

Twelf User's Guide

Version 1.2

Frank Pfenning and Carsten Schuermann

1 Introduction

Twelf is current version of a succession of implementations of the logical framework LF. Previous systems include Elf (which provided type reconstruction and the operational semantics reimplemented in Twelf) and MLF (which implemented module-level constructs loosely based on the signatures and functors of ML still missing from Twelf).

Twelf should be understood as research software. This means comments, suggestions, and bug reports are extremely welcome, but there are no guarantees regarding response times. The same remark applies to these notes which constitute the only documentation on the present Twelf implementation.

For current information including download instructions, publications, and mailing list, see the Twelf home page at <http://www.cs.cmu.edu/~twelf/>.

Below we state the typographic conventions in this manual.

<code>code</code>	for Twelf or ML code
<code>'samp'</code>	for characters and small code fragments
<i>metavar</i>	for placeholders in code
keyboard	for input in verbatim examples
KEY	for keystrokes
<i>math</i>	for mathematical expressions
<i>emph</i>	for emphasized phrases

File names for examples given in this guide are relative to the main directory of the Twelf installation. For example `'examples/guide/nd.elf'` may be found in `'/usr/local/twelf/examples/guide/nd.elf'` if Twelf was installed into the `'/usr/local/'` directory.

1.1 New Features

While the underlying type theory has not changed, the Twelf implementation differs from older Elf implementation in a few ways. Mostly, these are simplifications and improvements. The main

feature which has not yet been ported is the Elf server interface to Emacs. Also, while the type checker is more efficient now, the operational semantics does not yet incorporate some of the optimizations of the older Elf implementations and is therefore slower.

Syntax (see Chapter 3 [Syntax], page 7)

The quote ‘`’` character is no longer a special character in the lexer, and ‘`=`’ (equality) is now a reserved identifier. The syntax of `%name` declarations has changed by allowing only one preferred name to be specified. Also, `%name`, `%infix`, `%prefix` and `%postfix` declarations must be terminated by a period ‘`.`’ which previously was optional. Further, single line comments now must start with ‘`%whitespace`’ or ‘`%%`’ in order to avoid misspelled keywords of the form ‘`%keyword`’ to be ignored.

Type theory

Elf 1.5 had two experimental features which are not available in Twelf: polymorphism and the classification of *type* as a type.

Definitions (see Section 3.3 [Definitions], page 9)

Twelf offers definitions which were not available in Elf.

Searching for definitions (see Section 5.2 [Solve Declaration], page 20)

Elf had a special top-level query form `sigma [x:A] B` which searched for a solution `M : A` and then solved the result of substituting `M` for `x` in `B`. In Twelf this mechanism has been replaced by a declaration `%solve c : A` which searches for a solution `M : A` and then defines `c = M : A`, where the remaining free variables are implicitly universally quantified.

Query declarations (see Section 5.1 [Query Declaration], page 19)

Twelf allows queries in ordinary Elf files as ‘`%query`’ declarations. Queries are specified with an expected number of solutions, and the number of solutions to search for, which can be used to test implementations.

Operational semantics (see Section 5.5 [Operational Semantics], page 23)

Twelf eliminates the distinction between static and dynamic signatures. Instead, dependent function types `{x:A} B` where `x` occurs in the normal form of `B` are treated statically, while non-dependent function type `A -> B` or `B <- A` or `{x:A} B` where `x` does not occur in `B` are treated dynamically.

Modes (see Chapter 6 [Modes], page 27)

Twelf offers a mode checker which was only partially supported in Elf.

Termination (see Chapter 7 [Termination], page 31)

Twelf offers a termination checker which can verify that certain programs represent decision procedures.

Theorem prover (see Chapter 8 [Theorem Prover], page 39)

Although very limited at present, an experimental prover for theorems and meta-theorems (that is, properties of signatures) is now available. It does not yet support lemmas or meta-hypothetical reasoning, which are currently under development.

Emacs interface (see Chapter 11 [Emacs Interface], page 55)

The Elf mode has remained basically unchanged, but the Elf server interface has not yet been ported.

1.2 Quick Start

Assuming you are running on a Unix system with SML of New Jersey already installed (see Section 9.6 [Installation], page 50) you can build Twelf as follows. Here ‘%’ is assumed to be the shell prompt. You may need to edit the file ‘Makefile’ to give the proper location for `sml-cm`.

```
% gunzip twelf-1-2.tar.gz
% tar -xf twelf-1-2.tar
% cd twelf
% make
% bin/twelf-server
Twelf 1.2, Aug 27 1998
%% OK %%
```

You can now load the examples used in this guide and pose an example query as shown below. The prompt from the Twelf top-level is ‘?-’. To drop from the Twelf top-level to the ML top-level, type `C-c` (CTRL c). To exit the Twelf server you may issue the `quit` command or type `C-d` (CTRL c).

```
Config.read examples/guide/sources.cfg
Config.load
top
?- of (lam [x] x) T.
Solving...
T = arrow T1 T1.
More? y
No more solutions
?- C-c
interrupt
%% OK %%
quit
%
```


2 Lexical Conventions

Lexical analysis of Twelf has purposely been kept simple, with few reserved characters and identifiers. As a result one may need to use more whitespace to separate identifiers than in other languages. For example, ‘ $A \rightarrow B$ ’ or ‘ $A+B$ ’ are single identifiers, while ‘ $A \rightarrow B$ ’ and ‘ $A + B$ ’ both consist of 3 identifiers.

During parsing, identifiers are resolved as reserved identifiers, constants, bound variables, or free variables, following the usual rules of static scoping in lambda-calculi.

2.1 Reserved Characters

The following table lists the reserved characters in Twelf.

‘:’	colon, constant declaration or ascription
‘.’	period, terminates declarations
‘(’ ‘)’	parentheses, for grouping terms
‘[’ ‘]’	brackets, for lambda abstraction
‘{’ ‘}’	braces, for quantification (dependent function types)
<i>whitespace</i>	separates identifiers; one of space, newline, tab, carriage return, vertical tab or formfeed
‘%’	introduces comments or special keyword declarations
‘% <i>whitespace</i> ’ ‘%%’	comment terminated by the end of the line, may contain any characters
‘%{’ ‘}%’	delimited comment, nested ‘%{’ and ‘}%’ must match
‘% <i>keyword</i> ’	various declarations
‘%.’	end of input stream
‘”’	doublequote, disallowed
other printing characters	identifier constituents

2.2 Identifiers

All printing characters that are not reserved can be included in identifiers, which are separated by whitespace or reserved characters. In particular, $A \rightarrow B$ is an identifier, whereas $A \rightarrow B$ stands for the type of functions from A to B .

An uppercase identifier is one which begins with an underscore ‘_’ or a letter in the range ‘A’ through ‘Z’. A lowercase identifier begins with any other character except a reserved one. Numbers also count as lowercase identifiers and are not interpreted specially. Free variables in a declaration must be uppercase, bound variables and constants may be either uppercase or lowercase identifiers.

There are also four reserved identifiers with a predefined meaning which cannot be changed. Keep in mind that these can be constituents of other identifiers which are not interpreted specially.

‘ \rightarrow ’	function type
‘ \leftarrow ’	reverse function type
‘_’	hole, to be filled by term reconstruction
‘=’	definition
‘type’	the kind <i>type</i>

Constants have static scope, which means that they can be shadowed by subsequent declarations. A shadowed identifier (which can no longer be referred to in input) is printed as *%id%*. The printer for terms renames bound variables so they do not shadow constants.

Free uppercase identifiers in declarations represent schematic variables. In order to distinguish them from other kinds of variables and constants they are printed as ‘*id*’ (backquote, followed by the identifier name) in error messages.

3 Syntax

In LF, deductive systems are represented by signatures consisting of constant declarations. Twelf implements declarations in a straightforward way and generalizes signatures by also allowing definitions, which are semantically transparent. Twelf currently does not have module-level constructs so that, for example, signatures cannot be named. Instead, multiple signatures can be manipulated in the programming environment using configurations (see Section 9.1 [Configurations], page 45).

The LF type theory which underlies LF is stratified into three levels: objects M and N , types A and B , and kinds K . Twelf does not syntactically distinguish these levels and simply uses one syntactic category of term. Similarly, object-level constants c and type-level constants a as well as variables share one name space of identifiers.

In explanations and examples we will use letters following the mathematical conventions above to clarify the roles of various terms. We also use U and V to stand for arbitrary terms.

3.1 Grammar

The grammar below defines the non-terminals **sig**, **decl**, **term** and uses the terminal **id** which stands for identifiers (see Section 2.2 [Identifiers], page 6). The comments show the meaning in LF. There are various special declarations **%keyword** such as **%infix** or **%theorem** which are omitted here and detailed in the appropriate sections.

```

sig ::=                                % Empty signature
    | decl sig                         % Constant declaration

decl ::= id : term.                   %  $a : K$  or  $c : A$ 
    | id : term = term.               %  $d : A = M$ 
    | id = term.                     %  $d = M$ 
    | _ : term = term.               % anonymous definition, for type-checking
    | _ = term.                     % anonymous definition, for type-checking
    | %infix ixdecl.                 % operator declaration
    | %prefix pxdecl.               % operator declaration
    | %postfix pxdecl.              % operator declaration
    | %name id id.                   % name preference declaration
    | %query qdecl.                 % query declaration
    | %solve id : term.              % solve declaration
    | %mode mdecl.                  % mode declaration
    | %terminates tdecl.            % termination declaration
    | %theorem thdecl.              % theorem declaration
    | %prove pdecl.                 % prove declaration

```

```

term ::= type                % type
      | id                  % variable  $x$  or constant  $a$  or  $c$ 
      | term -> term        %  $A \rightarrow B$ 
      | term <- term        %  $A \leftarrow B$ , same as  $B \rightarrow A$ 
      | {id : term} term    %  $\Pi x : A. K$  or  $\Pi x : A. B$ 
      | [id : term] term    %  $\lambda x : A. B$  or  $\lambda x : A. M$ 
      | term term           %  $A M$  or  $M N$ 
      | term : term        % explicit type ascription
      | _                  % hole, to be filled by term reconstruction
      | {id} term          % same as {id:_} term
      | [id] term          % same as [id:_] term

```

The constructs $\{x:U\} V$ and $[x:U] V$ bind the identifier x in V , which may shadow other constants or bound variables. As usual in type theory, $U \rightarrow V$ is treated as an abbreviation for $\{x:U\} V$ where x does not appear in V . However, there is a subtlety in that the latter allows an implicit argument (see Section 4.2 [Implicit Arguments], page 14) to depend on x while the former does not.

In the order of precedence, we disambiguate the syntax as follows:

1. Juxtaposition (application) is left associative and has highest precedence.
2. User declared infix, prefix, or postfix operators (see below).
3. ‘ \rightarrow ’ is right and ‘ \leftarrow ’ left associative with equal precedence.
4. ‘ $:$ ’ is left associative.
5. ‘ $\{ \}$ ’ and ‘ $[\]$ ’ are weak prefix operators.

For example, the following are parsed identically:

```

d : a <- b <- {x} c x -> p x.
d : ({x} c x -> p x) -> b -> a.
d : ((a <- b) <- ({x:_} ((c x) -> (p x)))).

```

3.2 Constructor Declaration

New type families or object constructors can be introduced with

```

condec ::= id : term.      %  $a : K$  or  $c : A$ 

```

Here a stands for a type family and K for its kind, whereas c is an objects constructor and A its type. Identifiers are resolved as follows:

1. Any identifier x may be bound by the innermost enclosing binder for x of the form $\{x:A\}$ or $[x:A]$.
2. Any identifier which is not explicitly bound may be a declared or defined constant.
3. Any uppercase identifier, that is, identifier starting with ‘_’ (underscore) or an upper case letter, may be a free variable. Free variables are interpreted universally and their type is inferred from their occurrences (see Chapter 4 [Term Reconstruction], page 13).
4. Any other undeclared identifier is flagged as an error.

3.3 Definitions

Twelf supports notational definitions, currently employing a restriction to allow a simple and efficient internal treatment. Semantically, definitions are completely transparent, that is, both for type checking and the operational semantics definitions may be expanded.

```
defn ::= id : term = term.    % d : A = M
      | id = term.           % d = M
```

where the second is equivalent to `id : _ = term`. Definitions can only be made on the level of objects, not at the level of type families because the interaction of such definitions with logic programming search has not been fully investigated.

In order to avoid always expanding definitions, Twelf currently only permits strict definitions (see Section 4.4 [Strict Definitions], page 16). A definition of a constant c is strict if all arguments to c (implicit or explicit) have at least one strict occurrence (see Section 4.3 [Strict Occurrences], page 15) in the right-hand side of the definition, and the right-hand side contains at least one constant. In practice, most notational definitions are strict. For some examples, see Section 3.6 [Sample Signature], page 11 and Section 4.4 [Strict Definitions], page 16.

The power of definitions in Twelf, however, is severely limited by the lack of recursion. It should only be thought of as notational definition, not as a computational mechanism. Complex operations need to be defined as logic programs, taking advantage of the operational semantics assigned to signatures (see Chapter 5 [Logic Programming], page 19).

3.4 Operator Declaration

The user may declare constants to be infix, prefix, or postfix operators. Operator precedence properties are associated with constants, which must therefore already have been declared with a

type or kind and a possible definition. It is illegal to shadow an infix, prefix, or postfix operator with a bound variable. We use `nat` for the terminal natural numbers.

```

assoc ::= none    % not associative
       | left    % left associative
       | right   % right associative

prec ::= nat      % 0 <= prec < 10000
ixdecl ::= assoc prec id

pxdecl ::= prec id

decl ::= ...
       | %infix ixdecl.
       | %prefix pxdecl.
       | %postfix pxdecl.

```

During parsing, ambiguous successive operators of identical precedence such as `a <- b -> c` are flagged as errors. Note that it is not possible to declare an operator with equal or higher precedence than juxtaposition or equal or lower precedence than `'->'` and `'<-'`.

3.5 Name Preferences

During printing, Twelf frequently has to assign names to anonymous variables. In order to improve readability, the user can declare a name preference for anonymous variables based on their type. Thus name preferences are declared for type family constants. Note that name preferences are not used to disambiguate the types of identifiers during parsing.

```

decl ::= ...
       | %name id id.

```

Following our same conventions, a name preference declaration has the form `%name a id`, that is, the first identifier must be a type family already declared and the second is the name preference for variables of type `a`. The second identifier must be uppercase, that is, start with a letter from 'A' through 'Z' or an underscore '_'. Anonymous variables will then be named `id1`, `id2`, etc.

3.6 Sample Signature

Below is a signature for intuitionistic first-order logic over an unspecified domain of individuals and atomic propositions. It illustrates constant declarations and definitions and the use of operator precedence and name preference declarations. It may be found in the file ‘examples/guide/nd.elf’.

```

%%% Individuals
i : type.                                %name i T

%%% Propositions
o : type.                                %name o A

imp   : o -> o -> o.                    %infix right 10 imp
and   : o -> o -> o.                    %infix right 11 and
true  : o.
or    : o -> o -> o.                    %infix right 11 or
false : o.
forall : (i -> o) -> o.
exists : (i -> o) -> o.

not : o -> o = [A:o] A imp false.

%%% Natural Deductions
nd : o -> type.                          %name nd D

impi   : (nd A -> nd B) -> nd (A imp B).
impe   : nd (A imp B) -> nd A -> nd B.
andi   : nd A -> nd B -> nd (A and B).
ande1  : nd (A and B) -> nd A.
ande2  : nd (A and B) -> nd B.
truei  : nd (true).
% no truee
ori1   : nd A -> nd (A or B).
ori2   : nd B -> nd (A or B).
ore    : nd (A or B) -> (nd A -> nd C) -> (nd B -> nd C) -> nd C.
% no falsei
falsee : nd false -> nd C.
foralli : ({x:i} nd (A x)) -> nd (forall A).
foralle : nd (forall A) -> {T:i} nd (A T).
existsi : {T:i} nd (A T) -> nd (exists A).
existse : nd (exists A) -> ({x:i} nd (A x) -> nd C) -> nd C.

noti : (nd A -> nd false) -> nd (not A)
      = [D:nd A -> nd false] impi D.
note : nd (not A) -> nd A -> nd false
      = [D:nd (not A)] [E:nd A] impe D E.

```


4 Term Reconstruction

Representations of deductions in LF typically contain a lot of redundant information. In order to make LF practical, Twelf gives the user the opportunity to omit redundant information in declarations and reconstructs it from context. Unlike for functional languages, this requires recovering objects as well as types, so we refer to this phase as term reconstruction.

There are criteria which guarantee that the term reconstruction problem is decidable, but unfortunately these criteria are either very complicated or still force much redundant information to be supplied. Therefore, the Twelf implementation employs a reconstruction algorithm which always terminates and gives one of three answers:

1. yes, and here is the most general reconstruction;
2. no, and here is the problem; or
3. maybe.

The last characterizes the situations where there is insufficient information to guarantee a most general solution to the term reconstruction problem. Because of the decidable nature of type-checking in LF, the user can always annotate the term further until it falls into one of the definitive categories.

4.1 Implicit Quantifiers

The model of term reconstruction employed by Twelf is straightforward, although it employs a relatively complex algorithm. The basic principle is a duality between quantifiers omitted in a constant declaration and implicit arguments where the constant is used. Recall some definitions in the signature defining natural deductions (see Section 3.6 [Sample Signature], page 11).

```
o : type.
and : o -> o -> o.    %infix right 10 and
nd : o -> type.
andi : nd A -> nd B -> nd (A and B).
```

The last declaration contains *A* and *B* as free variables. Type reconstruction infers most general types for the free variables in a constant declaration and adds implicit quantifiers. In the example above, *A* and *B* must both be of type *o*. The internal form of the constant thus has one of the following two forms.

```

andi : {A:o} {B:o} nd A -> nd B -> nd (A and B).
andi : {B:o} {A:o} nd A -> nd B -> nd (A and B).

```

These forms are printed during type reconstruction, so the user can examine if the result of reconstruction matches his expectations.

4.2 Implicit Arguments

The quantifiers on *A* and *B* in the declaration

```

andi : nd A -> nd B -> nd (A and B).

```

were implicit. The corresponding arguments to `andi` are also implicit. In fact, since the order of the reconstructed quantifiers is arbitrary, we cannot know in which order to supply the arguments, so they must always be omitted. Thus a constant with *n* implicit quantifiers is supplied with *n* implicit arguments wherever it is seen. These implicit arguments are existential variables whose value may be determined from context by unification.

For example, using also

```

true : o.
truei: nd (true).

```

we have

```

(andi truei truei) : nd (true and true).

```

During parsing, the expression `(andi truei truei)` is interpreted as

```

(andi _ _ truei truei)

```

where the two underscores stand for the implicit *A* and *B* arguments to `andi`. They are replaced by existential variables whose value will be determined during type reconstruction. We call them *A1* and *A2* and reason as follows.

```

|- andi : {A:o} {B:o} nd A -> nd B -> nd (A and B)
|- andi A1 : {B:o} nd A1 -> nd B -> nd (A1 and B)
|- andi A1 A2 : nd A1 -> nd A2 -> nd (A1 and A2)

```


At this point, we need a to infer the type of the application `(andi A1 A2) truei`. This equates the actual type of the argument with the expected type of the argument.

```
|- andi A1 A2 : nd A1 -> nd A2 -> nd (A1 and A2)
|- truei : nd true
-----
|- andi A1 A2 truei : nd A2 -> nd (A1 and A2)
   where nd true = nd A1
```

The equation can be solved by instantiating `A1` to `true` and we continue:

```
|- andi true A2 truei : nd A2 -> nd (true and A2)
|- truei : nd true
-----
|- andi true A2 truei truei : nd (true and A2)
   where nd true = nd A2
|- andi true true truei truei : nd (true and true)
```

The last line is the expected result. In this way, term reconstruction can always be reduced to solving equations such that every solution to the set of equations leads to a valid typing and vice versa.

4.3 Strict Occurrences

Both for type reconstruction and the operational semantics, Twelf must solve equations between objects and types. Unfortunately, it is undecidable if a set of equations in the LF type theory has a solution. Worse yet, even if it has solutions, it may not have a most general solution. Therefore, Twelf postpones difficult equations as constraints and solves only those within the pattern fragment (see *Miller 1991, Journal of Logic and Computation* and *Pfenning 1991, Logical Frameworks*). In this fragment, principal solutions always exist and can be computed efficiently. If constraints remain after term reconstruction, the constant declaration is rejected as ambiguous which indicates that the user must supply more type information. We illustrate this phenomenon and a typical solution in our natural deduction example.

A central concept useful for understanding the finer details of type reconstruction is the notion of a *strict occurrence* of a free variable. We call a position in a term *rigid* if it is not in the argument of a free variable. We then call an occurrence of a free variable *strict* if the occurrence is in a rigid position and all its arguments (possibly none) are distinct bound variables.

If all free variable occurrences in all declarations in a signature are strict, then term reconstruction will always either fail or succeed with a principal solution, provided no further terms are omitted (that is, replaced by an underscore).

If a free variable in a declaration of a constant `c` has no strict occurrence at all, then its type can almost never be inferred and most uses of `c` will lead to a constraint.

If a free variable has strict and non-strict occurrences then in most cases term reconstruction will provide a definitive answer, but there is no guarantee. Mostly this is because most general answers simply do not exist, but sometimes because the algorithm generates, but cannot solve constraints with unique solutions.

We use some advanced examples from the natural deduction signature to illustrate these concepts and ideas. In the declarations

```
forallli : ({x:i} nd (A x)) -> nd (forall A).
forallle : nd (forall A) -> {T:i} nd (A T).
```

all free variables have a strict occurrence. However, if we had decided to leave `T` as an implicit argument,

```
forallle : nd (forall A) -> nd (A T).
```

then `T` has no strict occurrence. While this declaration is accepted as unambiguous (with `A:i -> o` and `T:i`), any future use of `forallle` most likely leads to constraints on `T` which cannot be solved.

4.4 Strict Definitions

Definitions are currently restricted so that each argument to the defined constant, may it be implicit or explicit, must have at least one strict occurrence in the right-hand side. For example, the definition of `not` in the signature for natural deduction (see Section 3.6 [Sample Signature], page 11)

```
not : o -> o = [A:o] A imp false.
```

is strict since the only argument `A` has a strict occurrence in `A imp false`. On the other hand, the definition

```

noti : ({p:o} nd A -> nd p) -> nd (not A)
      = [D] impi ([u:nd A] D false u).

```

which gives a possible derived introduction rule for negation is not strict: the argument `D` has only one occurrence, and this occurrence is not strict since the argument `false` is not a variable bound in the body, but a constant.

However, the definitions

```

noti : (nd A -> nd false) -> nd (not A)
      = [D:nd A -> nd false] impi D.
note : nd (not A) -> nd A -> nd false
      = [D:nd (not A)] [E:nd A] impe D E.

```

are both strict since arguments `D` and `E` both have strict occurrences. Type-checking these definitions requires that the definition of `not A` is expanded to `A imp false`.

Note that free variables in the type and the right-hand side of a definition are shared. In the above example, `A` occurs both in the types and the right hand side and it should be thought of as the same `A`. With the implicit quantifiers and abstractions restored, the definitions above have the following form.

```

noti : {A:o} (nd A -> nd false) -> nd (not A)
      = [A:o] [D:nd A -> nd false] impi D.
note : {A:o} nd (not A) -> nd A -> nd false
      = [A:o] [D:nd (not A)] [E:nd A] impe D E.

```

4.5 Type Ascription

In some circumstances it is useful to directly ascribe a type in order to disambiguate declarations. For example, the term `or1 true1` has principal type `nd (true or B)` for a free variable `B`. If we want to constrain this to a derivation of `nd (true or false)` we can write `or1 true1 : nd (true or false)`.

Explicit type ascription sometimes helps when the source of a type error is particularly hard to discern: we can ascribe an expected type to a subterm, thus verifying our intuition about constituent terms in a declaration.

4.6 Error Messages

When term reconstruction fails, Twelf issues an error message with the location of the declaration in which the problem occurred and the disagreement encountered. A typical message is

```
examples/nd/nd.elf:37.35-37.41 Error: Type mismatch
Expected: o
Found:    (i -> o) -> o
Expression clash
```

which points to an error in the file ‘`examples/nd/nd.elf`’, line 37, characters 35 through 41 where an argument to a function was expected to have type `o`, but was found to have type `(i -> o) -> o`.

If constraints remain, the error location is the whole declaration with the message

```
filename:location Error: Typing ambiguous -- unresolved constraints
```

The *filename* and *location* information can be used by Emacs (see Chapter 11 [Emacs Interface], page 55) to jump to the specified location in the given file for editing of the incorrect declaration for the constant `c`. The *location* has the form *line1.column1-line2.column2* and represent Twelf's best guess as to the source of the error. Due to the propagation of non-trivial constraints the source of a type reconstruction failure can sometimes not be pinpointed very precisely.

5 Logic Programming

Twelf gives an operational interpretation to signatures under the computation-as-proof-search paradigm. The fundamental idea is to fix a simple search strategy and then search for a derivation of a query according to this strategy. The result may be a substitution for the free variables in a query and a derivation, or explicit failure. It is also possible that the computation does not terminate.

A query can be posed in three different ways: as a `%query` declaration, as a `%solve` declaration, or interactively, using a top-level invoked from ML with `Twelf.top` which prompts with ‘?’ (see Section 5.3 [Interactive Queries], page 20).

```

query ::= id : term    % X : A, X a free variable
      | term          % A

bound ::= nat          % number of solutions
      | *             % unbounded number

qdecl ::= bound bound query % expected solutions, try limit, query

decl ::= ...
      | %query qdecl.      % execute query
      | %solve id : term.  % solve and name proof term

```

In all of these cases, the free variables in a query are interpreted existentially, which is in contrast to constant declarations where free variables are interpreted universally. In particular, free variables might be instantiated during type reconstruction and during execution of the query.

5.1 Query Declaration

The query form

```
%query expected try A.
```

will try to solve the query `A` and verify that it gives the *expected* number of solutions, but it will never try to find more than indicated by *try*. It succeeds and prints a message, whose precise form depends on the value of `Twelf.chatter` if `A` has the expected number of solutions; otherwise it either fails with an error message or does not terminate. ‘`%query`’ has no other effect on the state of Twelf. Here are some examples.

```

%query 1 * A.      % check that A has exactly one solution
%query 1 1 A.      % check that A has at least one solution
%query * 3 A.      % A has infinitely many solutions, check 3
%query * * A.      % fails if A has finitely many solutions
%query 1 0 A.      % skip this query

```

5.2 Solve Declaration

The query form

```
%solve c : A.
```

will search for the first solution M of A and then define

```
c : A = M.
```

If there are any free variables remaining in M or A after search, they will be implicitly quantified in the new definition. This form of definition is particularly useful to compute and name inputs to future queries. An example of this feature from the file ‘`examples/nd/lam.elf`’ can be found in Section 8.5 [Proof Realizations], page 42.

5.3 Interactive Queries

An interactive top-level can be invoked using the SML expression `Twelf.top ()`; . The top-level prompts with ‘?- ’ and awaits the input of a query, terminated by a period ‘.’ and a `RET`.

After the query has been parsed, Twelf reconstructs implicit type information, issuing a warning if constraints remain. The result is executed as a query. At any point during the processing of a query the user may interrupt with `C-c` (that is, `CTRL` and `c`) to drop back into ML’s interactive top-level.

When Twelf has found a solution, it prints the *answer substitution* for all free variables in the query, including the proof term variable if one was given. It also notes if there are remaining equational constraints, but currently does not print them.

The top-level then waits for input, which is interpreted as follows

y, Y, or ; backtrack and search for another solution

q or Q quit Twelf's top-level and return to ML

n, N, or anything else
 return to prompt for another query

5.4 Sample Trace

As an example we consider lists of propositions and some simple operations on them, as they might be used when programming a theorem prover.

```
list : type.
nil  : list.
cons : o -> list -> list.
```

First, we want to write a program to append two lists to obtain their concatenation. This is expressed as a relation between the three lists, which in turn is implemented as a type family

```
append : list -> list -> list -> type.

appNil  : append nil K K.
appCons : append (cons X L) K (cons X M)
           <- append L K M.
```

Here, we use the synonym $A \leftarrow B$ for $B \rightarrow A$ to improve readability. We say A if B .

The first sample query concatenates the singleton lists containing `true` and `false`. We proceed as if we had loaded the appropriate files and started a top-level with `Twelf.top ();`.

```
?- append (cons true nil) (cons false nil) L.
```

Here, L is a free existential variable. We search for an object M : `append (cons true nil) (cons false nil) L`, even though this object will not be shown in this form or query. Each constant declaration in the signature is tried in turn, unifying the head with the goal above. In this manner, we obtain the following sequence of goals and actions. Note that the intermediate forms and comments are not printed when this is run. They are added here to illustrate the behavior.

```

% original goal after parsing and type reconstruction
?- append (cons true nil) (cons false nil) L.
[try appNil:
  append nil K1 K1
  = append (cons true nil) (cons false nil) L
  unification fails with constant clash: nil <> cons
]
[try appCons:
  append (cons X1 L1) K2 (cons X1 M1)
  = append (cons true nil) (cons false nil) L
  unification succeeds with
  X1 = true, L1 = nil, K2 = cons false nil, L = cons true M1
]
% subgoal
?- append nil (cons false nil) M1.
[try appNil:
  append nil K3 K3
  = append nil (cons false nil) M1
  unification and subgoal succeeds with
  K3 = cons false nil, M1 = cons false nil
]

```

At this point the overall goal succeeds and we read off the answer substitution for the only free variable in the query, namely `L`. It was first determined to be `cons true M1` and then `M1` was instantiated to `cons false nil`, leading to the instantiation

```
L = cons true (cons false nil).
```

If instead we pose the query

```
?- X : append (cons true nil) (cons false nil) L.
```

we also obtain the proof term

```
L = cons true (cons false nil);
X = appCons appNil.
```

As another example we consider a query with several solutions which are enumerated when we ask for further results. This time we do not trace the steps of the execution, but show the interaction verbatim.


```

?- append L K (cons true (cons false nil)).
Solving...
K = cons true (cons false nil);
L = nil.
More? y
K = cons false nil;
L = cons true nil.
More? y
K = nil;
L = cons true (cons false nil).
More? y
No more solutions

```

5.5 Operational Semantics

The operational semantics of Twelf is a form of typed constraint logic programming. We will use standard terminology from this area. A type family which is used in a program or goal is called a *predicate*. A constant declaration in a signature which is available during search is called a *clause*. A clause typically has the form $c : a \ M1 \ \dots \ Mm \leftarrow A1 \leftarrow \dots \leftarrow An$, where $a \ M1 \ \dots \ Mm$ is the *head of the clause* and $A1$ through An are the *subgoals*. A clause is used to reduce a goal to subgoals by a process called *backchaining*. Backchaining unifies the head of the clause with the current goal to generate *subgoals*. Next, we *select* one of the subgoals as a current goal and continue the search process. Actually, instead of unification (which is undecidable in LF), Twelf employs *constraint simplification* and carries along equational constraints in a normal form.

A hypothesis which is introduced during search is a *local assumption*; a parameter is a *local parameter*. Parameters act like constants in unification, except that their occurrences might be restricted due to *parameter dependency*.

Without going into a formal description, here are the central ideas of the operational semantics.

Clause selection.

The clauses are tried in the following order: from most recent to least recent local assumption, then from first to last clause in the global signature.

Subgoal selection.

Subgoals are solved from the inside out. For example, when a clause $c : A \leftarrow B \leftarrow C$. is applied to solve the goal $?- A$. then the first subgoal is B and the second subgoal C . Truly dependent variables will only be subject to unification and never give rise to a subgoal. For example $c : \{X:b\} \ a \ X \leftarrow a \ c$ is a clause with head $a \ X$, subgoal $a \ c$, and existential variable X .

Unification.

An atomic goal is unified with the clause head using higher-order pattern unification. All equations outside this fragment are postponed and carried along as constraints.

Local assumptions.

A goal of the form $?- A \rightarrow B$. introduces a local assumption A and then solves B under this assumption. To solve atomic goals, local assumptions are tried before global clauses, using the most recently made assumption first. Note that this is different from Prolog `assert` in that A is available only for solving B .

Local parameters.

Parameters are introduced into proof search by goals of the form $?- \{x:A\} B$. which generates a new parameter a and then solves the result of substituting a for x in B . Parameters are also called *universal variables* since they are not subject to instantiation during unification. Local parameters will never be used as local assumptions during search.

5.6 Sample Program

As an example, we consider simple type inference for the pure lambda-calculus. An extension of this example to Mini-ML is given in the course notes *Pfenning 1992, Computation and Deduction*. The code below can be found in the file `'examples/guide/lam.elf'`.

```

% Simple types
tp : type.                                %name tp T.

arrow : tp -> tp -> tp.                    % T1 => T2

% Expressions
exp : type.                                %name exp E.

lam   : (exp -> exp) -> exp.                % lam x. E
app   : exp -> exp -> exp.                  % (E1 E2)

% Type inference
% |- E : T (expression E has type T)

of : exp -> tp -> type.                    %name of P.

tp_lam : of (lam E) (arrow T1 T2)           % |- lam x. E : T1 => T2
        <- ({x:exp} of x T1 -> of (E x) T2). % if x:T1 |- E : T2.

tp_app : of (app E1 E2) T1                  % |- E1 E2 : T1
        <- of E1 (arrow T2 T1)              % if |- E1 : T2 => T1
        <- of E2 T2.                        % and |- E2 : T2.

```

Again, we have used the notation $A \leftarrow B$ to emphasize the interpretation of constant declarations as clauses. We now trace the query which infers the most general type of the identity function, represented as $\text{lam } [x:\text{exp}] \ x$. We indicate the scope of hypotheses which are introduced during search by indentation.

```

% original query, T free
?- of (lam [x:exp] x) T.
% use tp_lam with E = ([x:exp] x) and T = arrow T1 T2
% subgoal
?- {x:exp} of x T1 -> of x T2.
% introduce parameter e
?- of e T1 -> of e T2.
% introduce hypothesis labelled p
p:of e T1
  ?- of e T2.
% succeed by hypothesis p with T1 = T2

```

At this point the query succeeds and prints the answer substitution.

```

T = arrow T1 T1.
More? y
No more solutions

```

We requested more solution by typing `y`, but there are no further possibilities. The free variable `T1` in the answer substitution means that every instance of `arrow T1 T1` provides a solution to the original query. In other words, `lam [x:exp] x` has type `arrow T1 T1` for all types `T1`.

As a second example we verify that self-application is not well-typed in the simply-typed lambda-calculus.

```
?- of (lam [x:exp] app x x) T.
% use tp_lam with T = arrow T1 T2
% subgoal
?- {x:exp} of x T1 -> of (app x x) T2.
% introduce parameter e
?- of e T1 -> of (app e e) T2.
% introduce hypothesis p:of a T1
p:of e T1
  ?- of (app e e) T2.
  % use tp_app
  % first subgoal
  ?- of e (arrow T3 T2).
  % succeed by hypothesis p with T1 = arrow T3 T2
  % second subgoal
  ?- of e T3.
  % fail, since T3 = arrow T3 T2 has no solution
```

At the point where the second subgoals fails we backtrack. However, all other alternatives fail immediately, since the clause head does not unify with the goal and the overall query fails.

6 Modes

In most cases, the correctness of the algorithmic interpretation of a signature as a logic program depends on a restriction to queries of a certain form. Often, this is a restriction of some arguments to *inputs* which must be given as *ground* objects, that is, objects not containing any existential variables. In return, one often obtains *outputs* which will also be ground. In the logic programming terminology, the information about which arguments to a predicate should be considered input and output is called *mode information*.

Twelf supports a simple system of modes. It checks explicit mode declarations by the programmer against the signature and signals errors if the prescribed information flow is violated. Currently, queries are not checked against the mode declaration.

Mode checking is useful to uncover certain types of errors which elude the type-checker. It can also be used to generate more efficient code, although the compiler currently does not take advantage of mode information.

There are two forms of mode declarations: a short form which is adequate and appropriate most of the time, and a long form which is sometimes necessary to ascribe the right modes to implicit arguments

```
mdecl ::= smdecl    % short mode declaration
        | fmdecl    % full mode declaration

decl ::= ...
        | %mode mdecl.
```

6.1 Short Mode Declaration

There are two forms of mode declarations: a short and a full form. The short form is an abbreviation which is expanded into the full form when it is unambiguous.

```

mode ::= +      % input
      | *      % unrestricted
      | -      % output

mid ::= mode id % named mode identifier, one token

smdecl ::= id      % type family a
        | smdecl mid % argument mode

```

Mode declarations for a predicate **a** must come before any clauses defining **a**. Note that the mode followed with the identifier must be one token, such as `+L` and not `+ L`. The short form is most convenient in typical situations. For example, we can declare that the `append` program (see Section 5.4 [Sample Trace], page 21) takes the first two arguments as input and produces the third as output.

```

append : list -> list -> list -> type.
%mode append +L +K -M.

```

If we wanted to use `append` to split a list into two sublists, we would instead declare

```

append : list -> list -> list -> type.
%mode append -L -K +M.

```

where the clauses `appNil` and `appCons` remain unchanged.

In the lambda-calculus type checker (see Section 5.6 [Sample Program], page 24), the type must be an unrestricted argument.

```

of : exp -> tp -> type.
%mode of +E *T.

```

If we declare it as an input argument, `%mode of +E +T`, we obtain an error pointing to the first occurrence of `T2` in the clause `tp_app` reproduced below.

```

examples/nd/lam.elf:27.20-27.22 Error:
Occurrence of variable T2 in input (+) argument not necessarily ground

tp_app : of (app E1 E2) T1
         <- of E1 (arrow T2 T1)
         <- of E2 T2.

```

If we declare it as an output argument, `%mode of +E -T`, we obtain an error pointing to the second occurrence of `T1` in the clause `tp_lam` reproduced below.

```
examples/nd/lam.elf:25.8-25.10 Error:
Occurrence of variable T1 in output (-) argument not necessarily ground

tp_lam : of (lam E) (arrow T1 T2)
        <- ({x:exp}
            of x T1 -> of (E x) T2).
```

In general, for a mode declaration in short form the arguments are specified exactly as they would look in the program. This means one cannot specify the modes of implicit arguments which are filled in by term reconstruction. These modes are reconstructed as follows: each implicit argument which appears in the type of an input argument is considered input '+', those among the remaining which appear in an output argument are considered output '-', the rest are unrestricted. The mode declaration is echoed in full form, so the user can verify the correctness of the modes assigned to implicit arguments. If the inferred full mode declaration is incorrect, or if one wants to be explicit about modes, one should use full mode declarations (see Section 6.2 [Full Mode Declaration], page 29).

6.2 Full Mode Declaration

To specify modes for implicit arguments one must use the full form of mode declaration. A **mode** can be one of '+', '*', or '-' (see Section 6.1 [Short Mode Declaration], page 27).

```
fmdecl ::= mode {id : term} fmdecl
        | mode {id} fmdecl
        | term
```

The term following the mode prefix in a full mode declaration must always have the form $a\ x_1 \dots x_n$ where x_1 through x_n are variables declared in the mode prefix. As an example, we give an alternative specification of the `append` predicate.

```
append : list -> list -> list -> type.
%mode +{L:list} +{K:list} -{M:list} append L K M.
```

6.3 Mode Checking

Mode checking for input, output, and unrestricted arguments examines each clause as it is encountered. The algorithm performs a kind of abstract interpretation of the clause, keeping track of a list of the existential variables for which it knows that they will be ground.

1. We assume each existential variable with a strict occurrence (see Section 4.3 [Strict Occurrences], page 15) in an input argument to the clause head to be ground.
2. We traverse the subgoals in evaluation order (see Section 5.5 [Operational Semantics], page 23). For each subgoal we first verify that all input arguments will be ground, using the information about the existential variables collected so far. If this check succeeds we add all variables which have a strict occurrence in an output argument of the subgoal to the list of variables with known ground instantiations.
3. After the last subgoal has been examined, we verify that the output arguments in the clause head are now also ground.

Arguments whose mode is unrestricted are ignored: they do not need to be checked, and they do not contribute any information about the instantiations of existential variables.

7 Termination

Besides checking types and modes, Twelf can also verify if a given type family, when interpreted as a logic program, always terminates on well-moded goals. In many cases this means that the program implements a decision procedure. Of course, in general termination is undecidable, so we only check a simple sufficient condition.

Checking termination presupposes that the program is well-typed and guarantees termination only when the arguments involved in the termination order are ground. This will always be true for well-moded goals, since mode and termination declarations must be consistent.

Termination is different from checking types and modes in that it is not checked incrementally as the signature is read. Instead, termination of a predicate is a global property of the program once it has been read. Thus termination declarations came after the predicate has been fully defined; further extensions of the predicate are not checked and may invalidate termination.

The termination checker is rather rudimentary in that it only allows lexicographic and simultaneous extensions of the subterm ordering. Moreover, it does not take into account if a result returned by a predicate is smaller than an input argument. Nonetheless, for the style of programs written in Twelf, the termination of many decision procedures can be verified.

7.1 Termination Declaration

The termination orders we construct are lexicographic or simultaneous extensions of the subterm ordering explained in Section 7.2 [Subterm Ordering], page 33. The termination declaration associates the termination order with argument positions of predicates via call patterns.

The case of mutually recursive predicates is particularly complex and requires mutual call patterns and mutual arguments. Their syntax is given below; they are explained in Section 7.5 [Mutual Recursion], page 36.

```

args ::=
  | id args      % named argument
  | _ args       % anonymous argument

callpat ::= id args      % a x1 ... xn

callpats ::=
  | (callpat) callpats
                    % mutual call patterns

ids ::=
  | id ids       % argument name

marg ::= id          % single argument
       | ( ids )    % mutual arguments

orders ::=
  | order orders % component order

order ::= marg        % subterm order
        | { orders } % lexicographic order
        | [ orders ] % simultaneous order

tdecl ::= order callpats % termination declaration

decl ::= ...
        | %terminates tdecl. % termination declaration

```

All identifiers in the order specification of a termination declaration must be upper case, must occur in the call patterns, and no variable may be repeated. Furthermore, all arguments participating in the termination order must occur in the call patterns in input positions.

The most frequent form of termination declaration is

```
%terminates Xi (a X1 ... Xn).
```

which expresses that predicate `a` terminates because recursive calls decrease the input argument `Xi` according to the subterm ordering (see Section 7.2 [Subterm Ordering], page 33).

As an example, we consider a proof that simple type inference (see Section 5.6 [Sample Program], page 24) terminates. Recall the relevant program fragment (see ‘`examples/guide/lam.elf`’).

```

of : exp -> tp -> type.                %name of P.
%mode of +E *T.

tp_lam : of (lam E) (arrow T1 T2)      % |- lam x. E : T1 => T2
      <- ({x:exp}                      % if x:T1 |- E : T2.
          of x T1 -> of (E x) T2).

tp_app : of (app E1 E2) T1             % |- E1 E2 : T1
      <- of E1 (arrow T2 T1)           % if |- E1 : T2 => T1
      <- of E2 T2.                    % and |- E2 : T2.

```

The typability of an expression is always reduced to the typability of its subexpressions. Therefore any call to the `of` predicate with a ground expression should terminate. In general, termination can only be checked for input arguments, and all calls must be well-moded (see Section 6.3 [Mode Checking], page 30). Twelf verifies termination with the declaration

```
%terminates E (of E T).
```

Here, `E` specifies the decreasing argument, namely the first argument of the typing judgment as expressed in the call pattern `(of E T)`.

A corresponding attempt to show that evaluation always terminates,

```
%terminates E (eval E V).
```

fails for the clause `ev_app` with the message

```

examples/guide/lam.elf:1053-1068 Error:
Termination violation:
(E1' V2) < (app E1 E2)

```

indicating that in a recursive call the term `E1' V2` could not be shown to be smaller than `app E1 E2`. In our example, of course, evaluation need not terminate for precisely this reason.

7.2 Subterm Ordering

On first-order terms, that is, terms not containing lambda-abstraction, the subterm ordering is familiar: $M < N$ if M is a strict subterm of N , that is, M is a subterm N and M is different from N .

On higher-order terms, the relation is slightly more complicated because we must allow the substitution of parameters for bound variables without destroying the subterm relation. Consider, for example, the case of the typing rule

```

of : exp -> tp -> type.                %name of P.
%mode of +E *T.

tp_lam : of (lam E) (arrow T1 T2)      % |- lam x. E : T1 => T2
        <- ({x:exp}                    % if x:T1 |- E : T2.
            of x T1 -> of (E x) T2).

```

from the signature for type inference (see Section 5.6 [Sample Program], page 24) in the file ‘example/guide/lam.elf’. We must recognize that

$(E\ x) < (lam\ E)$

according to the subterm ordering. This is because E stands for a term $[y:exp]\ E'$ and so $E\ x$ has the same structure as E' except that y (a bound variable) has been replaced by x (a parameter). This kind of pattern arises frequently in Twelf programs.

On the other hand, the restriction to parameter arguments of functions is critical. For example, the lax rule

```

tp_applam : of (app (lam E1) E2) T2
             <- of (E1 E2) T2.

```

which applies $E1$ to $E2$ which is not a parameter, is indeed not terminating. This can be seen from the query

```
?- of (app (lam [x:exp] app x x) (lam [y:exp] app y y)) T.
```

The restriction of the arguments to parameters can be lifted when the type of the argument is not mutually recursive with the result type of the function. For example, the signature for natural deduction (see Section 3.6 [Sample Signature], page 11, contains no constructor which allows propositions to occur inside individual terms. Therefore

$(A\ T) < (forall\ A)$

where $A : i \rightarrow o$ and $T : i$ is an arbitrary term (not just a parameter). Intuitively, this is correct because the number of quantifiers and logical connectives is smaller on the left, since T cannot contain such quantifiers or connectives.

This kind of precise analysis is important, for example, in the proof of cut elimination or the termination of polymorphic type reconstruction.

7.3 Lexicographic Orders

Lexicographic orders are specified as

$$\{O1 \dots On\}$$

Using vi and wi for corresponding argument structures whose order is already defined, we compare them lexicographically as follows:

$$\begin{aligned} \{v1 \dots vn\} < \{w1 \dots wn\}, \text{ if} \\ &v1 < w1, \text{ or} \\ &v1 = w1 \text{ and } v2 < w2, \text{ or} \\ &\dots \\ &v1 = w1, v2 = w2, \dots, \text{ and } vn < wn. \end{aligned}$$

A lexicographic order is needed, for example, to show termination of Ackermann's function, defined in `examples/arith/arith.elf` with the termination declaration in `examples/arith/arith.thm`.

7.4 Simultaneous Orders

Simultaneous orders require that one of its elements decreases while all others remain the same. This is strictly weaker than a lexicographic ordering built from the same components. Technically speaking it is therefore redundant for termination checking, since the corresponding lexicographic ordering could be used instead. However, for inductive theorem proving it is quite useful, since the search space for simultaneous induction is much smaller than for lexicographic induction.

Simultaneous orders are specified as

$[O1 \dots On]$

Using vi and wi for corresponding argument structures whose order is already defined, we compare them simultaneously as follows:

$[v1 \dots vn] < [w1 \dots wn]$, if
 $v1 < w1$, $v2 < = w2$, ..., and $vn < = wn$, or
 $v1 < = w1$, $v2 < w2$, ..., and $vn < = wn$, or
...
 $v1 < = w1$, $v2 < = w2$, ..., and $vn < wn$.

A combination of simultaneous and lexicographic order is used, for example, in the admissibility of cut found in ‘examples/cut-elim/int.thm’, where either the cut formula A gets smaller, or if A stays the same, either the derivation of the left or right premise get smaller while the other stays the same.

7.5 Mutual Recursion

Mutually recursive predicates present a challenge to termination checking, since decreasing arguments might appear in different positions. Moreover, mutually recursive predicates a and a' might be prioritized so that when a calls a' all termination arguments remain the same, but when a' calls a the arguments are smaller according to the termination order.

To handle the association of related argument in mutually recursive predicates, so-called *mutual arguments* can be specified in a termination order. They are given as

$(X1 \dots Xn)$

The priority between predicates is indicated by the order of the call patterns. If we analyze call patterns

$(a1 \text{ args1})$
 $(a2 \text{ args2})$
...
 $(an \text{ argsn})$

then a_i may call a_j for $i < j$ with equal termination arguments, but calls of a_i from a_j must decrease the termination order.

Mutual arguments are used, for example, in the proofs of soundness (file `'examples/lp-horn/uni-sound.thm'`) and completeness (file `'examples/lp-horn/uni-complete.thm'`) of uniform derivations for Horn logic.

8 Theorem Prover

Disclaimer: The theorem proving component of Twelf is in an even more experimental stage and currently under active development. There are two main restrictions which limit its utility: (1) it only support reasoning about closed objects, and (2) it cannot apply lemmas automatically.

Nonetheless, it can prove a number of interesting examples automatically which illustrate our approach the meta-theorem proving which is described in *Schuermann and Pfenning 1998, CADE*. These examples include type preservation for Mini-ML, one direction of compiler correctness for different abstract machines, soundness and completeness for logic programming interpreters, and the deduction theorem for Hilbert’s formulation of propositional logic. These and other examples can be found in the example directories of the Twelf distribution (see Chapter 12 [Examples], page 65).

A *theorem* in Twelf is, properly speaking, a meta-theorem: it expresses a property of objects constructed over a fixed LF signature. Theorems are stated in the meta-logic M2 whose quantifiers range over LF objects. In the simplest case, we may just be asserting the existence of an LF object of a given type. This only requires direct search for a proof term, using methods inspired by logic programming. More generally, we may claim and prove forall/exists statements which allow us to express meta-theorems which require structural induction, such as type preservation under evaluation in a simple functional language (see Section 5.6 [Sample Program], page 24).

8.1 Theorem Declaration

There are two forms of declarations related to the proving of theorems and meta-theorems. The first, `%theorem`, states a theorem as a meta-formula (`mform`) in the meta-logic M2 defined below. The second, `%prove`, gives a resource bound, a theorem, and an induction ordering and asks Twelf to search for a proof.

Note that a well-typed `%theorem` declaration always succeeds, while the `%prove` declaration only succeeds if Twelf can find a proof.

```

dec ::= {id:term}      % x : A
      | {id}           % x

decs ::= dec
       | dec decs

mform ::= forall* decs mform % implicit universal
        | forall decs mform  % universal
        | exists decs mform  % existential
        | true               % truth

thdecl ::= id : mform      % theorem name a, spec

pdecl ::= nat order callpats % bound, induction order, theorems

decl ::= ...
        | %theorem thdecl. % theorem declaration
        | %prove pdecl.    % prove declaration

```

The prover only accepts quantifier alternations of the form `forall* decs forall decs exists decs true`. Note that the implicit quantifiers (which will be suppressed when printing the proof terms) must all be collected in front.

The syntax and meaning of `order` and `callpats` are explained in Chapter 7 [Termination], page 31, since they are also critical notions in the simpler termination checker.

8.2 Sample Theorems

As a first example, we use the theorem prover to establish a simple theorem in first-order logic (namely that A implies A for any proposition A), using the signature for natural deduction (see Section 3.6 [Sample Signature], page 11).

```

%theorem
trivI : exists {D:{A:o} nd (A imp A)}
      true.

%prove 2 {} (trivI D).

```

The empty termination ordering `{}` instructs Twelf not to use induction to prove the theorem. The declarations above succeed, and with the default setting of 3 for `Twelf.chatter` we see

```

%theorem trivI : ({A:o} nd (A imp A)) -> type.
%prove 2 {} (trivI D).
%mode -{D:{A:o} nd (A imp A)} trivI D.
% -----
/trivI/: trivI ([A:o] impi ([D1:nd A] D1)).
% -----

```

The line starting with `%theorem` shows the way the theorem will be realized as a logic program predicate, the line starting with `/trivI/` gives the implementation, which, in this case, consists of just one line.

The second example is the type preservation theorem for evaluation in the lambda-calculus. This is a continuation of the example in Section 5.6 [Sample Program], page 24 in the file ‘examples/guide/lam.elf’. Type preservation states that if an expression E has type T and E evaluates to V , the V also has type T . This is expressed as the following `%theorem` declaration.

```

%theorem
tps : forall* {E:exp} {V:exp} {T:tp}
      forall {D:eval E V} {P:of E T}
      exists {Q:of V T}
      true.

```

The proof proceeds by structural induction on D , the evaluation from E to V . Therefore we can search for the proof with the following declaration (where the size bound of 5 on proof term size is somewhat arbitrary).

```

%prove 5 D (tps D P Q).

```

Twelf finds and displays the proof easily. The resulting program is installed in the global signature and can then be used to apply type preservation (see Section 8.5 [Proof Realizations], page 42).

8.3 Proof Steps

We expect the proof search component of Twelf to undergo major changes in the near future, so we only briefly review the current state.

Proving proceeds using three main kinds of steps:

- Filling** Using iterative deepening, Twelf searches directly for objects to fill the existential quantifiers, given all the constants in the signature and the universally quantified variables in the theorem. The number of constructors in the answer substitution for each existential quantifier is bounded by the size which is given as part of the `%prove` declaration, thus guaranteeing termination (in principle).
- Recursion** Based on the termination ordering, Twelf appeals to the induction hypothesis on smaller arguments. If there are several ways to use the induction hypothesis, Twelf non-deterministically picks one which has not yet been used. Since there may be infinitely many different ways to apply the induction hypothesis, the parameter `Twelf.Prover.maxRecurse` bounds the number of recursion steps in each case of the proof.
- Splitting** Based on the types of the universally quantified variables, Twelf distinguishes all possible cases by considering all constructors in the signatures. It never splits a variable which appears as an index in an input argument, and if there are several possibilities it picks the one with fewest resulting cases. Splitting can go on indefinitely, so the parameter `Twelf.Prover.maxSplit` bounds the number of times a variable may be split.

8.4 Search Strategies

The basic proof steps of filling, recursion, and splitting are sequentialized in a simple strategy which never backtracks. First we attempt to fill all existential variables simultaneously. If that fails we recurse by trying to find new ways to appeal to the induction hypothesis. If this is not possible, we pick a variable to distinguish cases and then prove each subgoal in turn. If none of the steps are possible we fail.

This behavior can be changed with the parameter `Twelf.Prover.strategy` which defaults to `Twelf.Prover.FRS` (which means Filling-Recursion-Splitting). When set to `Twelf.Prover.RFS` Twelf will first try recursion, then filling, followed by splitting. This is often faster, but fails in some cases where the default strategy succeeds.

8.5 Proof Realizations

Proofs of meta-theorems are realized as logic programs. Such a logic program is a relational representation of the constructive proof and can be executed to generate witness terms for the

existentials from given instances of the universal quantifiers. As an example, we consider once more type preservation (see Section 8.2 [Sample Theorems], page 40).

After the declarations,

```
%theorem
tps : forall* {E:exp} {V:exp} {T:tp}
      forall {D:eval E V} {P:of E T}
      exists {Q:of V T}
      true.

%prove 5 D (tps D P Q).
```

Twelf answers

```
/tps/tp_lam/ev_lam/:
  tps ev_lam (tp_lam ([x:exp] [P2:of x T1] P1 x P2))
    (tp_lam ([x:exp] [P3:of x T1] P1 x P3)).

/tps/tp_app/ev_app/tp_lam/:
  tps (ev_app D1 D2 D3) (tp_app P1 P2) P6
    <- tps D3 P2 (tp_lam ([x:exp] [P4:of x T2] P3 x P4))
    <- tps D2 P1 P5
    <- tps D1 (P3 E5 P5) P6.
```

which is the proof of type preservation expressed as a logic program with two clauses: one for evaluation of a lambda-abstraction, and one for application. Using the `%solve` declaration (see Section 5.2 [Solve Declaration], page 20) we can, for example, evaluate and type-check the identity applied to itself and then use type preservation to obtain a typing derivation for the resulting value.

```
e0 = (app (lam [x] x) (lam [y] y)).
%solve p0 : of e0 T.
%solve d0 : eval e0 V.
%solve tps0 : tps d0 p0 Q.
```

Recall that `%solve c : V` executes the query `V` and defines the constant `c` to abbreviate the resulting proof term.

9 ML Interface

The Twelf implementation defines a number of ML functions embedded in structures which can be called to load files, execute queries, and set environment parameters such as the verbosity level of the interaction. These functions and parameters are available in the `Twelf` structure. If you open the `Twelf` structure with

```
open Twelf
```

after compiling and loading Twelf, you do not have to type the ‘`Twelf.`’ to the functions shown below.

Previous implementations of Elf offered a stand-alone command interpreter but this has not yet been ported. To exit Twelf and ML call `Twelf.OS.exit ()`;

9.1 Configurations

Groups of Twelf files are managed in *configurations*. A configuration is defined by a file, by convention called ‘`sources.cfg`’, which resides in the same directory as the Twelf source files. The configuration file must contain at most one Twelf source file per line, and the files must be listed in dependency order. A configuration *config* can then be defined from the file by the ML declaration

```
val config = Twelf.Config.read "sources.cfg";
```

By convention, the filenames end in the extensions

‘ <code>.elf</code> ’	for constant declarations and definitions or mixed files,
‘ <code>.quy</code> ’	for files which contain query declarations,
‘ <code>.thm</code> ’	for files which contain <code>%theorem</code> and <code>%proof</code> declarations.

File names may not contain whitespace. They are interpreted relative to the current working directory of ML, but resolved into absolute path names when the configuration file is read. To change the current working directory call

```
Twelf.OS.getDir (); (* get working directory *)
Twelf.OS.chdir "directory"; (* change working directory *)
```

As an example, we show how the Mini-ML configuration is defined and loaded, assuming your current working directory is the root directory of Twelf.

```
val mini_ml = Twelf.Config.read "examples/mini-ml/sources.cfg";
Twelf.Config.load mini_ml;
```

Note that the identifier bound to the configuration (`mini_ml` in this example), must be a legal ML identifier, usually consisting only of alphanumeric characters and underscores. The call to `Twelf.Config.load` returns either `Twelf.OK` or `Twelf.ABORT`. It reads each file in turn, starting from an empty signature, printing the results of type reconstruction and search based on the value of the `Twelf.chatter` variable (see Section 9.3 [Environment Parameters], page 47). If another configuration or file has previously been read, all the declarations will first be deleted so that `Twelf.Config.load` always starts from the same state.

Loading a configuration will stop at the first error encountered, issue an appropriate message and return `Twelf.ABORT`. If there is an unexpected internal error (which indicates a bug in the Twelf implementation), it raises an uncaught exception instead and returns to the ML top-level.

To explore the behavior of programs interactively, you may call the Twelf top-level with

```
Twelf.top ();
```

which is explained in Section 5.3 [Interactive Queries], page 20.

9.2 Loading Files

Twelf also allows direct management of the signature by loading individual files. This is generally not recommended because successive declarations simply accumulate in the global signature which may lead to unexpected behavior. The relevant function calls are

```
Twelf.reset ();
Twelf.loadFile "file";
```

where `Twelf.reset ()` resets the current global signature to be empty and `Twelf.readFile "file"` loads the given *file* whose name is interpreted relative to the current working directory.

Caution: Reading a file twice will not replace the declarations of the first pass by the second, but simply add them to the current signature. If names are reused, old declarations will be shadowed,

but they are still in the global signature and might be used in the search for a solution to a query or in theorem proving, leading to unexpected behavior. When in doubt, use configurations (see Section 9.1 [Configurations], page 45) or call `Twelf.reset ()`.

9.3 Environment Parameters

Various flags and parameters can be used to modify the behavior of Twelf and the messages it issues. They are given below with the assignment of the default value.

`Twelf.chatter := 3;`

Controls the detail of the information which is printed when signatures are read.

0 Nothing.

1 Just file names.

2 File names and number of query solutions.

3 Each declarations after type reconstruction.

4 Debug information.

5 More debug information.

`Twelf.doubleCheck := false;`

If `true`, each declaration is checked again for type correctness after type reconstruction. This is expensive and useful only for your peace of mind, since type checking is significantly simpler than type reconstruction.

`Twelf.Print.implicit := false;`

If `true`, implicit arguments (normally elided) are printed. Sometimes this is useful to track particularly baffling errors.

`Twelf.Print.depth := NONE;`

If `SOME(d)` then terms deeper than level `d` are printed as ‘%%’.

`Twelf.Print.length := NONE;`

If `SOME(1)` then argument lists longer than 1 are truncated with ‘...’.

`Twelf.Print.indent := 3;`

Controls the amount of indentation for printing nested terms.

`Twelf.Print.width := 80;`

The value used to decide when to break lines during printing of terms.

`Twelf.Prover.strategy := Twelf.Prover.FRS;`

Determines the strategy, where `F`=Filling, `R`=Recursion, and `S`=Splitting. Can also be `Twelf.Prover.RFS`.

```
Twelf.Prover.maxSplit := 2;
```

The maximal number of generations of a variable introduced by splitting. Setting is to 0 will prohibit proof by cases.

```
Twelf.Prover.maxRecurse := 10;
```

The maximal number of appeals to the induction hypothesis in any case during a proof.

9.4 Timing Statistics

Twelf has a few utilities to collect run-time statistics which are useful mainly for the developers. They are collected in the structure `Timers`. Timing information is cumulative in an ML session.

```
Twelf.Timers.show ();
```

Show the value of timers and reset them to zero.

```
Twelf.Timers.reset ();
```

Simply reset all timers to zero.

```
Twelf.Timers.check ();
```

Display the value of timers, but do not reset them.

Caution: Normally, the various times are exclusive, except that the runtime includes the garbage collection time which is shown separately. However, there is a problem the time for printing the answer substitution to a query is charged both to `Printing` and `Solving`.

9.5 Twelf Signature

For reference, here is the ML signature `TWELF` of the `Twelf` structure which defines most functions and flags relevant to loading and executing Twelf programs.

```
signature TWELF =
sig
  structure Print :
  sig
    val implicit : bool ref          (* false, print implicit args *)
    val depth : int option ref      (* NONE, limit print depth *)
    val length : int option ref     (* NONE, limit argument length *)
    val indent : int ref            (* 3, indentation of subterms *)
    val width : int ref             (* 80, line width *)
  end
end
```

```

structure Timers :
sig
  val show : unit -> unit          (* show and reset timers *)
  val reset : unit -> unit         (* reset timers *)
  val check : unit -> unit         (* display, but not no reset *)
end

structure OS :
sig
  val chDir : string -> unit       (* change working directory *)
  val getDir : unit -> string      (* get working directory *)
  val exit : unit -> unit          (* exit Twelf and ML *)
end

structure Prover :
sig
  datatype Strategy = RFS | FRS    (* F=Fill, R=Recurse, S=Split *)
  val strategy : Strategy ref      (* FRS, strategy used for %prove *)
  val maxSplit : int ref           (* 2, bound on splitting *)
  val maxRecurse : int ref         (* 10, bound on recursion *)
end

val chatter : int ref              (* 3, chatter level *)
val doubleCheck : bool ref         (* false, check after reconstruction *)
datatype Status = OK | ABORT       (* return status *)

val reset : unit -> unit           (* reset global signature *)
val loadFile : string -> Status    (* load file *)
val readDecl : unit -> Status      (* read declaration interactively *)
val decl : string -> Status        (* print declaration of constant *)

val top : unit -> unit             (* top-level for interactive queries *)

structure Config :
sig
  type config                      (* configuration *)
  val read : string -> config       (* read config file *)
  val load : config -> Status       (* reset and load configuration *)
  val define : string list -> config (* explicitly define configuration *)
end

val version : string               (* Twelf version *)
end; (* signature TWELF *)

```


10 Twelf Server

The Twelf server is a stand-alone command interpreter which provides the functionality of the **Twelf** structure in ML (see Chapter 9 [ML Interface], page 45), but allows no ML definitions. It is significantly smaller than Standard ML and is the recommended way to interact with Twelf except for developers. Its behavior regarding configurations is slightly different in that the server maintains a current configuration, rather than allowing the binding of names to configurations. Configuration are defined with the **Config.read** command which takes a configuration filename as argument.

In Emacs, the Twelf server typically runs in a process buffer called ***twelf-server***. The user can select this buffer and directly type commands to the Twelf server. This style of interaction is inherited from the **comint** package for Emacs, but typically one works through advanced commands in Twelf mode (see Section 11.1 [Twelf Mode], page 55).

The Twelf server prompts with **%% OK %%** or **%% ABORT %%** depending on the success of failure of the previous operation. It accepts commands and their arguments on one line, except that additional Twelf declarations which may be required are read separately, following the command line. Reading declarations can be forcibly terminated with the end-of-file token **“%.’**.

10.1 Server Types

The server commands employ arguments of the following types.

file	The name of a file, relative to the current working directory.
id	A Twelf identifier
strategy	Either FRS or RFS (see Section 8.4 [Search Strategies], page 42)
bool	Either true or false
nat	A natural number (starting at 0)
limit	Either * (to indicate no limit) or a natural number

10.2 Server Commands

The Twelf server recognized the following commands.

set *parameter value*

Set *parameter* to *value*, where *parameter* is on of the following (explained in Section 9.3 [Environment Parameters], page 47).

chatter *nat*

doubleCheck *bool*

Print.implicit *bool*

Print.depth *limit*

Print.length *limit*

Print.indent *nat*

Print.width *nat*

Prover.strategy *strategy*

Prover.maxSplit *nat*

Prover.maxRecurse *nat*

get *parameter*

Print the current value of *parameter* (see table above).

Timers.show

Print and reset timers.

Timers.reset

Reset timers.

Timers.check

Print, but do not reset timers.

OS.chdir *file*

Change working directory to *file*.

OS.getDir

Print current working directory.

OS.exit Exit Twelf server.

quit Quit Twelf server (same as exit).

Config.read *file*

Read current configuration from *file*.

Config.load

Load current configuration

reset Reset global signature.

loadFile *file*

Load Twelf file *file*.

`decl id` Show constant declaration for *id*.

`top` Enter interactive query loop (see Section 5.3 [Interactive Queries], page 20)

11 Emacs Interface

The Twelf mode for Emacs provides some functions and utilities for editing Twelf source and for interacting with an inferior Twelf server process which can load configurations, files, and individual declarations and track the source location of errors. It also provides an interface to the tags package which allows simple editing of groups of files, constant name completion, and locating of constant declarations within the files of a configuration.

Note that in order to use the Emacs interface you need to include the line

```
(load "directory/emacs/twelf-init.el")
```

in your ‘.emacs’ file, where *directory* is the Twelf root directory.

11.1 Twelf Mode

The Twelf mode in Emacs provides support for editing and indentation, syntax highlighting (including colors) (see Section 11.11 [Syntax Highlighting], page 62), and communication commands for interacting with a Twelf server running as an inferior process to Emacs. It defines a menu which is added to the menu bar, usually at the top of each Emacs frame.

Many commands apply to the current declaration, which is the declaration in which we find the Emacs cursor (not the cursor of the window system). If the cursor is between declarations, the declaration after point is considered current. From the point of view of Emacs, single declarations never include consecutive blank lines, which provides some insulation against missing closing delimiters.

Normally, Twelf mode is entered automatically when a Twelf source file is edited (see Section 11.12 [Emacs Initialization], page 62), but it can also be switched on or off directly with **M-x twelf-mode**.

M-x twelf-mode

Toggle Twelf mode, the major mode for editing Twelf code.

11.2 Editing Commands

The editing commands in Twelf mode partially analyse the structure of the text at the cursor position as Twelf code and try to indent accordingly. This is not always perfect.

TAB

M-x twelf-indent-line

Indent current line as Twelf code. This recognizes comments, matching delimiters, and standard infix operators.

DEL

M-x backward-delete-char-untabify

Delete character backward, changing tabs into spaces.

M-C-q

M-x twelf-indent-decl

Indent each line of the current declaration.

M-x twelf-indent-region

Indent each line of the region as Twelf code.

11.3 Type Checking Commands

The Twelf mode provides simple commands which cause the server to load or reload the current configuration, the file edited in the current buffer, or just the declaration at point. Each of these command can be preceded by a prefix argument (for example, C-u C-c C-c) which will select the Twelf server buffer after completion of the command. The Twelf server buffer can also be forced to be shown with the C-c C-u Emacs command.

C-c C-c

M-x twelf-save-check-config

Save its modified buffers and then check the current Twelf configuration. With prefix argument also displays Twelf server buffer. If necessary, this will start up an Twelf server process.

C-c C-s

M-x twelf-save-check-file

Save buffer and then check it by giving a command to the Twelf server. In Twelf Config minor mode, it reconfigures the server. With prefix argument also displays Twelf server buffer.

C-c C-d

M-x twelf-check-declaration

Send the current declaration to the Twelf server process for checking. With prefix argument also displays Twelf server buffer.

C-c c

M-x twelf-type-const

Display the type of the constant before point. Note that the type of the constant will be ‘absolute’ rather than the type of the particular instance of the constant.

C-c C-u

M-x twelf-server-display

Display Twelf server buffer, moving to the end of output. With prefix argument also selects the Twelf server buffer.

11.4 Error Tracking

Error messages by the Twelf server are flagged with the filename and an educated guess as to the source of the error (see Section 4.6 [Error Messages], page 18). These can be interpreted by Emacs to jump directly to the suspected site.

Sometimes, the server buffer and the the server itself believe to have different working directories. In that case, error tracking may not be able to find the file, and an explicit call to `OS.chdir` or `M-x cd` in the server buffer may be required.

C-c ‘

M-x twelf-next-error

Find the next error by parsing the Twelf server or Twelf-SML buffer. Move the error message on the top line of the window; put the cursor at the beginning of the error source. If the error message specifies a range, the mark is placed at the end.

C-c =

M-x twelf-goto-error

Go to the error reported on the current line or below. Also updates the error cursor to the current line.

11.5 Server State

The server state consists of the current configuration and a number of parameters described in Chapter 10 [Twelf Server], page 51. The current configuration is often set implicitly, with the `C-c` command in a configuration buffer, but it can also be set explicitly.

`C-c <`

`M-x twelf-set`

Sets the Twelf parameter `PARAM` to `VALUE`. When called interactively, prompts for parameter and value, supporting completion.

`C-c >`

`M-x twelf-get`

Prints the value of the Twelf parameter `PARAM`. When called interactively, prompts for parameter, supporting completion.

`C-c C-i`

`M-x twelf-server-interrupt`

Interrupt the Twelf server process.

`M-x twelf-server`

Start an Twelf server process in a buffer named `*twelf-server*`. Any previously existing process is deleted after confirmation. Optional argument `PROGRAM` defaults to the value of the variable `twelf-server-program`. This locally re-binds `'twelf-server-timeout'` to 15 secs.

`M-x twelf-server-configure`

Initializes the Twelf server configuration from `CONFIG-FILE`. A configuration file is a list of relative file names in dependency order. Lines starting with `%` are treated as comments. Starts a Twelf servers if necessary.

`M-x twelf-reset`

Reset the global signature of Twelf maintained by the server.

`M-x twelf-server-quit`

Kill the Twelf server process.

`M-x twelf-server-restart`

Restarts server and re-initializes configuration. This is primarily useful during debugging of the Twelf server code or if the Twelf server is hopelessly wedged.

`M-x twelf-server-send-command`

Restarts server and re-initializes configuration. This is primarily useful during debugging of the Twelf server code or if the Twelf server is hopelessly wedged.

11.6 Info File

The content of this file in Info format can be visited directly and does not need to be tied into the Info tree. See the documentation for the Emacs info package for more info

C-c C-h

M-x twelf-info

Visit the Twelf User's Guide in info format in Emacs. With a prefix argument it prompts for the info file name, which defaults to the value of the `twelf-info-file` variable.

11.7 Tags Files

Tags files provide a convenient way to group files, such as Twelf configurations. See the documentation for the Emacs etags package for more information.

M-x twelf-tag

Create tags file for current configuration. If the current configuration is `sources.cfg`, the tags file is `TAGS`. If current configuration is named `FILE.cfg`, tags file will be named `FILE.tag` Errors are displayed in the Twelf server buffer.

M-.

M-x find-tag TAG

Selects the buffer that the tag is contained in and puts point at its definition.

C-x 4 .

M-x find-tag-other-window TAG

Selects the buffer that TAG is contained in in another window and puts point at its definition.

C-c q

M-x tags-query-replace FROM TO

Query-replace-regexp FROM with TO through all files listed in tags table.

C-c s

M-x tags-search REGEXP

Search through all files listed in tags table for match for REGEXP.

M-,

M-x tags-loop-continue

Continue last C-c s or C-c q command.

11.8 Twelf Timers

The following commands obtain the runtime statistics of the Twelf server.

M-x twelf-timers-reset

Reset the Twelf timers.

M-x twelf-timers-show

Show and reset the Twelf timers.

M-x twelf-timers-check

Show the Twelf timers without resetting them.

11.9 Twelf-SML Mode

There is some support for interacting with Twelf, even when it is run within ML, rather than as a stand-alone server. You can start an SML in which you intend to run Twelf with **M-x twelf-sml**; the buffer will then be in Twelf-SML mode.

If you intend to send command to a buffer running Twelf in SML (rather than the Twelf server), you can switch to a minor mode 2Twelf-SML with **M-x twelf-to-twelf-sml**.

M-x twelf-sml

Run an inferior Twelf-SML process in a buffer **twelf-sml**. If there is a process already running in **twelf-sml**, just switch to that buffer. With argument, allows you to change the program which defaults to the value of `twelf-sml-program`. Runs the hooks from `twelf-sml-mode-hook` (after the `comint-mode-hook` is run).

M-x twelf-to-twelf-sml-mode

Toggles minor mode for sending queries to Twelf-SML instead of Twelf server.

C-c C-e

M-x twelf-sml-send-query

Send the current declaration to the inferior Twelf-SML process as a query. Prefix argument means `switch-to-twelf-sml` afterwards.

C-c C-r

M-x twelf-sml-send-region

Send the current region to the inferior Twelf-SML process. Prefix argument means `switch-to-twelf-sml` afterwards.

C-c RETURN

M-x twelf-sml-send-newline

Send a newline to the inferior Twelf-SML process. If a prefix argument is given, switches to Twelf-SML buffer afterwards.

C-c ;

M-x twelf-sml-send-semicolon

Send a semi-colon to the inferior Twelf-SML process. If a prefix argument is given, switched to Twelf-SML buffer afterwards.

C-c d

M-x twelf-sml-cd DIR

Make DIR become the Twelf-SML process' buffer's default directory and furthermore issue an appropriate command to the inferior Twelf-SML process.

M-x twelf-sml-quit

Kill the Twelf-SML process.

11.10 Emacs Variables

A number of Emacs variables can be changed to customize the behavior of Twelf mode. The list below is not complete; please refer to the Emacs Lisp sources in `emacs/twelf.el` for additional information.

twelf-indent

Indent for Twelf expressions.

twelf-server-program

Default Twelf server program.

twelf-info-file

Default Twelf info file.

twelf-mode-hook

List of hook functions to run when switching to Twelf mode.

twelf-server-mode-hook

List of hook functions to run when switching to Twelf Server mode.

twelf-sml-program

Default Twelf-SML program.

twelf-sml-mode-hook

List of hook functions for Twelf-SML mode.

11.11 Syntax Highlighting

Twelf also provides syntax highlighting, which helps make Elf code more readable. This highlighting can use different colors and faces. Unfortunately, the necessary libraries are at present not standardized between XEmacs and FSF Emacs, which means that highlighting support is less general and less portable than the plain Twelf mode.

At present, highlighting has not been extensively tested in various versions of Emacs, but the font-lock mode provided in `'emacs/twelf-font.el'` seems to work at least in XEmacs version 19.16 and FSF Emacs version 19.34. The alternative highlight mode provided in `'emacs/twelf-hilit'` appears to work in FSF Emacs 19.34.

Unlike other font-lock modes, Twelf's fontification is not 'electric' in that it does not fontify as one types. One has to explicitly issue a command to fontify the current Twelf declaration or current buffer, since single-line highlighting is too error-prone and multi-line immediate highlighting is not well supported in current versions of font lock mode.

C-c C-l

M-x twelf-font-fontify-decl

Fontifies the current Twelf declaration.

C-c l

M-x twelf-font-fontify-buffer

Fontitifies the current buffer as Twelf code

M-x twelf-font-unfontify

Removes fontification from current buffer.

11.12 Emacs Initialization

If Twelf has been properly installed, you can use the Twelf's Emacs interface with the default settings simply by adding the line

```
(load "directory/emacs/twelf-init.el")
```

to your `'emacs'` file, where *directory* is the Twelf root directory. In order to customize the behavior, you might copy the file `'emacs/twelf-init.el'` or its contents and change it as appropriate.

11.13 Command Summary

```

--- Editing Commands ---
TAB          twelf-indent-line
DEL          backward-delete-char-untabify
M-C-q       twelf-indent-decl

--- Type Checking ---
C-c C-c     twelf-save-check-config
C-c C-s     twelf-save-check-file
C-c C-d     twelf-check-declaration
C-c c       twelf-type-const
C-c C-u     twelf-server-display

--- Error Tracking ---
C-c '       twelf-next-error
C-c =       twelf-goto-error

--- Syntax Highlighting ---
C-c C-l     twelf-font-fontify-decl
C-c l       twelf-font-fontify-buffer

--- Server State ---
C-c <       twelf-set
C-c >       twelf-get
C-c C-i     twelf-server-interrupt
M-x twelf-server
M-x twelf-server-configure
M-x twelf-server-quit
M-x twelf-server-restart
M-x twelf-server-send-command

--- Info ---
C-c C-h     twelf-info

--- Timers ---
M-x twelf-timers-reset
M-x twelf-timers-show
M-x twelf-timers-check

--- Tags (standard Emacs etags package) ---
M-x twelf-tag
M-.         find-tag (standard binding)
C-x 4 .     find-tag-other-window (standard binding)
C-c q       tags-query-replace (Twelf mode binding)
C-c s       tags-search (Twelf mode binding)
M-,         tags-loop-continue (standard binding)
            visit-tags-table, list-tags, tags-apropos

```

```
--- Communication with inferior Twelf-SML process (not Twelf Server) ---  
M-x twelf-sml  
C-c C-e      twelf-sml-send-query  
C-c C-r      twelf-sml-send-region  
C-c RET      twelf-sml-send-newline  
C-c ;        twelf-sml-send-semicolon  
C-c d        twelf-sml-cd  
M-x twelf-sml-quit  
--- Variables ---  
twelf-indent
```

12 Installation

At present, Twelf has been tested in SML of New Jersey (version 110 or higher) and MLWorks, both of which implement Standard ML (revised 1997) and the Standard ML Basis Library. The instructions below apply to a Unix system. For instructions for other architectures or updates please check the file ‘INSTALL’ at the Twelf home page and in the Twelf root directory after unpacking the distribution.

On a Unix system you unpack the sources with

```
gunzip twelf-1-2.tar.gz
tar -xf twelf-1-2.tar
cd twelf
make
```

This builds the Twelf server (see Chapter 10 [Twelf Server], page 51) for your current architecture and makes it accessible as ‘bin/twelf-server’. It also installs the Twelf Emacs interface (see Chapter 11 [Emacs Interface], page 55), but you must add a line

```
(load "directory/emacs/twelf-init.el")
```

to your ‘.emacs’ file, where *directory* is the root directory into which you installed Twelf. Note that the Twelf installation cannot be moved after it has been compiled with `make`, since absolute pathnames are built into the executable scripts.

Note that the Twelf server presently only works with Standard ML of New Jersey, since interrupt handling is implementation specific.

If you would like to use Twelf as a structure in SML, you can then call

```
make twelf-sml
```

which creates ‘bin/twelf-sml’ for the Twelf-SML mode (see Section 11.9 [Twelf-SML Mode], page 60). Calling `make clean` will remove temporary files created by the SML compiler, but not the executable file.

SML of New Jersey (free, version 110 or higher)

See <http://cm.bell-labs.com/cm/cs/what/smlnj/index.html>

MLWorks (commercial)

See <http://www.harlequin.com/products/ads/ml/ml.html>

In MLWorks, you can presently only directly load the Twelf sources, using the file 'load.sml'.

```
mlworks-basis  start MLWorks with basis library in Twelf root directory
use "load.sml";  compile and load Twelf
```

13 Examples

We give here only a brief reference to the examples in the ‘`examples/`’ subdirectory of the distribution. Each example comes in a separate subdirectory whose name is listed below.

‘ <code>arith</code> ’	Associativity and commutative of unary addition.
‘ <code>ccc</code> ’	Cartesian-closed categories (currently incomplete).
‘ <code>church-rosser</code> ’	The Church-Rosser theorem for untyped lambda-calculus.
‘ <code>compile</code> ’	Various compilers starting from Mini-ML.
‘ <code>cut-elim</code> ’	Cut elimination for intuitionistic and classical logic.
‘ <code>fol</code> ’	Simple theorems in first-order logic.
‘ <code>guide</code> ’	Examples from Users’ Guide.
‘ <code>lp</code> ’	Logic programming, uniform derivations.
‘ <code>lp-horn</code> ’	Horn fragment of logic programming.
‘ <code>mini-ml</code> ’	Mini-ML, type preservation and related theorems.
‘ <code>polylam</code> ’	Polymorphic lambda-calculus.
‘ <code>prop-calc</code> ’	Natural deduction and Hilbert propositional calculus
‘ <code>units</code> ’	Mini-ML extended with units (currently incomplete).

In each directory or subdirectory you can find a file ‘`sources.cfg`’ which defines the standard configuration, usually just the basic theory. The ‘`test.cfg`’ which also defines an extended configuration with some test queries and theorems. Most examples also have a ‘`README`’ file with a brief explanation and pointer to the literature.

Index

%

<code>%infix</code>	9
<code>%mode</code>	27, 29
<code>%postfix</code>	9
<code>%prefix</code>	9
<code>%prove</code>	39
<code>%query</code>	19
<code>%solve</code>	20
<code>%terminates</code>	31
<code>%theorem</code>	39

A

<code>add-hook</code>	62
ambiguity	17
arguments, implicit	14
arguments, mutual	36
arithmetic	65
assumptions	23
<code>auto-mode-alist</code>	62
<code>autoload</code>	62

B

backquote, before variables	6
<code>bool</code>	51
bound variables	8

C

call patterns	31
Cartesian-closed categories	65
case, upper and lower	6
characters, reserved	5
Church-Rosser theorem	65
clause selection	23
colors	62
commands, Emacs	63
commands, server	51
<code>Config.load</code>	52
<code>Config.read</code>	52
Configurations	45

current declaration	55
cut elimination	65

D

<code>decl</code>	53
declaration	7
declaration, currnt	55
declarations	7
declarations, mode	27
declarations, name preference	10
declarations, operator	9
declarations, termination	31
declarations, theorem	39
definitions	9
definitions, strict	16
display, of server buffer	56
documentation	58

E

editing	55
Emacs variables	61
environment parameters	47
error messages	18
error tracking	57
examples, from user's guide	65
executing proofs	42
existential quantifier	39

F

faces	62
<code>file</code>	51
file names	1
files, configuration	45
files, loading	46
filling	41
first-order logic	65
free variables	8

G

get 52

H

Hilbert calculus 65

Horn logic, theory 65

I

id 51

identifiers, reserved 6

implicit arguments 14

implicit quantifiers 13

indentation 55

info file 58

initializing Twelf mode 62

input mode 27

installation 50

interrupt 57

K

kinds 7

L

lambda calculus example 24

lambda-calculus, polymorphic 65

lambda-calculus, untyped 65

left 9

LF 1

limit 51

load-path 62

loadFile 52

loading files 46

local assumptions 23

local parameters 23

logic programming 19

logic programming, theory 65

logical framework 1

M

M-x backward-delete-char-untabify 56

M-x find-tag 59

M-x find-tag-other-window 59

M-x tags-loop-continue 59

M-x tags-query-replace 59

M-x tags-search 59

M-x twelf-check-declaration 56

M-x twelf-font-fontify-buffer 62

M-x twelf-font-fontify-decl 62

M-x twelf-font-unfontify 62

M-x twelf-get 58

M-x twelf-goto-error 57

M-x twelf-indent-decl 56

M-x twelf-indent-line 56

M-x twelf-indent-region 56

M-x twelf-info 59

M-x twelf-mode 55

M-x twelf-next-error 57

M-x twelf-reset 58

M-x twelf-save-check-config 56

M-x twelf-save-check-file 56

M-x twelf-server 58

M-x twelf-server-configure 58

M-x twelf-server-display 57

M-x twelf-server-interrupt 58

M-x twelf-server-quit 58

M-x twelf-server-restart 58

M-x twelf-server-send-command 58

M-x twelf-set 58

M-x twelf-sml 60

M-x twelf-sml-cd 61

M-x twelf-sml-quit 61

M-x twelf-sml-send-newline 60

M-x twelf-sml-send-query 60

M-x twelf-sml-send-region 60

M-x twelf-sml-send-semicolon 61

M-x twelf-tag 59

M-x twelf-timers-check 60

M-x twelf-timers-reset 60

M-x twelf-timers-show 60

M-x twelf-to-twelf-sml-mode 60

M-x twelf-type-const 57

meta-logic 39

Mini-ML, compilation 65

Mini-ML, theory 65

Mini-ML, with units 65

ML implementations 50

ML interface	45
MLWorks	50
mode checking	30
mode declaration, full form	29
mode declarations, short form	27
modes	27
mutual arguments	36
mutual recursion	36

N

name preferences	10
nat	51
natural deduction	11
none	9

O

objects	7
occurrences, rigid	15
occurrences, strict	15
open	45
operational semantics	23
operator declarations	9
order	31
order, lexicographic	35
order, simultaneous	35
order, subterm	33
OS.chDir	52
OS.exit	52
OS.getDir	52
output mode	27

P

parameters	23
parameters, environment	47
precedence	9
proof realizations	42

Q

quantifier, existential	39
quantifier, universal	39
quantifiers, implicit	13
queries	19
queries, interactive	20

quit	52
-------------------	----

R

recursion	41
reserved characters	5
reserved identifiers	6
reset	52
right	9
rigid occurrences	15
running time	48

S

search strategy	42
semantics, operational	23
server	51
server buffer	56
server commands	51
server parameters, setting	57
server state	57
server timers	59
server types	51
set	52
setting server parameters	57
signature	7
signature TWELF	48
solving queries	20
splitting	41
Standard ML of New Jersey	50
statistics	48
strategy	51
strict definitions	16
strict occurrences	15
structure Twelf	48
subgoal selection	23
subterm order	33
syntax highlighting	62

T

tagging configurations	59
tags file	59
term	7
term reconstruction	13
termination checking	31

termination declarations	31	<code>Twelf.Print.implicit</code>	47
termination order	31	<code>Twelf.Print.indent</code>	47
theorem declarations	39	<code>Twelf.Print.length</code>	47
theorem prover	39	<code>Twelf.Print.width</code>	47
<code>Timers.check</code>	52	<code>Twelf.Prover.FRS</code>	42
<code>Timers.reset</code>	52	<code>Twelf.Prover.maxRecurse</code>	41, 48
<code>Timers.show</code>	52	<code>Twelf.Prover.maxSplit</code>	41, 48
timing statistics	48	<code>Twelf.Prover.RFS</code>	42
<code>top</code>	53	<code>Twelf.Prover.strategy</code>	42, 47
top-level, query	20	<code>Twelf.reset</code>	46
tracking errors	57	<code>Twelf.Timers.check</code>	48
Twelf home page	1	<code>Twelf.Timers.reset</code>	48
Twelf mode in Emacs	55	<code>Twelf.Timers.show</code>	48
Twelf server	51	type ascription	17
<code>twelf-indent</code>	61	type checking, from Emacs	56
<code>twelf-info-file</code>	61	type families	7
<code>twelf-mode-hook</code>	61	type inference example	24
<code>twelf-server-mode-hook</code>	61	type reconstruction	13
<code>twelf-server-program</code>	61	types	7
Twelf-SML mode	60	types, server	51
<code>twelf-sml-mode-hook</code>	61	typographical conventions	1
<code>twelf-sml-program</code>	61		
<code>Twelf.ABORT</code>	45	U	
<code>Twelf.chatter</code>	47	unification	23
<code>Twelf.Config.load</code>	45	universal quantifier	39
<code>Twelf.Config.read</code>	45		
<code>Twelf.doubleCheck</code>	47	V	
<code>Twelf.loadFile</code>	46	variable naming	10
<code>Twelf.OK</code>	45	variable scope	8
<code>Twelf.OS.chDir</code>	45	variables, bound	8
<code>Twelf.OS.getDir</code>	45	variables, Emacs	61
<code>Twelf.Print.depth</code>	47	variables, free	8

Table of Contents

1	Introduction	1
1.1	New Features	1
1.2	Quick Start	3
2	Lexical Conventions	5
2.1	Reserved Characters	5
2.2	Identifiers	6
3	Syntax	7
3.1	Grammar	7
3.2	Constructor Declaration	8
3.3	Definitions	9
3.4	Operator Declaration	9
3.5	Name Preferences	10
3.6	Sample Signature	11
4	Term Reconstruction	13
4.1	Implicit Quantifiers	13
4.2	Implicit Arguments	14
4.3	Strict Occurrences	15
4.4	Strict Definitions	16
4.5	Type Ascription	17
4.6	Error Messages	18
5	Logic Programming	19
5.1	Query Declaration	19
5.2	Solve Declaration	20
5.3	Interactive Queries	20
5.4	Sample Trace	21
5.5	Operational Semantics	23
5.6	Sample Program	24
6	Modes	27
6.1	Short Mode Declaration	27
6.2	Full Mode Declaration	29
6.3	Mode Checking	30

7	Termination	31
7.1	Termination Declaration	31
7.2	Subterm Ordering	33
7.3	Lexicographic Orders	35
7.4	Simultaneous Orders	35
7.5	Mutual Recursion	36
8	Theorem Prover	39
8.1	Theorem Declaration	39
8.2	Sample Theorems	40
8.3	Proof Steps	41
8.4	Search Strategies	42
8.5	Proof Realizations	42
9	ML Interface	45
9.1	Configurations	45
9.2	Loading Files	46
9.3	Environment Parameters	47
9.4	Timing Statistics	48
9.5	Twelf Signature	48
10	Twelf Server	51
10.1	Server Types	51
10.2	Server Commands	51
11	Emacs Interface	55
11.1	Twelf Mode	55
11.2	Editing Commands	56
11.3	Type Checking Commands	56
11.4	Error Tracking	57
11.5	Server State	58
11.6	Info File	59
11.7	Tags Files	59
11.8	Twelf Timers	60
11.9	Twelf-SML Mode	60
11.10	Emacs Variables	61
11.11	Syntax Highlighting	62
11.12	Emacs Initialization	62
11.13	Command Summary	63
12	Installation	65

13	Examples.....	67
Index	69

