*, then the subgoal becomes  $\{d : \text{expr2}\} \text{expr3}$ ; to satisfy this subgoal, first the clause  $d : \text{expr2}$  is added to the fact database, then the subgoal *expr3* is tried. Once *expr3* succeeds or fails, the dynamic clause  $d : \text{expr2}$  is removed.

Our program has 16 clauses and 13 constructor declarations. Instead of showing the whole program, we will show just one clause to illustrate the use of dependent types. The following clause “executes” a command *find*  $X\ Y\ P$  in the case that  $X$  maps in exactly two steps to  $Y$ ; in this case, path-compression is performed:

$$\begin{aligned} \text{find\_tac2} : \text{find } X\ Y\ P, Hs \vdash H \leftarrow \\ \text{canon } X\ Z\ Pxz \leftarrow \\ \text{canon } Z\ Y\ Pzy \leftarrow \\ \text{canonical } Y \leftarrow ! \leftarrow \\ \{d : \text{canon } X\ Y\ (\text{trans } Pxz\ Pzy)\} Hs \vdash H. \end{aligned}$$

The first line matches the *find* command; lines 2 and 3 match the case that  $X$  links to  $Y$  in two steps, with proofs  $Pxz$  and  $Pzy$  respectively; line 4 checks that  $Y$  is its own canonical representative. Then there is a Prolog “cut” (!), to prevent other interpretations of the *find* command from matching<sup>5</sup>. Then a new atomic clause is

<sup>5</sup> We are using a version of Twelf with “cut”; the standard distribution does not have this operation.

added to the global database, stating that  $Y$  is the canonical representative of  $X$  with proof  $trans\ Pxz\ Pzy$ ; finally, the remaining command-list  $Hs$  is executed. The old clause  $canon\ X\ Z\ Pxz$  is still there, but by careful use of cuts, the algorithm will never have occasion to use it.

When *find\_tac2* adds a new clause to the global database, the dependent type of the *canon* constructor ensures that it must be with a valid proof. When a proof  $P$  is returned after a set of commands  $Hs \vdash A$  by  $P$ , the dependent type of  $\vdash$  ensures that it proves the theorem that is claimed. The correctness of *find\_tac2* and similar clauses is guaranteed statically.

## 6 Related Work

Using dependent types in proofs was not possible in the corresponding Lambda Prolog version of our tactic-style theorem prover. Lambda Prolog, however, has polymorphic types, which Twelf does not, and these types provide some advantages in a Lambda Prolog implementation of tactics and tacticals. For example, in Lambda Prolog, only one version of the goal constructor for universal quantification is needed:

$$all : (A \rightarrow goal) \rightarrow goal.$$

where  $A$  is a type variable that can be instantiated with any type. Thus, the implementation of the goal constructors and tacticals does not have to change when we change object logics. In contrast, in Twelf, one *all* constructor is needed for each type that needs to be quantified. Twelf also does not allow quantification over predicates. In Lambda Prolog, tactics can be implemented as predicates taking two goals as arguments, which means that tacticals would have predicate arguments. To illustrate, if this were possible in Twelf, there would no longer be a need for the *tactic* constructor and the type  $goal \rightarrow goal \rightarrow type$  would become the definition of the type *tac*. Some of the code would look like:

$$\begin{aligned} tac &= goal \rightarrow goal \rightarrow type. \\ t1 : & \quad initial\_tac (Hs \vdash A \text{ by } P) tt \leftarrow \\ & \quad \quad \quad nth\_item\ N\ (A \text{ by } P)\ Hs. \\ tactical2 : & \quad then\ T1\ T2\ InG\ OutG \leftarrow \\ & \quad \quad T1\ InG\ MidG \leftarrow T2\ MidG\ OutG. \end{aligned}$$

Pollack (1995) discusses the use of dependent types in LCF-style provers to avoid the need for validations. As a first step, a modification of the unforgeable abstract data type *theorem* is presented. The new data type makes the structure of the theorem explicit in the ML type, resulting in a more informative type. Then, a more expressive metalanguage with dependent types is proposed. When taking this step, the notion of tactic is modified; a tactic in this setting becomes the statement of a derived or admissible rule along with its proof in the LEGO system (Pollack, 1994). Applying the tactic means applying the new rule as a lemma. Programming decision procedures for proving subgoals is also mentioned, but example programs are not given.

McBride (2001) presents an implementation of first-order unification using a dependently typed functional language derived from the LEGO system. The language is a strongly normalizing type theory, so he is able to establish termination. Bove (1999) also programs unification in a dependently typed functional language. She uses Martin-Löf’s type theory as a programming language and works within the the ALF system (Altenkirch *et al.*, 1994). She also establishes termination. In addition, she provides a methodology for extracting a Haskell program from the type theory version. It would be interesting to compare these programs to a dependently typed logic-programming implementation of the same algorithm.

## 7 Conclusion

We have shown how dependent types can guarantee partial correctness of tactics in a tactic-style theorem prover written in Twelf. We have also shown that other proof strategies such as decision procedures can benefit similarly from the use of dependent types. In both of these examples, the fact that object-level proofs were constructed and returned as a result of proof search was a crucial element of the program. By using dependent types to represent such proofs, it is not possible to write tactics or other proof procedures that construct proofs that will not check when submitted to a proof checker.

Both Coq and Twelf contain dependently typed languages intended for describing object-logic terms. The designers of these systems didn’t really intend that large-scale programs written in these “little” languages would be executed within Coq or Twelf. We have demonstrated that there’s a significant software-engineering advantage to using the little language in Twelf instead of programming in ML, which is the surrounding implementation’s language. The same demonstration could probably have been done using Coq’s object language, a dependently typed functional language (as contrasted with Twelf’s dependently typed Prolog-like language).

Although the tactical prover discussed in this paper is just a prototype, we are confident that these techniques will scale to full-size provers and decision procedures. We have used similar techniques in other dependently typed proof-manipulation programs in Twelf, and the dependent types assist, not impede, program development.

## References

- Aho, Alfred V., Hopcroft, John E., & Ullman, Jeffrey D. (1974). *The design and analysis of computer algorithms*. Reading, MA: Addison-Wesley.
- Altenkirch, Thorsten, Gaspes, Veronica, Nordström, Bengt, & von Sydow, Björn. (1994). *A user’s guide to ALF*. Tech. rept. Chalmers University of Technology, Sweden.
- Appel, Andrew W. (2000). *Hints on proving theorems in Twelf*. [www.cs.princeton.edu/~appel/twelf-tutorial](http://www.cs.princeton.edu/~appel/twelf-tutorial).
- Appel, Andrew W., & Felty, Amy P. (2000). A semantic model of types and machine instructions for proof-carrying code. *Pages 243–253 of: POPL ’00: The 27th ACM SIGPLAN-SIGACT symposium on principles of programming languages*. New York: ACM Press.



- Barras, Bruno, *et al.* (1998). *The Coq Proof Assistant reference manual*. Tech. rept. INRIA.
- Bove, Ana. (1999). *Programming in Martin-Löf type theory: Unification, a non-trivial example*. Licentiate Thesis, Chalmers University of Technology and Göteborg University.
- Felty, Amy. (1993). Implementing tactics and tacticals in a higher-order logic programming language. *Journal of automated reasoning*, **11**(1), 43–81.
- Gordon, M. J. C., & Melham, T. F. (1993). *Introduction to HOL—a theorem proving environment for higher order logic*. Cambridge University Press.
- Gordon, Michael J., Milner, Robin, & Wadsworth, Christopher P. (1979). *Edinburgh LCF: A mechanised logic of computation*. Lecture Notes in Computer Science, vol. 78. New York: Springer-Verlag.
- Harper, Robert, Honsell, Furio, & Plotkin, Gordon. (1993). A framework for defining logics. *Journal of the ACM*, **40**(1), 143–184.
- McBride, Conor. (2001). First-order unification by structural recursion. *Journal of functional programming*. To appear.
- Nadathur, Gopalan, & Miller, Dale. (1988). An overview of lambda prolog. Bowen, K., & Kowalski, R. (eds), *Fifth international conference and symposium on logic programming*. MIT Press.
- Paulson, Lawrence C. (1994). *Isabelle: A generic theorem prover*. Lecture Notes in Computer Science, vol. 828. Springer-Verlag.
- Pfenning, Frank, & Schürmann, Carsten. (1999). System description: Twelf — a meta-logical framework for deductive systems. *The 16th international conference on automated deduction*. Berlin: Springer-Verlag.
- Pollack, Robert. (1994). *The theory of LEGO: A proof checker for the extended calculus of constructions*. Ph.D. thesis, University of Edinburgh.
- Pollack, Robert. (1995). On extensibility of proof checkers. *Pages 140–161 of: Dybjer, P., Nordstrom, B., & Smith, J. (eds), Types for proofs and programs: International workshop types'94, bästad, june 1994, selected papers*. LNCS, vol. 996. Springer Verlag Lecture Notes in Computer Science.