



University of Amsterdam

Dept. of Social Science Informatics (SWI)
Roeterstraat 15, 1018 WB Amsterdam
The Netherlands
Tel. (+31) 20 5256121

SWI-Prolog 5.2

Reference Manual

Updated for version 5.2.11, November 2003

Jan Wielemaker

jan@swi.psy.uva.nl
<http://www.swi-prolog.org>

SWI-Prolog is a Prolog implementation based on a subset of the WAM (Warren Abstract Machine). SWI-Prolog was developed as an *open* Prolog environment, providing a powerful and bi-directional interface to C in an era this was unknown to other Prolog implementations. This environment is required to deal with XPCE, an object-oriented GUI system developed at SWI. XPCE is used at SWI for the development of knowledge-intensive graphical applications.

As SWI-Prolog became more popular, a large user-community provided requirements that guided its development. Compatibility, portability, scalability, stability and providing a powerful development environment have been the most important requirements. Edinburgh, Quintus, SICStus and the ISO-standard guide the development of the SWI-Prolog primitives.

This document gives an overview of the features, system limits and built-in predicates.

Contents

1	Introduction	8
1.1	SWI-Prolog	8
1.1.1	Books about Prolog	8
1.2	Status	9
1.3	Compliance to the ISO standard	9
1.4	Should you be using SWI-Prolog?	9
1.5	The XPCE GUI system for Prolog	10
1.6	Release Notes	11
1.6.1	Version 1.8 Release Notes	11
1.6.2	Version 1.9 Release Notes	12
1.6.3	Version 2.0 Release Notes	12
1.6.4	Version 2.5 Release Notes	12
1.6.5	Version 2.6 Release Notes	13
1.6.6	Version 2.7 Release Notes	13
1.6.7	Version 2.8 Release Notes	13
1.6.8	Version 2.9 Release Notes	13
1.6.9	Version 3.0 Release Notes	14
1.6.10	Version 3.1 Release Notes	14
1.6.11	Version 3.3 Release Notes	14
1.6.12	Version 3.4 Release Notes	15
1.6.13	Version 4.0 Release Notes	16
1.6.14	Version 5.0 Release Notes	16
1.6.15	Version 5.1 Release Notes	16
1.6.16	Version 5.2 Release Notes	16
1.7	Donate to the SWI-Prolog project	16
1.8	Acknowledgements	17
2	Overview	18
2.1	Getting started quickly	18
2.1.1	Starting SWI-Prolog	18
2.1.2	Executing a query	19
2.2	The user's initialisation file	19
2.3	Initialisation files and goals	20
2.4	Command line options	20
2.5	GNU Emacs Interface	23
2.6	Online Help	23
2.7	Query Substitutions	24
2.7.1	Limitations of the History System	25
2.8	Reuse of toplevel bindings	25
2.9	Overview of the Debugger	26

2.10	Compilation	29
2.10.1	During program development	29
2.10.2	For running the result	29
2.11	Environment Control (Prolog flags)	32
2.12	An overview of hook predicates	38
2.13	Automatic loading of libraries	39
2.14	Garbage Collection	41
2.15	Syntax Notes	41
2.15.1	ISO Syntax Support	41
2.16	System limits	43
2.16.1	Limits on memory areas	43
2.16.2	Other Limits	43
2.16.3	Reserved Names	45
3	Initialising and Managing a Prolog Project	46
3.1	The project source-files	46
3.1.1	File Names and Locations	46
3.1.2	Project Special Files	47
3.2	Using modules	48
3.3	The test-edit-reload cycle	49
3.3.1	Locating things to edit	49
3.3.2	Editing and incremental compilation	49
3.4	Using the PceEmacs built-in editor	50
3.4.1	Activating PceEmacs	50
3.4.2	Bluffing through PceEmacs	50
3.4.3	Prolog Mode	52
3.5	The Graphical Debugger	54
3.5.1	Invoking the window-based debugger	54
3.6	The Prolog Navigator	54
3.7	Accessing the IDE from your program	55
3.8	Summary of the iDE	55
4	Built-in predicates	57
4.1	Notation of Predicate Descriptions	57
4.2	Character representation	57
4.3	Loading Prolog source files	58
4.3.1	Quick load files	64
4.4	Listing and Editor Interface	65
4.5	Verify Type of a Term	67
4.6	Comparison and Unification of Terms	67
4.6.1	Standard Order of Terms	67
4.7	Control Predicates	69
4.8	Meta-Call Predicates	71
4.9	ISO compliant Exception handling	73
4.9.1	Debugging and exceptions	74
4.9.2	The exception term	74
4.9.3	Printing messages	75

4.10	Handling signals	76
4.10.1	Notes on signal handling	78
4.11	The 'block' control-structure	78
4.12	DCG Grammar rules	79
4.13	Database	80
4.13.1	Update view	82
4.13.2	Indexing databases	82
4.14	Declaring predicates properties	83
4.15	Examining the program	84
4.16	Input and output	87
4.16.1	ISO Input and Output Streams	87
4.16.2	Edinburgh-style I/O	91
4.16.3	Switching Between Edinburgh and ISO I/O	93
4.17	Status of streams	93
4.18	Primitive character I/O	94
4.19	Term reading and writing	97
4.20	Analysing and Constructing Terms	101
4.21	Analysing and Constructing Atoms	103
4.22	Classifying characters	105
4.22.1	Case conversion	107
4.23	Representing text in strings	107
4.24	Operators	108
4.25	Character Conversion	109
4.26	Arithmetic	110
4.27	Arithmetic Functions	111
4.28	Adding Arithmetic Functions	115
4.29	Built-in list operations	115
4.30	Finding all Solutions to a Goal	117
4.31	Invoking Predicates on all Members of a List	118
4.32	Forall	118
4.33	Formatted Write	118
4.33.1	Writef	118
4.33.2	Format	120
4.33.3	Programming Format	122
4.34	Terminal Control	122
4.35	Operating System Interaction	123
4.35.1	Dealing with time and date	125
4.35.2	Controlling the PLWIN.EXE console window	125
4.36	File System Interaction	126
4.37	User Toplevel Manipulation	130
4.38	Creating a Protocol of the User Interaction	131
4.39	Debugging and Tracing Programs	131
4.40	Obtaining Runtime Statistics	133
4.41	Execution profiling	135
4.41.1	Profiling predicates	135
4.41.2	Visualizing profiling data	136
4.41.3	Information gathering	137

4.42	Memory Management	137
4.43	Windows DDE interface	139
4.43.1	DDE client interface	139
4.43.2	DDE server mode	140
4.44	Miscellaneous	141
5	Using Modules	143
5.1	Why Using Modules?	143
5.2	Name-based versus Predicate-based Modules	143
5.3	Defining a Module	144
5.4	Importing Predicates into a Module	144
5.4.1	Reserved Modules	145
5.5	Using the Module System	145
5.5.1	Object Oriented Programming	146
5.6	Meta-Predicates in Modules	147
5.6.1	Definition and Context Module	147
5.6.2	Overruling Module Boundaries	148
5.7	Dynamic Modules	148
5.8	Module Handling Predicates	149
5.9	Compatibility of the Module System	150
5.9.1	Emulating <code>meta_predicate/1</code>	151
6	Multi-threaded applications	153
6.1	Creating and destroying Prolog threads	153
6.2	Monitoring threads	155
6.3	Thread communication	156
6.3.1	Message queues	156
6.3.2	Signalling threads	158
6.3.3	Threads and dynamic predicates	159
6.4	Thread synchronisation	159
6.5	Thread-support library(<code>threadutil</code>)	161
6.6	Multi-threaded mixed C and Prolog applications	161
6.6.1	A Prolog thread for each native thread (<i>one-to-one</i>)	162
6.6.2	Pooling Prolog engines (<i>many-to-many</i>)	163
6.7	Multithreading and the XPCE graphics system	164
6.8	Threaded Prolog: Status and plans	165
7	Foreign Language Interface	166
7.1	Overview of the Interface	166
7.2	Linking Foreign Modules	166
7.2.1	What linking is provided?	167
7.2.2	What kind of loading should I be using?	167
7.3	Dynamic Linking of shared libraries	167
7.4	Using the library <code>shlib</code> for <code>.DLL</code> and <code>.so</code> files	168
7.4.1	Static Linking	169
7.5	Interface Data types	170
7.5.1	Type <code>term.t</code> : a reference to a Prolog term	170

7.5.2	Other foreign interface types	172
7.6	The Foreign Include File	173
7.6.1	Argument Passing and Control	173
7.6.2	Atoms and functors	174
7.6.3	Analysing Terms via the Foreign Interface	176
7.6.4	Constructing Terms	183
7.6.5	Unifying data	185
7.6.6	Calling Prolog from C	189
7.6.7	Discarding Data	192
7.6.8	Foreign Code and Modules	192
7.6.9	Prolog exceptions in foreign code	193
7.6.10	Catching Signals (Software Interrupts)	195
7.6.11	Miscellaneous	196
7.6.12	Errors and warnings	198
7.6.13	Environment Control from Foreign Code	199
7.6.14	Querying Prolog	199
7.6.15	Registering Foreign Predicates	200
7.6.16	Foreign Code Hooks	201
7.6.17	Storing foreign data	202
7.6.18	Embedding SWI-Prolog in other applications	206
7.7	Linking embedded applications using plld	209
7.7.1	A simple example	211
7.8	The Prolog ‘home’ directory	211
7.9	Example of Using the Foreign Interface	213
7.10	Notes on Using Foreign Code	215
7.10.1	Memory Allocation	215
7.10.2	Compatibility between Prolog versions	215
7.10.3	Debugging Foreign Code	215
7.10.4	Name Conflicts in C modules	216
7.10.5	Compatibility of the Foreign Interface	216
8	Generating Runtime Applications	217
8.1	Limitations of <code>qsave_program</code>	219
8.2	Runtimes and Foreign Code	219
8.3	Using program resources	220
8.3.1	Predicates Definitions	221
8.3.2	The <code>plrc</code> program	222
8.4	Finding Application files	222
8.4.1	Passing a path to the application	223
8.5	The Runtime Environment	223
8.5.1	The Runtime Emulator	223
A	The SWI-Prolog library	225
A.1	<code>lists</code> : List Manipulation	225
A.1.1	Set Manipulation	226
A.2	<code>check</code> : Elementary completeness checks	227
A.3	<code>debug</code> : Some reusable code to help debugging applications	228

A.4	readutil: Reading lines, streams and files	229
A.5	netscape: Activating your Web-browser	230
A.6	registry: Manipulating the Windows registry	230
A.7	url: Analysing and constructing URL	231
B	Hackers corner	234
B.1	Examining the Environment Stack	234
B.2	Intercepting the Tracer	236
B.3	Hooks using the <code>exception/3</code> predicate	237
B.4	Hooks for integrating libraries	237
B.5	Hooks for loading files	238
B.6	Readline Interaction	239
C	Glossary of Terms	240
D	SWI-Prolog License Conditions and Tools	245
D.1	The SWI-Prolog kernel and foreign libraries	245
D.1.1	The SWI-Prolog Prolog libraries	245
D.2	Contributing to the SWI-Prolog project	246
D.3	Software support to keep track of license conditions	246
E	Summary	248
E.1	Predicates	248
E.2	Library predicates	259
E.2.1	<code>lists</code>	259
E.2.2	<code>check</code>	259
E.2.3	<code>readutil</code>	259
E.2.4	<code>netscape</code>	259
E.2.5	<code>registry</code>	259
E.2.6	<code>url</code>	260
E.3	Arithmetic Functions	261
E.4	Operators	263

1

Introduction

1.1 SWI-Prolog

SWI-Prolog started back in 1986 with the requirement for a Prolog that could handle recursive interaction with the C-language: Prolog calling C and C calling Prolog recursively. Those days Prolog systems were very aware of its environment and we needed such a system to support interactive applications. Since then, SWI-Prolog's development has been guided by requests from the user community, especially focussing on (in arbitrary order) interaction with the environment, scalability, (I/O) performance, standard compliance, teaching and the program development environment.

SWI-Prolog is based on a very restricted form of the WAM (Warren Abstract Machine) described in [Bowen & Byrd, 1983] which defines only 7 instructions. Prolog can easily be compiled into this language and the abstract machine code is easily decompiled back into Prolog. As it is also possible to wire a standard 4-port debugger in the WAM interpreter there is no need for a distinction between compiled and interpreted code. Besides simplifying the design of the Prolog system itself this approach has advantages for program development: the compiler is simple and fast, the user does not have to decide in advance whether debugging is required and the system only runs slightly slower when in debug mode. The price we have to pay is some performance degradation (taking out the debugger from the WAM interpreter improves performance by about 20%) and somewhat additional memory usage to help the decompiler and debugger.

SWI-Prolog extends the minimal set of instructions described in [Bowen & Byrd, 1983] to improve performance. While extending this set care has been taken to maintain the advantages of decompilation and tracing of compiled code. The extensions include specialised instructions for unification, predicate invocation, some frequently used built-in predicates, arithmetic, and control (`;/2`, `|/2`), if-then (`->/2`) and negation-by-failure (`\+/1`).

1.1.1 Books about Prolog

This manual does not describe the full syntax and semantics of Prolog, nor how one should write a program in Prolog. These subjects have been described extensively in the literature. See [Bratko, 1986], [Sterling & Shapiro, 1986], and [Clocksin & Melish, 1987]. For more advanced Prolog material see [O'Keefe, 1990]. Syntax and standard operator declarations conform to the 'Edinburgh standard'. Most built in predicates are compatible with those described in [Clocksin & Melish, 1987]. SWI-Prolog also offers a number of primitive predicates compatible with Quintus Prolog¹ [Qui, 1997] and BIM_Prolog² [BIM, 1989].

ISO compliant predicates are based on "Prolog: The Standard", [Deransart *et al.*, 1996], validated using [Hodgson, 1998].

¹Quintus is a trademark of Quintus Computer Systems Inc., USA

²BIM is a trademark of BIM sa/nv., Belgium

1.2 Status

This manual describes version 5.2 of SWI-Prolog. SWI-Prolog has been used now for many years. The application range includes Prolog course material, meta-interpreters, simulation of parallel Prolog, learning systems, natural language processing, complex interactive systems, web-server and web-server components. Although we experienced rather obvious and critical bugs can remain unnoticed for a remarkable long period, we assume the basic Prolog system is fairly stable. Bugs can be expected in infrequently used built-in predicates.

Some bugs are known to the author. They are described as footnotes in this manual.

1.3 Compliance to the ISO standard

SWI-Prolog 3.3.0 implements all predicates described in “Prolog: The Standard” [Deransart *et al.*, 1996].

Exceptions and warning are still weak. Some SWI-Prolog predicates silently fail on conditions where the ISO specification requires an exception (`functor/3` for example). Some predicates print warnings rather than raising an exception. All predicates where exceptions may be caused due to a correct program operating in an imperfect world (I/O, arithmetic, resource overflows) should behave according to the ISO standard. In other words: SWI-Prolog should be able to execute any program conforming to [Deransart *et al.*, 1996] that does not rely on exceptions generated by errors in the program.

1.4 Should you be using SWI-Prolog?

There are a number of reasons why you better choose a commercial Prolog system, or another academic product:

- *SWI-Prolog is not supported*
Although I usually fix bugs shortly after a bug report arrives, I cannot promise anything. Now that the sources are provided, you can always dig into them yourself.
- *Memory requirements and performance are your first concerns*
A number of commercial compilers are more keen on memory and performance than SWI-Prolog. I do not wish to sacrifice some of the nice features of the system, nor its portability to compete on raw performance.
- *You need features not offered by SWI-Prolog*
In this case you may wish to give me suggestions for extensions. If you have great plans, please contact me (you might have to implement them yourself however).

On the other hand, SWI-Prolog offers some nice facilities:

- *Nice environment*
This includes ‘Do What I Mean’, automatic completion of atom names, history mechanism and a tracer that operates on single key-strokes. Interfaces to some standard editors are provided (and can be extended), as well as a facility to maintain programs (see `make/0`).

- *Very fast compiler*
Even very large applications can be loaded in seconds on most machines. If this is not enough, there is a Quick Load Format that is slightly more compact and loading is almost always I/O bound.
- *Transparent compiled code*
SWI-Prolog compiled code can be treated just as interpreted code: you can list it, trace it, etc. This implies you do not have to decide beforehand whether a module should be loaded for debugging or not. Also, performance is much better than the performance of most interpreters.
- *Profiling*
SWI-Prolog offers tools for performance analysis, which can be very useful to optimise programs. Unless you are very familiar with Prolog and Prolog performance considerations this might be more helpful than a better compiler without these facilities.
- *Flexibility*
SWI-Prolog can easily be integrated with C, supporting non-determinism in Prolog calling C as well as C calling Prolog (see section 7. It can also be *embedded* embedded in external programs (see section 7.7). System predicates can be redefined locally to provide compatibility with other Prolog systems.
- *Integration with XPCE*
SWI-Prolog offers a tight integration to the Object Oriented Package for User Interface Development, called XPCE [Anjewierden & Wielemaker, 1989]. XPCE allows you to implement graphical user interfaces that are source-code compatible over Unix/X11, Win32 (Windows 95/98/ME and NT/2000/XP) and MacOS X (darwin).

1.5 The XPCE GUI system for Prolog

The XPCE GUI system for dynamically typed languages has been with SWI-Prolog for a long time. It is developed by Anjo Anjewierden and Jan Wielemaker from the department of SWI, University of Amsterdam. It aims at a high-productive development environment for graphical applications based on Prolog.

Object oriented technology has proven to be a suitable model for implementing GUIs, which typically deal with things Prolog is not very good at: event-driven control and global state. With XPCE, we designed a system that has similar characteristics that make Prolog such a powerful tool: dynamic typing, meta-programming and dynamic modification of the running system.

XPCE is an object-system written in the C-language. It provides for the implementation of methods in multiple languages. New XPCE classes may be defined from Prolog using a simple, natural syntax. The body of the method is executed by Prolog itself, providing a natural interface between the two systems. Below is a very simple class definition.

```
:- pce_begin_class(prolog_lister, frame,
                  "List Prolog predicates").

initialise(Self) :->
    "As the C++ constructor"::
    send_super(Self, initialise, 'Prolog Lister'),
```

```

send(Self, append, new(D, dialog)),
send(D, append,
      text_item(predicate, message(Self, list, @arg1))),
send(new(view), below, D).

list(Self, From:name) :->
  "List predicates from specification":
  ( catch(term_to_atom(Term, From), _, fail)
  -> get(Self, member, view, V),
      current_output(Old),
      pce_open(V, write, Fd),
      set_output(Fd),
      listing(Term),
      close(Fd),
      set_output(Old)
  ; send(Self, report, error, 'Syntax error')
  ).

:- pce_end_class.

test :- send(new(prolog_lister), open).

```

Its 165 built-in classes deal with the meta-environment, data-representation and—of course—graphics. The graphics classes concentrate on direct-manipulation of diagrammatic representations.

Availability. XPCE runs on most Unixtm platforms, Windows 95/98/ME, Windows NT/2000/XP and MacOS X (using X11). In the past, versions for Quintus- and SICStus Prolog as well as some Lisp dialects have existed. After discontinuing active Lisp development at SWI the Lisp versions have died. Active development on the Quintus and SICStus versions has been stopped due to lack of standardisation in the the Prolog community. If adequate standards emerge we are happy to actively support other Prolog implementations.

Info. further information is available from <http://www.swi-prolog.org/packages/xpce/> or by E-mail to info@www.swi-prolog.org.

1.6 Release Notes

Collected release-notes. This section only contains some highlights. Smaller changes to especially older releases have been removed. For a complete log, see the file `ChangeLog` from the distribution.

1.6.1 Version 1.8 Release Notes

Version 1.8 offers a stack-shifter to provide dynamically expanding stacks on machines that do not offer operating-system support for implementing dynamic stacks.

1.6.2 Version 1.9 Release Notes

Version 1.9 offers better portability including an MS-Windows 3.1 version. Changes to the Prolog system include:

- *Redefinition of system predicates*
Redefinition of system predicates was allowed silently in older versions. Version 1.9 only allows it if the new definition is headed by a `:- redefine_system_predicate/1` directive.
- *'Answer' reuse*
The toplevel maintains a table of bindings returned by toplevel goals and allows for reuse of these bindings by prefixing the variables with the \$ sign. See section 2.8.
- *Better source code administration*
Allows for proper updating of multifile predicates and finding the sources of individual clauses.

1.6.3 Version 2.0 Release Notes

New features offered:

- *32-bit Virtual Machine*
Removes various limits and improves performance.
- *Inline foreign functions*
'Simple' foreign predicates no longer build a Prolog stack-frame, but are directly called from the VM. Notably provides a speedup for the test predicates such as `var/1`, etc.
- *Various compatibility improvements*
- *Stream based I/O library*
All SWI-Prolog's I/O is now handled by the stream-package defined in the foreign include file `SWI-Stream.h`. Physical I/O of Prolog streams may be redefined through the foreign language interface, facilitating much simpler integration in window environments.

1.6.4 Version 2.5 Release Notes

Version 2.5 is an intermediate release on the path from 2.1 to 3.0. All changes are to the foreign-language interface, both to user- and system-predicates implemented in the C-language. The aim is twofold. First of all to make garbage-collection and stack-expansion (stack-shifts) possible while foreign code is active without the C-programmer having to worry about locking and unlocking C-variables pointing to Prolog terms. The new approach is closely compatible to the Quintus and SIC-Stus Prolog foreign interface using the `+term` argument specification (see their respective manuals). This allows for writing foreign interfaces that are easily portable over these three Prolog platforms.

Apart from various bug fixes listed in the Changelog file, these are the main changes since 2.1.0:

- *ISO compatibility*
Many ISO compatibility features have been added: `open/4`, arithmetic functions, syntax, etc.
- *Win32*
Many fixes for the Win32 (NT, '95 and win32s) platforms. Notably many problems related to pathnames and a problem in the garbage collector.

- *Performance*
Many changes to the clause indexing system: added hash-tables, lazy computation of the index information, etc.
- *Portable saved-states*
The predicate `qsave_program/[1,2]` allows for the creating of machine independent saved-states that load very quickly.

1.6.5 Version 2.6 Release Notes

Version 2.6 provides a stable implementation of the features added in the 2.5.x releases, but at the same time implements a number of new features that may have impact on the system stability.

- *32-bit integer and double float arithmetic*
The biggest change is the support for full 32-bit signed integers and raw machine-format double precision floats. The internal data representation as well as the arithmetic instruction set and interface to the arithmetic functions has been changed for this.
- *Embedding for Win32 applications*
The Win32 version has been reorganised. The Prolog kernel is now implemented as Win32 DLL that may be embedded in C-applications. Two front ends are provided, one for window-based operation and one to run as a Win32 console application.
- *Creating stand-alone executables*
Version 2.6.0 can create stand-alone executables by attaching the saved-state to the emulator. See `qsave_program/2`.

1.6.6 Version 2.7 Release Notes

Version 2.7 reorganises the entire data-representation of the Prolog data itself. The aim is to remove most of the assumption on the machine's memory layout to improve portability in general and enable embedding on systems where the memory layout may depend on invocation or on how the executable is linked. The latter is notably a problem on the Win32 platforms. Porting to 64-bit architectures is feasible now.

Furthermore, 2.7 lifts the limits on arity of predicates and number of variables in a clause considerably and allow for further expansion at minimal cost.

1.6.7 Version 2.8 Release Notes

With version 2.8, we declare the data-representation changes of 2.7.x stable. Version 2.8 exploits the changes of 2.7 to support 64-bit processors like the DEC Alpha. As of version 2.8.5, the representation of recorded terms has changed, and terms on the heap are now represented in a compiled format. SWI-Prolog no longer limits the use of `malloc()` or uses assumptions on the addresses returned by this function.

1.6.8 Version 2.9 Release Notes

Version 2.9 is the next step towards version 3.0, improving ISO compliance and introducing ISO compliant exception handling. New are `catch/3`, `throw/1`, `abolish/1`, `write_term/[2,3]`,

`write_canonical/[1,2]` and the C-functions `PL_exception()` and `PL_throw()`. The predicates `display/[1,2]` and `displayq/[1,2]` have been moved to `backcomp`, so old code referring to them will autoload them.

The interface to `PL_open_query()` has changed. The *debug* argument is replaced by a bitwise or'ed *flags* argument. The values `FALSE` and `TRUE` have their familiar meaning, making old code using these constants compatible. Non-zero values other than `TRUE` (1) will be interpreted different.

1.6.9 Version 3.0 Release Notes

Complete redesign of the saved-state mechanism, providing the possibility of ‘program resources’. See `resource/3`, `open_resource/3`, and `qsave_program/[1,2]`.

1.6.10 Version 3.1 Release Notes

Improvements on exception-handling. Allows relating software interrupts (signals) to exceptions, handling signals in Prolog and C (see `on_signal/3` and `PL_signal()`). Prolog stack overflows now raise the `resource_error` exception and thus can be handled in Prolog using `catch/3`.

1.6.11 Version 3.3 Release Notes

Version 3.3 is a major release, changing many things internally and externally. The highlights are a complete redesign of the high-level I/O system, which is now based on explicit streams rather than current input/output. The old Edinburgh predicates (`see/1`, `tell/1`, etc.) are now defined on top of this layer instead of the other way around. This fixes various internal problems and removes Prolog limits on the number of streams.

Much progress has been made to improve ISO compliance: handling strings as lists of one-character atoms is now supported (next to character codes as integers). Many more exceptions have been added and printing of exceptions and messages is rationalised using Quintus and SICStus Prolog compatible `print_message/2`, `message_hook/3` and `print_message_lines/3`. All predicates described in [Deransart *et al.*, 1996] are now implemented.

As of version 3.3, SWI-Prolog adheres the ISO *logical update view* for dynamic predicates. See section 4.13.1 for details.

SWI-Prolog 3.3 includes garbage collection on atoms, removing the last serious memory leak especially in text-manipulation applications. See section 7.6.2. In addition, both the user-level and foreign interface supports atoms holding *0-bytes*.

Finally, an alpha version of a multi-threaded SWI-Prolog for Linux is added. This version is still much slower than the single-threaded version due to frequent access to ‘thread-local-data’ as well as some too detailed mutex locks. The basic thread API is ready for serious use and testing however. See section 6.

Incompatible changes

A number of incompatible changes result from this upgrade. They are all easily fixed however.

- `!/0`, `call/1`

The cut now behaves according to the ISO standard. This implies it works in compound goals passed to `call/1` and is local to the *condition* part of if-then-else as well as the argument of `\+/1`.

- `atom_chars/2`
This predicate is now ISO compliant and thus generates a list of one-character atoms. The behaviour of the old predicate is available in the —also ISO compliant— `atom_codes/2` predicate. Safest repair is a replacement of all `atom_chars` into `atom_codes`. If you do not want to change any source-code, you might want to use

```
user:goal_expansion(atom_chars(A,B), atom_codes(A,B)).
```
- `number_chars/2`
Same applies for `number_chars/2` and `number_codes/2`.
- `feature/2`, `set_feature/2`
These are replaced by the ISO compliant `current_prolog_flag/2` and `set_prolog_flag/2`. The library `backcomp` provides definitions for `feature/2` and `set_feature/2`, so no source **has** to be updated.
- *Accessing command-line arguments*
This used to be provided by the undocumented `'$argv'/1` and Quintus compatible library `unix/1`. Now there is also documented `current_prolog_flag(argv, Argv)`.
- `dup_stream/2`
Has been deleted. New stream-aliases can deal with most of the problems for which `dup_stream/2` was designed and `dup/2` from the `clib` package can with most others.
- `op/3`
Operators are now **local to modules**. This implies any modification of the operator-table does not influence other modules. This is consistent with the proposed ISO behaviour and a necessity to have any usable handling of operators in a multi-threaded environment.
- `set_prolog_flag(character_escapes, Bool)`
This prolog flag is now an interface to changing attributes on the current source-module, effectively making this flag module-local as well. This is required for consistent handling of sources written with ISO (obligatory) character-escape sequences together with old Edinburgh code.
- `current_stream/3` and `stream_position`
These predicates have been moved to `quintus`.

1.6.12 Version 3.4 Release Notes

The 3.4 release is a consolidation release. It consolidates the improvements and standard conformance of the 3.3 releases. This version is closely compatible with the 3.3 version except for one important change:

- *Argument order in select/3*
The list-processing predicate `select/3` somehow got into a very early version of SWI-Prolog with the wrong argument order. This has been fixed in 3.4.0. The correct order is `select(?Elem, ?List, ?Rest)`.

As `select/3` has no error conditions, runtime checking cannot be done. To simplify debugging, the library module `checkselect` will print references to `select/3` in your source

code and install a version of `select` that enters the debugger if `select` is called and the second argument is not a list.

This library can be loaded explicitly or by calling `check_old_select/0`.

1.6.13 Version 4.0 Release Notes

As of version 4.0 the standard distribution of SWI-Prolog is bundled with a number of its popular extension packages, among which the now open source XPCE GUI toolkit (see section 1.5). No significant changes have been made to the basic SWI-Prolog engine.

Some useful tricks in the integrated environment:

- *Register the GUI tracer*
Using a call to `gui_tracer/0`, hooks are installed that replace the normal command-line driven tracer with a graphical frontend.
- *Register PceEmacs for editing files*
From your initialisation file, you can load `emacs/swi_prolog` that cause `edit/1` to use the built-in PceEmacs editor.

1.6.14 Version 5.0 Release Notes

Version 5.0 marks a breakpoint in the philosophy, where SWI-Prolog moves from a dual GPL/proprietary to a uniform LGPL (Lesser GNU Public Licence) schema, providing a widely usable Free Source Prolog implementation.

On the technical site the development environment, consisting of source-level debugger, integrated editor and various analysis and navigation tools progress steadily towards a mature set of tools.

Many portability issues have been improved, including a port to MacOS X (Darwin).

For details, please visit the new website at <http://www.swi-prolog.org>

1.6.15 Version 5.1 Release Notes

Version 5.1 is a beta-series introducing portable multi-threading. See chapter 6. In addition it introduces many new facilities to support server applications, such as the new `rlimit` library to limit system resources and the possibility to set timeouts on input streams.

1.6.16 Version 5.2 Release Notes

Version 5.2 consolidates the 5.1.x beta series that introduced threading and many related modifications to the kernel.

1.7 Donate to the SWI-Prolog project

If you are happy with SWI-Prolog, you care it to be around for much longer while it becomes faster, more stable and with more features you should consider to donate to the SWI-Prolog foundation. Please visit the page below.

<http://www.swi-prolog.org/donate.html>

1.8 Acknowledgements

Some small parts of the Prolog code of SWI-Prolog are modified versions of the corresponding Edinburgh C-Prolog code: grammar rule compilation and `writeln/2`. Also some of the C-code originates from C-Prolog: finding the path of the currently running executable and the code underlying `absolute_file_name/2`. Ideas on programming style and techniques originate from C-Prolog and Richard O’Keefe’s *thief* editor. An important source of inspiration are the programming techniques introduced by Anjo Anjewierden in PCE version 1 and 2.

I also would like to thank those who had the fate of using the early versions of this system, suggested extensions or reported bugs. Among them are Anjo Anjewierden, Huub Knops, Bob Wielinga, Wouter Jansweijer, Luc Peerdeman, Eric Nombden, Frank van Harmelen, Bert Rengel.

Martin Jansche (`jansche@novell1.gs.uni-heidelberg.de`) has been so kind to reorganise the sources for version 2.1.3 of this manual.

Horst von Brand has been so kind to fix many typos in the 2.7.14 manual. Thanks!

Overview

2

2.1 Getting started quickly

2.1.1 Starting SWI-Prolog

Starting SWI-Prolog on Unix

By default, SWI-Prolog is installed as ‘pl’, though some administrators call it ‘swipl’ or ‘swi-prolog’. The command-line arguments of SWI-Prolog itself and its utility programs are documented using standard Unix man pages. SWI-Prolog is normally operated as an interactive application simply by starting the program:

```
machine% pl
Welcome to SWI-Prolog (Version 5.2.0)
Copyright (c) 1990-2003 University of Amsterdam.
SWI-Prolog comes with ABSOLUTELY NO WARRANTY. This is free software,
and you are welcome to redistribute it under certain conditions.
Please visit http://www.swi-prolog.org for details.
```

```
For help, use ?- help(Topic). or ?- apropos(Word).
```

```
1 ?-
```

After starting Prolog, one normally loads a program into it using `consult/1`, which—for historical reasons—may be abbreviated by putting the name of the program file between square brackets. The following goal loads the file `likes.pl` containing clauses for the predicates `likes/2`:

```
?- [likes].
% likes compiled, 0.00 sec, 596 bytes.
```

```
Yes
```

```
?-
```

After this point, Unix and Windows users unite, so if you are using Unix please continue at section [2.1.2](#).

Starting SWI-Prolog on Windows

After SWI-Prolog has been installed on a Windows system, the following important new things are available to the user:

- A folder (called *directory* in the remainder of this document) called `pl` containing the executables, libraries, etc. of the system. No files are installed outside this directory.
- A program `plwin.exe`, providing a window for interaction with Prolog. The program `plcon.exe` is a version of SWI-Prolog that runs in a DOS-box.
- The file-extension `.pl` is associated with the program `plwin.exe`. Opening a `.pl` file will cause `plwin.exe` to start, change directory to the directory in which the file-to-open resides and load this file.

The normal way to start with the `likes.pl` file mentioned in section 2.1.1 is by simply double-clicking this file in the Windows explorer.

2.1.2 Executing a query

After loading a program, one can ask Prolog queries about the program. The query below asks Prolog what food 'sam' likes. The system responds with `X = <value>` if it can prove the goal for a certain `X`. The user can type the semi-colon (;)¹ if (s)he wants another solution, or `RETURN` if (s)he is satisfied, after which Prolog will say **Yes**. If Prolog answers **No**, it indicates it cannot find any (more) answers to the query. Finally, Prolog can answer using an error message to indicate the query or program contains an error.

```
?- likes(sam, X).
```

```
X = dahl ;
```

```
X = tandoori ;
```

```
...
```

```
X = chips ;
```

```
No
```

```
?-
```

2.2 The user's initialisation file

After the necessary system initialisation the system consults (see `consult/1`) the user's startup file. The base-name of this file follows conventions of the operating system. On MS-Windows, it is the file `pl.ini` and on Unix systems `.plrc`. The file is searched using the `file_search_path/2` clauses for `user_profile`. The table below shows the default value for this search-path.

	Unix	Windows
local	<code>.</code>	<code>.</code>
home	<code>~</code>	<code>%HOME%</code> or <code>%HOMEDRIVE%\%HOMEPATH%</code>
global		SWI-Home directory or <code>%WINDIR%</code> or <code>%SYSTEMROOT%</code>

¹On most installations, single-character commands are executed without waiting for the `RETURN` key.

After the first startup file is found it is loaded and Prolog stops looking for further startup files. The name of the startup file can be changed with the ‘`-f file`’ option. If *File* denotes an absolute path, this file is loaded, otherwise the file is searched for using the same conventions as for the default startup file. Finally, if *file* is none, no file is loaded.

2.3 Initialisation files and goals

Using commandline arguments (see section 2.4), SWI-Prolog can be forced to load files and execute queries for initialisation purposes or non-interactive operation. The most commonly used options are `-f file` or `-s file` to make Prolog load a file, `-g goal` to define an initialisation goal and `-t goal` to define the *toplevel goal*. The following is a typical example for starting an application directly from the commandline.

```
machine% pl -s load.pl -g go -t halt
```

It tells SWI-Prolog to load `load.pl`, start the application using the *entry-point* `go/0` and —instead of entering the interactive *toplevel*— exit after completing `go/0`. The `-q` may be used to suppress all informational messages.

In MS-Windows, the same can be achieved using a short-cut with appropriately defined commandline arguments. A typically seen alternative is to write a file `run.pl` with content as illustrated below. Double-clicking `run.pl` will start the application.

```
:- [load].           % load program
:- go.              % run it
:- halt.           % and exit
```

Section 2.10.2 discusses further scripting options and chapter 8 discusses the generation of runtime executables. Runtime executables are a mean to deliver executables that do not require the Prolog system.

2.4 Command line options

The full set of command line options is given below:

-help

When given as the only option, it summarises the most important options.

-v

When given as the only option, it summarises the version and the architecture identifier.

-arch

When given as the only option, it prints the architecture identifier (see `current_prolog_flag(arch, Arch)`) and exits. See also `-dump-runtime-variables`.

-dump-runtime-variables

When given as the only option, it prints a sequence of variable settings that can be used in shell-scripts to deal with Prolog parameters. This feature is also used by `plld` (see section 7.7). Below is a typical example of using this feature.

```
eval `pl -dump-runtime-variables`
cc -I$PLBASE/include -L$PLBASE/runtime/$PLARCH ...
```

-q

Set the prolog-flag `verbose` to `silent`, suppressing informational and banner messages.

-Lsize[km]

Give local stack limit (2 Mbytes default). Note that there is no space between the size option and its argument. By default, the argument is interpreted in Kbytes. Postfixing the argument with `m` causes the argument to be interpreted in Mbytes. The following example specifies 32 Mbytes local stack.

```
% pl -L32m
```

A maximum is useful to stop buggy programs from claiming all memory resources. `-L0` sets the limit to the highest possible value. See section 2.16.

-Gsize[km]

Give global stack limit (4 Mbytes default). See `-L` for more details.

-Tsize[km]

Give trail stack limit (4 Mbytes default). This limit is relatively high because trail-stack overflows are not often caused by program bugs. See `-L` for more details.

-Asize[km]

Give argument stack limit (1 Mbytes default). The argument stack limits the maximum nesting of terms that can be compiled and executed. SWI-Prolog does ‘last-argument optimisation’ to avoid many deeply nested structure using this stack. Enlarging this limit is only necessary in extreme cases. See `-L` for more details.

-c file ...

Compile files into an ‘intermediate code file’. See section 2.10.

-o output

Used in combination with `-c` or `-b` to determine output file for compilation.

-O

Optimised compilation. See `current_prolog_flag/2`.

-s file

Use *file* as a script-file. The script file is loaded after the initialisation file specified with the `-f file` option. Unlike `-f file`, using `-s` does not stop Prolog from loading the personal initialisation file.

-f file

Use *file* as initialisation file instead of the default `.plrc` (Unix) or `pl.ini` (Windows). ‘`-f none`’ stops SWI-Prolog from searching for a startup file. This option can be used as an alternative to `-s file` that stops Prolog from loading the personal initialisation file. See also section 2.2.

-F *script*

Selects a startup-script from the SWI-Prolog home directory. The script-file is named `<script>.rc`. The default *script* name is deduced from the executable, taking the leading alphanumerical characters (letters, digits and underscore) from the program-name. `-F none` stops looking for a script. Intended for simple management of slightly different versions. One could for example write a script `iso.rc` and then select ISO compatibility mode using `pl -F iso` or make a link from `iso-pl` to `pl`.

-g *goal*

Goal is executed just before entering the top level. Default is a predicate which prints the welcome message. The welcome message can thus be suppressed by giving `-g true`. *goal* can be a complex term. In this case quotes are normally needed to protect it from being expanded by the Unix shell.

-t *goal*

Use *goal* as interactive toplevel instead of the default goal `prolog/0`. *goal* can be a complex term. If the toplevel goal succeeds SWI-Prolog exits with status 0. If it fails the exit status is 1. This flag also determines the goal started by `break/0` and `abort/0`. If you want to stop the user from entering interactive mode start the application with `'-g goal'` and give `'halt'` as toplevel.

-tty

Unix only. Switches controlling the terminal for allowing single-character commands to the tracer and `get_single_char/1`. By default manipulating the terminal is enabled unless the system detects it is not connected to a terminal or it is running as a GNU-Emacs inferior process. This flag is sometimes required for smooth interaction with other applications.

-nosignals

Inhibit any signal handling by Prolog, a property that is sometimes desirable for embedded applications. This option sets the flag `signals` to `false`. See section 7.6.18 for details.

-x *bootfile*

Boot from *bootfile* instead of the system's default boot file. A bootfile is a file resulting from a Prolog compilation using the `-b` or `-c` option or a program saved using `qsave_program/[1,2]`.

-p *alias=path1[:path2...]*

Define a path alias for `file_search_path`. *alias* is the name of the alias, *path1 ...* is a list of values for the alias. On Windows the list-separator is `;`. On other systems it is `:`. A value is either a term of the form `alias(value)` or `pathname`. The computed aliases are added to `file_search_path/2` using `asserta/1`, so they precede predefined values for the alias. See `file_search_path/2` for details on using this file-location mechanism.

--

Stops scanning for more arguments, so you can pass arguments for your application after this one. See `current_prolog_flag/2` using the flag `argv` for obtaining the commandline arguments.

The following options are for system maintenance. They are given for reference only.

-b *initfile ... -c file ...*

Boot compilation. *initfile ...* are compiled by the C-written bootstrap compiler, *file ...* by the normal Prolog compiler. System maintenance only.

-d *level*

Set debug level to *level*. Only has effect if the system is compiled with the `-DO_DEBUG` flag. System maintenance only.

2.5 GNU Emacs Interface

The default Prolog mode for GNU-Emacs can be activated by adding the following rules to your Emacs initialisation file:

```
(setq auto-mode-alist
      (append
        '(("\.pl" . prolog-mode))
        auto-mode-alist))
(setq prolog-program-name "pl")
(setq prolog-consult-string "[user].\n")
;If you want this. Indentation is either poor or I don't use
;it as intended.
;(setq prolog-indent-width 8)
```

Unfortunately the default Prolog mode of GNU-Emacs is not very good. An alternative `prolog.el` file for GNU-Emacs 20 is available from <http://www.freesoft.cz/pdm/software/emacs/prolog-mode/> and for GNU-Emacs 19 from <http://w1.858.telia.com/u85810764/Prolog-mode/index.html>

2.6 Online Help

Online help provides a fast lookup and browsing facility to this manual. The online manual can show predicate definitions as well as entire sections of the manual.

The online help is displayed from the file `'MANUAL'`. The file `helpidx` provides an index into this file. `'MANUAL'` is created from the \LaTeX sources with a modified version of `dvitty`, using `overstrike` for printing bold text and `underlining` for rendering italic text. XPCe is shipped with `swi_help`, presenting the information from the online help in a hypertext window. The `prolog`-flag `write_help_with_overstrike` controls whether or not `help/1` writes its output using `overstrike` to realise bold and underlined output or not. If this `prolog`-flag is not set it is initialised by the help library to `true` if the `TERM` variable equals `xterm` and `false` otherwise. If this default does not satisfy you, add the following line to your personal startup file (see section 2.2):

```
:- set_prolog_flag(write_help_with_overstrike, true).
```

help

Equivalent to `help(help/1)`.

help(+What)

Show specified part of the manual. *What* is one of:

$\langle Name \rangle / \langle Arity \rangle$	Give help on specified predicate
$\langle Name \rangle$	Give help on named predicate with any arity or C interface function with that name
$\langle Section \rangle$	Display specified section. Section numbers are dash-separated numbers: 2-3 refers to section 2.3 of the manual. Section numbers are obtained using <code>apropos/1</code> .

Examples:

```
?- help(assert).      Give help on predicate assert
?- help(3-4).        Display section 3.4 of the manual
?- help('PL_retry'). Give help on interface function PL_retry()
```

See also `apropos/1`, and the SWI-Prolog home page at <http://www.swi-prolog.org>, which provides a FAQ, an HTML version of manual for online browsing and HTML and PDF versions for downloading.

apropos(+Pattern)

Display all predicates, functions and sections that have *Pattern* in their name or summary description. Lowercase letters in *Pattern* also match a corresponding uppercase letter. Example:

```
?- apropos(file).   Display predicates, functions and sections that have 'file'
                   (or 'File', etc.) in their summary description.
```

explain(+ToExplain)

Give an explanation on the given 'object'. The argument may be any Prolog data object. If the argument is an atom, a term of the form *Name/Arity* or a term of the form *Module:Name/Arity*, `explain` will try to explain the predicate as well as possible references to it.

explain(+ToExplain, -Explanation)

Unify *Explanation* with an explanation for *ToExplain*. Backtracking yields further explanations.

2.7 Query Substitutions

SWI-Prolog offers a query substitution mechanism similar to that of Unix `cs` (`cs(1)`), called 'history'. The availability of this feature is controlled by `set_prolog_flag/2`, using the `history` `prolog-flag`. By default, `history` is available if the `prolog-flag` `readline` is `false`. To enable this feature, remembering the last 50 commands, put the following into your startup file (see section 2.2:

```
:- set_prolog_flag(history, 50).
```

The history system allows the user to compose new queries from those typed before and remembered by the system. It also allows to correct queries and syntax errors. SWI-Prolog does not offer the Unix `cs` capabilities to include arguments. This is omitted as it is unclear how the first, second, etc. argument should be defined.²

The available history commands are shown in table 2.1.

²One could choose words, defining words as a sequence of alpha-numeric characters and the word separators as anything else, but one could also choose Prolog arguments

!!.	Repeat last query
!nr.	Repeat query numbered $\langle nr \rangle$
!str.	Repeat last query starting with $\langle str \rangle$
!?str.	Repeat last query holding $\langle str \rangle$
^old^new.	Substitute $\langle old \rangle$ into $\langle new \rangle$ in last query
!nr^old^new.	Substitute in query numbered $\langle nr \rangle$
!str^old^new.	Substitute in query starting with $\langle str \rangle$
!?str^old^new.	Substitute in query holding $\langle str \rangle$
h.	Show history list
!h.	Show this list

Table 2.1: History commands

```

1 ?- maplist(plus(1), "hello", X).

X = [105,102,109,109,112]

Yes
2 ?- format('~s~n', [$X]).
ifmmp

Yes
3 ?-

```

Figure 2.1: Reusing toplevel bindings

2.7.1 Limitations of the History System

History expansion is executed after *raw-reading*. This is the first stage of `read_term/2` and friends, reading the term into a string while deleting comment and canonising blank. This makes it hard to use it for correcting syntax errors. Command-line editing as provided using the GNU-readline library is more suitable for this. History expansion is first of all useful for executing or combining commands from long ago.

2.8 Reuse of toplevel bindings

Bindings resulting from the successful execution of a toplevel goal are asserted in a database. These values may be reused in further toplevel queries as `$Var`. Only the latest binding is available. Example:

Note that variables may be set by executing `=/2`:

```

6 ?- X = statistics.

X = statistics

Yes
7 ?- $X.

```

```

28.00 seconds cpu time for 183,128 inferences
4,016 atoms, 1,904 functors, 2,042 predicates, 52 modules
55,915 byte codes; 11,239 external references

```

	Limit	Allocated	In use
Heap :			624,820 Bytes
Local stack :	2,048,000	8,192	404 Bytes
Global stack :	4,096,000	16,384	968 Bytes
Trail stack :	4,096,000	8,192	432 Bytes

```

Yes
8 ?-

```

2.9 Overview of the Debugger

SWI-Prolog has a 6-port tracer, extending the standard 4-port tracer [Clocksin & Melish, 1987] with two additional ports. The optional *unify* port allows the user to inspect the result after unification of the head. The *exception* port shows exceptions raised by `throw/1` or one of the built-in predicates. See section 4.9.

The standard ports are called `call`, `exit`, `redo`, `fail` and `unify`. The tracer is started by the `trace/0` command, when a spy point is reached and the system is in debugging mode (see `spy/1` and `debug/0`) or when an exception is raised.

The interactive toplevel goal `trace/0` means “trace the next query”. The tracer shows the port, displaying the port name, the current depth of the recursion and the goal. The goal is printed using the Prolog predicate `write_term/2`. The style is defined by the prolog-flag `debugger_print_options` and can be modified using this flag or using the `w`, `p` and `d` commands of the tracer.

On *leashed ports* (set with the predicate `leash/1`, default are `call`, `exit`, `redo` and `fail`) the user is prompted for an action. All actions are single character commands which are executed **without** waiting for a return, unless the command line option `-tty` is active. Tracer options:

+ (Spy)

Set a spy point (see `spy/1`) on the current predicate.

- (No spy)

Remove the spy point (see `nospy/1`) from the current predicate.

/ (Find)

Search for a port. After the `/`, the user can enter a line to specify the port to search for. This line consists of a set of letters indicating the port type, followed by an optional term, that should unify with the goal run by the port. If no term is specified it is taken as a variable, searching for any port of the specified type. If an atom is given, any goal whose functor has a name equal to that atom matches. Examples:

```

1 ?- visible(+all), leash(-exit).

Yes
2 ?- trace, min([3, 2], X).
   Call: ( 3) min([3, 2], G235) ? creep
   Unify: ( 3) min([3, 2], G235)
   Call: ( 4) min([2], G244) ? creep
   Unify: ( 4) min([2], 2)
   Exit: ( 4) min([2], 2)
   Call: ( 4) min(3, 2, G235) ? creep
   Unify: ( 4) min(3, 2, G235)
   Call: ( 5) 3 < 2 ? creep
   Fail: ( 5) 3 < 2 ? creep
   Redo: ( 4) min(3, 2, G235) ? creep
   Exit: ( 4) min(3, 2, 2)
   Exit: ( 3) min([3, 2], 2)

Yes
[trace] 3 ?-

```

Figure 2.2: Example trace

/f	Search for any fail port
/fe solve	Search for a fail or exit port of any goal with name solve
/c solve(a, _)	Search for a call to solve/2 whose first argument is a variable or the atom a
/a member(_, _)	Search for any port on member/2. This is equivalent to setting a spy point on member/2.

. (Repeat find)

Repeat the last find command (see ‘/’).

A (Alternatives)

Show all goals that have alternatives.

C (Context)

Toggle ‘Show Context’. If on the context module of the goal is displayed between square brackets (see section 5). Default is off.

L (Listing)

List the current predicate with `listing/1`.

a (Abort)

Abort Prolog execution (see `abort/0`).

b (Break)

Enter a Prolog break environment (see `break/0`).

c (Creep)

Continue execution, stop at next port. (Also return, space).

d (Display)

Set the `max_depth(Depth)` option of `debugger_print_options`, limiting the depth to which terms are printed. See also the `w` and `p` options.

e (Exit)

Terminate Prolog (see `halt/0`).

f (Fail)

Force failure of the current goal.

g (Goals)

Show the list of parent goals (the execution stack). Note that due to tail recursion optimization a number of parent goals might not exist any more.

h (Help)

Show available options (also ‘?’).

i (Ignore)

Ignore the current goal, pretending it succeeded.

l (Leap)

Continue execution, stop at next spy point.

n (No debug)

Continue execution in ‘no debug’ mode.

p (Print)

Set the `prolog-flag debugger_print_options` to `[quoted(true), portray(true), max_depth(10)]`. This is the default.

r (Retry)

Undo all actions (except for database and i/o actions) back to the call port of the current goal and resume execution at the call port.

s (Skip)

Continue execution, stop at the next port of **this** goal (thus skipping all calls to children of this goal).

u (Up)

Continue execution, stop at the next port of **the parent** goal (thus skipping this goal and all calls to children of this goal). This option is useful to stop tracing a failure driven loop.

w (Write)

Set the `prolog-flag debugger_print_options` to `[quoted(true)]`, bypassing `portray/1`, etc.

The ideal 4 port model as described in many Prolog books [[Clocksin & Melish, 1987](#)] is not visible in many Prolog implementations because code optimisation removes part of the choice- and

exit-points. Backtrack points are not shown if either the goal succeeded deterministically or its alternatives were removed using the cut. When running in debug mode (`debug/0`) choice points are only destroyed when removed by the cut. In debug mode, tail recursion optimisation is switched off.³

Reference information to all predicates available for manipulating the debugger is in section 4.39.

2.10 Compilation

2.10.1 During program development

During program development, programs are normally loaded using `consult/1`, or the list abbreviation. It is common practice to organise a project as a collection of source-files and a *load-file*, a Prolog file containing only `use_module/[1,2]` or `ensure_loaded/1` directives, possibly with a definition of the *entry-point* of the program, the predicate that is normally used to start the program. This file is often called `load.pl`. If the entry-point is called `go`, a typical session starts as:

```
% pl
<banner>

1 ?- [load].
<compilation messages>

Yes
2 ?- go.
<program interaction>
```

When using Windows, the user may open `load.pl` from the Windows explorer, which will cause `plwin.exe` to be started in the directory holding `load.pl`. Prolog loads `load.pl` before entering the toplevel.

2.10.2 For running the result

There are various options if you want to make your program ready for real usage. The best choice depends on whether the program is to be used only on machines holding the SWI-Prolog development system, the size of the program and the operating system (Unix vs. Windows).

Using PrologScript

New in version 4.0.5 is the possibility to use a Prolog source file directly as a Unix script-file. the same mechanism is useful to specify additional parameters for running a Prolog file on Windows.

If the first letter of a Prolog file is `#`, the first line is treated as comment.⁴ To create a Prolog script, make the first line start like this:

```
#!/path/to/pl <options> -s
```

³This implies the system can run out of local stack in debug mode, while no problems arise when running in non-debug mode.

⁴The `#`-sign can be the legal start of a normal Prolog clause. In the unlikely case this is required, leave the first line blank or add a header-comment.

Prolog recognises this starting sequence and causes the interpreter to receive the following argument-list:

```
/path/to/pl <options> -s <script> -- <ScriptArguments>
```

Instead of `-s`, the user may use `-f` to stop Prolog from looking for a personal initialisation file. Here is a simple script doing expression evaluation:

```
#!/usr/bin/pl -q -t main -f

eval :-
    current_prolog_flag(argv, Argv),
    append(_, [--|Args], Argv),
    concat_atom(Args, ' ', SingleArg),
    term_to_atom(Term, SingleArg),
    Val is Term,
    format('~w~n', [Val]).

main :-
    catch(eval, E, (print_message(error, E), fail)),
    halt.

main :-
    halt(1).
```

And here are two example runs:

```
% eval 1+2
3
% eval foo
ERROR: Arithmetic: `foo/0' is not a function
%
```

The Windows version supports the `#!` construct too, but here it serves a rather different role. The Windows shell already allows the user to start Prolog source-files directly through the Windows file-type association. Windows however makes it rather complicated to provide additional parameters, such as the required stack-size for an individual Prolog file. The `#!` line provides for this, providing a more flexible approach than changing the global defaults. The following starts Prolog with unlimited stack-size on the given source-file:

```
#!/usr/bin/pl -L0 -T0 -G0 -s
```

....

Note the use of `/usr/bin/pl`, which specifies the interpreter. This argument is ignored in the Windows version, but required to ensure best cross-platform compatibility.

Creating a shell-script

With the introduction of *PrologScript* (see section 2.10.2), using shell-scripts as explained in this section has become redundant for most applications.

Especially on Unix systems and not-too-large applications, writing a shell-script that simply loads your application and calls the entry-point is often a good choice. A skeleton for the script is given below, followed by the Prolog code to obtain the program arguments.

```
#!/bin/sh

base=<absolute-path-to-source>
PL=pl

exec $PL -f none -g "load_files(['$base/load'],[silent(true)])" \
    -t go -- $*

go :-
    current_prolog_flag(argv, Arguments),
    append(_SystemArgs, [--|Args], Arguments), !,
    go(Args).

go(Args) :-
    ...
```

On Windows systems, similar behaviour can be achieved by creating a shortcut to Prolog, passing the proper options or writing a .bat file.

Creating a saved-state

For larger programs, as well as for programs that are required to run on systems that do not have the SWI-Prolog development system installed, creating a saved state is the best solution. A saved state is created using `qsave_program/[1,2]` or using the linker `plld(1)`. A saved state is a file containing machine-independent intermediate code in a format dedicated for fast loading. Optionally, the emulator may be integrated in the saved state, creating a single-file, but machine-dependent, executable. This process is described in chapter 8.

Compilation using the -c commandline option

This mechanism loads a series of Prolog source files and then creates a saved-state as `qsave_program/2` does. The command syntax is:

```
% pl [option ...] [-o output] -c file ...
```

The *options* argument are options to `qsave_program/2` written in the format below. The option-names and their values are described with `qsave_program/2`.

```
--option-name=option-value
```

For example, to create a stand-alone executable that starts by executing `main/0` and for which the source is loaded through `load.pl`, use the command

```
% pl --goal=main --stand_alone=true -o myprog -c load.pl
```

This performs exactly the same as executing

```
% pl
<banner>
?- [load].
?- qsave_program(myprog,
                 [ goal(main),
                   stand_alone(true)
                 ]).
?- halt.
```

2.11 Environment Control (Prolog flags)

The predicates `current_prolog_flag/2` and `set_prolog_flag/2` allow the user to examine and modify the execution environment. It provides access to whether optional features are available on this version, operating system, foreign-code environment, command-line arguments, version, as well as runtime flags to control the runtime behaviour of certain predicates to achieve compatibility with other Prolog environments.

current_prolog_flag(?Key, -Value)

The predicate `current_prolog_flag/2` defines an interface to installation features: options compiled in, version, home, etc. With both arguments unbound, it will generate all defined prolog-flags. With the 'Key' instantiated it unify the value of the prolog-flag. Features come in three types: boolean prolog-flags, prolog-flags with an atom value and prolog-flags with an integer value. A boolean prolog-flag is true iff the prolog-flag is present **and** the *Value* is the atom `true`. Currently defined keys in alphabetical order:

abort_with_exception (*bool, changeable*)

Determines how `abort/0` is realised. See the description of `abort/0` for details.

agc_margin (*integer, changeable*)

If this amount of atoms has been created since the last atom-garbage collection, perform atom garbage collection at the first opportunity. Initial value is 10,000. May be changed. A value of 0 (zero) disables atom garbage collection. See also `PL_register_atom()`.

allow_variable_name_as_functor (*bool, changeable*)

If true (default is false), `Functor(arg)` is read as if it was written '`Functor'(arg)`'. Some applications use the Prolog `read/1` predicate for reading an application defined script language. In these cases, it is often difficult to explain to non-Prolog users of the application that constants and functions can only start with a lowercase letter. Variables can be turned into atoms starting with an uppercase atom by calling `read_term/2` using the option `variable_names` and binding the variables to their name. Using this feature, `F(x)` can be turned into valid syntax for such script languages. Suggested by Robert van Engelen. SWI-Prolog specific.

argv (*list*)

List is a list of atoms representing the command-line arguments used to invoke SWI-Prolog. Please note that **all** arguments are included in the list returned.

arch (*atom*)

Identifier for the hardware and operating system SWI-Prolog is running on. Used to select foreign files for the right architecture. See also section 7.4 and `file_search_path/2`.

associate (*atom, changeable*)

On Windows systems, this is set to the filename-extension (e.g. `pl` or `pro` associated with `plwin.exe`).

autoload (*bool, changeable*)

If `true` (default) autoloading of library functions is enabled. See section 2.13.

backquoted_string (*bool, changeable*)

If `true` (default `false`), `read` translates text between backquotes into a string object (see section 4.23). This flag is mainly for compatibility to LPA Prolog.

bounded (*true*)

ISO prolog-flag describing integer representation is bound by `min_integer` and `max_integer`.

c_cc (*atom*)

Name of the C-compiler used to compile SWI-Prolog. Normally either `gcc` or `cc`. See section 7.7.

c_ldflags (*atom*)

Special linker flags passed to link SWI-Prolog. See section 7.7.

c_libs (*atom*)

Libraries passed to the C-linker when SWI-Prolog was linked. May be used to determine the libraries needed to create statically linked extensions for SWI-Prolog. See section 7.7.

char_conversion (*bool, changeable*)

Determines whether character-conversion takes place while reading terms. See also `char_conversion/2`.

character_escapes (*bool, changeable*)

If `true` (default), `read/1` interprets `\` escape sequences in quoted atoms and strings. May be changed. This flag is local to the module in which it is changed.

compiled_at (*atom*)

Describes when the system has been compiled. Only available if the C-compiler used to compile SWI-Prolog provides the `__DATE__` and `__TIME__` macros.

console_menu (*bool*)

Set to `true` in `plwin.exe` to indicate the console supports menus. See also section 4.35.2.

dde (*bool*)

Set to `true` if this instance of Prolog supports DDE as described in section 4.43.

debug (*bool, changeable*)

Switch debugging mode on/off. If debug mode is activated the system traps encountered spy-points (see `spy/1`) and trace-points (see `trace/1`). In addition, tail-recursion optimisation is disabled and the system is more conservative in destroying choice-points to simplify debugging.

Disabling these optimisations can cause the system to run out of memory on programs that behave correctly if debug mode is off.

debug_on_error (*bool, changeable*)

If `true`, start the tracer after an error is detected. Otherwise just continue execution. The goal that raised the error will normally fail. See also `fileerrors/2` and the prolog-flag `report_error`. May be changed. Default is `true`, except for the runtime version.

debugger_print_options (*term, changeable*)

This argument is given as `option-list` to `write_term/2` for printing goals by the debugger. Modified by the ‘w’, ‘p’ and ‘<N> d’ commands of the debugger. Default is `[quoted(true), portray(true), max_depth(10)]`.

debugger_show_context (*bool, changeable*)

If `true`, show the context module while printing a stack-frame in the tracer. Normally controlled using the ‘C’ option of the tracer.

double_quotes (*codes,chars,atom,string, changeable*)

This flag determines how double-quotes strings are read by Prolog and is —like `character_escapes`— maintained for each module. If `codes` (default), a list of character-codes is returned, if `chars` a list of one-character atoms, if `atom` double quotes are the same as single-quotes and finally, `string` reads the text into a Prolog string (see section 4.23). See also `atom_chars/2` and `atom_codes/2`.

dynamic_stacks (*bool*)

If `true`, the system uses some form of ‘sparse-memory management’ to realise the stacks. If `false`, `malloc()/realloc()` are used for the stacks. In earlier days this had consequences for foreign code. As of version 2.5, this is no longer the case.

Systems using ‘sparse-memory management’ are a bit faster as there is no stack-shifter, and checking the stack-boundary is often realised by the hardware using a ‘guard-page’. Also, memory is actually returned to the system after a garbage collection or call to `trim_stacks/0` (called by `prolog/0` after finishing a user-query).

editor (*atom, changeable*)

Determines the editor used by `edit/1`. See section 4.4 for details on selecting the editor used.

emacs_inferior_process (*bool*)

If `true`, SWI-Prolog is running as an *inferior process* of (GNU/X-)Emacs. SWI-Prolog assumes this is the case if the environment variable `EMACS` is `t` and `INFERIOR` is `yes`.

executable (*atom*)

Path-name of the running executable. Used by `qsave_program/2` as default emulator.

file_name_variables (*bool, changeable*)

If `true` (default `false`), expand `$varname` and `~` in arguments of builtin-predicates that accept a file name (`open/3`, `exists_file/1`, `access_file/2`, etc.). The predicate `expand_file_name/2` should be used to expand environment variables and wildcard patterns. This prolog-flag is intended for backward compatibility with older versions of SWI-Prolog.

float_format (*atom, changeable*)

C-library `printf()` format specification used by `write/1` and friends to determine how floating point numbers are printed. The default is `%g`. The specified value is passed to

printf() without further checking. For example, if you want more digits printed, `% .12g` will print all floats using 12 digits instead of the default 6.

When using quoted-write, the output is guaranteed to contain a decimal dot or exponent, so `read/1` reads a floating point number. See also `format/[1,2]`, `write_term/[2,3]`.

gc (*bool, changeable*)

If true (default), the garbage collector is active. If false, neither garbage-collection, nor stack-shifts will take place, even not on explicit request. May be changed.

gui (*bool*)

Set to `true` if XPCE is around and can be used for graphics.

history (*integer, changeable*)

If *integer* > 0, support Unix `cs(1)` like history as described in section 2.7. Otherwise, only support reusing commands through the commandline editor. The default is to set this prolog-flag to 0 if a commandline editor is provided (see prolog-flag `readline`) and 15 otherwise.

home (*atom*)

SWI-Prolog's notion of the home-directory. SWI-Prolog uses its home directory to find its startup file as `<home>/boot32.prc` (32-bit machines) or `<home>/boot64.prc` (64-bit machines) and to find its library as `<home>/library`.

hwnd (*integer*)

In `plwin.exe`, this refers to the MS-Windows window-handle of the console window.

integer_rounding_function (*down,toward_zero*)

ISO prolog-flag describing rounding by `//` and `rem` arithmetic functions. Value depends on the C-compiler used.

iso (*bool, changeable*)

Include some weird ISO compatibility that is incompatible to normal SWI-Prolog behaviour. Currently it has the following effect:

- `is/2` and evaluation under `flag/3` do not automatically convert floats to integers if the float represents an integer.
- The `//2` (float division) *always* return a float, even if applied to integers that can be divided.
- In the standard order of terms (see section 4.6.1), all floats are before all integers.
- `atom_length/2` yields an instantiation error if the first argument is a number.
- `clause/[2,3]` raises a permission error when accessing static predicates.
- `abolish/[1,2]` raises a permission error when accessing static predicates.

max_arity (*unbounded*)

ISO prolog-flag describing there is no maximum arity to compound terms.

max_integer (*integer*)

Maximum integer value. Most arithmetic operations will automatically convert to floats if integer values above this are returned.

max_tagged_integer (*integer*)

Maximum integer value represented as a 'tagged' value. Tagged integers require 4-bytes storage and are used for indexing. Larger integers are represented as 'indirect data' and require 16-bytes on the stacks (though a copy requires only 4 additional bytes).

min_integer (*integer*)

Minimum integer value.

min_tagged_integer (*integer*)

Start of the tagged-integer value range.

open_shared_object (*bool*)

If true, `open_shared_object/2` and friends are implemented, providing access to shared libraries (`.so` files) or dynamic link libraries (`.DLL` files).

optimise (*bool, changeable*)

If true, compile in optimised mode. The initial value is true if Prolog was started with the `-O` commandline option.

Currently optimise compilation implies compilation of arithmetic, and deletion of redundant `true/0` that may result from `expand_goal/2`.

Later versions might imply various other optimisations such as integrating small predicates into their callers, eliminating constant expressions and other predictable constructs. Source code optimisation is never applied to predicates that are declared dynamic (see `dynamic/1`).

pid (*int*)

Process identifier of the running Prolog process. Existence of this flag is implementation dependent.

pipe (*bool, changeable*)

If true, `open(pipe(command), mode, Stream)`, etc. are supported. Can be changed to disable the use of pipes in applications testing this feature. Not recommended.

readline (*bool*)

If true, SWI-Prolog is linked with the readline library. This is done by default if you have this library installed on your system. It is also true for the Win32 `plwin.exe` version of SWI-Prolog, which realises a subset of the readline functionality.

resource_database (*atom*)

Set to the absolute-filename of the attached state. Typically this is the file `boot32.prc`, the file specified with `-x` or the running executable. See also `resource/3`.

report_error (*bool, changeable*)

If true, print error messages, otherwise suppress them. May be changed. See also the `debug_on_error` prolog-flag. Default is true, except for the runtime version.

runtime (*bool*)

If true, SWI-Prolog is compiled with `-DO_RUNTIME`, disabling various useful development features (currently the tracer and profiler).

saved_program (*bool*)

If true, Prolog is started from a state saved with `qsave_program/[1,2]`.

shared_object_extension (*atom*)

Extension used by the operating system for shared objects. `.so` for most Unix systems and `.dll` for Windows. Used for locating files using the `file_type` executable. See also `absolute_file_name/3`.

signals (*bool*)

Determine whether Prolog is handling signals (software interrupts). This flag is

`false` if the hosting OS does not support signal handling or the command-line option `-nosignals` is active. See section 7.6.18 for details.

system_thread_id (*int*)

On MT systems (section 6, refers to the thread-identifier used by the system for the calling thread. See also `thread_self/1`.

tail_recursion_optimisation (*bool, changeable*)

Determines whether or not tail-recursion optimisation is enabled. Normally the value of this flag is equal to the `debug` flag. As programs may run out of stack if tail-recursion optimisation is omitted, it is sometimes necessary to enable it during debugging.

toplevel_print_anon (*bool, changeable*)

If `true`, `toplevel` variables starting with an underscore (`_`) are printed normally. If `false` they are hidden. This may be used to hide bindings in complex queries from the `toplevel`.

toplevel_print_options (*term, changeable*)

This argument is given as `option-list` to `write_term/2` for printing results of queries. Default is `[quoted(true), portray(true), max_depth(10)]`.

toplevel_var_size (*int, changeable*)

Maximum size counted in literals of a term returned as a binding for a variable in a `toplevel` query that is saved for re-use using the `$` variable reference. See section 2.8.

trace_gc (*bool, changeable*)

If `true` (`false` is the default), garbage collections and stack-shifts will be reported on the terminal. May be changed.

tty_control (*bool*)

Determines whether the terminal is switched to raw mode for `get_single_char/1`, which also reads the user-actions for the trace. May be set. See also the `+/-tty` command-line option.

unix (*bool*)

If `true`, the operating system is some version of Unix. Defined if the C-compiler used to compile this version of SWI-Prolog either defines `__unix__` or `unix`.

unknown (*fail,warning,error, changeable*)

Determines the behaviour if an undefined procedure is encountered. If `fail`, the predicates fails silently. If `warn`, a warning is printed, and execution continues as if the predicate was not defined and if `error` (default), an `existence_error` exception is raised. This flag is local to each module.

verbose (*Atom, changeable*)

This flag is used by `print_message/2`. If its value is `silent`, messages of type `informational` and `banner` are suppressed. The `-q` switches the value from the initial `normal` to `silent`.

verbose_autoload (*bool, changeable*)

If `true` the normal `consult` message will be printed if a library is autoloaded. By default this message is suppressed. Intended to be used for debugging purposes.

verbose_file_search (*bool, changeable*)

If `true` (default `false`), print messages indicating the progress of `absolute_file_name/[2,3]` in locating files. Intended for debugging complicated file-search paths. See also `file_search_path/2`.

version (*integer*)

The version identifier is an integer with value:

$$10000 \times Major + 100 \times Minor + Patch$$

Note that in releases up to 2.7.10 this prolog-flag yielded an atom holding the three numbers separated by dots. The current representation is much easier for implementing version-conditional statements.

windows (*bool*)

If `true`, the operating system is an implementation of Microsoft Windows (3.1, 95/98/ME, NT/2000/XP, etc.).

write_help_with_overstrike (*bool*)

Internal flag used by `help/1` when writing to a terminal. If `true` it prints bold and underlined text using *overstrike*.

xpce (*bool*)

Set to `true` if the XPCE graphics system is loaded.

xpce_version (*atom*)

Set to the version of the loaded XPCE system.

set_prolog_flag(+Key, +Value)

Define a new prolog-flag or change its value. *Key* is an atom. If the flag is a system-defined flag that is not marked *changeable* above, an attempt to modify the flag yields a `permission_error`. If the provided *Value* does not match the type of the flag, a `type_error` is raised.

In addition to ISO, SWI-Prolog allows for user-defined prolog flags. The type of the flag is determined from the initial value and cannot be changed afterwards.

2.12 An overview of hook predicates

SWI-Prolog provides a large number of hooks, mainly to control handling messages, debugging, startup, shut-down, macro-expansion, etc. Below is a summary of all defined hooks with an indication of their portability.

- `portray/1`
Hook into `write_term/3` to alter the way terms are printed (ISO).
- `message_hook/3`
Hook into `print_message/2` to alter the way system messages are printed (Quintus/SICStus).
- `library_directory/1`
Hook into `absolute_file_name/3` to define new library directories. (most Prolog system).
- `file_search_path/2`
Hook into `absolute_file_name/3` to define new search-paths (Quintus/SICStus).

- `term_expansion/2`
Hook into `load_files/1` to modify read terms before they are compiled (macro-processing) (most Prolog system).
- `goal_expansion/2`
Same as `term_expansion/2` for individual goals (SICStus).
- `prolog_load_file/2`
Hook into `load_files/2` to load other data-formats for Prolog sources from ‘non-file’ resources. The `load_files/2` predicate is the ancestor of `consult/1`, `use_module/1`, etc.
- `prolog_edit:locate/3`
Hook into `edit/1` to locate objects (SWI).
- `prolog_edit:edit_source/1`
Hook into `edit/1` to call some internal editor (SWI).
- `prolog_edit:edit_command/2`
Hook into `edit/1` to define the external editor to use (SWI).
- `prolog_list_goal/1`
Hook into the tracer to list the code associated to a particular goal (SWI).
- `prolog_trace_interception/4`
Hook into the tracer to handle trace-events (SWI).
- `prolog:debug_control_hook/1`
Hook in `spy/1`, `nosp/1`, `nosp/0` and `debugging/0` to extend these control-predicates to higher-level libraries.
- `prolog:help_hook/1`
Hook in `help/0`, `help/1` and `apropos/1` to extend the help-system.
- `resource/3`
Defines a new resource (not really a hook, but similar) (SWI).
- `exception/3`
Old attempt to a generic hook mechanism. Handles undefined predicates (SWI).

2.13 Automatic loading of libraries

If —at runtime— an undefined predicate is trapped the system will first try to import the predicate from the module’s default module. If this fails the *auto loader* is activated. On first activation an index to all library files in all library directories is loaded in core (see `library_directory/1` and `file_search_path/2`). If the undefined predicate can be located in the one of the libraries that library file is automatically loaded and the call to the (previously undefined) predicate is restarted. By default this mechanism loads the file silently. The `current_prolog_flag/2` `verbose_autoload` is provided to get verbose loading. The `prolog-flag` `autoload` can be used to enable/disable the entire auto load system.

The auto-loader only works if the unknown flag (see `unknown/2`) is set to `trace` (default). A more appropriate interaction with this flag should be considered.

Autoloading only handles (library) source files that use the module mechanism described in chapter 5. The files are loaded with `use_module/2` and only the trapped undefined predicate will be imported to the module where the undefined predicate was called. Each library directory must hold a file `INDEX.pl` that contains an index to all library files in the directory. This file consists of lines of the following format:

```
index(Name, Arity, Module, File).
```

The predicate `make/0` updates the autoload index. It searches for all library directories (see `library_directory/1` and `file_search_path/2`) holding the file `MKINDEX.pl` or `INDEX.pl`. If the current user can write or create the file `INDEX.pl` and it does not exist or is older than the directory or one of its files, the index for this directory is updated. If the file `MKINDEX.pl` exists updating is achieved by loading this file, normally containing a directive calling `make_library_index/2`. Otherwise `make_library_index/1` is called, creating an index for all `*.pl` files containing a module.

Below is an example creating a completely indexed library directory.

```
% mkdir ~/lib/prolog
% cd !$
% pl -g true -t 'make_library_index(.)'
```

If there are more than one library files containing the desired predicate the following search schema is followed:

1. If there is a library file that defines the module in which the undefined predicate is trapped, this file is used.
2. Otherwise library files are considered in the order they appear in the `library_directory/1` predicate and within the directory alphabetically.

make_library_index(+Directory)

Create an index for this directory. The index is written to the file `'INDEX.pl'` in the specified directory. Fails with a warning if the directory does not exist or is write protected.

make_library_index(+Directory, +ListOfPatterns)

Normally used in `MKINDEX.pl`, this predicate creates `INDEX.pl` for *Directory*, indexing all files that match one of the file-patterns in *ListOfPatterns*.

Sometimes library packages consist of one public load file and a number of files used by this load-file, exporting predicates that should not be used directly by the end-user. Such a library can be placed in a sub-directory of the library and the files containing public functionality can be added to the index of the library. As an example we give the XPCE library's `MKINDEX.pl`, including the public functionality of `trace/browse.pl` to the autoloadable predicates for the XPCE package.

```
:- make_library_index('.',
    [ '*.pl',
      'trace/browse.pl'
    ]).
```


reload_library_index

Force reloading the index after modifying the set of library directories by changing the rules for `library_directory/2`, `file_search_path/2`, adding or deleting `INDEX.pl` files. This predicate does *not* update the `INDEX.pl` files. Check `make_library_index/[1,2]` and `make/0` for updating the index files.

Normally, the index is reloaded automatically if a predicate cannot be found in the index and the set of library directories has changed. Using `reload_library_index/0` is necessary if directories are removed or the order of the library directories is changed.

2.14 Garbage Collection

SWI-Prolog provides garbage-collection, last-call optimization and atom garbage collection. These features are controlled using prolog flags (see `current_prolog_flag/2`).

2.15 Syntax Notes

SWI-Prolog uses standard ‘Edinburgh’ syntax. A description of this syntax can be found in the Prolog books referenced in the introduction. Below are some non-standard or non-common constructs that are accepted by SWI-Prolog:

- `0'⟨char⟩`
This construct is not accepted by all Prolog systems that claim to have Edinburgh compatible syntax. It describes the character code of `⟨char⟩`. To test whether `C` is a lower case character one can use `between(0'a, 0'z, C)`.
- `/* ... */ ... */`
The `/* ... */` comment statement can be nested. This is useful if some code with `/* ... */` comment statements in it should be commented out.

2.15.1 ISO Syntax Support

SWI-Prolog offers ISO compatible extensions to the Edinburgh syntax.

Character Escape Syntax

Within quoted atoms (using single quotes: `'⟨atom⟩'`) special characters are represented using escape-sequences. An escape sequence is lead in by the backslash (`\`) character. The list of escape sequences is compatible with the ISO standard, but contains one extension and the interpretation of numerically specified characters is slightly more flexible to improve compatibility.

`\a`
Alert character. Normally the ASCII character 7 (beep).

`\b`
Backspace character.

`\c`

No output. All input characters up to but not including the first non-layout character are skipped. This allows for the specification of pretty-looking long lines. For compatibility with Quintus Prolog. Not supported by ISO. Example:

```
format('This is a long line that would look better if it was \c
      split across multiple physical lines in the input')
```

`\<RETURN>`

No output. Skips input till the next non-layout character or to the end of the next line. Same intention as `\c` but ISO compatible.

`\f`

Form-feed character.

`\n`

Next-line character.

`\r`

Carriage-return only (i.e. go back to the start of the line).

`\t`

Horizontal tab-character.

`\v`

Vertical tab-character (ASCII 11).

`\x23`

Hexadecimal specification of a character. 23 is just an example. The 'x' may be followed by a maximum of 2 hexadecimal digits. The closing `\` is optional. The code `\xa\3` emits the character 10 (hexadecimal 'a') followed by '3'. The code `\x201` emits 32 (hexadecimal '20') followed by '1'. According to ISO, the closing `\` is obligatory and the number of digits is unlimited. The SWI-Prolog definition allows for ISO compatible specification, but is compatible with other implementations.

`\40`

Octal character specification. The rules and remarks for hexadecimal specifications apply to octal specifications too, but the maximum allowed number of octal digits is 3.

`\<character>`

Any character immediately preceded by a `\` and not covered by the above escape sequences is copied verbatim. Thus, `'\\'` is an atom consisting of a single `\` and `'\''` and `'\''` both describe the atom with a single `'`.

Character escaping is only available if the `current_prolog_flag(character_escapes, true)` is active (default). See `current_prolog_flag/2`. Character escapes conflict with `writeln/2` in two ways: `\40` is interpreted as decimal 40 by `writeln/2`, but character escapes handling by `read` has already interpreted as 32 (40 octal). Also, `\1` is translated to a single '1'. It is advised to use the more widely supported `format/[2,3]` predicate instead. If you insist upon using `writeln/2`, either switch `character_escapes` to `false`, or use double `\\`, as in `writeln('\\1')`.

Syntax for non-decimal numbers

SWI-Prolog implements both Edinburgh and ISO representations for non-decimal numbers. According to Edinburgh syntax, such numbers are written as $\langle radix \rangle' \langle number \rangle$, where $\langle radix \rangle$ is a number between 2 and 36. ISO defines binary, octal and hexadecimal numbers using $0[base] \langle number \rangle$. For example: `A is 0b100 \ / 0xf00` is a valid expression. Such numbers are always unsigned.

2.16 System limits

2.16.1 Limits on memory areas

SWI-Prolog has a number of memory areas which are only enlarged to a certain limit. The default sizes for these areas should suffice for most applications, but big applications may require larger ones. They are modified by command line options. The table below shows these areas. The first column gives the option name to modify the size of the area. The option character is immediately followed by a number and optionally by a k or m. With k or no unit indicator, the value is interpreted in Kbytes (1024 bytes), with m, the value is interpreted in Mbytes (1024 × 1024 bytes).

The local-, global- and trail-stack are limited to 128 Mbytes on 32 bit processors, or more generally to $2^{\text{bits-per-long}-5}$ bytes.

The PrologScript facility described in section 2.10.2 provides a mechanism for specifying options with the load-file. On Windows the default stack-sizes are controlled using the Windows registry on the key `HKEY_CURRENT_USER\Software\SWI\Prolog` using the names `localSize`, `globalSize` and `trailSize`. The value is a `DWORD` expressing the default stack size in Kbytes. A GUI for modifying these values is provided using the XPCE package. To use this, start the XPCE manual tools using `manpce/0`, after which you find *Preferences* in the *File* menu.

The heap

With the heap, we refer to the memory area used by `malloc()` and friends. SWI-Prolog uses the area to store atoms, functors, predicates and their clauses, records and other dynamic data. As of SWI-Prolog 2.8.5, no limits are imposed on the addresses returned by `malloc()` and friends.

On some machines, the runtime stacks described above are allocated using ‘sparse allocation’. Virtual space up to the limit is claimed at startup and committed and released while the area grows and shrinks. On Win32 platform this is realised using `VirtualAlloc()` and friends. On Unix systems this is realised using `mmap()`.

2.16.2 Other Limits

Clauses The only limit on clauses is their arity (the number of arguments to the head), which is limited to 1024. Raising this limit is easy and relatively cheap, removing it is harder.

Atoms and Strings SWI-Prolog has no limits on the sizes of atoms and strings. `read/1` and its derivatives however normally limit the number of newlines in an atom or string to 5 to improve error detection and recovery. This can be switched off with `style_check/1`.

The number of atoms is limited to 16777216 (16M) on 32-bit machines. On 64-bit machines this is virtually unlimited. See also section 7.6.2.

Option	Default	Area name	Description
-L	2M	local stack	The local stack is used to store the execution environments of procedure invocations. The space for an environment is reclaimed when it fails, exits without leaving choice points, the alternatives are cut off with the !/0 predicate or no choice points have been created since the invocation and the last subclause is started (tail recursion optimisation).
-G	4M	global stack	The global stack is used to store terms created during Prolog's execution. Terms on this stack will be reclaimed by backtracking to a point before the term was created or by garbage collection (provided the term is no longer referenced).
-T	4M	trail stack	The trail stack is used to store assignments during execution. Entries on this stack remain alive until backtracking before the point of creation or the garbage collector determines they are no longer needed any longer.
-A	1M	argument stack	The argument stack is used to store one of the intermediate code interpreter's registers. The amount of space needed on this stack is determined entirely by the depth in which terms are nested in the clauses that constitute the program. Overflow is most likely when using long strings in a clause.

Table 2.2: Memory areas

Address space SWI-Prolog data is packed in a 32-bit word, which contains both type and value information. The size of the various memory areas is limited to 128 Mb for each of the areas, except for the program heap, which is not limited.

Integers Integers are 32-bit (64 on 64-bit machines) to the user, but integers up to the value of the `max_tagged_integer` prolog-flag are represented more efficiently.

Floats Floating point numbers are represented as C-native double precision floats, 64 bit IEEE on most machines.

2.16.3 Reserved Names

The boot compiler (see `-b` option) does not support the module system. As large parts of the system are written in Prolog itself we need some way to avoid name clashes with the user's predicates, database keys, etc. Like Edinburgh C-Prolog [Pereira, 1986] all predicates, database keys, etc. that should be hidden from the user start with a dollar (\$) sign (see `style_check/1`).

Initialising and Managing a Prolog Project

3

Prolog text-books give you an overview of the Prolog language. The manual tells you what predicates are provided in the system and what they do. This chapter wants to explain how to run a project. There is no ultimate ‘right’ way to do this. Over the years we developed some practice in this area and SWI-Prolog’s commands are there to support this practice. This chapter describes the conventions and supporting commands.

The first two sections (section 3.1 and section 3.2 only require plain Prolog. The remainder discusses the use of the built-in graphical tools that require the XPCE graphical library installed on your system.

3.1 The project source-files

Organisation of source-files depends largely on the size of your project. If you are doing exercises for a Prolog course you’ll normally use one file for each exercise. If you have a small project you’ll work with one directory holding a couple of files and some files to link it all together. Even bigger projects will be organised in sub-projects each using their own directory.

3.1.1 File Names and Locations

File Name Extensions

The first consideration is what extension to use for the source-files. Tradition calls for `.pl`, but conflicts with Perl force the use of another extension on systems where extensions have global meaning, such as MS-Windows. On such systems `.pro` is the common alternative.¹

All versions of SWI-Prolog load files with the extension `.pl` as well as with the registered alternative extension without explicitly specifying the extension. For portability reasons we propose the following convention:

If there is no conflict because you do not use a conflicting application or the system does not force a unique relation between extension and application, use `.pl`.

With a conflict choose `.pro` and use this extension for the files you want to load through your file-manager. Use `.pl` for all other files for maximal portability.

Project Directories

Large projects are generally composed of sub-projects, each using their own directory or directory-structure. If nobody else will ever touch your files and you use only one computer there is little to

¹On MS-Windows, the alternative extension is stored in the registry-key `HKEY_CURRENT_USER/Software/SWI/Prolog/fileExtension` or `HKEY_LOCAL_MACHINE/Software/SWI/Prolog/fileExtension`.

worry about, but this is rarely the case with a large project.

To improve portability, SWI-Prolog uses the POSIX notation for filenames, which uses the forward slash (/) to separate directories. Just before hitting the file-system it uses `prolog_to_os_filename/2` to convert the filename to the conventions used by the hosting operating system. It is *strongly* advised to write paths using the /, especially on systems using the \ for this purpose (MS-Windows). Using \ violates the portability rules and requires you to *double* the \ due to the Prolog quoted-atom escape rules.

Portable code should use `prolog_to_os_filename/2` to convert computed paths into system-paths when constructing commands for `shell/1` and friends.

Sub-projects using search-paths

Thanks to Quintus, Prolog adapted an extensible mechanism for searching files using `file_search_path/2`. This mechanism allows for comfortable and readable specifications.

Suppose you have extensive library packages on graph-algorithms, set-operations and ui-primitives. These sub-projects are likely candidates for re-use in future projects. A good choice is to create a directory with sub-directories for each of these sub-projects.

Next, there are three options. One is to add the sub-projects to the directory-hierarchy of the current project. Another is to use a completely dislocated directory and finally the sub-project can be added to the SWI-Prolog hierarchy. Using local installation, a typical `file_search_path/2` is:

```
:- prolog_load_context(directory, Dir),
   asserta(user:file_search_path(myapp, Dir)).

user:file_search_path(graph, myapp(graph)).
user:file_search_path(ui, myapp(ui)).
```

For using sub-projects in the SWI-Prolog hierarchy one should use the path-alias `swi` as basis. For a system-wide installation use an absolute-path.

Extensive sub-projects with a small well-defined API should define a load-file using `use_module/1` calls to import the various library-components and export the API.

3.1.2 Project Special Files

There are a number of tasks you typically carry out on your project, such as loading it, creating a saved-state, debugging it, etc. Good practice on large projects is to define small files that hold the commands to execute such a task, name this file after the task and give it a file-extension that makes starting easy (see section 3.1.1). The task *load* is generally central to these tasks. Here is a tentative list.

- *load.pl*
Use this file to set up the environment (prolog flags and file search paths) and load the sources. Quite commonly this file also provides convenient predicates to parse command-line options and start the application.
- *run.pl*
Use this file to start the application. Normally it loads `load.pl` in silent-mode, and calls one of the starting predicates from `load.pl`.

- *save.pl*
Use this file to create a saved-state of the application by loading `load.pl` and call `qsave_program/2` to generate a saved-state with the proper options.
- *debug.pl*
Loads the program for debugging. In addition to loading `load.pl` this file defines rules for `portray/1` to modify printing rules for complex terms and customisation rules for the debugger and editing environment. It may start some of these tools.

3.2 Using modules

Modules have been debated fiercely in the Prolog world. Despite all counter-arguments we feel they are extremely useful because

- *They hide local predicates*
This is the reason they have been invented in the first place. Hiding provides two features. They allow for short predicate names without worrying about conflicts. Given the flat name-space introduced by modules, they still require meaningful module names as well as meaningful names for exported predicates.
- *They document the interface*
Possibly more important than avoiding name-conflicts is their role in documenting which part of the file is for public usage and which is private. When editing a module you may assume you can reorganise anything but the name and semantics of the exported predicates without worrying.
- *They help the editor*
The PceEmacs built-in editor does on-the-fly cross-referencing of the current module, colouring predicates based on their origin and usage. Using modules, the editor can quickly find out what is provided by the imported modules by reading just the first term. This allows it to indicate real-time which predicates are not used or not defined.

Using modules is generally easy. Only if you write meta-predicates (predicates reasoning about other predicates) that are exported from a module good understanding of resolution of terms to predicates inside a module is required. Here is a typical example from `readutil`.

```
:- module(read_util,
    [ read_line_to_codes/2,      % +Fd, -Codes
      read_line_to_codes/3,      % +Fd, -Codes, ?Tail
      read_stream_to_codes/2,    % +Fd, -Codes
      read_stream_to_codes/3,    % +Fd, -Codes, ?Tail
      read_file_to_codes/3,      % +File, -Codes, +Options
      read_file_to_terms/3       % +File, -Terms, +Options
    ]).
```


3.3 The test-edit-reload cycle

SWI-Prolog does not enforce the use of a particular editor for writing down Prolog source code. Editors are complicated programs that must be mastered in detail for real productive programming and if you are familiar with a specific editor you should not be forced to change. You may specify your favourite editor using the prolog flag `editor`, the environment variable `EDITOR` or by defining rules for `prolog_edit:edit_source/1` (see section 4.4).

The use of a built-in editor, which is selected by setting the prolog-flag `editor` to `pce_emacs`, has advantages. The XPCE *editor* object around which the built-in PceEmacs is built can be opened as a Prolog stream allowing analysis of your source by the real Prolog system.

3.3.1 Locating things to edit

The central predicate for editing something is `edit/1`, an extensible front-end that searches for objects (files, predicates, modules as well as XPCE classes and methods) in the Prolog database. If multiple matches are found it provides a choice. Together with the built-in completion on atoms bound to the TAB key this provides a quick way to edit objects:

```
?- edit(country).
Please select item to edit:

1 chat:country/10      '/staff/jan/lib/prolog/chat/countr.pl':16
2 chat:country/1      '/staff/jan/lib/prolog/chat/world0.pl':72
```

Your choice?

3.3.2 Editing and incremental compilation

One of the nice features of Prolog is that the code can be modified while the program is running. Using pure Prolog you can trace a program, find it is misbehaving, enter a *break environment*, modify the source code, reload it and finally do *retry* on the misbehaving predicate and try again. This sequence is not uncommon for long-running programs. For faster programs one normally aborts after understanding the misbehaviour, edit the source, reload it and try again.

One of the nice features of SWI-Prolog is the availability of `make/0`, a simple predicate that checks all loaded source files to see which ones you have modified. It then reloads these files, considering the module from which the file was loaded originally. This greatly simplifies the trace-edit-verify development cycle. After the tracer reveals there is something wrong with `prove/3`, you do:

```
?- edit(prove).
```

Now edit the source, possibly switching to other files and making multiple changes. After finishing invoke `make/0`, either through the editor UI (**Compile/Make** (`Control-C Control-M`)) or on the toplevel and watch the files being reloaded.²

```
?- make.
% show compiled into photo_gallery 0.03 sec, 3,360 bytes
```

²Watching these files is a good habit. If expected files are not reloaded you may have forgotten to save them from the editor or you may have been editing the wrong file (wrong directory).

3.4 Using the PceEmacs built-in editor

3.4.1 Activating PceEmacs

Initially `edit/1` uses the editor specified in the `EDITOR` environment variable. There are two ways to force it to use the built-in editor. One is to set the prolog-flag `editor` to `pce_emacs` and the other is by starting the editor explicitly using the `emacs/[0,1]` predicates.

3.4.2 Bluffing through PceEmacs

PceEmacs closely mimics Richard Stallman's GNU-Emacs commands, adding features from modern window-based editors to make it more acceptable for beginners.³

At the basis, PceEmacs maps keyboard sequences to methods defined on the extended *editor* object. Some frequently used commands are, with their key-binding, presented in the menu-bar above each editor window. A complete overview of the bindings for the current *mode* is provided through Help/Show key bindings (`Control-h Control-b`).

Edit modes

Modes are the heart of (Pce)Emacs. Modes define dedicated editing support for a particular kind of (source-)text. For our purpose we want *Prolog mode*. There are various ways to make PceEmacs use Prolog mode for a file.

- *Using the proper extension*
If the file ends in `.pl` or the selected alternative (e.g. `.pro`) extension, Prolog mode is selected.
- *Using `#!/path/to/pl`*
If the file is a *Prolog Script* file, starting with the line `#!/path/to/pl options -s`, Prolog mode is selected regardless of the extension
- *Using `-- Prolog --`*
If the above sequence appears in the first line of the file (inside a Prolog comment) Prolog mode is sequence
- *Explicit selection*
Finally, using File/Mode/Prolog (*y*)ou can switch to Prolog mode explicitly.

Frequently used editor commands

Below we list a few important commands and how to activate them.

- *Cut/Copy/Paste*
These commands follow Unix/X11 traditions. You're best suited with a three-button mouse. After selecting using the left-mouse (double-click uses word-mode and triple line-mode), the selected text is *automatically* copied to the clipboard (X11 primary selection on Unix). *Cut* is achieved using the `DEL` key or by typing something else at the location. *Paste* is achieved using the middle-mouse (or wheel) button. If you don't have a middle mouse-button, pressing the

³Decent merging with MS-Windows control-key conventions is difficult as many conflict with GNU-Emacs. Especially the cut/copy/paste commands conflict with important GNU-Emacs commands.

left- and right-button at the same time is interpreted as a middle-button click. If nothing helps there is the **Edit/Paste** menu-entry. Text is pasted at the caret-location.

- *Undo*
Undo is bound to the GNU-Emacs **Control-`_`** as well as the MS-Windows **Control-Z** sequence.
- *Abort*
Multi-key sequences can be aborted at any stage using **Control-G**.

- *Find*
Find (Search) is started using **Control-S** (forward) or **Control-R** (backward). PceEmacs implements *incremental search*. This is difficult to use for novices, but very powerful once you get the clue. After one of the above start-keys the system indicates search mode in the status line. As you are typing the search-string, the system searches for it, extending the search with every character you type. It illustrates the current match using a green background.

If the target cannot be found, PceEmacs warns you and no longer extends the search-string.⁴ During search some characters have special meaning. Typing anything but these characters commits the search, re-starting normal edit mode. Special commands are:

Control-S

Search for next forwards.

Control-R

Search for next backwards.

Control-W

Extend search to next word-boundary.

Control-G

Cancel search, go back to where it started.

ESC

Commit search, leaving caret at found location.

Backspace

Remove a character from the search string.

- *Dynamic Abbreviation*
Also called *dabbrev* is an important feature of Emacs clones to support programming. After typing the first few letters of an identifier you may hit **Alt-/**, causing PceEmacs to search backwards for identifiers that start the same and using it to complete the text you typed. A second **Alt-/** searches further backwards. If there are no hits before the caret it starts searching forwards. With some practice, this system allows for very fast entering code with nice and readable identifiers (or other difficult long words).
- *Open (a file)*
Is called **File/Find file** (**Control-x Control-f**). By default the file is loaded into the current window. If you want to keep this window, Hit **Alt-s** or click the little icon at the bottom-left to make the window *sticky*.

⁴GNU-Emacs keeps extending the string, but why? Adding more text will not make it match.

- *Split view*

Sometimes you want to look at two places of the same file. To do this, use **Control-x 2** to create a new window pointing to the same file. Do not worry, you can edit as well as move around in both. **Control-x 1** kills all other windows running on the same file.

These were the most commonly used commands. In section section 3.4.3 we discuss specific support for dealing with Prolog source code.

3.4.3 Prolog Mode

In the previous section (section 3.4.2) we explained the basics of PceEmacs. Here we continue with Prolog specific functionality. Possibly the most interesting is *Syntax highlighting*. Unlike most editors where this is based on simple patterns, PceEmacs syntax highlighting is achieved by Prolog itself actually reading and interpreting the source as you type it. There are three moments at which PceEmacs checks (part of) the syntax.

- *After typing a .*

After typing a `.` that is not preceded by a *symbol* character the system assumes you completed a clause, tries to find the start of this clause and verifies the syntax. If this process succeeds it colours the elements of the clause according to the rules given below. Colouring is done using information from the last full check on this file. If it fails, the syntax error is displayed in the status line and the clause is not coloured.

- *After the command Control-c Control-s*

Acronym for **C**heck **S**yntax it performs the same checks as above for the clause surrounding the caret. On a syntax error however, the caret is moved to the expected location of the error.⁵

- *After pausing for two seconds*

After a short pause (2 seconds), PceEmacs opens the edit-buffer and reads it as a whole, creating an index of defined, called, dynamic, imported and exported predicates. After completing this, it re-reads the file and colours all clauses and calls with valid syntax.

- *After typing Control-l Control-l*

The **Control-l** commands re-centers the window (scrolls the window to make the caret the center of the window). Hitting this command twice starts the same process as above.

The colour schema itself is defined in `emacs/prolog_colour`. The colouring can be extended and modified using multifile predicates. Please check this source-file for details. In general, underlined objects have a popup (right-mouse button) associated for common commands such as viewing the documentation or source. **Bold** text is used to indicate the definition of objects (typically predicates when using plain Prolog). Other colours follow intuitive conventions. See table 3.4.3.

Layout support Layout is not ‘just nice’, it is *essential* for writing readable code. There is much debate on the proper layout of Prolog. PceEmacs, being a rather small project supports only one particular style for layout.⁶ Below are examples of typical constructs.

⁵In most cases the location where the parser cannot proceed is further down the file than the actual error-location.

⁶Defined in Prolog in the file `emacs/prolog_mode`, you may wish to extend this. Please contribute your extensions!

Clauses	
Blue bold	Head of an exported predicate
Red bold	Head of a predicate that is not called
Black Bold	Head of remaining predicates
Calls in the clause-body	
Blue	Call to built-in or imported predicate
Red	Call to not-defined predicate
Purple	Call to dynamic predicate
Other entities	
Dark green	Comment
Dark blue	Quoted atom or string
Brown	Variable

Table 3.1: Colour conventions

```

head(arg1, arg2).

head(arg1, arg2) :- !.

head(Arg1, arg2) :- !,
    call1(Arg1).

head(Arg1, arg2) :-
    (   if(Arg1)
    -> then
    ;   else
    ).

head(Arg1) :-
    (   a
    ;   b
    ).

head :-
    a(many,
      long,
      arguments(with,
                many,
                more),
      and([ a,
            long,
            list,
            with,
            a,
            | tail
            ])).

```

PceEmacs uses the same conventions as GNU-Emacs. The `TAB` key indents the current line according to the syntax rules. `Alt-q` indents all lines of the current clause. It provides support for head, calls (indented 1 tab), if-then-else, disjunction and argument-lists broken across multiple lines as illustrated above.

Finding your way around

The command `Alt-` extracts name and arity from the caret location and jumps (after conformation or edit) to the definition of the predicate. It does so based on the source-location database of loaded predicates also used by `edit/1`. This makes locating predicates reliable if all sources are loaded and up-to-date (see `make/0`).

In addition, references to files in `use_module/[1,2]`, `consult/1`, etc. are red if the file cannot be found and underlined blue if the file can be loaded. A popup allows for opening the referenced file.

3.5 The Graphical Debugger

SWI-Prolog offers two debuggers. One is the traditional text-console based 4-port Prolog tracer and the other is a window-based source-level debugger. The window-based debugger requires XPCE installed. It operates based on the `prolog_trace_interception/4` hook and other low-level functionality described in chapter B.

Window-based tracing provides much better overview due to the eminent relation to your source-code, a clear list of named variables and their bindings as well as a graphical overview of the call and choice-point stack. There are some drawbacks though. Using a textual trace on the console one can scroll back and examine the past, while the graphical debugger just presents a (much better) overview of the current state.

3.5.1 Invoking the window-based debugger

Whether the text-based or window-based debugger is used is controlled using the predicates `guitracer/0` and `noguitracer/0`. Entering debug mode is controlled using the normal predicates for this: `trace/0` and `spy/1`. In addition, PceEmacs prolog mode provides the command `Prolog/Break at (Control-c b)` to insert a break-point at a specific location in the source-code.

guitracer

This predicate installs the above-mentioned hooks that redirect tracing to the window-based environment. No window appears. The debugger window appears as actual tracing is started through `trace/0`, by hitting a spy-point defined by `spy/1` or a break-point defined using PceEmacs command `Prolog/Break at (Control-c b)`.

noguitracer

Disable the hooks installed by `guitracer/0`, reverting to normal text-console based tracing.

3.6 The Prolog Navigator

Another tool is the *Prolog Navigator*. This tool can be started from PceEmacs using the command `Browse/Prolog navigator`, from the GUI debugger or using the programmatic IDE interface described in section 3.7.

3.7 Accessing the IDE from your program

Over the years a collection of IDE components have been developed, each with their own interface. In addition, some of these components require each other and loading IDE components must be on demand to avoid the IDE being part of a saved-state (see `qsave_program/2`). For this reason, access to the IDE will be concentrated on a single interface called `prolog_ide/1`:

prolog_ide(+Action)

This predicate ensures the IDE enabling XPCE component is loaded, creates the XPCE class `prolog_ide` and sends *Action* to its one and only instance `\index{@prolog_ide}\objectname{prolog_ide}`. *Action* is one of the following:

open_navigator(+Directory)

Open the Prolog Navigator (see section 3.6) in the given *Directory*.

open_debug_status

Open a window to edit spy- and trace-points.

open_query_window

Opens a little window to run Prolog queries from a GUI component.

3.8 Summary of the iDE

The SWI-Prolog development environment consists of a number of interrelated but not (yet) integrated tools. Here is a list of the most important features and tips.

- *Atom completion*
The console⁷ completes a partial atom on the TAB key and shows alternatives on the command Alt-?.
- *Use edit/1 to finding locations*
The command `edit/1` takes the name of a file, module, predicate or other entity registered through extensions and starts the users preferred editor at the right location.
- *Select editor*
External editors are selected using the EDITOR environment variable, by setting the prolog flag `editor` or by defining the hook `prolog_edit:edit_source/1`.
- *Update Prolog after editing*
Using `make/0`, all files you have edited are re-loaded.
- *PceEmacs*
Offers syntax-highlighting and checking based on real-time parsing of the editor's buffer, layout-support and navigation support.

⁷On Windows this is realised by `plwin.exe`, on Unix through the GNU readline library, which is included automatically when found by `configure`.

- *Using the graphical debugger*
The predicates `guitracer/0` and `noguitracer/0` switch between traditional text-based and window-based debugging. The tracer is activated using the `trace/0`, `spy/1` or menu-items from PceEmacs or the PrologNavigator.
- *The Prolog Navigator*
Shows the file-structure and structure inside the file. It allows for loading files, editing, setting spy-points, etc.

4

Built-in predicates

4.1 Notation of Predicate Descriptions

We have tried to keep the predicate descriptions clear and concise. First the predicate name is printed in bold face, followed by the arguments in italics. Arguments are preceded by a ‘+’, ‘-’ or ‘?’ sign. ‘+’ indicates the argument is input to the predicate, ‘-’ denotes output and ‘?’ denotes ‘either input or output’.¹ Constructs like ‘*op/3*’ refer to the predicate ‘*op*’ with arity ‘3’. Finally, arguments may have the ‘:’ specifier, which implies the argument is module-sensitive. Normally the argument is a *callable* term referring to a predicate in a specific module. See section 5 for more information on module-handling.

4.2 Character representation

In traditional (Edinburgh-) Prolog, characters are represented using *character-codes*. Character codes are integer indices into a specific character set. Traditionally the character set was 7-bits US-ASCII. 8-bit character sets have been allowed for a long time, providing support for national character sets, of which iso-latin-1 (ISO 8859-1) is applicable to many western languages. Text-files are supposed to represent a sequence of character-codes.

ISO Prolog introduces three types, two of which are used for characters and one for accessing binary streams (see *open/4*). These types are:

- *code*
A *character-code* is an integer representing a single character. As files may use multi-byte encoding for supporting different character sets (utf-8 encoding for example), reading a code from a text-file is in general not the same as reading a byte.
- *char*
Alternatively, characters may be represented as *one-character-atoms*. This is a very natural representation, hiding encoding problems from the programmer as well as providing much easier debugging.
- *byte*
Bytes are used for accessing binary-streams.

The current version of SWI-Prolog does not provide support for multi-byte character encoding. This implies for example that it is not capable of breaking a multi-byte encoded atom into characters. For SWI-Prolog, bytes and codes are the same and one-character-atoms are simple atoms containing one byte.

¹These marks do **not** suggest instantiation (e.g. *var(+Var)*).

To ease the pain of these multiple representations, SWI-Prolog's built-in predicates dealing with character-data work as flexible as possible: they accept data in any of these formats as long as the interpretation is unambiguous. In addition, for output arguments that are instantiated, the character is extracted before unification. This implies that the following two calls are identical, both testing whether the next input characters is an a.

```
peek_code(Stream, a).
peek_code(Stream, 97).
```

These multiple-representations are handled by a large number of built-in predicates, all of which are ISO-compatible. For converting between code and character there is `char_code/2`. For breaking atoms and numbers into characters are `atom_chars/2`, `atom_codes/2`, `number_codes/2` and `number_chars/2`. For character I/O on streams there is `get_char/[1,2]`, `get_code/[1,2]`, `get_byte/[1,2]`, `peek_char/[1,2]`, `peek_code/[1,2]`, `peek_byte/[1,2]`, `put_code/[1,2]`, `put_char/[1,2]` and `put_byte/[1,2]`. The prolog-flag `double_quotes` (see `current_prolog_flag/2`) controls how text between double-quotes is interpreted.

4.3 Loading Prolog source files

This section deals with loading Prolog source-files. A Prolog source file is a plain text file containing a Prolog program or part thereof. Prolog source files come in three flavours:

A traditional Prolog source file contains a Prolog clauses and directives, but no *module-declaration*. They are normally loaded using `consult/1` or `ensure_loaded/1`.

A module Prolog source file starts with a module declaration. The subsequent Prolog code is loaded into the specified module and only the *public* predicates are made available to the context loading the module. Module files are normally loaded using `use_module/[1,2]`. See chapter 5 for details.

An include Prolog source file is loaded using the `include/1` directive and normally contains only directives.

Prolog source-files are located using `absolute_file_name/3` with the following options:

```
locate_prolog_file(Spec, Path) :-
    absolute_file_name(Spec,
                      [ file_type(prolog),
                        access(read)
                      ],
                      Path).
```

The `file_type(prolog)` option is used to determine the extension of the file using `prolog_file_type/2`. The default extension is `.pl`. `Spec` allows for the *path-alias* construct defined by `absolute_file_name/3`. The most commonly used path-alias is `library(LibraryFile)`. The example below loads the library file `oset.pl` (containing predicates for manipulating ordered sets).

```
:- use_module(library(ose)).
```

SWI-Prolog recognises grammar rules (DCG) as defined in [Clocksin & Melish, 1987]. The user may define additional compilation of the source file by defining the dynamic predicates `term_expansion/2` and `goal_expansion/2`. Transformations by `term_expansion/2` overrule the systems grammar rule transformations. It is not allowed to use `assert/1`, `retract/1` or any other database predicate in `term_expansion/2` other than for local computational purposes.²

Directives may be placed anywhere in a source file, invoking any predicate. They are executed when encountered. If the directive fails, a warning is printed. Directives are specified by `:-/1` or `?-/1`. There is no difference between the two.

SWI-Prolog does not have a separate `reconsult/1` predicate. Reconsulting is implied automatically by the fact that a file is consulted which is already loaded.

load_files(+Files, +Options)

The predicate `load_files/2` is the parent of all the other loading predicates except for `include/2`. It currently supports a subset of the options of Quintus `load_files/2`. *Files* is either specifies a single, or a list of source-files. The specification for a source-file is handled `absolute_file_name/2`. See this predicate for the supported expansions. *Options* is a list of options using the format

OptionName(OptionValue)

The following options are currently supported:

if(Condition)

Load the file only if the specified condition is satisfied. The value `true` loads the file unconditionally, `changed` loads the file if it was not loaded before, or has been modified since it was loaded the last time, `not_loaded` loads the file if it was not loaded before.

must_be_module(Bool)

If `true`, raise an error if the file is not a module file. Used by `use_module/[1, 2]`.

imports(ListOrAll)

If `all` and the file is a module file, import all public predicates. Otherwise import only the named predicates. Each predicate is referred to as `<name>/<arity>`. This option has no effect if the file is not a module file.

silent(Bool)

If `true`, load the file without printing a message. The specified value is the default for all files loaded as a result of loading the specified files.

autoload(Bool)

If `true` (default `false`), indicate this load is a *demand* load. This implies that, depending on the setting of the prolog-flag `verbose_autoload` the load-action is printed at level `informational` or `silent`. See also `print_message/2` and `current_prolog_flag/2`.

²It does work for normal loading, but not for `qcompile/1`.

qcompile(Bool)

If this call appears in a directive of a file that is compiled into Quick Load Format using `qcompile/1` and this flag is `true`, the contents of the argument files are included in the `.qlf` file instead of the loading directive.

stream(Input)

This SWI-Prolog extension compiles the data from the stream *Input*. If this option is used, *Files* must be a single atom which is used to identify the source-location of the loaded clauses as well as remove all clauses if the data is re-consulted.

This option is added to allow compiling from non-file locations such as databases, the web, the *user* (see `consult/1`) or other servers.

The `load_files/2` predicate can be hooked to load other data or data from other objects than files. See `prolog_load_file/2` for a description and `http_load` for an example.

consult(+File)

Read *File* as a Prolog source file. *File* may be a list of files, in which case all members are consulted in turn. *File* may start with the Unix shell special sequences `~`, `<user>` and `$(var)`. *File* may also be `library(Name)`, in which case the libraries are searched for a file with the specified name. See also `library_directory/1` and `file_search_path/2`. `consult/1` may be abbreviated by just typing a number of file names in a list. Examples:

```
?- consult(load).           % consult load or load.pl
?- [library(quintus)].     % load Quintus compatibility library
?- [user].
```

`Consult/1` is equivalent to `load_files(Files, [])`, except for handling the special file *user*, which reads clauses from the terminal. See also the `stream(Input)` option of `load_files/2`.

ensure_loaded(+File)

If the file is not already loaded, this is equivalent to `consult/1`. Otherwise, if the file defines a module, import all public predicates. Finally, if the file is already loaded, is not a module file and the context module is not the global user module, `ensure_loaded/1` will call `consult/1`.

With the semantics, we hope to get as closely possible to the clear semantics without the presence of a module system. Applications using modules should consider using `use_module/[1, 2]`.

Equivalent to `load_files(Files, [if(changed)])`.

include(+File)

Pretend the terms in *File* are in the source-file in which `:- include(File)` appears. The `include` construct is only honoured if it appears as a directive in a source-file. Normally *File* contains a sequence of directives.

require(+ListOfNameAndArity)

Declare that this file/module requires the specified predicates to be defined “with their commonly accepted definition”. This predicate originates from the Prolog portability layer for XPCE. It is intended to provide a portable mechanism for specifying that this module requires the specified predicates.

The implementation normally first verifies whether the predicate is already defined. If not, it will search the libraries and load the required library.

SWI-Prolog, having autoloading, does **not** load the library. Instead it creates a procedure header for the predicate if it does not exist. This will flag the predicate as ‘undefined’. See also `check/0` and `autoload/0`.

make

Consult all source files that have been changed since they were consulted. It checks *all* loaded source files: files loaded into a compiled state using `pl -c ...` and files loaded using `consult` or one of its derivatives. The predicate `make/0` is called after `edit/1`, automatically reloading all modified files. If the user uses an external editor (in a separate window), `make/0` is normally used to update the program after editing. In addition, `make/0` updates the `autoload` indices (see section 2.13) and runs `list_undefined/0` from the `check` library to report on undefined predicates.

library_directory(?Atom)

Dynamic predicate used to specify library directories. Default `./lib`, `~/lib/prolog` and the system’s library (in this order) are defined. The user may add library directories using `assert/1`, `asserta/1` or remove system defaults using `retract/1`.

file_search_path(+Alias, ?Path)

Dynamic predicate used to specify ‘path-aliases’. This feature is best described using an example. Given the definition

```
file_search_path(demo, '/usr/lib/prolog/demo').
```

the file specification `demo(myfile)` will be expanded to `/usr/lib/prolog/demo/myfile`. The second argument of `file_search_path/2` may be another alias.

Below is the initial definition of the file search path. This path implies `swi(<Path>)` refers to a file in the SWI-Prolog home directory. The alias `foreign(<Path>)` is intended for storing shared libraries (`.so` or `.DLL` files). See also `load_foreign_library/[1,2]`.

```
user:file_search_path(library, X) :-
    library_directory(X).
user:file_search_path(swi, Home) :-
    current_prolog_flag(home, Home).
user:file_search_path(foreign, swi(ArchLib)) :-
    current_prolog_flag(arch, Arch),
    atom_concat('lib/', Arch, ArchLib).
user:file_search_path(foreign, swi(lib)).
```

The `file_search_path/2` expansion is used by all loading predicates as well as by `absolute_file_name/[2,3]`.

The prolog-flag `verbose_file_search` can be set to `true` to help debugging Prolog’s search for files.

expand_file_search_path(+Spec, -Path)

Unifies *Path* with all possible expansions of the file name specification *Spec*. See also `absolute_file_name/3`.

prolog_file_type(?Extension, ?Type)

This dynamic multifile predicate defined in module `user` determines the extensions considered by `file_search_path/2`. *Extension* is the filename extension without the leading dot, *Type* denotes the type as used by the `file_type(Type)` option of `file_search_path/2`. Here is the initial definition of `prolog_file_type/2`:

```
user:prolog_file_type(pl,      prolog).
user:prolog_file_type(Ext,    prolog) :-
    current_prolog_flag(associate, Ext),
    Ext \== pl.
user:prolog_file_type(qlf,    qlf).
user:prolog_file_type(Ext,    executable) :-
    current_prolog_flag(shared_object_extension, Ext).
```

Users may wish to change the extension used for Prolog source files to avoid conflicts (for example with `perl`) as well as to be compatible with some specific implementation. The preferred alternative extension is `.pro`.

source_file(?File)

Succeeds if *File* is a loaded Prolog source file. *File* is the absolute and canonical path to the source-file.

source_file(?Pred, ?File)

Is true if the predicate specified by *Pred* was loaded from file *File*, where *File* is an absolute path name (see `absolute_file_name/2`). Can be used with any instantiation pattern, but the database only maintains the source file for each predicate. See also `clause_property/2`.

prolog_load_context(?Key, ?Value)

Determine loading context. The following keys are defined:

Key	Description
<code>module</code>	Module into which file is loaded
<code>file</code>	File loaded
<code>stream</code>	Stream identifier (see <code>current_input/1</code>)
<code>directory</code>	Directory in which <i>File</i> lives.
<code>term_position</code>	Position of last term read. Term of the form ' <code>\$stream_position</code> ' (0, <i>Line</i> , 0, 0, 0)

Quintus compatibility predicate. See also `source_location/2`.

source_location(-File, -Line)

If the last term has been read from a physical file (i.e., not from the file `user` or a string), unify *File* with an absolute path to the file and *Line* with the line-number in the file. New code should use `prolog_load_context/2`.

term_expansion(+Term1, -Term2)

Dynamic predicate, normally not defined. When defined by the user all terms read during consulting that are given to this predicate. If the predicate succeeds Prolog will assert *Term2* in the database rather than the read term (*Term1*). *Term2* may be a term of the form ‘?- *Goal*’ or ‘:- *Goal*’. *Goal* is then treated as a directive. If *Term2* is a list all terms of the list are stored in the database or called (for directives). If *Term2* is of the form below, the system will assert *Clause* and record the indicated source-location with it.

```
'$source_location' (<File>, <Line>):<Clause>
```

When compiling a module (see chapter 5 and the directive `module/2`), `expand_term/2` will first try `term_expansion/2` in the module being compiled to allow for term-expansion rules that are local to a module. If there is no local definition, or the local definition fails to translate the term, `expand_term/2` will try `term_expansion/2` in module `user`. For compatibility with SICStus and Quintus Prolog, this feature should not be used. See also `expand_term/2`, `goal_expansion/2` and `expand_goal/2`.

expand_term(+Term1, -Term2)

This predicate is normally called by the compiler to perform preprocessing. First it calls `term_expansion/2`. If this predicate fails it performs a grammar-rule translation. If this fails it returns the first argument.

goal_expansion(+Goal1, -Goal2)

Like `term_expansion/2`, `goal_expansion/2` provides for macro-expansion of Prolog source-code. Between `expand_term/2` and the actual compilation, the body of clauses analysed and the goals are handed to `expand_goal/2`, which uses the `goal_expansion/2` hook to do user-defined expansion.

The predicate `goal_expansion/2` is first called in the module that is being compiled, and then on the user module.

Only goals appearing in the body of clauses when reading a source-file are expanded using mechanism, and only if they appear literally in the clause, or as an argument to the meta-predicates `not/1`, `call/1`, `once/1`, `ignore/1`, `findall/3`, `bagof/3`, `setof/3` or `forall/2`. A real predicate definition is required to deal with dynamically constructed calls.

expand_goal(+Goal1, -Goal2)

This predicate is normally called by the compiler to perform preprocessing. First it calls `goal_expansion/2`. If this fails it returns the first argument.

at_initialization(+Goal)

Register *Goal* to be run when the system initialises. Initialisation takes place after reloading a `.qlf` (formerly `.wic`) file as well as after reloading a saved-state. The hooks are run in the order they were registered. A warning message is issued if *Goal* fails, but execution continues. See also `at_halt/1`

at_halt(+Goal)

Register *Goal* to be run from `PL_cleanup()`, which is called when the system halts. The hooks are run in the reverse order they were registered (FIFO). Success or failure executing a hook is ignored. If the hook raises an exception this is printed using `print_message/2`. An attempt to call `halt/[0,1]` from a hook is ignored.

initialization(+Goal)

Call *Goal* and register it using `at_initialization/1`. Directives that do other things than creating clauses, records, flags or setting predicate attributes should normally be written using this tag to ensure the initialisation is executed when a saved system starts. See also `qsave_program/[1,2]`.

compiling

Succeeds if the system is compiling source files with the `-c` option or `qcompile/1` into an intermediate code file. Can be used to perform conditional code optimisations in `term_expansion/2` (see also the `-O` option) or to omit execution of directives during compilation.

preprocessor(-Old, +New)

Read the input file via a Unix process that acts as preprocessor. A preprocessor is specified as an atom. The first occurrence of the string ‘%f’ is replaced by the name of the file to be loaded. The resulting atom is called as a Unix command and the standard output of this command is loaded. To use the Unix C preprocessor one should define:

```
?- preprocessor(Old, '/lib/cpp -C -P %f'), consult(...).
```

```
Old = none
```

4.3.1 Quick load files

SWI-Prolog supports compilation of individual or multiple Prolog source files into ‘Quick Load Files’. A ‘Quick Load Files’ (`.qlf` file) stores the contents of the file in a precompiled format.

These files load considerably faster than source files and are normally more compact. They are machine independent and may thus be loaded on any implementation of SWI-Prolog. Note however that clauses are stored as virtual machine instructions. Changes to the compiler will generally make old compiled files unusable.

Quick Load Files are created using `qcompile/1`. They are loaded using `consult/1` or one of the other file-loading predicates described in section 4.3. If `consult` is given the explicit `.pl` file, it will load the Prolog source. When given the `.qlf` file, it will load the file. When no extension is specified, it will load the `.qlf` file when present and the `.pl` file otherwise.

qcompile(+File)

Takes a single file specification like `consult/1` (i.e., accepts constructs like `library(LibFile)`) and creates a Quick Load File from *File*. The file-extension of this file is `.qlf`. The base name of the Quick Load File is the same as the input file.

If the file contains ‘:- `consult(+File)`’, ‘:- `[+File]`’ or ‘:- `load_files(+File, [qcompile(true), ...])`’ statements, the referred files are compiled into the same `.qlf` file. Other directives will be stored in the `.qlf` file and executed in the same fashion as when loading the `.pl` file.

For `term_expansion/2`, the same rules as described in section 2.10 apply.

Conditional execution or optimisation may test the predicate `compiling/0`.

Source references (`source_file/2`) in the Quick Load File refer to the Prolog source file from which the compiled code originates.

4.4 Listing and Editor Interface

SWI-Prolog offers an extensible interface which allows the user to edit objects of the program: predicates, modules, files, etc. The editor interface is implemented by `edit/1` and consists of three parts: *locating*, *selecting* and *starting the editor*.

Any of these parts may be extended or redefined by adding clauses to various multi-file (see `multifile/1`) predicates defined in the module `prolog_edit`.

The built-in edit specifications for `edit/1` (see `prolog_edit:locate/3`) are described below.

Fully specified objects	
<code><Module>:<Name>/<Arity></code>	Refers a predicate
<code>module(<Module>)</code>	Refers to a module
<code>file(<Path>)</code>	Refers to a file
<code>source_file(<Path>)</code>	Refers to a loaded source-file
Ambiguous specifications	
<code><Name>/<Arity></code>	Refers this predicate in any module
<code><Name></code>	Refers to (1) named predicate in any module with any arity, (2) a (source) file or (3) a module.

`edit(+Specification)`

First exploits `prolog_edit:locate/3` to translate *Specification* into a list of *Locations*. If there is more than one ‘hit’, the user is asked to select from the locations found. Finally, `prolog_edit:edit_source/1` is used to invoke the user’s preferred editor. Typically, `edit/1` can be handed the name of a predicate, module, basename of a file, XPCE class, XPCE method, etc.

`edit`

Edit the ‘default’ file using `edit/1`. The default file is the file loaded with the commandline option `-s` or, in windows, the file loaded by double-clicking from the Windows shell.

`prolog_edit:locate(+Spec, -FullSpec, -Location)`

Where *Spec* is the specification provided through `edit/1`. This multifile predicate is used to enumerate locations at which an object satisfying the given *Spec* can be found. *FullSpec* is unified with the complete specification for the object. This distinction is used to allow for ambiguous specifications. For example, if *Spec* is an atom, which appears as the base-name of a loaded file and as the name of a predicate, *FullSpec* will be bound to `file(Path)` or `Name/Arity`.

Location is a list of attributes of the location. Normally, this list will contain the term `file(File)` and —if available— the term `line(Line)`.

`prolog_edit:locate(+Spec, -Location)`

Same as `prolog_edit:locate/3`, but only deals with fully-specified objects.

`prolog_edit:edit_source(+Location)`

Start editor on *Location*. See `prolog_edit:locate/3` for the format of a location term. This multi-file predicate is normally not defined. If it succeeds, `edit/1` assumes the editor is started.

If it fails, `edit/1` uses its internal defaults, which are defined by the prolog-flag `editor` and/or the environment variable `EDITOR`. The following rules apply. If the prolog-flag `editor` is of the format `$(name)`, the editor is determined by the environment variable `(name)`. Else, if this flag is `pce_emacs` or `built_in` and `XPCE` is loaded or can be loaded, the built-in Emacs clone is used. Else, if the environment `EDITOR` is set, this editor is used. Finally, `vi` is used as default on Unix systems and `notepad` on Windows.

See the default user preferences file `dotfiles/dotplrc` for examples.

prolog_edit:edit_command(+Editor, -Command)

Determines how *Editor* is to be invoked using `shell/1`. *Editor* is the determined editor (see `edit_source/1`), without the full path specification, and without possible (exe) extension. *Command* is an atom describing the command. The pattern `%f` is replaced by the full file-name of the location, and `%d` by the line number. If the editor can deal with starting at a specified line, two clauses should be provided, one holding only the `%f` pattern, and one holding both patterns.

The default contains definitions for `vi`, `emacs`, `emacsclient`, `vim` and `notepad` (latter without line-number version).

Please contribute your specifications to `jan@swi.psy.uva.nl`.

prolog_edit:load

Normally not-defined multifile predicate. This predicate may be defined to provide loading hooks for user-extensions to the `edit` module. For example, `XPCE` provides the code below to load `swi_edit`, containing definitions to locate classes and methods as well as to bind this package to the `PceEmacs` built-in editor.

```
:- multifile prolog_edit:load/0.

prolog_edit:load :-
    ensure_loaded(library(swi_edit)).
```

listing(+Pred)

List specified predicates (when an atom is given all predicates with this name will be listed). The listing is produced on the basis of the internal representation, thus losing user's layout and variable name information. See also `portray_clause/1`.

listing

List all predicates of the database using `listing/1`.

portray_clause(+Clause)

Pretty print a clause. A clause should be specified as a term `'(Head) :- (Body)'`. Facts are represented as `'(Head) :- true'` or simply `(Head)`. See also `portray_clause/2`.

portray_clause(+Stream, +Clause)

Pretty print a clause to *Stream*. See `portray_clause/1` for details.

4.5 Verify Type of a Term

var(+Term)

Succeeds if *Term* currently is a free variable.

nonvar(+Term)

Succeeds if *Term* currently is not a free variable.

integer(+Term)

Succeeds if *Term* is bound to an integer.

float(+Term)

Succeeds if *Term* is bound to a floating point number.

number(+Term)

Succeeds if *Term* is bound to an integer or a floating point number.

atom(+Term)

Succeeds if *Term* is bound to an atom.

string(+Term)

Succeeds if *Term* is bound to a string.

atomic(+Term)

Succeeds if *Term* is bound to an atom, string, integer or floating point number.

compound(+Term)

Succeeds if *Term* is bound to a compound term. See also `functor/3` and `=../2`.

callable(+Term)

Succeeds if *Term* is bound to an atom or a compound term, so it can be handed without type-error to `call/1`, `functor/3` and `=../2`.

ground(+Term)

Succeeds if *Term* holds no free variables.

4.6 Comparison and Unification of Terms

4.6.1 Standard Order of Terms

Comparison and unification of arbitrary terms. Terms are ordered in the so called “standard order”. This order is defined as follows:

1. *Variables* < *Atoms* < *Strings* < *Numbers* < *Terms*³
2. *Old Variable* < *New Variable*⁴
3. *Atoms* are compared alphabetically.

³Strings might be considered atoms in future versions. See also section 4.23

⁴In fact the variables are compared on their (dereferenced) addresses. Variables living on the global stack are always < than variables on the local stack. Programs should not rely on the order in which variables are sorted.

4. *Strings* are compared alphabetically.
5. *Numbers* are compared by value. Integers and floats are treated identically.
6. *Compound* terms are first checked on their arity, then on their functor-name (alphabetically) and finally recursively on their arguments, leftmost argument first.

If the prolog flag (see `current_prolog_flag/2`) `iso` is defined, all floating point numbers precede all integers.

`+Term1 == +Term2`

Succeeds if *Term1* is equivalent to *Term2*. A variable is only identical to a sharing variable.

`+Term1 \== +Term2`

Equivalent to `\+Term1 == Term2`.

`+Term1 = +Term2`

Unify *Term1* with *Term2*. Succeeds if the unification succeeds.

`unify_with_occurs_check(+Term1, +Term2)`

As `=/2`, but using *sound-unification*. That is, a variable only unifies to a term if this term does not contain the variable itself. To illustrate this, consider the two goals below:

```
1 ?- A = f(A).
```

```
A = f(f(f(f(f(f(f(f(f(...))))))))
```

```
2 ?- unify_with_occurs_check(A, f(A)).
```

No

I.e. the first creates a *cyclic-term*, which is printed as an infinitely nested `f/1` term (see the `max_depth` option of `write_term/2`). The second executes logically sound unification and thus fails.

`+Term1 \= +Term2`

Equivalent to `\+Term1 = Term2`.

`+Term1 ==@ +Term2`

Succeeds if *Term1* is ‘structurally equal’ to *Term2*. Structural equivalence is weaker than equivalence (`==/2`), but stronger than unification (`=/2`). Two terms are structurally equal if their tree representation is identical and they have the same ‘pattern’ of variables. Examples:

<code>a</code>	<code>==@</code>	<code>A</code>	<code>false</code>
<code>A</code>	<code>==@</code>	<code>B</code>	<code>true</code>
<code>x(A,A)</code>	<code>==@</code>	<code>x(B,C)</code>	<code>false</code>
<code>x(A,A)</code>	<code>==@</code>	<code>x(B,B)</code>	<code>true</code>
<code>x(A,B)</code>	<code>==@</code>	<code>x(C,D)</code>	<code>true</code>

`+Term1 \==@ +Term2`

Equivalent to `'\+Term1 ==@ Term2'`.

`+Term1 @< +Term2`

Succeeds if *Term1* is before *Term2* in the standard order of terms.

`+Term1 @=< +Term2`

Succeeds if both terms are equal (`==/2`) or *Term1* is before *Term2* in the standard order of terms.

`+Term1 @> +Term2`

Succeeds if *Term1* is after *Term2* in the standard order of terms.

`+Term1 @>= +Term2`

Succeeds if both terms are equal (`==/2`) or *Term1* is after *Term2* in the standard order of terms.

compare(?*Order*, *+Term1*, *+Term2*)

Determine or test the *Order* between two terms in the standard order of terms. *Order* is one of `<`, `>` or `=`, with the obvious meaning.

4.7 Control Predicates

The predicates of this section implement control structures. Normally these constructs are translated into virtual machine instructions by the compiler. It is still necessary to implement these constructs as true predicates to support meta-calls, as demonstrated in the example below. The predicate finds all currently defined atoms of 1 character long. Note that the cut has no effect when called via one of these predicates (see `!/0`).

```
one_character_atoms(As) :-
    findall(A, (current_atom(A), atom_length(A, 1)), As).
```

fail

Always fail. The predicate `fail/0` is translated into a single virtual machine instruction.

true

Always succeed. The predicate `true/0` is translated into a single virtual machine instruction.

repeat

Always succeed, provide an infinite number of choice points.

!

Cut. Discard choice points of parent frame and frames created after the parent frame. As of SWI-Prolog 3.3, the semantics of the cut are compliant with the ISO standard. This implies that the cut is transparent to `;/2`, `->/2` and `*->/2`. Cuts appearing in the *condition* part of `->/2` and `*->/2` as well as in `\+/1` are local to the condition.⁵

```
t1 :- (a, !, fail ; b).           % cuts a/0 and t1/0
t2 :- (a -> b, ! ; c).           % cuts b/0 and t2/0
t3 :- call((a, !, fail ; b)).    % cuts a/0
t4 :- \+(a, !, fail ; b).       % cuts a/0
```

⁵Up to version 4.0.6, the sequence `X=!`, `X` acted as a true cut. This feature has been deleted for ISO compliance.

+Goal1 , +Goal2

Conjunction. Succeeds if both ‘Goal1’ and ‘Goal2’ can be proved. It is defined as (this definition does not lead to a loop as the second comma is handled by the compiler):

```
Goal1, Goal2 :- Goal1, Goal2.
```

+Goal1 ; +Goal2

The ‘or’ predicate is defined as:

```
Goal1 ; _Goal2 :- Goal1.
_Goal1 ; Goal2 :- Goal2.
```

+Goal1 | +Goal2

Equivalent to `;/2`. Retained for compatibility only. New code should use `;/2`.

+Condition -> +Action

If-then and If-Then-Else. The `->/2` construct commits to the choices made at its left-hand side, destroying choice-points created inside the clause (by `;/2`), or by goals called by this clause. Unlike `!/0`, the choicepoint of the predicate as a whole (due to multiple clauses) is **not** destroyed. The combination `;/2` and `->/2` acts as if defines by:

```
If -> Then; _Else :- If, !, Then.
If -> _Then; Else :- !, Else.
If -> Then :- If, !, Then.
```

Please note that `(If -> Then)` acts as `(If -> Then ; fail)`, making the construct *fail* if the condition fails. This unusual semantics is part of the ISO and all de-facto Prolog standards.

+Condition *-> +Action ; +Else

This construct implements the so-called ‘soft-cut’. The control is defined as follows: If *Condition* succeeds at least once, the semantics is the same as `(Condition, Action)`. If *Condition* does not succeed, the semantics is that of `(\+ Condition, Else)`. In other words, If *Condition* succeeds at least once, simply behave as the conjunction of *Condition* and *Action*, otherwise execute *Else*.

The construct `A *-> B`, i.e. without an *Else* branche, is translated as the normal conjunction `A, B`.⁶

\+ +Goal

Succeeds if ‘Goal’ cannot be proven (mnemonic: + refers to *provable* and the backslash (\) is normally used to indicate negation in Prolog).

⁶BUG: The decompiler implemented by `clause/2` returns this construct as a normal conjunction too.

4.8 Meta-Call Predicates

Meta call predicates are used to call terms constructed at run time. The basic meta-call mechanism offered by SWI-Prolog is to use variables as a subclause (which should of course be bound to a valid goal at runtime). A meta-call is slower than a normal call as it involves actually searching the database at runtime for the predicate, while for normal calls this search is done at compile time.

call(+Goal)

Invoke *Goal* as a goal. Note that clauses may have variables as subclauses, which is identical to `call/1`.

call(+Goal, +ExtraArg1, ...)

Append *ExtraArg1*, *ExtraArg2*, ... to the argument list of *Goal* and call the result. For example, `call(plus(1), 2, X)` will call `plus/3`, binding *X* to 3.

The `call/[2..]` construct is handled by the compiler, which implies that redefinition as a predicate has no effect. The predicates `call/[2-6]` are defined as true predicates, so they can be handled by interpreted code.

apply(+Term, +List)

Append the members of *List* to the arguments of *Term* and call the resulting term. For example: `apply(plus(1), [2, X])` will call `plus(1, 2, X)`. `apply/2` is incorporated in the virtual machine of SWI-Prolog. This implies that the overhead can be compared to the overhead of `call/1`. New code should use `call/[2..]` if the length of *List* is fixed, which is more widely supported and faster because there is no need to build and examine the argument list.

not(+Goal)

Succeeds when *Goal* cannot be proven. Retained for compatibility only. New code should use `\+/1`.

once(+Goal)

Defined as:

```
once(Goal) :-
    Goal, !.
```

`once/1` can in many cases be replaced with `->/2`. The only difference is how the cut behaves (see `!/0`). The following two clauses are identical:

```
1) a :- once((b, c)), d.
2) a :- b, c -> d.
```

ignore(+Goal)

Calls *Goal* as `once/1`, but succeeds, regardless of whether *Goal* succeeded or not. Defined as:

```
ignore(Goal) :-
    Goal, !.
ignore(_).
```

call_with_depth_limit(+Goal, +Limit, -Result)

If *Goal* can be proven without recursion deeper than *Limit* levels, `call_with_depth_limit/3` succeeds, binding *Result* to the deepest recursion level used during the proof. Otherwise, *Result* is unified with `depth_limit_exceeded` if the limit was exceeded during the proof, or the entire predicate fails if *Goal* fails without exceeding *Limit*.

The depth-limit is guarded by the internal machinery. This may differ from the depth computed based on a theoretical model. For example, `true/0` is translated into an inlined virtual machine instruction. Also, `repeat/0` is not implemented as below, but as a non-deterministic foreign predicate.

```
repeat .
repeat :-
    repeat .
```

As a result, `call_with_depth_limit/3` may still loop infinitely on programs that should theoretically finish in finite time. This problem can be cured by using Prolog equivalents to such built-in predicates.

This predicate may be used for theorem-provers to realise techniques like *iterative deepening*. It was implemented after discussion with Steve Moyle `smoyle@ermine.ox.ac.uk`.

call_cleanup(:Goal, +Catcher, :Cleanup)

Calls *Goal*. If *Goal* is completely finished, either by deterministic success, failure, its choice-point being cut or raising an exception and *Catcher* unifies to the termination code (see below), *Cleanup* is called. Success or failure of *Cleanup* is ignored and possibly choicepoints it created are destroyed (as `once/1`). If *cleanup* throws an exception this is executed as normal.⁷

Catcher is unified with a term describing how the call has finished. If this unification fails, *Cleanup* is *not* called.

exit

Goal succeeded without leaving any choicepoints.

fail

Goal failed.

!

Goal succeeded with choicepoints and these are now discarded by the execution of a cut (or other pruning of the search tree such as if-then-else).

exception(Exception)

Goal raised the given *Exception*.

Typical use of this predicate is cleanup of permanent data storage required to execute *Goal*, close file-descriptors, etc. The example below provides a non-deterministic search for a term in a file, closing the stream as needed.

⁷BUG: During the execution of *Cleanup*, garbage collection and stack-shifts are disabled.


```

term_in_file(Term, File) :-
    open(File, read, In),
    call_cleanup(term_in_stream(Term, In), _, close(In)).

term_in_stream(Term, In) :-
    repeat,
    read(In, T),
    (   T == end_of_file
    ->  !, fail
    ;   T = Term
    ).

```

Note that this predicate is impossible to implement in Prolog other than reading all terms into a list, close the file and call `member/2` because without `call_cleanup/3` there is no way to gain control if the choicepoint left by `repeat` is killed by a cut.

This predicate is a SWI-Prolog extension. See also `call_cleanup/2` for compatibility to other Prolog implementations.

call_cleanup(:Goal, :Cleanup)

This predicate is equivalent to `call_cleanup(Goal, _, Cleanup)`, calling *Cleanup* regardless of the reason for termination and without providing information. This predicate provides compatibility to a number of other Prolog implementations.

4.9 ISO compliant Exception handling

SWI-Prolog defines the predicates `catch/3` and `throw/1` for ISO compliant raising and catching of exceptions. In the current implementation (4.0.6), most of the built-in predicates generate exceptions, but some obscure predicates merely print a message, start the debugger and fail, which was the normal behaviour before the introduction of exceptions.

catch(:Goal, +Catcher, :Recover)

Behaves as `call/1` if no exception is raised when executing *Goal*. If an exception is raised using `throw/1` while *Goal* executes, and the *Goal* is the innermost goal for which *Catcher* unifies with the argument of `throw/1`, all choicepoints generated by *Goal* are cut, the system backtracks to the start of `catch/3` while preserving the thrown exception term and *Recover* is called as in `call/1`.

The overhead of calling a goal through `catch/3` is very comparable to `call/1`. Recovery from an exception is much slower, especially if the exception-term is large due to the copying thereof.

throw(+Exception)

Raise an exception. The system looks for the innermost `catch/3` ancestor for which *Exception* unifies with the *Catcher* argument of the `catch/3` call. See `catch/3` for details.

ISO demands `throw/1` to make a copy of *Exception*, walk up the stack to a `catch/3` call, backtrack and try to unify the copy of *Exception* with *Catcher*. SWI-Prolog delays making a copy of *Exception* and backtracking until it actually found a matching `catch/3` goal. The

advantage is that we can start the debugger at the first possible location while preserving the entire exception context if there is no matching `catch/3` goal. This approach can lead to different behaviour if *Goal* and *Catcher* of `catch/3` call share variables. We assume this to be highly unlikely and could not think of a scenario where this is useful.⁸

If an exception is raised in a callback from C (see chapter 7) and not caught in the same call-back, `PL_next_solution()` fails and the exception context can be retrieved using `PL_exception()`.

4.9.1 Debugging and exceptions

Before the introduction of exceptions in SWI-Prolog a runtime error was handled by printing an error message, after which the predicate failed. If the `prolog_flag` (see `current_prolog_flag/2`) `debug_on_error` was in effect (default), the tracer was switched on. The combination of the error message and trace information is generally sufficient to locate the error.

With exception handling, things are different. A programmer may wish to trap an exception using `catch/3` to avoid it reaching the user. If the exception is not handled by user-code, the interactive toplevel will trap it to prevent termination.

If we do not take special precautions, the context information associated with an unexpected exception (i.e, a programming error) is lost. Therefore, if an exception is raised, which is not caught using `catch/3` and the toplevel is running, the error will be printed, and the system will enter trace mode.

If the system is in a non-interactive callback from foreign code and there is no `catch/3` active in the current context, it cannot determine whether or not the exception will be caught by the external routine calling Prolog. It will then base its behaviour on the `prolog_flag debug_on_error`:

- `current_prolog_flag(debug_on_error, false)`
The exception does not trap the debugger and is returned to the foreign routine calling Prolog, where it can be accessed using `PL_exception()`. This is the default.
- `current_prolog_flag(debug_on_error, true)`
If the exception is not caught by Prolog in the current context, it will trap the tracer to help analysing the context of the error.

While looking for the context in which an exception takes place, it is advised to switch on debug mode using the predicate `debug/0`.

4.9.2 The exception term

Builtin predicates generates exceptions using a term `error(Formal, Context)`. The first argument is the ‘formal’ description of the error, specifying the class and generic defined context information. When applicable, the ISO error-term definition is used. The second part describes some additional context to help the programmer while debugging. In its most generic form this is a term of the form `context(Name/Arity, Message)`, where *Name/Arity* describes the built-in predicate that raised the error, and *Message* provides an additional description of the error. Any part of this structure may be a variable if no information was present.

⁸I’d like to acknowledge Bart Demoen for his clarifications on these matters.

4.9.3 Printing messages

The predicate `print_message/2` may be used to print a message term in a human readable format. The other predicates from this section allow the user to refine and extend the message system. The most common usage of `print_message/2` is to print error messages from exceptions. The code below prints errors encountered during the execution of *Goal*, without further propagating the exception and without starting the debugger.

```
... ,
catch(Goal, E,
      ( print_message(error, E),
        fail
      )),
... 
```

Another common use is to defined `message_hook/3` for printing messages that are normally *silent*, suppressing messages, redirecting messages or make something happen in addition to printing the message.

print_message(+Kind, +Term)

The predicate `print_message/2` is used to print messages, notably from exceptions in a human-readable format. *Kind* is one of `informational`, `banner`, `warning`, `error`, `help` or `silent`. A human-readable message is printed to the stream `user_error`.

If the prolog flag (see `current_prolog_flag/2`) `verbose` is `silent`, messages with *Kind* `informational`, or `banner` are treated as `silent`. See `-q`.

This predicate first translates the *Term* into a list of ‘message lines’ (see `print_message_lines/3` for details). Next it will call the hook `message_hook/3` to allow the user intercepting the message. If `message_hook/3` fails it will print the message unless *Kind* is `silent`.

The `print_message/2` predicate and its rules are in the file `<plhome>/boot/messages.pl`, which may be inspected for more information on the error messages and related error terms.

See also `message_to_string/2`.

print_message_lines(+Stream, +Prefix, +Lines)

Print a message (see `print_message/2`) that has been translated to a list of message elements. The elements of this list are:

`<Format>-<Args>`

Where *Format* is an atom and *Args* is a list of format argument. Handed to `format/3`.

`flush`

If this appears as the last element, *Stream* is flushed (see `flush_output/1`) and no final newline is generated.

`at_same_line`

If this appears as first element, no prefix is printed for the first line and the line-position is not forced to 0 (see `format/1`, `~N`).

<Format>

Handed to `format/3` as `format(Stream, Format, [])`.

nl

A new line is started and if the message is not complete the *Prefix* is printed too.

See also `print_message/2` and `message_hook/3`.

message_hook(+Term, +Kind, +Lines)

Hook predicate that may be define in the module `user` to intercept messages from `print_message/2`. *Term* and *Kind* are the same as passed to `print_message/2`. *Lines* is a list of format statements as described with `print_message_lines/3`. See also `message_to_string/2`.

This predicate should be defined dynamic and multifile to allow other modules defining clauses for it too.

message_to_string(+Term, -String)

Translates a message-term into a string object (see section 4.23. Primarily intended to write messages to Windows in XPCE (see section 1.5) or other GUI environments.

4.10 Handling signals

As of version 3.1.0, SWI-Prolog is capable to handle software interrupts (signals) in Prolog as well as in foreign (C) code (see section 7.6.10).

Signals are used to handle internal errors (execution of a non-existing CPU instruction, arithmetic domain errors, illegal memory access, resource overflow, etc.), as well as for dealing asynchronous inter-process communication.

Signals are defined by the POSIX standard and part of all Unix machines. The MS-Windows Win32 provides a subset of the signal handling routines, lacking the vital functionality to raise a signal in another thread for achieving asynchronous inter-process (or inter-thread) communication (Unix `kill()` function).

on_signal(+Signal, -Old, :New)

Determines the reaction on *Signal*. *Old* is unified with the old behaviour, while the behaviour is switched to *New*. As with similar environment-control predicates, the current value is retrieved using `on_signal(Signal, Current, Current)`.

The action description is an atom denoting the name of the predicate that will be called if *Signal* arrives. `on_signal/3` is a meta predicate, which implies that `<Module>:<Name>` refers the `<Name>/1` in the module `<Module>`.

Two predicate-names have special meaning. `throw` implies Prolog will map the signal onto a Prolog exception as described in section 4.9. `default` resets the handler to the settings active before SWI-Prolog manipulated the handler.

Signals bound to a foreign function through `PL_signal()` are reported using the term `$foreign_function(Address)`.

After receiving a signal mapped to `throw`, the exception raised has the structure

```
error(signal(<SigName>, <SigNum>), <Context>)
```

One possible usage of this is, for example, to limit the time spent on proving a goal. This requires a little C-code for setting the alarm timer (see chapter 7):

```
#include <SWI-Prolog.h>
#include <unistd.h>

foreign_t
pl_alarm(term_t time)
{ double t;

  if ( PL_get_float(time, &t) )
    { alarm((long)(t+0.5));

      PL_succeed;
    }

  PL_fail;
}

install_t
install()
{ PL_register_foreign("alarm", 1, pl_alarm, 0);
}
```

Next, we can define the Prolog below. This will run *Goal* just as once/1, throwing the exception `error(signal(alm, _), _)` if a timeout occurs.⁹

```
:- load_foreign_library(alarm).

:- on_signal(alm, _, throw).

:- module_transparent
    call_with_time_limit/2.

call_with_time_limit(Goal, MaxTime) :-
    alarm(MaxTime),
    call_cleanup(Goal, _, alarm(0)), !.
```

The signal names are defined by the C-Posix standards as symbols of the form `SIG_⟨SIGNAME⟩`. The Prolog name for a signal is the lowercase version of `⟨SIGNAME⟩`. The predicate `current_signal/3` may be used to map between names and signals.

Initially, some signals are mapped to `throw`, while all other signals are default. The following signals throw an exception: `ill`, `fpe`, `segv`, `pipe`, `alm`, `bus`, `xcpu`, `xfsz` and `vtalm`.

⁹Note that `call_with_depth_limit/2` is defined in `time`, part of the ‘clib’ package. The version provided in this library runs on POSIX systems as well as MS-Windows and can schedule multiple concurrent alarms.

current_signal(?Name, ?Id, ?Handler)

Enumerate the currently defined signal handling. *Name* is the signal name, *Id* is the numerical identifier and *Handler* is the currently defined handler (see `on_signal/3`).

4.10.1 Notes on signal handling

Before deciding to deal with signals in your application, please consider the following:

- *Portability*
On MS-Windows, the signal interface is severely limited. Different Unix brands support different sets of signals, and the relation between signal name and number may vary.
- *Safety*
Signal handling is not completely safe in the current implementation, especially if `throw` is used in combination with external foreign code. The system will use the C `longjmp()` construct to direct control to the innermost `PL_next_solution()`, thus forcing an external procedure to be abandoned at an arbitrary moment. Most likely not all SWI-Prologs own foreign code is (yet) safe too. For the multi-threaded versions this is even worse: signals can easily violate thread synchronisation consistency.

The C-interface described in section 7.6.10 provides the option `PL_SIGSYNC` for registering a signal handler that delays delivery of signals to a safe point. Unfortunately this may cause signals to be delayed for a long time if Prolog is executing foreign code.
- *Garbage Collection*
The garbage collector will block all signals that are handled by Prolog. While handling a signal, the garbage-collector is disabled.
- *Time of delivery*
Normally delivery is immediate (or as defined by the operating system used). Signals are blocked when the garbage collector is active, and internally delayed if they occur within in a ‘critical section’. The critical sections are generally very short.

4.11 The ‘block’ control-structure

The `block/3` predicate and friends have been introduced before ISO compatible `catch/3` exception handling for compatibility with some Prolog implementation. The only feature not covered by `catch/3` and `throw/1` is the possibility to execute global cuts. New code should use `catch/3` and `throw/1` to deal with exceptions.

block(+Label, +Goal, -ExitValue)

Execute *Goal* in a *block*. *Label* is the name of the block. *Label* is normally an atom, but the system imposes no type constraints and may even be a variable. *ExitValue* is normally unified to the second argument of an `exit/2` call invoked by *Goal*.

exit(+Label, +Value)

Calling `exit/2` makes the innermost *block* which *Label* unifies exit. The block’s *ExitValue* is unified with *Value*. If this unification fails the block fails.

fail(+Label)

Calling `fail/1` makes the innermost *block* which *Label* unifies fail immediately. Implemented as

```
fail(Label) :- !(Label), fail.
```

!(+Label)

Cut all choice-points created since the entry of the innermost *block* which *Label* unifies.

4.12 DCG Grammar rules

Grammar rules form a comfortable interface to *difference-lists*. They are designed both to support writing parsers that build a parse-tree from a list as for generating a flat list from a term. Unfortunately, Definite Clause Grammar (DCG) handling is not part of the Prolog standard. Most Prolog engines implement DCG, but the details differ slightly.

Grammar rules look like ordinary clauses using `-->/2` for separating the head and body rather than `:-/2`. Expanding grammar rules is done by `expand_term/2`, which adds two additional arguments to each term for representing the difference list. We will illustrate the behaviour by defining a rule-set for parsing an integer.

```
integer(I) -->
    digit(D0),
    digits(D),
    { number_chars(I, [D0|D])
    }.

digits([D|T]) -->
    digit(D), !,
    digits(T).
digits([]) -->
    [].

digit(D) -->
    [D],
    { code_type(D, digit)
    }.
```

The body of a grammar rule can contain three types of terms. A compound term interpreted as a reference to a grammar-rule. Code between `{...}` is interpreted as a reference to ordinary Prolog code and finally, a list is interpreted as a sequence of literals. The Prolog control-constructs (`\+/1`, `->/2`, `://2`, `/2` and `!/0`) can be used in grammar rules.

Grammar rule-sets are called using the builtin predicates `phrase/2` and `phrase/3`:

phrase(+RuleSet, +InputList)

Equivalent to `phrase(RuleSet, InputList, [])`.

phrase(+RuleSet, +InputList, -Rest)

Activate the rule-set with given name. ‘InputList’ is the list of tokens to parse, ‘Rest’ is unified with the remaining tokens if the sentence is parsed correctly. The example below calls the rule-set ‘integer’ defined above.

```
?- phrase(integer(X), "42 times", Rest).
```

```
X = 42
```

```
Rest = [32, 116, 105, 109, 101, 115]
```

4.13 Database

SWI-Prolog offers three different database mechanisms. The first one is the common assert/retract mechanism for manipulating the clause database. As facts and clauses asserted using `assert/1` or one of its derivatives become part of the program these predicates compile the term given to them. `retract/1` and `retractall/1` have to unify a term and therefore have to decompile the program. For these reasons the assert/retract mechanism is expensive. On the other hand, once compiled, queries to the database are faster than querying the recorded database discussed below. See also `dynamic/1`.

The second way of storing arbitrary terms in the database is using the “recorded database”. In this database terms are associated with a *key*. A key can be an atom, integer or term. In the last case only the functor and arity determine the key. Each key has a chain of terms associated with it. New terms can be added either at the head or at the tail of this chain. This mechanism is considerably faster than the assert/retract mechanism as terms are not compiled, but just copied into the heap.

The third mechanism is a special purpose one. It associates an integer or atom with a key, which is an atom, integer or term. Each key can only have one atom or integer associated with it. It is faster than the mechanisms described above, but can only be used to store simple status information like counters, etc.

abolish(:PredicateIndicator)

Removes all clauses of a predicate with functor *Functor* and arity *Arity* from the database. All predicate attributes (dynamic, multifile, index, etc.) are reset to their defaults. Abolishing an imported predicate only removes the import link; the predicate will keep its old definition in its definition module.

According to the ISO standard, `abolish/1` can only be applied to dynamic procedures. This is odd, as for dealing with dynamic procedures there is already `retract/1` and `retractall/1`. The `abolish/1` predicate has been introduced in DEC-10 Prolog precisely for dealing with static procedures. In SWI-Prolog, `abolish/1` works on static procedures, unless the prolog flag `iso` is set to `true`.

It is advised to use `retractall/1` for erasing all clauses of a dynamic predicate.

abolish(+Name, +Arity)

Same as `abolish(Name/Arity)`. The predicate `abolish/2` conforms to the Edinburgh standard, while `abolish/1` is ISO compliant.

redefine_system_predicate(+Head)

This directive may be used both in module `user` and in normal modules to redefine any system predicate. If the system definition is redefined in module `user`, the new definition is the default definition for all sub-modules. Otherwise the redefinition is local to the module. The system definition remains in the module `system`.

Redefining system predicate facilitates the definition of compatibility packages. Use in other context is discouraged.

retract(+Term)

When *Term* is an atom or a term it is unified with the first unifying fact or clause in the database. The fact or clause is removed from the database.

retractall(+Head)

All facts or clauses in the database for which the *head* unifies with *Head* are removed.

assert(+Term)

Assert a fact or clause in the database. *Term* is asserted as the last fact or clause of the corresponding predicate.

asserta(+Term)

Equivalent to `assert/1`, but *Term* is asserted as first clause or fact of the predicate.

assertz(+Term)

Equivalent to `assert/1`.

assert(+Term, -Reference)

Equivalent to `assert/1`, but *Reference* is unified with a unique reference to the asserted clause. This key can later be used with `clause/3` or `erase/1`.

asserta(+Term, -Reference)

Equivalent to `assert/2`, but *Term* is asserted as first clause or fact of the predicate.

assertz(+Term, -Reference)

Equivalent to `assert/2`.

recorda(+Key, +Term, -Reference)

Assert *Term* in the recorded database under key *Key*. *Key* is an integer, atom or term. *Reference* is unified with a unique reference to the record (see `erase/1`).

recorda(+Key, +Term)

Equivalent to `recorda(Key, Value, -)`.

recordz(+Key, +Term, -Reference)

Equivalent to `recorda/3`, but puts the *Term* at the tail of the terms recorded under *Key*.

recordz(+Key, +Term)

Equivalent to `recordz(Key, Value, -)`.

recorded(+Key, -Value, -Reference)

Unify *Value* with the first term recorded under *Key* which does unify. *Reference* is unified with the memory location of the record.

recorded(+Key, -Value)

Equivalent to `recorded(Key, Value, -)`.

erase(+Reference)

Erase a record or clause from the database. *Reference* is an integer returned by `recorda/3` or `recorded/3`, `clause/3`, `assert/2`, `asserta/2` or `assertz/2`. Other integers might conflict with the internal consistency of the system. Erase can only be called once on a record or clause. A second call also might conflict with the internal consistency of the system.¹⁰

flag(+Key, -Old, +New)

Key is an atom, integer or term. Unify *Old* with the old value associated with *Key*. If the key is used for the first time *Old* is unified with the integer 0. Then store the value of *New*, which should be an integer, float, atom or arithmetic expression, under *Key*. `flag/3` is a very fast mechanism for storing simple facts in the database. Example:

```
:- module_transparent succeeds_n_times/2.

succeeds_n_times(Goal, Times) :-
    (   flag(succeeds_n_times, Old, 0),
        Goal,
        flag(succeeds_n_times, N, N+1),
        fail
    ;   flag(succeeds_n_times, Times, Old)
    ).
```

4.13.1 Update view

Traditionally, Prolog systems used the *immediate update view*: new clauses became visible to predicates backtracking over dynamic predicates immediately and retracted clauses became invisible immediately.

Starting with SWI-Prolog 3.3.0 we adhere the *logical update view*, where backtrackable predicates that enter the definition of a predicate will not see any changes (either caused by `assert/1` or `retract/1`) to the predicate. This view is the ISO standard, the most commonly used and the most ‘safe’.¹¹ Logical updates are realised by keeping reference-counts on predicates and *generation* information on clauses. Each change to the database causes an increment of the generation of the database. Each goal is tagged with the generation in which it was started. Each clause is flagged with the generation it was created as well as the generation it was erased. Only clauses with ‘created’ ... ‘erased’ interval that encloses the generation of the current goal are considered visible.

4.13.2 Indexing databases

By default, SWI-Prolog, as most other implementations, indexes predicates on their first argument. SWI-Prolog allows indexing on other and multiple arguments using the declaration `index/1`.

For advanced database indexing, it defines `hash_term/2`:

¹⁰BUG: The system should have a special type for pointers, thus avoiding the Prolog user having to worry about consistency matters. Currently some simple heuristics are used to determine whether a reference is valid.

¹¹For example, using the immediate update view, no call to a dynamic predicate is deterministic.

hash_term(+Term, -HashKey)

If *Term* is a ground term (see `ground/1`), *HashKey* is unified with a positive integer value that may be used as a hash-key to the value. If *Term* is not ground, the predicate succeeds immediately, leaving *HashKey* an unbound variable.

This predicate may be used to build hash-tables as well as to exploit argument-indexing to find complex terms more quickly.

The hash-key does not rely on temporary information like addresses of atoms and may be assumed constant over different invocations of SWI-Prolog.

4.14 Declaring predicates properties

This section describes directives which manipulate attributes of predicate definitions. The functors `dynamic/1`, `multifile/1` and `discontiguous/1` are operators of priority 1150 (see `op/3`), which implies the list of predicates they involve can just be a comma separated list:

```
:- dynamic
    foo/0,
    baz/2.
```

On SWI-Prolog all these directives are just predicates. This implies they can also be called by a program. Do not rely on this feature if you want to maintain portability to other Prolog implementations.

dynamic +Name/+Arity, ...

Informs the interpreter that the definition of the predicate(s) may change during execution (using `assert/1` and/or `retract/1`). In the multi-threaded version, the clauses of dynamic predicates are shared between the threads. The directive `thread_local/1` provides an alternative where each threads has its own clause-list for the predicate. Dynamic predicates can be turned into static ones using `compile_predicates/1`.

compile_predicates(:ListOfNameArity)

Compile a list of specified dynamic predicates (see `dynamic/1` and `assert/1`) into normal static predicates. This call tells the Prolog environment the definition will not change anymore and further calls to `assert/1` or `retract/1` on the named predicates raise a permission error. This predicate is designed to deal with parts of the program that is generated at runtime but does not change during the remainder of the program execution.¹²

multifile +Name/+Arity, ...

Informs the system that the specified predicate(s) may be defined over more than one file. This stops `consult/1` from redefining a predicate when a new definition is found.

discontiguous +Name/+Arity, ...

Informs the system that the clauses of the specified predicate(s) might not be together in the source file. See also `style_check/1`.

¹²The specification of this predicate is from Richard O'Keefe. The implementation is allowed to optimise the predicate. This is not yet implemented. In multi-threaded Prolog however, static code runs faster as it does not require synchronization. This is particularly true on SMP hardware.

index(+Head)

Index the clauses of the predicate with the same name and arity as *Head* on the specified arguments. *Head* is a term of which all arguments are either ‘1’ (denoting ‘index this argument’) or ‘0’ (denoting ‘do not index this argument’). Indexing has no implications for the semantics of a predicate, only on its performance. If indexing is enabled on a predicate a special purpose algorithm is used to select candidate clauses based on the actual arguments of the goal. This algorithm checks whether indexed arguments might unify in the clause head. Only atoms, integers and compound terms are considered. Compound terms are indexed on the combination of their name and arity. Indexing is very useful for predicates with many clauses representing facts.

Due to the representation technique used at most 4 arguments can be indexed. All indexed arguments should be in the first 32 arguments of the predicate. If more than 4 arguments are specified for indexing only the first 4 will be accepted. Arguments above 32 are ignored for indexing.

By default all predicates with $\langle \text{arity} \rangle \geq 1$ are indexed on their first argument. It is possible to redefine indexing on predicates that already have clauses attached to them. This will initiate a scan through the predicates clause list to update the index summary information stored with each clause.

If—for example—one wants to represent sub-types using a fact list ‘sub_type(Sub, Super)’ that should be used both to determine sub- and super types one should declare sub_type/2 as follows:

```
:- index(sub_type(1, 1)).

sub_type(horse, animal).
...
...
```

4.15 Examining the program**current_atom(-Atom)**

Successively unifies *Atom* with all atoms known to the system. Note that current_atom/1 always succeeds if *Atom* is instantiated to an atom.

current_functor(?Name, ?Arity)

Successively unifies *Name* with the name and *Arity* with the arity of functors known to the system.

current_flag(-FlagKey)

Successively unifies *FlagKey* with all keys used for flags (see flag/3).

current_key(-Key)

Successively unifies *Key* with all keys used for records (see recorda/3, etc.).

current_predicate(?Name, ?Head)

Successively unifies *Name* with the name of predicates currently defined and *Head* with the most general term built from *Name* and the arity of the predicate. This predicate succeeds for

all predicates defined in the specified module, imported to it, or in one of the modules from which the predicate will be imported if it is called.

current_predicate(:*Name/Arity*)

ISO compliant implementation of `current_predicate/2`. Unlike `current_predicate/2`, the current implementation of `current_predicate/1` does not consider predicates that can be autoloaded 'current'.

predicate_property(?*Head*, ?*Property*)

Succeeds if *Head* refers to a predicate that has property *Property*. Can be used to test whether a predicate has a certain property, obtain all properties known for *Head*, find all predicates having *property* or even obtaining all information available about the current program. *Property* is one of:

interpreted

Is true if the predicate is defined in Prolog. We return true on this because, although the code is actually compiled, it is completely transparent, just like interpreted code.

built_in

Is true if the predicate is locked as a built-in predicate. This implies it cannot be redefined in its definition module and it can normally not be seen in the tracer.

foreign

Is true if the predicate is defined in the C language.

dynamic

Is true if `assert/1` and `retract/1` may be used to modify the predicate. This property is set using `dynamic/1`.

volatile

If true, the clauses are not saved into a saved-state by `qsave_program/[1,2]`. This property is set using `volatile/1`.

thread_local

If true (only possible on the multi-threaded version) each thread has its own clauses for the predicate. This property is set using `thread_local/1`.

multifile

Is true there may be multiple (or no) file providing clauses for the predicate. This property is set using `multifile/1`.

undefined

Is true if a procedure definition block for the predicate exists, but there are no clauses for it and it is not declared dynamic or multifile. This is true if the predicate occurs in the body of a loaded predicate, an attempt to call it has been made via one of the meta-call predicates or the predicate had a definition in the past. See the library package `check` for example usage.

transparent

Is true if the predicate is declared transparent using the `module.transparent/1` declaration.

exported

Is true if the predicate is in the public list of the context module.

imported_from(*Module*)

Is true if the predicate is imported into the context module from module *Module*.

indexed(*Head*)

Predicate is indexed (see `index/1`) according to *Head*. *Head* is a term whose name and arity are identical to the predicate. The arguments are unified with ‘1’ for indexed arguments, ‘0’ otherwise.

file(*FileName*)

Unify *FileName* with the name of the source file in which the predicate is defined. See also `source_file/2`.

line_count(*LineNumber*)

Unify *LineNumber* with the line number of the first clause of the predicate. Fails if the predicate is not associated with a file. See also `source_file/2`.

number_of_clauses(*ClauseCount*)

Unify *ClauseCount* to the number of clauses associated with the predicate. Fails for foreign predicates.

dwim_predicate(+*Term*, -*Dwim*)

‘Do What I Mean’ (‘dwim’) support predicate. *Term* is a term, which name and arity are used as a predicate specification. *Dwim* is instantiated with the most general term built from *Name* and the arity of a defined predicate that matches the predicate specified by *Term* in the ‘Do What I Mean’ sense. See `dwim_match/2` for ‘Do What I Mean’ string matching. Internal system predicates are not generated, unless `style_check(+dollar)` is active. Backtracking provides all alternative matches.

clause(?*Head*, ?*Body*)

Succeeds when *Head* can be unified with a clause head and *Body* with the corresponding clause body. Gives alternative clauses on backtracking. For facts *Body* is unified with the atom *true*. Normally `clause/2` is used to find clause definitions for a predicate, but it can also be used to find clause heads for some body template.

clause(?*Head*, ?*Body*, ?*Reference*)

Equivalent to `clause/2`, but unifies *Reference* with a unique reference to the clause (see also `assert/2`, `erase/1`). If *Reference* is instantiated to a reference the clause’s head and body will be unified with *Head* and *Body*.

nth_clause(?*Pred*, ?*Index*, ?*Reference*)

Provides access to the clauses of a predicate using their index number. Counting starts at 1. If *Reference* is specified it unifies *Pred* with the most general term with the same name/arity as the predicate and *Index* with the index-number of the clause. Otherwise the name and arity of *Pred* are used to determine the predicate. If *Index* is provided *Reference* will be unified with the clause reference. If *Index* is unbound, backtracking will yield both the indices and the references of all clauses of the predicate. The following example finds the 2nd clause of `member/2`:

```
?- nth_clause(member(_,_), 2, Ref), clause(Head, Body, Ref).
```

```
Ref = 160088
```

```
Head = system : member(G575, [G578|G579])
Body = member(G575, G579)
```

clause_property(+ClauseRef, -Property)

Queries properties of a clause. *ClauseRef* is a reference to a clause as produced by `clause/3`, `nth_clause/3` or `prolog_frame_attribute/3`. *Property* is one of the following:

file(FileName)

Unify *FileName* with the name of the source file in which the clause is defined. Fails if the clause is not associated to a file.

line_count(LineNumber)

Unify *LineNumber* with the line number of the clause. Fails if the clause is not associated to a file.

fact

True if the clause has no body.

erased

True if the clause has been erased, but not yet reclaimed because it is referenced.

4.16 Input and output

SWI-Prolog provides two different packages for input and output. The native I/O system is based on the ISO standard predicates `open/3`, `close/1` and `friends`.¹³ Being more widely portable and equipped with a clearer and more robust specification, new code is encouraged to use these predicates for manipulation of I/O streams.

Section 4.16.2 describes `tell/1`, `see/1` and `friends`, providing I/O in the spirit of the outdated Edinburgh standard. These predicates are layered on top of the ISO predicates. Both packages are fully integrated; the user may switch freely between them.

4.16.1 ISO Input and Output Streams

The predicates described in this section provide ISO compliant I/O, where streams are explicitly created using the predicate `open/3`. The resulting stream identifier is then passed as a parameter to the reading and writing predicates to specify the source or destination of the data.

This schema is not vulnerable to filename and stream ambiguities as well as changes to the working directory. New code is advised to use these predicates to manage input and output streams.

open(+SrcDest, +Mode, -Stream, +Options)

ISO compliant predicate to open a stream. *SrcDes* is either an atom, specifying a Unix file, or a term `'pipe(Command)'`, like `see/1` and `tell/1`. *Mode* is one of `read`, `write`, `append` or `update`. Mode `append` opens the file for writing, positioning the file-pointer at the end. Mode `update` opens the file for writing, positioning the file-pointer at the beginning of the file without truncating the file. *Stream* is either a variable, in which case it is bound to an integer identifying the stream, or an atom, in which case this atom will be the stream identifier.¹⁴ The *Options* list can contain the following options:

¹³Actually based on Quintus Prolog, providing this interface before the ISO standard existed.

¹⁴New code should use the `alias(Alias)` option for compatibility to the ISO standard

type(*Type*)

Using type `text` (default), Prolog will write a text-file in an operating-system compatible way. Using type `binary` the bytes will be read or written without any translation. Note there is no difference between the two on Unix systems.

alias(*Atom*)

Gives the stream a name. Below is an example. Be careful with this option as stream-names are global. See also `set_stream/2`.

```
?- open(data, read, Fd, [alias(input)]).

    ...
    read(input, Term),
    ...
```

eof_action(*Action*)

Defines what happens if the end of the input stream is reached. Action `eof_code` makes `get0/1` and friends return `-1` and `read/1` and friends return the atom `end_of_file`. Repetitive reading keeps yielding the same result. Action `error` is like `eof_code`, but repetitive reading will raise an error. With action `reset`, Prolog will examine the file again and return more data if the file has grown.

buffer(*Buffering*)

Defines output buffering. The atom `full` (default) defines full buffering, `line` buffering by line, and `false` implies the stream is fully unbuffered. Smaller buffering is useful if another process or the user is waiting for the output as it is being produced. See also `flush_output/[0,1]`. This option is not an ISO option.

close_on_abort(*Bool*)

If `true` (default), the stream is closed on an abort (see `abort/0`). If `false`, the stream is not closed. If it is an output stream, it will be flushed however. Useful for logfiles and if the stream is associated to a process (using the `pipe/1` construct).

lock(*LockingMode*)

Try to obtain a lock on the open file. Default is `none`, which does not lock the file. The value `read` or `shared` means other processes may read the file, but not write it. The value `write` or `exclusive` means no other process may read or write the file.

Locks are acquired through the POSIX function `fcntl()` using the command `F_SETLKW`, which makes a blocked call wait for the lock to be released. Please note that `fcntl()` locks are *advisory* and therefore only other applications using the same advisory locks honour your lock. As there are many issues around locking in Unix, especially related to NFS (network file system), please study the `fcntl()` manual page before trusting your locks!

The `lock` option is a SWI-Prolog extension.

The option `reposition` is not supported in SWI-Prolog. All streams connected to a file may be repositioned.

open(+*SrcDest*, +*Mode*, ?*Stream*)

Equivalent to `open/4` with an empty option-list.

open_null_stream(?Stream)

Open a stream that produces no output. All counting functions are enabled on such a stream. An attempt to read from a null-stream will immediately signal end-of-file. Similar to Unix `/dev/null`. *Stream* can be an atom, giving the null-stream an alias name.

close(+Stream)

Close the specified stream. If *Stream* is not open an error message is displayed. If the closed stream is the current input or output stream the terminal is made the current input or output.

close(+Stream, +Options)

Provides `close(Stream, [force(true)])` as the only option. Called this way, any resource error (such as write-errors while flushing the output buffer) are ignored.

stream_property(?Stream, ?StreamProperty)

ISO compatible predicate for querying status of open I/O streams. *StreamProperty* is one of:

file_name(Atom)

If *Stream* is associated to a file, unify *Atom* to the name of this file.

mode(IOMode)

Unify *IOMode* to the mode given to `open/4` for opening the stream. Values are: `read`, `write`, `append` and the SWI-Prolog extension `update`.

input

True if *Stream* has mode `read`.

output

True if *Stream* has mode `write`, `append` or `update`.

alias(Atom)

If *Atom* is bound, test if the stream has the specified alias. Otherwise unify *Atom* with the first alias of the stream.¹⁵

position(Term)

Unify *Term* with the current stream-position. A stream-position is a term of format `$stream_position(CharIndex, LineNo, LinePos)`. See also `set_stream_position/2`.

end_of_stream(E)

If *Stream* is an input stream, unify *E* with one of the atoms `not`, `at` or `past`. See also `at_end_of_stream/[0,1]`.

eof_action(A)

Unify *A* with one of `eof_code`, `reset` or `error`. See `open/4` for details.

reposition(Bool)

Unify *Bool* with `true` if the position of the stream can be set (see `seek/4`). It is assumed the position can be set if the stream has a *seek-function* and is not based on a POSIX file-descriptor that is not associated to a regular file.

type(T)

Unify *Bool* with `text` or `binary`.

¹⁵BUG: Backtracking does not give other aliases.

file_no(*Integer*)

If the stream is associated with a POSIX file-descriptor, unify *Integer* with the descriptor number. SWI-Prolog extension used primarily for integration with foreign code. See also `Sfileno()` from `SWI-Stream.h`.

buffer(*Buffering*)

SWI-Prolog extension to query the buffering mode of this stream. *Buffering* is one of `full`, `line` or `false`. See also `open/4`.

current_stream(?*Object*, ?*Mode*, ?*Stream*)

The predicate `current_stream/3` is used to access the status of a stream as well as to generate all open streams. *Object* is the name of the file opened if the stream refers to an open file, an integer file-descriptor if the stream encapsulates an operating-system stream or the atom `[]` if the stream refers to some other object. *Mode* is one of `read` or `write`.

set_stream_position(+*Stream*, +*Pos*)

Set the current position of *Stream* to *Pos*. *Pos* is a term as returned by `stream_property/2` using the `position(Pos)` property. See also `seek/4`.

seek(+*Stream*, +*Offset*, +*Method*, -*NewLocation*)

Reposition the current point of the given *Stream*. *Method* is one of `bof`, `current` or `eof`, indicating positioning relative to the start, current point or end of the underlying object. *NewLocation* is unified with the new offset, relative to the start of the stream.

If the seek modifies the current location, the line number and character position in the line are set to 0.

If the stream cannot be repositioned, a `reposition` error is raised. The predicate `seek/4` is compatible to Quintus Prolog, though the error conditions and signalling is ISO compliant. See also `stream_property/2` and `set_stream_position/2`.

set_stream(+*Stream*, +*Attribute*)

Modify an attribute of an existing stream. *Attribute* specifies the stream property to set. See also `stream_property/2` and `open/4`.

alias(*AliasName*)

Set the alias of an already created stream. If *AliasName* is the name of one of the standard streams is used, this stream is rebound. Thus, `set_stream(S, current_input)` is the same as `set_input/1` and by setting the alias of a stream to `user_input`, etc. all user terminal input is read from this stream. See also `interactor/0`.

buffer(*Buffering*)

Set the buffering mode of an already created stream. *Buffering* is one of `full`, `line` or `false`.

eof_action(*Action*)

Set end-of-file handling to one of `eof_code`, `reset` or `error`.

close_on_abort(*Bool*)

Determine whether or not the stream is closed by `abort/0`. By default streams are closed.

timeout(Seconds)

This option can be used to make streams generate an exception if it takes longer than *Seconds* before any new data arrives at the stream. The value *infinite* (default) makes the stream block indefinitely. Like `wait_for_input/3`, this call only applies to streams that support the `select()` system call. For further information about timeout handling, see `wait_for_input/3`. The exception is of the form

```
error(timeout_error(read, Stream), _)
```

record_position(Bool)

Do/do not record the line-count and line-position (see `line_count/2` and `line_position/2`).

file_name(FileName)

Set the file name associated to this stream. This call can be used to set the file for error-locations if *Stream* corresponds to *FileName* and is not obtained by opening the file directly but, for example, through a network service.

4.16.2 Edinburgh-style I/O

The package for implicit input and output destination is (almost) compatible to Edinburgh DEC-10 and C-Prolog. The reading and writing predicates refer to resp. the *current* input- and output stream. Initially these streams are connected to the terminal. The current output stream is changed using `tell/1` or `append/1`. The current input stream is changed using `see/1`. The streams current value can be obtained using `telling/1` for output- and `seeing/1` for input streams.

Source and destination are either a file, `user`, or a term `'pipe(Command)'`. The reserved stream name `user` refers to the terminal.¹⁶ In the predicate descriptions below we will call the source/destination argument *'SrcDest'*. Below are some examples of source/destination specifications.

```
?- see(data).           % Start reading from file 'data'.
?- tell(user).         % Start writing to the terminal.
?- tell(pipe(lpr)).    % Start writing to the printer.
```

Another example of using the `pipe/1` construct is shown below. Note that the `pipe/1` construct is not part of Prolog's standard I/O repertoire.

```
getwd(Wd) :-
    seeing(Old), see(pipe(pwd)),
    collect_wd(String),
    seen, see(Old),
    atom_codes(Wd, String).
```

```
collect_wd([C|R]) :-
    get0(C), C \== -1, !,
    collect_wd(R).
```

```
collect_wd([]).
```

¹⁶The ISO I/O layer uses `user_input`, `user_output` and `user_error`.

Compatibility notes

Unlike Edinburgh Prolog systems, `telling/1` and `seeing/1` do not return the filename of the current input/output, but the stream-identifier, to ensure the design pattern below works under all circumstances.¹⁷

```

... ,
telling(Old), tell(x),
... ,
told, tell(Old),
... ,

```

The predicates `tell/1` and `see/1` first check for `user`, the `pipe(command)` and a stream-handle. Otherwise, if the argument is an atom it is first compared to open streams associated to a file with *exactly* the same name. If such a stream, created using `tell/1` or `see/1` exists, output (input) is switch to the open stream. Otherwise a file with the specified name is opened.

The behaviour is compatible to Edinburgh Prolog. This is not without problems. Changing directory, non-file streams, multiple names referring to the same file easily lead to unexpected behaviour. New code, especially when managing multiple I/O channels should consider using the ISO I/O predicates defined in section 4.16.1.

see(+SrcDest)

Open *SrcDest* for reading and make it the current input (see `set_input/1`). If *SrcDest* is a stream-handle, just makes this stream the current input. See the introduction of section 4.16.2 for details.

tell(+SrcDest)

Open *SrcDest* for writing and make it the current output (see `set_output/1`). If *SrcDest* is a stream-handle, just makes this stream the current output. See the introduction of section 4.16.2 for details.

append(+File)

Similar to `tell/1`, but positions the file pointer at the end of *File* rather than truncating an existing file. The pipe construct is not accepted by this predicate.

seeing(?SrcDest)

Same as `current_input/1`, except that `user` is returned if the current input is the stream `user_input` to improve compatibility with traditional Edinburgh I/O. See the introduction of section 4.16.2 for details.

telling(?SrcDest)

Same as `current_output/1`, except that `user` is returned if the current output is the stream `user_output` to improve compatibility with traditional Edinburgh I/O. See the introduction of section 4.16.2 for details.

seen

Close the current input stream. The new input stream becomes *user_input*.

told

Close the current output stream. The new output stream becomes *user_output*.

¹⁷Filenames can be ambiguous and SWI-Prolog streams can refer to much more than just files.

4.16.3 Switching Between Edinburgh and ISO I/O

The predicates below can be used for switching between the implicit- and the explicit stream based I/O predicates.

set_input(+Stream)

Set the current input stream to become *Stream*. Thus, `open(file, read, Stream), set_input(Stream)` is equivalent to `see(file)`.

set_output(+Stream)

Set the current output stream to become *Stream*.

current_input(-Stream)

Get the current input stream. Useful to get access to the status predicates associated with streams.

current_output(-Stream)

Get the current output stream.

4.17 Status of streams

wait_for_input(+ListOfStreams, -ReadyList, +Timeout)

Wait for input on one of the streams in *ListOfStreams* and return a list of streams on which input is available in *ReadyList*. `wait_for_input/3` waits for at most *Timeout* seconds. *Timeout* may be specified as a floating point number to specify fractions of a second. If *Timeout* equals infinite, `wait_for_input/3` waits indefinitely.¹⁸

This predicate can be used to implement timeout while reading and to handle input from multiple sources. The following example will wait for input from the user and an explicitly opened second terminal. On return, *Inputs* may hold `user` or `P4` or both.

```
?- open('/dev/tty4', read, P4),
   wait_for_input([user, P4], Inputs, 0).
```

This predicate relies on the `select()` call on most operating systems. On Unix this call is implemented for any stream referring to a file-handle, which implies all OS-based streams: sockets, terminals, pipes, etc. On non-Unix systems `select()` is generally only implemented for socket-based streams. See also `socket` from the `clib` package.

Note that `wait_for_input/3` returns streams that have data waiting. This does not mean you can, for example, call `read/2` on the stream without blocking as the stream might hold an incomplete term. The predicate `set_stream/2` using the option `timeout(Seconds)` can be used to make the stream generate an exception if no new data arrives for within the timeout. Suppose two processes communicate by exchanging Prolog terms. The following code makes the server immune for clients that write an incomplete term:

¹⁸For compatibility reasons, a *Timeout* value of 0 (integer) also waits indefinitely. To call `select()` without giving up the CPU pass the float 0.0. If compatibility with versions older than 5.1.3 is desired pass the float value 1e-7.

```

    . . . ,
    tcp_accept(Server, Socket, _Peer),
    tcp_open(Socket, In, Out),
    set_stream(In, timeout(10)),
    catch(read(In, Term), _, (close(Out), close(In), fail)),
    . . . ,

```

character_count(+Stream, -Count)

Unify *Count* with the current character index. For input streams this is the number of characters read since the open, for output streams this is the number of characters written. Counting starts at 0.

line_count(+Stream, -Count)

Unify *Count* with the number of lines read or written. Counting starts at 1.

line_position(+Stream, -Count)

Unify *Count* with the position on the current line. Note that this assumes the position is 0 after the open. Tabs are assumed to be defined on each 8-th character and backspaces are assumed to reduce the count by one, provided it is positive.

fileerrors(-Old, +New)

Define error behaviour on errors when opening a file for reading or writing. Valid values are the atoms `on` (default) and `off`. First *Old* is unified with the current value. Then the new value is set to *New*.¹⁹

With the introduction of exception-handling, it is advised to use `catch/3` to catch possibly file-errors and act accordingly. Note that if *fileerrors* is `off`, no exception is generated.

4.18 Primitive character I/O

See section 4.2 for an overview of supported character representations.

nl

Write a newline character to the current output stream. On Unix systems `nl/0` is equivalent to `put(10)`.

nl(+Stream)

Write a newline to *Stream*.

put(+Char)

Write *Char* to the current output stream, *Char* is either an integer-expression evaluating to a character code ($0 \leq Char \leq 255$) or an atom of one character.

put(+Stream, +Char)

Write *Char* to *Stream*.

¹⁹Note that Edinburgh Prolog defines `fileerrors/0` and `noerrors/0`. As this does not allow you to switch back to the old mode I think this definition is better.

put_byte(+Byte)

Alias for `put / 1`.

put_byte(+Stream, +Byte)

Alias for `put / 2`

put_char(+Char)

Alias for `put / 1`.

put_char(+Stream, +Char)

Alias for `put / 2`

put_code(+Code)

Alias for `put / 1`.

put_code(+Stream, +Code)

Alias for `put / 2`

tab(+Amount)

Writes *Amount* spaces on the current output stream. *Amount* should be an expression that evaluates to a positive integer (see section 4.26).

tab(+Stream, +Amount)

Writes *Amount* spaces to *Stream*.

flush_output

Flush pending output on current output stream. `flush_output / 0` is automatically generated by `read / 1` and derivatives if the current input stream is `user` and the cursor is not at the left margin.

flush_output(+Stream)

Flush output on the specified stream. The stream must be open for writing.

ttyflush

Flush pending output on stream *user*. See also `flush_output / [0, 1]`.

get_byte(-Byte)

Read the current input stream and unify the next byte with *Byte* (an integer between 0 and 255). *Byte* is unified with -1 on end of file.

get_byte(+Stream, -Byte)

Read the next byte from *Stream*.

get_code(-Code)

Read the current input stream and unify *Code* with the character code of the next character. *Code* is unified with -1 on end of file. See also `get_char / 1`.

get_code(+Stream, -Code)

Read the next character-code from *Stream*.

get_char(-Char)

Read the current input stream and unify *Char* with the next character as a one-character-atom. See also `atom_chars / 2`. On end-of-file, *Char* is unified to the atom `end_of_file`.

get_char(+Stream, -Char)

Unify *Char* with the next character from *Stream* as a one-character-atom. See also `get_char/2`, `get_byte/2` and `get_code/2`.

get0(-Char)

Edinburgh version of the ISO `get_byte/1` predicate.

get0(+Stream, -Char)

Edinburgh version of the ISO `get_byte/2` predicate.

get(-Char)

Read the current input stream and unify the next non-blank character with *Char*. *Char* is unified with -1 on end of file.

get(+Stream, -Char)

Read the next non-blank character from *Stream*.

peek_byte(-Byte)

Reads the next input byte like `get_byte/1`, but does not remove it from the input stream.

peek_byte(+Stream, -Byte)

Reads the next input byte like `get_byte/2`, but does not remove it from the stream.

peek_code(-Code)

Reads the next input code like `get_code/1`, but does not remove it from the input stream.

peek_code(+Stream, -Code)

Reads the next input code like `get_code/2`, but does not remove it from the stream.

peek_char(-Char)

Reads the next input character like `get_char/1`, but does not remove it from the input stream.

peek_char(+Stream, -Char)

Reads the next input character like `get_char/2`, but does not remove it from the stream.

skip(+Char)

Read the input until *Char* or the end of the file is encountered. A subsequent call to `get0/1` will read the first character after *Char*.

skip(+Stream, +Char)

Skip input (as `skip/1`) on *Stream*.

get_single_char(-Char)

Get a single character from input stream 'user' (regardless of the current input stream). Unlike `get0/1` this predicate does not wait for a return. The character is not echoed to the user's terminal. This predicate is meant for keyboard menu selection etc. If SWI-Prolog was started with the `-tty` option this predicate reads an entire line of input and returns the first non-blank character on this line, or the character code of the newline (10) if the entire line consisted of blank characters.

at_end_of_stream

Succeeds after the last character of the current input stream has been read. Also succeeds if there is no valid current input stream.

at_end_of_stream(+Stream)

Succeeds after the last character of the named stream is read, or *Stream* is not a valid input stream. The end-of-stream test is only available on buffered input stream (unbuffered input streams are rarely used, see `open/4`).

copy_stream_data(+StreamIn, +StreamOut, +Len)

Copy *Len* bytes from stream *StreamIn* to *StreamOut*.

copy_stream_data(+StreamIn, +StreamOut)

Copy data all (remaining) data from stream *StreamIn* to *StreamOut*.

4.19 Term reading and writing

This section describes the basic term reading and writing predicates. The predicates `term_to_atom/2`, `atom_to_term/3` and `sformat/3` provide means for translating atoms and strings to terms. The predicates `format/[1,2]` and `writeln/2` provide formatted output.

There are two ways to manipulate the output format. The predicate `print/[1,2]` may be programmed using `portray/1`. The format of floating point numbers may be manipulated using the `prolog_flag` (see `current_prolog_flag/2`) `float_format`.

Reading is sensitive to the `prolog_flag` `character_escapes`, which controls the interpretation of the `\` character in quoted atoms and strings.

write_term(+Term, +Options)

The predicate `write_term/2` is the generic form of all Prolog term-write predicates. Valid options are:

quoted(Bool)

If `true`, atoms and functors that needs quotes will be quoted. The default is `false`.

backquoted_string(Bool)

If `true`, write a string object (see section 4.23) as `'...'`. The default depends on the `prolog_flag` with the same name.

character_escapes(Bool)

If `true`, and `quoted(true)` is active, special characters in quoted atoms and strings are emitted as ISO escape-sequences. Default is taken from the reference module (see below).

ignore_ops(Bool)

If `true`, the generic term-representation (`<functor>(<args>...)`) will be used for all terms, Otherwise (default), operators, list-notation and `{ }/1` will be written using their special syntax.

module(Module)

Define the reference module (default `user`). This defines the default value for the `character_escapes` option as well as the operator definitions to use. See also `op/3`.

numbervars(Bool)

If `true`, terms of the format `$VAR(N)`, where `<N>` is a positive integer, will be written as a variable name. The default is `false`.

portray(*Bool*)

If true, the hook `portray/1` is called before printing a term that is not a variable. If `portray/1` succeeds, the term is considered printed. See also `print/1`. The default is false. This option is an extension to the ISO `write_term` options.

max_depth(*Integer*)

If the term is nested deeper than *Integer*, print the remainder as eclipse (...). A 0 (zero) value (default) imposes no depth limit. This option also delimits the number of printed for a list. Example:

```
?- write_term(a(s(s(s(s(0))))), [a,b,c,d,e,f]), [max_depth(3)].
a(s(s(...)), [a, b|...])
```

Yes

Used by the toplevel and debugger to limit screen output. See also the `prolog-flags` `toplevel_print_options` and `debugger_print_options`.

write_term(+*Stream*, +*Term*, +*Options*)

As `write_term/2`, but output is sent to *Stream* rather than the current output.

write_canonical(+*Term*)

Write *Term* on the current output stream using standard parenthesised prefix notation (i.e., ignoring operator declarations). Atoms that need quotes are quoted. Terms written with this predicate can always be read back, regardless of current operator declarations. Equivalent to `write_term/2` using the options `ignore_ops` and `quoted`.

write_canonical(+*Stream*, +*Term*)

Write *Term* in canonical form on *Stream*.

write(+*Term*)

Write *Term* to the current output, using brackets and operators where appropriate. See `current_prolog_flag/2` for controlling floating point output format.

write(+*Stream*, +*Term*)

Write *Term* to *Stream*.

writeq(+*Term*)

Write *Term* to the current output, using brackets and operators where appropriate. Atoms that need quotes are quoted. Terms written with this predicate can be read back with `read/1` provided the currently active operator declarations are identical.

writeq(+*Stream*, +*Term*)

Write *Term* to *Stream*, inserting quotes.

print(+*Term*)

Prints *Term* on the current output stream similar to `write/1`, but for each (sub)term of *Term* first the dynamic predicate `portray/1` is called. If this predicate succeeds *print* assumes the (sub)term has been written. This allows for user defined term writing.

print(+*Stream*, +*Term*)

Print *Term* to *Stream*.

portray(+Term)

A dynamic predicate, which can be defined by the user to change the behaviour of `print/1` on (sub)terms. For each subterm encountered that is not a variable `print/1` first calls `portray/1` using the term as argument. For lists only the list as a whole is given to `portray/1`. If `portray` succeeds `print/1` assumes the term has been written.

read(-Term)

Read the next Prolog term from the current input stream and unify it with *Term*. On a syntax error `read/1` displays an error message, attempts to skip the erroneous term and fails. On reaching end-of-file *Term* is unified with the atom `end_of_file`.

read(+Stream, -Term)

Read *Term* from *Stream*.

read_clause(-Term)

Equivalent to `read/1`, but warns the user for variables only occurring once in a term (singleton variables) which do not start with an underscore if `style_check(singleton)` is active (default). Used to read Prolog source files (see `consult/1`). New code should use `read_term/2` with the option `singletons(warning)`.

read_clause(+Stream, -Term)

Read a clause from *Stream*. See `read_clause/1`.

read_term(-Term, +Options)

Read a term from the current input stream and unify the term with *Term*. The reading is controlled by options from the list of *Options*. If this list is empty, the behaviour is the same as for `read/1`. The options are upward compatible to Quintus Prolog. The argument order is according to the ISO standard. Syntax-errors are always reported using exception-handling (see `catch/3`). Options:

variables(Vars)

Unify *Vars* with a list of variables in the term. The variables appear in the order they have been read. See also `free_variables/2`. (ISO).

variable_names(Vars)

Unify *Vars* with a list of '*Name = Var*', where *Name* is an atom describing the variable name and *Var* is a variable that shares with the corresponding variable in *Term*. (ISO).

singletons(Vars)

As `variable_names`, but only reports the variables occurring only once in the *Term* read. Variables starting with an underscore ('_') are not included in this list. (ISO).

syntax_errors(Atom)

If `error` (default), throw an exception on a syntax error. Other values are `fail`, which causes a message to be printed using `print_message/2`, after which the predicate fails, `quiet` which causes the predicate to fail silently and `dec10` which causes syntax errors to be printed, after which `read_term/[2,3]` continues reading the next term. Using `dec10`, `read_term/[2,3]` never fails. (Quintus, SICStus).

module(Module)

Specify *Module* for operators, `character_escapes` flag and `double_quotes` flag.

The value of the latter two is overruled if the corresponding `read_term/3` option is provided. If no module is specified, the current ‘source-module’ is used. (SWI-Prolog).

character_escapes(*Bool*)

Defines how to read \ escape-sequences in quoted atoms. See the prolog-flags `character_escapes`, `current_prolog_flag/2`. (SWI-Prolog).

double_quotes(*Bool*)

Defines how to read ”...” strings. See the prolog-flags `double_quotes`, `current_prolog_flag/2`. (SWI-Prolog).

backquoted_string(*Bool*)

If `true`, read ‘...’ to a string object (see section 4.23). The default depends on the prolog flag with the same name.

term_position(*Pos*)

Unifies *Pos* with the starting position of the term read. *Pos* if of the same format as use by `stream_property/2`.

subterm_positions(*TermPos*)

Describes the detailed layout of the term. The formats for the various types of terms if given below. All positions are character positions. If the input is related to a normal stream, these positions are relative to the start of the input, when reading from the terminal, they are relative to the start of the term.

From-To

Used for primitive types (atoms, numbers, variables).

string_position(From, To)

Used to indicate the position of a string enclosed in double quotes (”).

brace_term_position(From, To, Arg)

Term of the form { . . }, as used in DCG rules. *Arg* describes the argument.

list_position(From, To, Elms, Tail)

A list. *Elms* describes the positions of the elements. If the list specifies the tail as | <TailTerm>, *Tail* is unified with the term-position of the tail, otherwise with the atom `none`.

term_position(From, To, FFrom, FTo, SubPos)

Used for a compound term not matching one of the above. *FFrom* and *FTo* describe the position of the functor. *SubPos* is a list, each element of which describes the term-position of the corresponding subterm.

read_term(+Stream, -Term, +Options)

Read term with options from *Stream*. See `read_term/2`.

read_history(+Show, +Help, +Special, +Prompt, -Term, -Bindings)

Similar to `read_term/2` using the option `variable_names`, but allows for history substitutions. `read_history/6` is used by the top level to read the user’s actions. *Show* is the command the user should type to show the saved events. *Help* is the command to get an overview of the capabilities. *Special* is a list of commands that are not saved in the history. *Prompt* is the first prompt given. Continuation prompts for more lines are determined by `prompt/2`. A %w in the prompt is substituted by the event number. See section 2.7 for available substitutions.

SWI-Prolog calls `read_history/6` as follows:

```
read_history(h, '!h', [trace], '%w ?- ', Goal, Bindings)
```

prompt(-Old, +New)

Set prompt associated with `read/1` and its derivatives. *Old* is first unified with the current prompt. On success the prompt will be set to *New* if this is an atom. Otherwise an error message is displayed. A prompt is printed if one of the read predicates is called and the cursor is at the left margin. It is also printed whenever a newline is given and the term has not been terminated. Prompts are only printed when the current input stream is *user*.

prompt1(+Prompt)

Sets the prompt for the next line to be read. Continuation lines will be read using the prompt defined by `prompt/2`.

4.20 Analysing and Constructing Terms

functor(?Term, ?Functor, ?Arity)

Succeeds if *Term* is a term with functor *Functor* and arity *Arity*. If *Term* is a variable it is unified with a new term holding only variables. `functor/3` silently fails on instantiation faults²⁰ If *Term* is an atom or number, *Functor* will be unified with *Term* and arity will be unified with the integer 0 (zero).

arg(?Arg, ?Term, ?Value)

Term should be instantiated to a term, *Arg* to an integer between 1 and the arity of *Term*. *Value* is unified with the *Arg*-th argument of *Term*. *Arg* may also be unbound. In this case *Value* will be unified with the successive arguments of the term. On successful unification, *Arg* is unified with the argument number. Backtracking yields alternative solutions.²¹ The predicate `arg/3` fails silently if $Arg = 0$ or $Arg > arity$ and raises the exception `domain_error(not_less_than_zero, Arg)` if $Arg < 0$.

setarg(+Arg, +Term, +Value)

Extra-logical predicate. Assigns the *Arg*-th argument of the compound term *Term* with the given *Value*. The assignment is undone if backtracking brings the state back into a position before the `setarg/3` call.

This predicate may be used for destructive assignment to terms, using them as an extra-logical storage bin. Always try hard to avoid the use of `setarg/3` as it is not supported by many Prolog systems and one has to be very careful about unexpected copying as well as unexpected not copying of terms.

?Term = .. ?List

List is a list which head is the functor of *Term* and the remaining arguments are the arguments of the term. Each of the arguments may be a variable, but not both. This predicate is called 'Univ'. Examples:

²⁰In version 1.2 instantiation faults led to error messages. The new version can be used to do type testing without the need to catch illegal instantiations first.

²¹The instantiation pattern `(-, +, ?)` is an extension to 'standard' Prolog.

```
?- foo(hello, X) =.. List.

List = [foo, hello, X]

?- Term =.. [baz, foo(1)]

Term = baz(foo(1))
```

numbervars(+Term, +Start, -End)

Unify the free variables of *Term* with a term $\$VAR(N)$, where N is the number of the variable. Counting starts at *Start*. *End* is unified with the number that should be given to the next variable. Example:

```
?- numbervars(foo(A, B, A), 0, End).

A = '$VAR'(0)
B = '$VAR'(1)
End = 2
```

See also the `numbervars` option to `write_term/3` and `numbervars/4`.

numbervars(+Term, +FunctionSymbol, +Start, -End)

As `numbervars/3`, but instead of using the fixed $\$VAR$ the function symbol is determined by extra *FunctionSymbol* argument.

free_variables(+Term, -List)

Unify *List* with a list of variables, each sharing with a unique variable of *Term*. For example:

```
?- free_variables(a(X, b(Y, X), Z), L).

L = [G367, G366, G371]
X = G367
Y = G366
Z = G371
```

copy_term(+In, -Out)

Make a copy of term *In* and unify the result with *Out*. Ground parts of *In* are shared by *Out*. Provided *In* and *Out* have no sharing variables before this call they will have no sharing variables afterwards. `copy_term/2` is semantically equivalent to:

```
copy_term(In, Out) :-
    recorda(copy_key, In, Ref),
    recorded(copy_key, Out, Ref),
    erase(Ref).
```

Note that as `copy_term/2` does not create copies for ground terms `setarg/3` affects both the copy and original.

4.21 Analysing and Constructing Atoms

These predicates convert between Prolog constants and lists of character codes. The predicates `atom_codes/2`, `number_codes/2` and `name/2` behave the same when converting from a constant to a list of character codes. When converting the other way around, `atom_codes/2` will generate an atom, `number_codes/2` will generate a number or exception and `name/2` will return a number if possible and an atom otherwise.

The ISO standard defines `atom_chars/2` to describe the ‘broken-up’ atom as a list of one-character atoms instead of a list of codes. Upto version 3.2.x, SWI-Prolog’s `atom_chars/2` behaved, compatible to Quintus and SICStus Prolog, like `atom_codes`. As of 3.3.x SWI-Prolog `atom_codes/2` and `atom_chars/2` are compliant to the ISO standard.

To ease the pain of all variations in the Prolog community, all SWI-Prolog predicates behave as flexible as possible. This implies the ‘list-side’ accepts either a code-list or a char-list and the ‘atom-side’ accept all atomic types (atom, number and string).

atom_codes(?Atom, ?String)

Convert between an atom and a list of character codes. If *Atom* is instantiated, it will be translated into a list of character codes and the result is unified with *String*. If *Atom* is unbound and *String* is a list of character codes, it will *Atom* will be unified with an atom constructed from this list.

atom_chars(?Atom, ?CharList)

As `atom_codes/2`, but *CharList* is a list of one-character atoms rather than a list of character codes²².

```
?- atom_chars(hello, X).
```

```
X = [h, e, l, l, o]
```

char_code(?Atom, ?Code)

Convert between character and character code for a single character.²³

number_chars(?Number, ?CharList)

Similar to `atom_chars/2`, but converts between a number and its representation as a list of one-character atoms. Fails with a `representation_error` if *Number* is unbound and *CharList* does not describe a number.

number_codes(?Number, ?CodeList)

As `number_chars/2`, but converts to a list of character codes rather than one-character atoms. In the mode -, +, both predicates behave identically to improve handling of non-ISO source.

atom_number(?Atom, ?Number)

Realises the popular combination of `atom_codes/2` and `number_codes/2` to convert between atom and number (integer or float) in one predicate, avoiding the intermediate list.

²²Upto version 3.2.x, `atom_chars/2` behaved as the current `atom_codes/2`. The current definition is compliant with the ISO standard

²³This is also called `atom_char/2` in older versions of SWI-Prolog as well as some other Prolog implementations. `atom_char/2` is available from the library `backcomp.pl`

name(?AtomOrInt, ?String)

String is a list of character codes representing the same text as *Atom*. Each of the arguments may be a variable, but not both. When *String* is bound to a character code list describing an integer and *Atom* is a variable *Atom* will be unified with the integer value described by *String* (e.g. `'name(N, "300"), 400 is N + 100'` succeeds).

int_to_atom(+Int, +Base, -Atom)

Convert *Int* to a textual representation using base *Base* and unify the result with *Atom*. If $Base \neq 10$ the base will be prepended to *Atom*. $Base = 0$ will try to interpret *Int* as a character code and return `0'<c>`. Otherwise $2 \leq Base \leq 36$. Some examples are given below.

```
int_to_atom(45, 2, A)   → A = 2'101101
int_to_atom(97, 0, A)  → A = 0'a
int_to_atom(56, 10, A) → A = 56
```

int_to_atom(+Int, -Atom)

Equivalent to `int_to_atom(Int, 10, Atom)`.

term_to_atom(?Term, ?Atom)

Succeeds if *Atom* describes a term that unifies with *Term*. When *Atom* is instantiated *Atom* is converted and then unified with *Term*. If *Atom* has no valid syntax, a `syntax_error` exception is raised. Otherwise *Term* is “written” on *Atom* using `write/1`.

atom_to_term(+Atom, -Term, -Bindings)

Use *Atom* as input to `read_term/2` using the option `variable_names` and return the read term in *Term* and the variable bindings in *Bindings*. *Bindings* is a list of *Name = Var* couples, thus providing access to the actual variable names. See also `read_term/2`. If *Atom* has no valid syntax, a `syntax_error` exception is raised.

atom_concat(?Atom1, ?Atom2, ?Atom3)

Atom3 forms the concatenation of *Atom1* and *Atom2*. At least two of the arguments must be instantiated to atoms, integers or floating point numbers. For ISO compliance, the instantiation-pattern `-, -, +` is allowed too, non-deterministically splitting the 3-th argument into two parts (as `append/3` does for lists). See also `string_concat/3`.

concat_atom(+List, -Atom)

List is a list of atoms, integers or floating point numbers. Succeeds if *Atom* can be unified with the concatenated elements of *List*. If *List* has exactly 2 elements it is equivalent to `atom_concat/3`, allowing for variables in the list.

concat_atom(?List, +Separator, ?Atom)

Creates an atom just like `concat_atom/2`, but inserts *Separator* between each pair of atoms. For example:

```
?- concat_atom([gnu, gnat], ', ', A).
```

```
A = 'gnu, gnat'
```

This predicate can also be used to split atoms by instantiating *Separator* and *Atom*:


```
?- concat_atom(L, -, 'gnu-gnat').
```

```
L = [gnu, gnat]
```

atom_length(+Atom, -Length)

Succeeds if *Atom* is an atom of *Length* characters long. This predicate also works for strings (see section 4.23). If the prolog flag `iso` is *not* set, it also accepts integers and floats, expressing the number of characters output when given to `write/1` as well as code-lists and character-lists, expressing the length of the list.²⁴

atom_prefix(+Atom, +Prefix)

Succeeds if *Atom* starts with the characters from *Prefix*. Its behaviour is equivalent to `?- concat(Prefix, -, Atom)`, but avoids the construction of an atom for the ‘remainder’.

sub_atom(+Atom, ?Before, ?Len, ?After, ?Sub)

ISO predicate for breaking atoms. It maintains the following relation: *Sub* is a sub-atom of *Atom* that starts at *Before*, has *Len* characters and *Atom* contains *After* characters after the match.

```
?- sub_atom(abc, 1, 1, A, S).
```

```
A = 1, S = b
```

The implementation minimalises non-determinism and creation of atoms. This is a very flexible predicate that can do search, prefix- and suffix-matching, etc.

4.22 Classifying characters

SWI-Prolog offers two comprehensive predicates for classifying characters and character-codes. These predicates are defined as built-in predicates to exploit the C-character classification’s handling of *locale* (handling of local character-sets). These predicates are fast, logical and deterministic if applicable.

In addition, there is the library `ctype` providing compatibility to some other Prolog systems. The predicates of this library are defined in terms of `code_type/2`.

char_type(?Char, ?Type)

Tests or generates alternative *Types* or *Chars*. The character-types are inspired by the standard C `<ctype.h>` primitives.

alnum

Char is a letter (upper- or lowercase) or digit.

alpha

Char is a letter (upper- or lowercase).

²⁴BUG: Note that `[]` is both an atom an empty code/character list. The predicate `atom_length/2` returns 2 for this atom.

csym

Char is a letter (upper- or lowercase), digit or the underscore (`_`). These are valid C- and Prolog symbol characters.

csymf

Char is a letter (upper- or lowercase) or the underscore (`_`). These are valid first characters for C- and Prolog symbols

ascii

Char is a 7-bits ASCII character (0..127).

white

Char is a space or tab. E.i. white space inside a line.

cntrl

Char is an ASCII control-character (0..31).

digit

Char is a digit.

digit(*Weight*)

Char is a digit with value *Weight*. I.e. `char_type(X, digit(6))` yields `X = '6'`. Useful for parsing numbers.

xdigit(*Weight*)

Char is a hexe-decimal digit with value *Weight*. I.e. `char_type(a, xdigit(X))` yields `X = '10'`. Useful for parsing numbers.

graph

Char produces a visible mark on a page when printed. Note that the space is not included!

lower

Char is a lower-case letter.

lower(*Upper*)

Char is a lower-case version of *Upper*. Only true if *Char* is lowercase and *Upper* uppercase.

to_lower(*Upper*)

Char is a lower-case version of *Upper*. For non-letters, or letter without case, *Char* and *Lower* are the same. See also `upcase_atom/2` and `downcase_atom/2`.

upper

Char is an upper-case letter.

upper(*Lower*)

Char is an upper-case version of *Lower*. Only true if *Char* is uppercase and *Lower* lowercase.

to_upper(*Lower*)

Char is an upper-case version of *Lower*. For non-letters, or letter without case, *Char* and *Lower* are the same. See also `upcase_atom/2` and `downcase_atom/2`.

punct

Char is a punctuation character. This is a `graph` character that is not a letter or digit.

space

Char is some form of layout character (tab, vertical-tab, newline, etc.).

end_of_file

Char is -1.

end_of_line

Char ends a line (ASCII: 10..13).

newline

Char is a the newline character (10).

period

Char counts as the end of a sentence (.,!,?).

quote

Char is a quote-character (" , ' , `).

paren(*Close*)

Char is an open-parenthesis and *Close* is the corresponding close-parenthesis.

code_type(?*Code*, ?*Type*)

As `char_type/2`, but uses character-codes rather than one-character atoms. Please note that both predicates are as flexible as possible. They handle either representation if the argument is instantiated and only will instantiate with an integer code or one-character atom depending of the version used. See also the prolog-flag `double_quotes`, `atom_chars/2` and `atom_codes/2`.

4.22.1 Case conversion

There is nothing in the Prolog standard for converting case in textual data. The SWI-Prolog predicates `code_type/2` and `char_type/2` can be used to test and convert individual characters. We have started some additional support:

downcase_atom(+*AnyCase*, -*LowerCase*)

Converts the characters of *AnyCase* into lowercase as `char_type/2` does (i.e. based on the defined *locale* if Prolog provides locale support on the hosting platform) and unifies the lowercase atom with *LowerCase*.

upcase_atom(+*AnyCase*, -*UpperCase*)

Converts, similar to `downcase_atom/2`, an atom to upper-case.

4.23 Representing text in strings

SWI-Prolog supports the data type *string*. Strings are a time and space efficient mechanism to handle text in Prolog. Strings are stored as a byte array on the global (term) stack and thus destroyed on backtracking and reclaimed by the garbage collector.

Strings were added to SWI-Prolog based on an early draft of the ISO standard, offering a mechanism to represent temporary character data efficiently. As SWI-Prolog strings can handle 0-bytes, they are frequently used through the foreign language interface (section 7) for storing arbitrary byte-sequences.

Starting with version 3.3, SWI-Prolog offers garbage collection on the atom-space as well as representing 0-bytes in atoms. Although strings and atoms still have different features, new code should consider using atoms to avoid too many representations for text as well as for compatibility to other Prolog implementations. Below are some of the differences:

- *creation*
Creating strings is fast, as the data is simply copied to the global stack. Atoms are unique and therefore more expensive in terms of memory and time to create. On the other hand, if the same text has to be represented multiple times, atoms are more efficient.
- *destruction*
Backtracking destroys strings at no cost. They are cheap to handle by the garbage collector, but it should be noted that extensive use of strings will cause many garbage collections. Atom garbage collection is generally faster.

String objects by default have no lexical representation and thus can only be created using the predicates below or through the foreign language interface (See chapter 7). There are two ways to make `read/1` read text into strings, both controlled through Prolog flags. One is by setting the `double_quotes` flag to `string` and the other is by setting the `backquoted_string` flag to `true`. In latter case, `'Hello world'` is read into a string and `write_term/2` prints strings between backquotes if `quoted` is `true`. This flag provides compatibility to LPA Prolog string handling.

string_to_atom(?String, ?Atom)

Logical conversion between a string and an atom. At least one of the two arguments must be instantiated. *Atom* can also be an integer or floating point number.

string_to_list(?String, ?List)

Logical conversion between a string and a list of character codes characters. At least one of the two arguments must be instantiated.

string_length(+String, -Length)

Unify *Length* with the number of characters in *String*. This predicate is functionally equivalent to `atom_length/2` and also accepts atoms, integers and floats as its first argument.

string_concat(?String1, ?String2, ?String3)

Similar to `atom_concat/3`, but the unbound argument will be unified with a string object rather than an atom. Also, if both *String1* and *String2* are unbound and *String3* is bound to text, it breaks *String3*, unifying the start with *String1* and the end with *String2* as `append` does with lists. Note that this is not particularly fast on long strings as for each redo the system has to create two entirely new strings, while the list equivalent only creates a single new list-cell and moves some pointers around.

sub_string(+String, ?Start, ?Length, ?After, ?Sub)

Sub is a substring of *String* starting at *Start*, with length *Length* and *String* has *After* characters left after the match. See also `sub_atom/5`.

4.24 Operators

Operators are defined to improve the readability of source-code. For example, without operators, to write $2*3+4*5$ one would have to write `+(*(2 , 3) , *(4 , 5))`. In Prolog, a number of operators have been predefined. All operators, except for the comma (`,`) can be redefined by the user.

Some care has to be taken before defining new operators. Defining too many operators might make your source 'natural' looking, but at the same time lead to hard to understand the limits of your

syntax. To ease the pain, as of SWI-Prolog 3.3.0, operators are local to the module in which they are defined. The module-table of the module `user` acts as default table for all modules. This global table can be modified explicitly from inside a module:

```
:- module(prove,
        [ prove/1
        ]).

:- op(900, xfx, user:(=>)).
```

Unlike what many users think, operators and quoted atoms have no relation: defining a atom as an operator does **not** influence parsing characters into atoms and quoting an atom does **not** stop it from acting as an operator. To stop an atom acting as an operator, enclose it in braces like this: `(myop)`.

op(+Precedence, +Type, :Name)

Declare *Name* to be an operator of type *Type* with precedence *Precedence*. *Name* can also be a list of names, in which case all elements of the list are declared to be identical operators. *Precedence* is an integer between 0 and 1200. Precedence 0 removes the declaration. *Type* is one of: `xf`, `yf`, `xfx`, `xfy`, `yfx`, `yfy`, `fy` or `fx`. The ‘f’ indicates the position of the functor, while `x` and `y` indicate the position of the arguments. ‘y’ should be interpreted as “on this position a term with precedence lower or equal to the precedence of the functor should occur”. For ‘x’ the precedence of the argument must be strictly lower. The precedence of a term is 0, unless its principal functor is an operator, in which case the precedence is the precedence of this operator. A term enclosed in brackets `(. . .)` has precedence 0.

The predefined operators are shown in table 4.1. Note that all operators can be redefined by the user.

current_op(?Precedence, ?Type, ?Name)

Succeeds when *Name* is currently defined as an operator of type *Type* with precedence *Precedence*. See also `op/3`.

4.25 Character Conversion

Although I wouldn’t really know for what you would like to use these features, they are provided for ISO complacency.

char_conversion(+CharIn, +CharOut)

Define that term-input (see `read_term/3`) maps each character read as *CharIn* to the character *CharOut*. Character conversion is only executed if the prolog-flag `char_conversion` is set to `true` and not inside quoted atoms or strings. The initial table maps each character onto itself. See also `current_char_conversion/2`.

current_char_conversion(?CharIn, ?CharOut)

Queries the current character conversion-table. See `char_conversion/2` for details.

1200	<i>xfx</i>	-->, :-
1200	<i>fx</i>	:-, ?-
1150	<i>fx</i>	dynamic, multifile, module.transparent, discontinuous, volatile, initialization
1100	<i>xfy</i>	i,
1050	<i>xfy</i>	->
1000	<i>xfy</i>	,
954	<i>xfy</i>	\
900	<i>fy</i>	\+
900	<i>fx</i>	~
700	<i>xfx</i>	<, =, =.., =@=, =:=, =<, ==, =\=, >, >=, @<, @=<, @>, @>=, \=, \==, is
600	<i>xfy</i>	:
500	<i>yfx</i>	+, -, /\, \/, xor
500	<i>fx</i>	+, -, ?, \
400	<i>yfx</i>	*, /, //, <<, >>, mod, rem
200	<i>xfx</i>	**
200	<i>xfy</i>	^

Table 4.1: System operators

4.26 Arithmetic

Arithmetic can be divided into some special purpose integer predicates and a series of general predicates for floating point and integer arithmetic as appropriate. The integer predicates are as “logical” as possible. Their usage is recommended whenever applicable, resulting in faster and more “logical” programs.

The general arithmetic predicates are optionally compiled now (see `set_prolog_flag/2` and the `-O` command line option). Compiled arithmetic reduces global stack requirements and improves performance. Unfortunately compiled arithmetic cannot be traced, which is why it is optional.

The general arithmetic predicates all handle *expressions*. An expression is either a simple number or a *function*. The arguments of a function are expressions. The functions are described in section 4.27.

between(+Low, +High, ?Value)

Low and *High* are integers, $High \geq Low$. If *Value* is an integer, $Low \leq Value \leq High$. When *Value* is a variable it is successively bound to all integers between *Low* and *High*.

succ(?Int1, ?Int2)

Succeeds if $Int2 = Int1 + 1$ and $Int1 \geq 0$. At least one of the arguments must be instantiated to a natural number. This predicate raises the domain-error `not_less_than_zero` if called with a negative integer. E.g. `succ(X, 0)` fails silently and `succ(X, -1)` raises a domain-error.²⁵

plus(?Int1, ?Int2, ?Int3)

Succeeds if $Int3 = Int1 + Int2$. At least two of the three arguments must be instantiated to integers.

²⁵The behaviour to deal with natural numbers only was defined by Richard O’Keefe to support the common count-down-to-zero in a natural way. Upto 5.1.8 `succ/2` also accepted negative integers.

$+Expr1 > +Expr2$

Succeeds when expression $Expr1$ evaluates to a larger number than $Expr2$.

$+Expr1 < +Expr2$

Succeeds when expression $Expr1$ evaluates to a smaller number than $Expr2$.

$+Expr1 = < +Expr2$

Succeeds when expression $Expr1$ evaluates to a smaller or equal number to $Expr2$.

$+Expr1 = > +Expr2$

Succeeds when expression $Expr1$ evaluates to a larger or equal number to $Expr2$.

$+Expr1 = \neq +Expr2$

Succeeds when expression $Expr1$ evaluates to a number non-equal to $Expr2$.

$+Expr1 = := +Expr2$

Succeeds when expression $Expr1$ evaluates to a number equal to $Expr2$.

$-Number \text{ is } +Expr$

Succeeds when $Number$ has successfully been unified with the number $Expr$ evaluates to. If $Expr$ evaluates to a float that can be represented using an integer (i.e, the value is integer and within the range that can be described by Prolog's integer representation), $Expr$ is unified with the integer value.

Note that normally, `is/2` will be used with unbound left operand. If equality is to be tested, `:=/2` should be used. For example:

<code>?- 1.0 is sin(pi/2).</code>	Fails!. <code>sin(pi/2)</code> evaluates to 1.0, but <code>is/2</code> will represent this as the integer 1, after which unify will fail.
<code>?- 1.0 is float(sin(pi/2)).</code>	Succeeds, as the <code>float/1</code> function forces the result to be float.
<code>?- 1.0 := sin(pi/2).</code>	Succeeds as expected.

4.27 Arithmetic Functions

Arithmetic functions are terms which are evaluated by the arithmetic predicates described above. SWI-Prolog tries to hide the difference between integer arithmetic and floating point arithmetic from the Prolog user. Arithmetic is done as integer arithmetic as long as possible and converted to floating point arithmetic whenever one of the arguments or the combination of them requires it. If a function returns a floating point value which is whole it is automatically transformed into an integer. There are three types of arguments to functions:

<i>Expr</i>	Arbitrary expression, returning either a floating point value or an integer.
<i>IntExpr</i>	Arbitrary expression that should evaluate into an integer.
<i>Int</i>	An integer.

In case integer addition, subtraction and multiplication would lead to an integer overflow the operands are automatically converted to floating point numbers. The floating point functions (`sin/1`, `exp/1`, etc.) form a direct interface to the corresponding C library functions used to compile SWI-Prolog. Please refer to the C library documentation for details on precision, error handling, etc.

- +Expr
Result = $-Expr$

+Expr1 + +Expr2
Result = $Expr1 + Expr2$

+Expr1 - +Expr2
Result = $Expr1 - Expr2$

+Expr1 * +Expr2
Result = $Expr1 \times Expr2$

+Expr1 / +Expr2
Result = $\frac{Expr1}{Expr2}$

+IntExpr1 mod +IntExpr2
 Modulo: *Result* = $IntExpr1 - (IntExpr1 // IntExpr2) \times IntExpr2$ The function `mod/2` is implemented using the C `%` operator. It's behaviour with negative values is illustrated in the table below.

2	=	17	mod	5
2	=	17	mod	-5
-2	=	-17	mod	5
-2	=	-17	mod	5

+IntExpr1 rem +IntExpr2
 Remainder of division: *Result* = `float_fractional_part(IntExpr1/IntExpr2)`

+IntExpr1 // +IntExpr2
 Integer division: *Result* = `truncate(Expr1/Expr2)`

abs(+Expr)
 Evaluate *Expr* and return the absolute value of it.

sign(+Expr)
 Evaluate to -1 if *Expr* < 0, 1 if *Expr* > 0 and 0 if *Expr* = 0.

max(+Expr1, +Expr2)
 Evaluates to the largest of both *Expr1* and *Expr2*.

min(+Expr1, +Expr2)
 Evaluates to the smallest of both *Expr1* and *Expr2*.

.(+Int, [I])

A list of one element evaluates to the element. This implies "a" evaluates to the character code of the letter 'a' (97). This option is available for compatibility only. It will not work if 'style_check(+string)' is active as "a" will then be transformed into a string object. The recommended way to specify the character code of the letter 'a' is 0'a.

random(+Int)

Evaluates to a random integer i for which $0 \leq i < Int$. The seed of this random generator is determined by the system clock when SWI-Prolog was started.

round(+Expr)

Evaluates *Expr* and rounds the result to the nearest integer.

integer(+Expr)

Same as round/1 (backward compatibility).

float(+Expr)

Translate the result to a floating point number. Normally, Prolog will use integers whenever possible. When used around the 2nd argument of is/2, the result will be returned as a floating point number. In other contexts, the operation has no effect.

float_fractional_part(+Expr)

Fractional part of a floating-point number. Negative if *Expr* is negative, 0 if *Expr* is integer.

float_integer_part(+Expr)

Integer part of floating-point number. Negative if *Expr* is negative, *Expr* if *Expr* is integer.

truncate(+Expr)

Truncate *Expr* to an integer. Same as float_integer_part/1.

floor(+Expr)

Evaluates *Expr* and returns the largest integer smaller or equal to the result of the evaluation.

ceiling(+Expr)

Evaluates *Expr* and returns the smallest integer larger or equal to the result of the evaluation.

ceil(+Expr)

Same as ceiling/1 (backward compatibility).

+IntExpr >> +IntExpr

Bitwise shift *IntExpr1* by *IntExpr2* bits to the right.

+IntExpr << +IntExpr

Bitwise shift *IntExpr1* by *IntExpr2* bits to the left.

+IntExpr \/ +IntExpr

Bitwise 'or' *IntExpr1* and *IntExpr2*.

+IntExpr /\ +IntExpr

Bitwise 'and' *IntExpr1* and *IntExpr2*.

+IntExpr xor +IntExpr

Bitwise 'exclusive or' *IntExpr1* and *IntExpr2*.

**** *+IntExpr*

Bitwise negation.

sqrt(*+Expr*)

Result = \sqrt{Expr}

sin(*+Expr*)

Result = $\sin Expr$. *Expr* is the angle in radians.

cos(*+Expr*)

Result = $\cos Expr$. *Expr* is the angle in radians.

tan(*+Expr*)

Result = $\tan Expr$. *Expr* is the angle in radians.

asin(*+Expr*)

Result = $\arcsin Expr$. *Result* is the angle in radians.

acos(*+Expr*)

Result = $\arccos Expr$. *Result* is the angle in radians.

atan(*+Expr*)

Result = $\arctan Expr$. *Result* is the angle in radians.

atan(*+YExpr*, *+XExpr*)

Result = $\arctan \frac{YExpr}{XExpr}$. *Result* is the angle in radians. The return value is in the range $[-\pi \dots \pi]$. Used to convert between rectangular and polar coordinate system.

log(*+Expr*)

Result = $\ln Expr$

log10(*+Expr*)

Result = $\lg Expr$

exp(*+Expr*)

Result = e^{Expr}

+Expr1 ****** *+Expr2*

Result = $Expr1^{Expr2}$

+Expr1 **^** *+Expr2*

Same as ******/2. (backward compatibility).

pi

Evaluates to the mathematical constant π (3.141593).

e

Evaluates to the mathematical constant e (2.718282).

cputime

Evaluates to a floating point number expressing the CPU time (in seconds) used by Prolog up till now. See also `statistics/2` and `time/1`.

4.28 Adding Arithmetic Functions

Prolog predicates can be given the role of arithmetic function. The last argument is used to return the result, the arguments before the last are the inputs. Arithmetic functions are added using the predicate `arithmetic_function/1`, which takes the head as its argument. Arithmetic functions are module sensitive, that is they are only visible from the module in which the function is defined and declared. Global arithmetic functions should be defined and registered from module `user`. Global definitions can be overruled locally in modules. The builtin functions described above can be redefined as well.

`arithmetic_function(+Head)`

Register a Prolog predicate as an arithmetic function (see `is/2`, `>/2`, etc.). The Prolog predicate should have one more argument than specified by *Head*, which is either a term *Name/Arity*, an atom or a complex term. This last argument is an unbound variable at call time and should be instantiated to an integer or floating point number. The other arguments are the parameters. This predicate is module sensitive and will declare the arithmetic function only for the context module, unless declared from module `user`. Example:

```
1 ?- [user].
:- arithmetic_function(mean/2).

mean(A, B, C) :-
    C is (A+B)/2.
user compiled, 0.07 sec, 440 bytes.

Yes
2 ?- A is mean(4, 5).

A = 4.500000
```

`current_arithmetic_function(?Head)`

Successively unifies all arithmetic functions that are visible from the context module with *Head*.

4.29 Built-in list operations

Most list operations are defined in the library `lists` described in section [A.1](#). Some that are implemented with more low-level primitives are built-in and described here.

`is_list(+Term)`

Succeeds if *Term* is bound to the empty list (`[]`) or a term with functor `'.'` and arity 2 and the second argument is a list.²⁶ This predicate acts as if defined by the following definition:

²⁶In versions before 5.0.1, `is_list/1` just checked for `[]` or `[_|_]` and `proper_list/1` had the role of the current `is_list/1`. The current definition is conform to the de-facto standard. Assuming proper coding standards, there should only be very few cases where a quick-and-dirty `is_list/1` is a good choice. Richard O'Keefe pointed at this issue.

```

is_list(X) :-
    var(X), !,
    fail.
is_list([]).
is_list(_|T) :-
    is_list(T).

```

memberchk(?Elem, +List)

Equivalent to `member/2`, but leaves no choice point.

length(?List, ?Int)

Succeeds if *Int* represents the number of elements of list *List*. Can be used to create a list holding only variables.

sort(+List, -Sorted)

Succeeds if *Sorted* can be unified with a list holding the elements of *List*, sorted to the standard order of terms (see section 4.6). Duplicates are removed. The implementation is in C, using *natural merge sort*²⁷

msort(+List, -Sorted)

Equivalent to `sort/2`, but does not remove duplicates.

keysort(+List, -Sorted)

List is a proper list whose elements are *Key-Value*, that is, terms whose principal functor is `(-)/2`, whose first argument is the sorting key, and whose second argument is the satellite data to be carried along with the key. `keysort/2` sorts *List* like `msort/2`, but only compares the keys. It is used to sort terms not on standard order, but on any criterion that can be expressed on a multi-dimensional scale. Sorting on more than one criterion can be done using terms as keys, putting the first criterion as argument 1, the second as argument 2, etc. The order of multiple elements that have the same *Key* is not changed. The implementation is in C, using *natural merge sort*.

predsort(+Pred, +List, -Sorted)

Sorts similar to `sort/2`, but determines the order of two terms by calling `Pred(-Delta, +E1, +E2)`. This call must unify *Delta* with one of `<`, `const>` or `=`. If built-in predicate `compare/3` is used, the result is the same as `sort/2`. See also `keysort/2`.²⁸

merge(+List1, +List2, -List3)

List1 and *List2* are lists, sorted to the standard order of terms (see section 4.6). *List3* will be unified with an ordered list holding both the elements of *List1* and *List2*. Duplicates are **not** removed.

merge_set(+Set1, +Set2, -Set3)

Set1 and *Set2* are lists without duplicates, sorted to the standard order of terms. *Set3* is unified with an ordered list without duplicates holding the union of the elements of *Set1* and *Set2*.

²⁷Contributed by Richard O'Keefe.

²⁸Please note that the semantics have changed between 3.1.1 and 3.1.2

4.30 Finding all Solutions to a Goal

findall(+Template, +Goal, -Bag)

Creates a list of the instantiations *Template* gets successively on backtracking over *Goal* and unifies the result with *Bag*. Succeeds with an empty list if *Goal* has no solutions. `findall/3` is equivalent to `bagof/3` with all free variables bound with the existence operator (\wedge), except that `bagof/3` fails when goal has no solutions.

bagof(+Template, +Goal, -Bag)

Unify *Bag* with the alternatives of *Template*, if *Goal* has free variables besides the one sharing with *Template* `bagof` will backtrack over the alternatives of these free variables, unifying *Bag* with the corresponding alternatives of *Template*. The construct `+Var \wedge Goal` tells `bagof` not to bind *Var* in *Goal*. `bagof/3` fails if *Goal* has no solutions.

The example below illustrates `bagof/3` and the \wedge operator. The variable bindings are printed together on one line to save paper.

```
2 ?- listing(foo).
```

```
foo(a, b, c).
```

```
foo(a, b, d).
```

```
foo(b, c, e).
```

```
foo(b, c, f).
```

```
foo(c, c, g).
```

Yes

```
3 ?- bagof(C, foo(A, B, C), Cs).
```

```
A = a, B = b, C = G308, Cs = [c, d] ;
```

```
A = b, B = c, C = G308, Cs = [e, f] ;
```

```
A = c, B = c, C = G308, Cs = [g] ;
```

No

```
4 ?- bagof(C, A $\wedge$ foo(A, B, C), Cs).
```

```
A = G324, B = b, C = G326, Cs = [c, d] ;
```

```
A = G324, B = c, C = G326, Cs = [e, f, g] ;
```

No

```
5 ?-
```

setof(+Template, +Goal, -Set)

Equivalent to `bagof/3`, but sorts the result using `sort/2` to get a sorted list of alternatives without duplicates.

4.31 Invoking Predicates on all Members of a List

All the predicates in this section call a predicate on all members of a list or until the predicate called fails. The predicate is called via `call/[2..]`, which implies common arguments can be put in front of the arguments obtained from the list(s). For example:

```
?- maplist(plus(1), [0, 1, 2], X).
```

```
X = [1, 2, 3]
```

we will phrase this as “*Predicate* is applied on ...”

checklist(+Pred, +List)

Pred is applied successively on each element of *List* until the end of the list or *Pred* fails. In the latter case the `checklist/2` fails.

maplist(+Pred, ?List1, ?List2)

Apply *Pred* on all successive pairs of elements from *List1* and *List2*. Fails if *Pred* can not be applied to a pair. See the example above.

sublist(+Pred, +List1, ?List2)

Unify *List2* with a list of all elements of *List1* to which *Pred* applies.

4.32 Forall

forall(+Cond, +Action)

For all alternative bindings of *Cond* *Action* can be proven. The example verifies that all arithmetic statements in the list *L* are correct. It does not say which is wrong if one proves wrong.

```
?- forall(member(Result = Formula, [2 = 1 + 1, 4 = 2 * 2]),
          Result == Formula).
```

4.33 Formatted Write

The current version of SWI-Prolog provides two formatted write predicates. The first is `writeln/[1,2]`, which is compatible with Edinburgh C-Prolog. The second is `writef/[1,2]`, which is compatible with Quintus Prolog. We hope the Prolog community will once define a standard formatted write predicate. If you want performance use `format/[1,2]` as this predicate is defined in C. Otherwise compatibility reasons might tell you which predicate to use.

4.33.1 Writeln

writeln(+Term)

Equivalent to `write(Term), nl`.

writef(+Atom)

Equivalent to `writef(Atom, [])`.

writeln(+Format, +Arguments)

Formatted write. *Format* is an atom whose characters will be printed. *Format* may contain certain special character sequences which specify certain formatting and substitution actions. *Arguments* then provides all the terms required to be output.

Escape sequences to generate a single special character:

<code>\n</code>	Output a newline character (see also <code>nl/[0,1]</code>)
<code>\l</code>	Output a line separator (same as <code>\n</code>)
<code>\r</code>	Output a carriage-return character (ASCII 13)
<code>\t</code>	Output the ASCII character TAB (9)
<code>\\</code>	The character <code>\</code> is output
<code>\%</code>	The character <code>%</code> is output
<code>\nnn</code>	where <i>nnn</i> is an integer (1-3 digits) the character with character code <i>nnn</i> is output (NB : <i>nnn</i> is read as decimal)

Note that `\l`, `\nnn` and `\\` are interpreted differently when character-escapes are in effect. See section 2.15.1.

Escape sequences to include arguments from *Arguments*. Each time a `%` escape sequence is found in *Format* the next argument from *Arguments* is formatted according to the specification.

<code>%t</code>	<code>print/1</code> the next item (mnemonic: term)
<code>%w</code>	<code>write/1</code> the next item
<code>%q</code>	<code>writeln/1</code> the next item
<code>%d</code>	Write the term, ignoring operators. See also <code>write_term/2</code> . Mnemonic: old Edinburgh display/1.
<code>%p</code>	<code>print/1</code> the next item (identical to <code>%t</code>)
<code>%n</code>	Put the next item as a character (i.e., it is a character code)
<code>%r</code>	Write the next item <i>N</i> times where <i>N</i> is the second item (an integer)
<code>%s</code>	Write the next item as a String (so it must be a list of characters)
<code>%f</code>	Perform a <code>ttyflush/0</code> (no items used)
<code>%Nc</code>	Write the next item Centered in <i>N</i> columns.
<code>%Nl</code>	Write the next item Left justified in <i>N</i> columns.
<code>%Nr</code>	Write the next item Right justified in <i>N</i> columns. <i>N</i> is a decimal number with at least one digit. The item must be an atom, integer, float or string.

swriteln(-String, +Format, +Arguments)

Equivalent to `writeln/2`, but “writes” the result on *String* instead of the current output stream.

Example:

```
?- swritef(S, '%15L%w', ['Hello', 'World']).
```

```
S = "Hello           World"
```

swritef(-String, +Format)

Equivalent to `swritef(String, Format, [])`.

4.33.2 Format

format(+Format)

Defined as `'format(Format) :- format(Format, []).'`

format(+Format, +Arguments)

Format is an atom, list of character codes, or a Prolog string. *Arguments* provides the arguments required by the format specification. If only one argument is required and this is not a list of character codes the argument need not be put in a list. Otherwise the arguments are put in a list.

Special sequences start with the tilde (~), followed by an optional numeric argument, followed by a character describing the action to be undertaken. A numeric argument is either a sequence of digits, representing a positive decimal number, a sequence `\langle character \rangle`, representing the character code value of the character (only useful for `~t`) or an asterisk (*), in when the numeric argument is taken from the next argument of the argument list, which should be a positive integer. Actions are:

- ~ Output the tilde itself.
- a Output the next argument, which should be an atom. This option is equivalent to **w**. Compatibility reasons only.
- c Interpret the next argument as an character code and add it to the output. This argument should be an integer in the range [0, ..., 255] (including 0 and 255).
- d Output next argument as a decimal number. It should be an integer. If a numeric argument is specified a dot is inserted *argument* positions from the right (useful for doing fixed point arithmetic with integers, such as handling amounts of money).
- D Same as **d**, but makes large values easier to read by inserting a comma every three digits left to the dot or right.
- e Output next argument as a floating point number in exponential notation. The numeric argument specifies the precision. Default is 6 digits. Exact representation depends on the C library function `printf()`. This function is invoked with the format `% .(precision)e`.
- E Equivalent to **e**, but outputs a capital E to indicate the exponent.
- f Floating point in non-exponential notation. See C library function `printf()`.
- g Floating point in **e** or **f** notation, whichever is shorter.
- G Floating point in **E** or **f** notation, whichever is shorter.
- i Ignore next argument of the argument list. Produces no output.
- k Give the next argument to `displayq/1` (canonical write).
- n Output a newline character.

- N Only output a newline if the last character output on this stream was not a newline. Not properly implemented yet.
- p Give the next argument to `print/1`.
- q Give the next argument to `writeln/1`.
- r Print integer in radix the numeric argument notation. Thus `~16r` prints its argument hexadecimal. The argument should be in the range $[2, \dots, 36]$. Lower case letters are used for digits above 9.
- R Same as `r`, but uses upper case letters for digits above 9.
- s Output text from a list of character codes or a string (see `string/1` and section 4.23) from the next argument.
- t All remaining space between 2 tabs tops is distributed equally over `~t` statements between the tabs tops. This space is padded with spaces by default. If an argument is supplied this is taken to be the character code of the character used for padding. This can be used to do left or right alignment, centering, distributing, etc. See also `~|` and `~+` to set tab stops. A tabs top is assumed at the start of each line.
 - | Set a tabs top on the current position. If an argument is supplied set a tabs top on the position of that argument. This will cause all `~t`'s to be distributed between the previous and this tabs top.
 - + Set a tabs top relative to the current position. Further the same as `~|`.
- w Give the next argument to `write/1`.
- W Give the next two argument to `write_term/2`. This option is SWI-Prolog specific.

Example:

```
simple_statistics :-
    <obtain statistics>           % left to the user
    format('~tStatistics~t~72|~n~n'),
    format('Runtime: ~\t ~2f~34| Inferences: ~\t ~D~72|~n',
           [RunT, Inf]),
    ....
```

Will output

```
Statistics

Runtime: ..... 3.45 Inferences: ..... 60,345
```

format(+Stream, +Format, +Arguments)

As `format/2`, but write the output on the given *Stream*.

sformat(-String, +Format, +Arguments)

Equivalent to `format/2`, but “writes” the result on *String* instead of the current output stream.

Example:

```
?- sformat(S, '~w~t~15|~w', ['Hello', 'World']).

S = "Hello           World"
```

sformat(-String, +Format)

Equivalent to `'sformat(String, Format, [])'`.

4.33.3 Programming Format

format_predicate(+Char, +Head)

If a sequence `~c` (tilde, followed by some character) is found, the format derivatives will first check whether the user has defined a predicate to handle the format. If not, the built in formatting rules described above are used. *Char* is either an ASCII value, or a one character atom, specifying the letter to be (re)defined. *Head* is a term, whose name and arity are used to determine the predicate to call for the redefined formatting character. The first argument to the predicate is the numeric argument of the format command, or the atom `default` if no argument is specified. The remaining arguments are filled from the argument list. The example below redefines `~n` to produce *Arg* times return followed by linefeed (so a (Grr.) DOS machine is happy with the output).

```
:- format_predicate(n, dos_newline(_Arg)).

dos_newline(Arg) :-
    between(1, Ar, _), put(13), put(10), fail ; true.
```

current_format_predicate(?Code, ?:Head)

Enumerates all user-defined format predicates. *Code* is the character code of the format character. *Head* is unified with a term with the same name and arity as the predicate. If the predicate does not reside in module `user`, *Head* is qualified with the definition module of the predicate.

4.34 Terminal Control

The following predicates form a simple access mechanism to the Unix termcap library to provide terminal independent I/O for screen terminals. These predicates are only available on Unix machines. The SWI-Prolog Windows consoles accepts the ANSI escape sequences.

tty_get_capability(+Name, +Type, -Result)

Get the capability named *Name* from the termcap library. See `termcap(5)` for the capability names. *Type* specifies the type of the expected result, and is one of `string`, `number` or `bool`. String results are returned as an atom, number result as an integer and `bool` results as the atom `on` or `off`. If an option cannot be found this predicate fails silently. The results are only computed once. Successive queries on the same capability are fast.

tty_goto(+X, +Y)

Goto position (X, Y) on the screen. Note that the predicates `line_count/2` and `line_position/2` will not have a well defined behaviour while using this predicate.

tty_put(+Atom, +Lines)

Put an atom via the termcap library function `tputs()`. This function decodes padding information in the strings returned by `tty_get_capability/3` and should be used to output these strings. *Lines* is the number of lines affected by the operation, or 1 if not applicable (as in almost all cases).

set_tty(-OldStream, +NewStream)

Set the output stream, used by `tty_put/2` and `tty_goto/2` to a specific stream. Default is `user_output`.

tty_size(-Rows, -Columns)

Determine the size of the terminal. Platforms:

Unix If the system provides *ioctl* calls for this, these are used and `tty_size/2` properly reflects the actual size after a user resize of the window. As a fallback, the system uses `tty_get_capability/2` using `li` and `co` capabilities. In this case the reported size reflects the size at the first call and is not updated after a user-initiated resize of the terminal.

Windows Getting the size of the terminal is provided for `plwin.exe`. The requested value reflects the current size. For the multi-threaded version the console that is associated with the `user_input` stream is used.

4.35 Operating System Interaction

shell(+Command, -Status)

Execute *Command* on the operating system. *Command* is given to the Bourne shell (`/bin/sh`). *Status* is unified with the exit status of the command.

On *Win32* systems, `shell/[1, 2]` executes the command using the `CreateProcess()` API and waits for the command to terminate. If the command ends with a `&` sign, the command is handed to the `WinExec()` API, which does not wait for the new task to terminate. See also `win_exec/2` and `win_shell/2`. Please note that the `CreateProcess()` API does **not** imply the Windows command interpreter (`command.exe` on Windows 95/98 and `cmd.exe` on Windows-NT) and therefore commands built-in to the command-interpreter can only be activated using the command interpreter. For example: `'command.exe /C copy file1.txt file2.txt'`

shell(+Command)

Equivalent to `'shell(Command, 0)'`.

shell

Start an interactive Unix shell. Default is `/bin/sh`, the environment variable `SHELL` overrides this default. Not available for *Win32* platforms.

win_exec(+Command, +Show)

Win32 systems only. Spawns a Windows task without waiting for its completion. *Show* is either `iconic` or `normal` and dictates the initial status of the window. The `iconic` option is notably handy to start (DDE) servers.

win_shell(+Operation, +File)

Win32 systems only. Opens the document *File* using the windows shell-rules for doing so. *Operation* is one of `open`, `print` or `explore` or another operation registered with the shell for the given document-type. On modern systems it is also possible to pass a URL as *File*, opening the URL in Windows default browser. This call interfaces to the Win32 API `ShellExecute()`.

win_registry_get_value(+Key, +Name, -Value)

Win32 systems only. Fetches the value of a Win32 registry key. *Key* is an atom formed as a path-name describing the desired registry key. *Name* is the desired attribute name of the key. *Value* is unified with the value. If the value is of type `DWORD`, the value is returned as an integer. If the value is a string it is returned as a Prolog atom. Other types are currently not supported. The default 'root' is `HKEY_CURRENT_USER`. Other roots can be specified explicitly as `HKEY_CLASSES_ROOT`, `HKEY_CURRENT_USER`, `HKEY_LOCAL_MACHINE` or `HKEY_USERS`. The example below fetches the extension to use for Prolog files (see `README.TXT` on the Windows version):

```
?- win_registry_get_value('HKEY_LOCAL_MACHINE/Software/SWI/Prolog',
                        fileExtension,
                        Ext).
```

```
Ext = pl
```

getenv(+Name, -Value)

Get environment variable. Fails silently if the variable does not exist. Please note that environment variable names are case-sensitive on Unix systems and case-insensitive on Windows.

setenv(+Name, +Value)

Set environment variable. *Name* and *Value* should be instantiated to atoms or integers. The environment variable will be passed to `shell/[0-2]` and can be requested using `getenv/2`. They also influence `expand_file_name/2`.

unsetenv(+Name)

Remove environment variable from the environment.

unix(+Command)

This predicate comes from the Quintus compatibility library and provides a partial implementation thereof. It provides access to some operating system features and unlike the name suggests, is not operating system specific. Defined *Command*'s are below.

system(+Command)

Equivalent to calling `shell/1`. Use for compatibility only.

shell(+Command)

Equivalent to calling `shell/1`. Use for compatibility only.

shell

Equivalent to calling `shell/0`. Use for compatibility only.

cd

Equivalent to calling `working_directory/2` to the expansion (see `expand_file_name/2`) of `~`. For compatibility only.

cd(+Directory)

Equivalent to calling `working_directory/2`. Use for compatibility only.

argv(-Argv)

Unify *Argv* with the list of commandline arguments provides to this Prolog run. Please note that Prolog system-arguments and application arguments are separated by `--`. Integer arguments are passed as Prolog integers, float arguments and Prolog floating point numbers and all other arguments as Prolog atoms. New applications should use the prolog-flag `argv`. See also `prolog prolog-flag argv`.

A stand-alone program could use the following skeleton to handle command-line arguments. See also section 6.

```
main :-
    unix(argv(Argv)),
    append(_PrologArgs, [--|AppArgs], Argv), !,
    main(AppArgs).
```

4.35.1 Dealing with time and date

There is no standard for time-representation in the Prolog community. SWI-Prolog represents it as a floating-point number using the same basic representation as the POSIX standard, seconds elapsed since the January 1970, 0 hours. This format is also used for predicates accessing time-information from files (see `time_file/2`).

get_time(-Time)

Return the number of seconds that elapsed since the epoch of the POSIX, tim representation: January 1970, 0 hours. *Time* is a floating point number. The granularity is system dependent.

convert_time(+Time, -Year, -Month, -Day, -Hour, -Minute, -Second, -MilliSeconds)

Convert a time stamp, provided by `get_time/1`, `time_file/2`, etc. *Year* is unified with the year, *Month* with the month number (January is 1), *Day* with the day of the month (starting with 1), *Hour* with the hour of the day (0–23), *Minute* with the minute (0–59). *Second* with the second (0–59) and *MilliSecond* with the milliseconds (0–999). Note that the latter might not be accurate or might always be 0, depending on the timing capabilities of the system. See also `convert_time/2`.

convert_time(+Time, -String)

Convert a time-stamp as obtained though `get_time/1` into a textual representation using the C-library function `ctime()`. The value is returned as a SWI-Prolog string object (see section 4.23). See also `convert_time/8`.

4.35.2 Controlling the PLWIN.EXE console window

The Windows executable `PLWIN.EXE` console has a number of predicates to control the appearance of the console. Being totally non-portable, we do not advice using it for your own application, but use `XPCE` or another portable GUI platform instead. We give the predicates for reference here.

window_title(-Old, +New)

Unify *Old* with the title displayed in the console and change the title to *New*.²⁹

²⁹BUG: This predicate should have been called `win_window_title` for consistent naming.

win_window_pos(+ListOfOptions)

Interface to the MS-Windows SetWindowPos() function, controlling size, position and stacking order of the window. *ListOfOptions* is a list that may hold any number of the terms below.

size(*W*, *H*)

Change the size of the window. *W* and *H* are expressed in character-units.

position(*X*, *Y*)

Change the top-left corner of the window. The values are expressed in pixel units.

zorder(*ZOrder*)

Change the location in the window stacking order. Values are *bottom*, *top*, *topmost* and *notopmost*. *Topmost* windows are displayed above all other windows.

show(*Bool*)

If *true*, show the window, if *false* hide the window.

activate

If present, activate the window.

win_has_menu

Succeeds if *win_insert_menu/2* and *win_insert_menu_item/4* are present.

win_insert_menu(+Label, +Before)

Insert a new entry (pulldown) in the menu. If the menu already contains this entry, nothing is done. The *Label* is the label and using the Windows conventions, a letter prefixed with & is underlined and defines the associated accelerator key. *Before* is the label before which this one must be inserted. Using - adds the new entry at the end (right). For example, the call below adds a *Application* entry just before the *Help* menu.

```
win_insert_menu('&Application', '&Help')
```

win_insert_menu_item(+Pulldown, +Label, +Before, :Goal)

Add an item to the named *Pulldown* menu. *Label* and *Before* are handled as in *win_insert_menu/2*, but the label - inserts a *separator*. *Goal* is called if the user selects the item.

4.36 File System Interaction

access_file(+File, +Mode)

Succeeds if *File* exists and can be accessed by this prolog process under mode *Mode*. *Mode* is one of the atoms *read*, *write*, *append*, *exist*, *none* or *execute*. *File* may also be the name of a directory. Fails silently otherwise. *access_file(File, none)* simply succeeds without testing anything.

If 'Mode' is *write* or *append*, this predicate also succeeds if the file does not exist and the user has write-access to the directory of the specified location.

exists_file(+File)

Succeeds when *File* exists and is a regular file. This does not imply the user has read and/or write permission for the file.

file_directory_name(+File, -Directory)

Extracts the directory-part of *File*. The returned *Directory* name does not end in /. There are two special cases. The directory-name of / is / itself and the directory-name if *File* does not contain any / characters is ..

file_base_name(+File, -BaseName)

Extracts the filename part from a path specification. If *File* does not contain any directory separators, *File* is returned.

same_file(+File1, +File2)

Succeeds if both filenames refer to the same physical file. That is, if *File1* and *File2* are the same string or both names exist and point to the same file (due to hard or symbolic links and/or relative vs. absolute paths).

exists_directory(+Directory)

Succeeds if *Directory* exists and is a directory. This does not imply the user has read, search and or write permission for the directory.

delete_file(+File)

Remove *File* from the file system.

rename_file(+File1, +File2)

Rename *File1* into *File2*. Currently files cannot be moved across devices.

size_file(+File, -Size)

Unify *Size* with the size of *File* in characters.

time_file(+File, -Time)

Unify the last modification time of *File* with *Time*. *Time* is a floating point number expressing the seconds elapsed since Jan 1, 1970. See also `convert_time/[2,8]` and `get_time/1`.

absolute_file_name(+File, -Absolute)

Expand a local file-name into an absolute path. The absolute path is canonised: references to . and .. are deleted. This predicate ensures that expanding a file-name it returns the same absolute path regardless of how the file is addressed. SWI-Prolog uses absolute file names to register source files independent of the current working directory. See also `absolute_file_name/3`. See also `absolute_file_name/3` and `expand_file_name/2`.

absolute_file_name(+Spec, +Options, -Absolute)

Converts the given file specification into an absolute path. *Option* is a list of options to guide the conversion:

extensions(ListOfExtensions)

List of file-extensions to try. Default is ''. For each extension, `absolute_file_name/3` will first add the extension and then verify the conditions imposed by the other options. If the condition fails, the next extension of the list is tried. Extensions may be specified both as ..ext or plain ext.

relative_to(+FileOrDir)

Resolve the path relative to the given directory or directory the holding the given file. Without this option, paths are resolved relative to the working directory (see `working_directory/2`) or, if *Spec* is atomic and `absolute_file_name/[2,3]` is executed in a directive, it uses the current source-file as reference.

access(Mode)

Imposes the condition `access_file(File, Mode)`. *Mode* is one of read, write, append, exist or none. See also `access_file/2`.

file_type(Type)

Defines extensions. Current mapping: `txt` implies [' '], `prolog` implies ['.pl', ' '], `executable` implies ['.so', ' '], `qlf` implies ['.qlf', ' '] and `directory` implies [' ']. The file-type source is an alias for `prolog` for compatibility to SICStus Prolog. See also `prolog_file_type/2`.

file_errors(fail/error)

If `error` (default), throw and `existence_error` exception if the file cannot be found. If `fail`, stay silent.³⁰

solutions(first/all)

If `first` (default), the predicates leaves no choice-point. Otherwise a choice-point will be left and backtracking may yield more solutions.

expand(true/false)

If `true` (default is `false`) and *Spec* is atomic, call `expand_file_name/2` followed by `member/2` on *Spec* before proceeding. This is a SWI-Prolog extension.

The prolog-flag `verbose_file_search` can be set to `true` to help debugging Prolog's search for files.

Compatibility considerations to common argument-order in ISO as well as SICStus `absolute_file_name/3` forced us to be flexible here. If the last argument is a list and the 2nd not, the arguments are swapped, making the call `absolute_file_name(+Spec, -Path, +Options)` valid as well.

is_absolute_file_name(+File)

True if *File* specifies and absolute path-name. On Unix systems, this implies the path starts with a '/'. For Microsoft based systems this implies the path starts with `<letter>:`. This predicate is intended to provide platform-independent checking for absolute paths. See also `absolute_file_name/2` and `prolog_to_os_filename/2`.

file_name_extension(?Base, ?Extension, ?Name)

This predicate is used to add, remove or test filename extensions. The main reason for its introduction is to deal with different filename properties in a portable manner. If the file system is case-insensitive, testing for an extension will be done case-insensitive too. *Extension* may be specified with or without a leading dot (.). If an *Extension* is generated, it will not have a leading dot.

expand_file_name(+Wildcard, -List)

Unify *List* with a sorted list of files or directories matching *Wildcard*. The normal Unix wildcard constructs '?', '*', '[...]' and '{...}' are recognised. The interpretation of '{...}'

³⁰Silent operation was the default up to version 3.2.6.

is interpreted slightly different from the C shell (`csh(1)`). The comma separated argument can be arbitrary patterns, including ‘{ . . . }’ patterns. The empty pattern is legal as well: ‘\{ . p l , \}’ matches either ‘. p l’ or the empty string.

If the pattern does contains wildcard characters, only existing files and directories are returned. Expanding a ‘pattern’ without wildcard characters returns the argument, regardless on whether or not it exists.

Before expanding wildchards, the construct `$var` is expanded to the value of the environment variable `var` and a possible leading `~` character is expanded to the user’s home directory.³¹

prolog_to_os_filename(?PrologPath, ?OsPath)

Converts between the internal Prolog pathname conventions and the operating-system pathname conventions. The internal conventions are Unix and this predicates is equivalent to `=/2` (unify) on Unix systems. On DOS systems it will change the directory-separator, limit the filename length map dots, except for the last one, onto underscores.

read_link(+File, -Link, -Target)

If `File` points to a symbolic link, unify `Link` with the value of the link and `Target` to the file the link is pointing to. `Target` points to a file, directory or non-existing entry in the file system, but never to a link. Fails if `File` is not a link. Fails always on systems that do not support symbolic links.

tmp_file(+Base, -TmpName)

Create a name for a temporary file. `Base` is an identifier for the category of file. The `TmpName` is guaranteed to be unique. If the system halts, it will automatically remove all created temporary files.

make_directory(+Directory)

Create a new directory (folder) on the filesystem. Raises an exception on failure. On Unix systems, the directory is created with default permissions (defined by the process `umask` setting).

delete_directory(+Directory)

Delete directory (folder) from the filesystem. Raises an exception on failure. Please note that in general it will not be possible to delete a non-empty directory.

working_directory(-Old, +New)

Unify `Old` with an absolute path to the current working directory and change working directory to `New`. Use the pattern `working_directory(CWD, CWD)` to get the current directory. See also `absolute_file_name/2` and `chdir/1`.³²

chdir(+Path)

Compatibility predicate. New code should use `working_directory/2`.

³¹On Windows, the home directory is determined as follows: if the environment variable `HOME` exists, this is used. If the variables `HOMEDRIVE` and `HOMEPATH` exist (Windows-NT), these are used. At initialisation, the system will set the environment variable `HOME` to point to the SWI-Prolog home directory if neither `HOME` nor `HOMEPATH` and `HOMEDRIVE` are defined

³²BUG: Some of the file-I/O predicates use local filenames. Changing directory while file-bound streams are open causes wrong results on `telling/1`, `seeing/1` and `current_stream/3`

4.37 User Toplevel Manipulation

break

Recursively start a new Prolog top level. This Prolog top level has its own stacks, but shares the heap with all break environments and the top level. Debugging is switched off on entering a break and restored on leaving one. The break environment is terminated by typing the system's end-of-file character (control-D). If the `-t toplevel` command line option is given this goal is started instead of entering the default interactive top level (`prolog/0`).

abort

Abort the Prolog execution and restart the top level. If the `-t toplevel` command line options is given this goal is started instead of entering the default interactive top level.

There are two implementations of `abort/0`. The default one uses the exception mechanism (see `throw/1`), throwing the exception `$aborted`. The other one uses the C-construct `longjmp()` to discard the entire environment and rebuild a new one. Using exceptions allows for proper recovery of predicates exploiting exceptions. Rebuilding the environment is safer if the Prolog stacks are corrupt. Therefore the system will use the rebuild-strategy if the abort was generated by an internal consistency check and the exception mechanism otherwise. Prolog can be forced to use the rebuild-strategy setting the prolog flag `abort_with_exception` to `false`.

halt

Terminate Prolog execution. Open files are closed and if the command line option `-tty` is not active the terminal status (see Unix `stty(1)`) is restored. Hooks may be registered both in Prolog and in foreign code. Prolog hooks are registered using `at_halt/1`. `halt/0` is equivalent to `halt(0)`.³³

halt(+Status)

Terminate Prolog execution with given status. Status is an integer. See also `halt/0`.

prolog

This goal starts the default interactive top level. Queries are read from the stream `user_input`. See also the `history` `prolog_flag` (`current_prolog_flag/2`). The `prolog/0` predicate is terminated (succeeds) by typing the end-of-file character (On most systems control-D).

The following two hooks allow for expanding queries and handling the result of a query. These hooks are used by the toplevel variable expansion mechanism described in section 2.8.

expand_query(+Query, -Expanded, +Bindings, -ExpandedBindings)

Hook in module `user`, normally not defined. *Query* and *Bindings* represents the query read from the user and the names of the free variables as obtained using `read_term/3`. If this predicate succeeds, it should bind *Expanded* and *ExpandedBindings* to the query and bindings to be executed by the toplevel. This predicate is used by the toplevel (`prolog/0`). See also `expand_answer/2` and `term_expansion/2`.

expand_answer(+Bindings, -ExpandedBindings)

Hook in module `user`, normally not defined. Expand the result of a successfully executed

³³BUG: In the multi-threaded version, `halt/0` does not work when not called from the *main* thread. In the current system a `permission_error` exception is raised. Future versions may enable `halt/0` from any thread.

toplevel query. *Bindings* is the query $\langle Name \rangle = \langle Value \rangle$ binding list from the query. *ExpandedBindings* must be unified with the bindings the toplevel should print.

4.38 Creating a Protocol of the User Interaction

SWI-Prolog offers the possibility to log the interaction with the user on a file.³⁴ All Prolog interaction, including warnings and tracer output, are written on the protocol file.

protocol(+File)

Start protocolling on file *File*. If there is already a protocol file open then close it first. If *File* exists it is truncated.

protocola(+File)

Equivalent to `protocol/1`, but does not truncate the *File* if it exists.

noprotoocol

Stop making a protocol of the user interaction. Pending output is flushed on the file.

protocolling(-File)

Succeeds if a protocol was started with `protocol/1` or `protocola/1` and unifies *File* with the current protocol output file.

4.39 Debugging and Tracing Programs

This section is a reference to the debugger interaction predicates. A more use-oriented overview of the debugger is in section 2.9.

If you have installed XPCE, you can use the graphical frontend of the tracer. This frontend is installed using the predicate `guitracer/0`.

trace

Start the tracer. `trace/0` itself cannot be seen in the tracer. Note that the Prolog toplevel treats `trace/0` special; it means ‘trace the next goal’.

tracing

Succeeds when the tracer is currently switched on. `tracing/0` itself can not be seen in the tracer.

notrace

Stop the tracer. `notrace/0` itself cannot be seen in the tracer.

guitracer

Installs hooks (see `prolog_trace_interception/4`) into the system that redirects tracing information to a GUI frontend providing structured access to variable-bindings, graphical overview of the stack and highlighting of relevant source-code.

noguitracer

Reverts back to the textual tracer.

³⁴A similar facility was added to Edinburgh C-Prolog by Wouter Jansweijer.

trace(+Pred)

Equivalent to `trace(Pred, +all)`.

trace(+Pred, +Ports)

Put a trace-point on all predicates satisfying the predicate specification *Pred*. *Ports* is a list of portnames (`call`, `redo`, `exit`, `fail`). The atom `all` refers to all ports. If the port is preceded by a `-` sign the trace-point is cleared for the port. If it is preceded by a `+` the trace-point is set.

The predicate `trace/2` activates debug mode (see `debug/0`). Each time a port (of the 4-port model) is passed that has a trace-point set the goal is printed as with `trace/0`. Unlike `trace/0` however, the execution is continued without asking for further information. Examples:

```
?- trace(hello).           Trace all ports of hello with any arity in any module.
?- trace(foo/2, +fail).    Trace failures of foo/2 in any module.
?- trace(bar/1, -all).     Stop tracing bar/1.
```

The predicate `debugging/0` shows all currently defined trace-points.

notrace(+Goal)

Call *Goal*, but suspend the debugger while *Goal* is executing. The current implementation cuts the choicepoints of *Goal* after successful completion. See `once/1`. Later implementations may have the same semantics as `call/1`.

debug

Start debugger. In debug mode, Prolog stops at spy- and trace-points, disables tail-recursion optimisation and aggressive destruction of choice-points to make debugging information accessible. Implemented by the Prolog flag `debug`.

nodebug

Stop debugger. Implemented by the prolog flag `debug`. See also `debug/0`.

debugging

Print debug status and spy points on current output stream. See also the prolog flag `debug`.

spy(+Pred)

Put a spy point on all predicates meeting the predicate specification *Pred*. See section 4.4.

nospys(+Pred)

Remove spy point from all predicates meeting the predicate specification *Pred*.

nospysall

Remove all spy points from the entire program.

leash(?Ports)

Set/query leashing (ports which allow for user interaction). *Ports* is one of `+Name`, `-Name`, `?Name` or a list of these. `+Name` enables leashing on that port, `-Name` disables it and `?Name` succeeds or fails according to the current setting. Recognised ports are: `call`, `redo`, `exit`, `fail` and `unify`. The special shorthand `all` refers to all ports, `full` refers to all ports except for the `unify` port (default). `half` refers to the `call`, `redo` and `fail` port.

visible(+Ports)

Set the ports shown by the debugger. See `leash/1` for a description of the port specification. Default is `full`.

unknown(-Old, +New)

Edinburgh-prolog compatibility predicate, interfacing to the ISO prolog flag `unknown`. Values are `trace` (meaning `error`) and `fail`. If the `unknown` flag is set to `warning`, `unknown/2` reports the value as `trace`.

style_check(+Spec)

Set style checking options. *Spec* is either `+<option>`, `-<option>`, `?<option>` or a list of such options. `+<option>` sets a style checking option, `-<option>` clears it and `?<option>` succeeds or fails according to the current setting. `consult/1` and `derivatives` resets the style checking options to their value before loading the file. If—for example—a file containing long atoms should be loaded the user can start the file with:

```
:- style_check(-atom).
```

Currently available options are:

Name	Default	Description
<code>singleton</code>	<code>on</code>	<code>read_clause/1</code> (used by <code>consult/1</code>) warns on variables only appearing once in a term (clause) which have a name not starting with an underscore.
<code>atom</code>	<code>on</code>	<code>read/1</code> and <code>derivatives</code> will produce an error message on quoted atoms or strings longer than 5 lines.
<code>dollar</code>	<code>off</code>	Accept dollar as a lower case character, thus avoiding the need for quoting atoms with dollar signs. System maintenance use only.
<code>discontiguous</code>	<code>on</code>	Warn if the clauses for a predicate are not together in the same source file.
<code>string</code>	<code>off</code>	Backward compatibility. See the prolog-flag <code>double_quotes</code> (<code>current_prolog_flag/2</code>).

4.40 Obtaining Runtime Statistics

statistics(+Key, -Value)

Unify system statistics determined by *Key* with *Value*. The possible keys are given in the table 4.2. The last part of the table contains keys for compatibility to other Prolog implementations (Quintus) for improved portability. Note that the ISO standard does not define methods to collect system statistics.

statistics

Display a table of system statistics on the current output stream.

time(+Goal)

Execute *Goal* just like `once/1` (i.e., leaving no choice points), but print used time, number

agc	Number of atom garbage-collections performed
agc_gained	Number of atoms removed
agc_time	Time spent in atom garbage-collections
cputime	(User) CPU time since Prolog was started in seconds
inferences	Total number of passes via the call and redo ports since Prolog was started.
heap	Estimated total size of the heap (see section 2.16.1)
heapused	Bytes heap in use by Prolog.
heaplimit	Maximum size of the heap (see section 2.16.1)
local	Allocated size of the local stack in bytes.
localused	Number of bytes in use on the local stack.
locallimit	Size to which the local stack is allowed to grow
global	Allocated size of the global stack in bytes.
globalused	Number of bytes in use on the global stack.
globallimit	Size to which the global stack is allowed to grow
trail	Allocated size of the trail stack in bytes.
trailused	Number of bytes in use on the trail stack.
traillimit	Size to which the trail stack is allowed to grow
atoms	Total number of defined atoms.
functors	Total number of defined name/arity pairs.
predicates	Total number of predicate definitions.
modules	Total number of module definitions.
codes	Total amount of byte codes in all clauses.
threads	MT-version: number of active threads
threads_created	MT-version: number of created threads
thread_cputime	MT-version: seconds CPU time used by finished threads. Supported on Windows-NT and later, Linux and possibly a few more. Verify it gives plausible results before using.
Compatibility keys	
runtime	[CPU time, CPU time since last] (milliseconds)
system_time	[System CPU time, System CPU time since last] (milliseconds)
real_time	[Wall time, Wall time since last] (seconds since 1970)
memory	[Total unshared data, free memory] (Uses getrusage() if available, otherwise incomplete own statistics.
stacks	[global use, local use]
program	[heap, 0]
global_stack	[global use, global free]
local_stack	[local use, local free]
trail	[trail use, 0]
garbage_collection	[number of GC, bytes gained, time spent]
stack_shifts	[global shifts, local shifts, time spent] (fails if no shifter in this version)
atoms	[number, memory use, 0]
atom_garbage_collection	[number of AGC, bytes gained, time spent]
core	Same as memory

Table 4.2: Keys for `statistics/2`

of logical inferences and the average number of *lips* (logical inferences per second). Note that SWI-Prolog counts the actual executed number of inferences rather than the number of passes through the call- and redo ports of the theoretical 4-port model.

4.41 Execution profiling

This section describes the hierarchical execution profiler introduced in SWI-Prolog 5.1.10. This profiler is based on ideas from `gprof` described in [Graham *et al.*, 1982]. The profiler consists of two parts: the information-gathering is built into the kernel,³⁵ and a presentation component which is defined in the `statistics` library. The latter can be hooked, which is used by the XPCE module `swi/pce_profile` to provide an interactive graphical representation of results.

4.41.1 Profiling predicates

Currently, the interface is kept compatible with the old profiler. As experience grows, it is likely that the old interface is replaced with one that better reflects the new capabilities. Feel free to examine the internal interfaces and report useful application thereof.

profile(:*Goal*)

Execute *Goal* just like `time/1`, collecting profiling statistics and call `show_profile(plain, 25)`. With XPCE installed this opens a graphical interface to the collected profiling data.

profile(:*Goal*, +*Style*, +*Number*)

Execute *Goal* just like `time/1`. Collect profiling statistics and show the top *Number* procedures on the current output stream (see `show_profile/1`) using *Style*. The results are kept in the database until `reset_profiler/0` or `profile/3` is called and can be displayed again with `show_profile/1`. The `profile/1` predicate is a backward compatibility interface to `profile/1`. The other predicates in this section are low-level predicates for special cases.

show_profile(+*Style*, +*Number*)

Show the collected results of the profiler. It shows the top *Number* predicates according the percentage CPU-time used. If *Style* is `plain` the time spent in the predicates itself is displayed. If *Style* is `cumulative` the time spent in its siblings (callees) is added to the predicate.

This predicate first calls `prolog:show_profile_hook/2`. If XPCE is loaded this hook is used to activate a GUI interface to visualise the profile results.

show_profile(+*Number*)

Compatibility. Same as `show_profile(plain, Number)`.

profiler(-*Old*, +*New*)

Query or change the status of the profiler. The status is a boolean (`true` or `false`) stating whether or not the profiler is collecting data. It can be used to enable or disable profiling certain parts of the program.

reset_profiler

Switches the profiler to `false` and clears all collected statistics.

³⁵There are two implementations; one based on `setitimer()` using the `SIGPROF` signal and one using Windows Multi Media (MM) timers. On other systems the profiler is not provided.

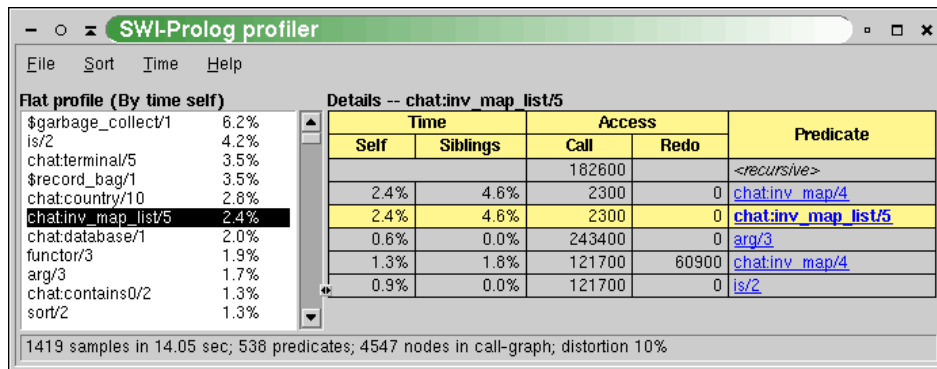


Figure 4.1: Execution profiler showing the activity of the predicate `chat:inv_map_list/5`.

`noprofile(+Name/+Arity, ...)`

Declares the predicate *Name/Arity* to be invisible to the profiler. The time spend in the named predicate is added to the caller and the callees are linked directly to the caller. This is particularly useful for simple meta-predicates such as `call/1`, `ignore/1`, `catch/3`, etc.

4.41.2 Visualizing profiling data

Browsing the annotated call-tree as described in section 4.41.3 itself is not very attractive. Therefore, the results are combined per predicate, collecting all *callers* and *callees* as well as the propagation of time and activations in both directions. Figure 4.1 illustrates this. The central yellowish line is the ‘current’ predicate with counts for time spent in the predicate (‘Self’), time spent in its children (‘Siblings’), activations through the call and redo ports. Above that are the *callers*. Here, the two time fields indicate how much time is spent serving each of the callers. The columns sum to the time in the yellowish line. The caller `<recursive>` are the number of recursive calls. Below the yellowish lines are the *callees*, with the time spent in the callee itself for serving the current predicate and the time spent in the callees of the callee (‘Siblings’), so the whole time-block adds up to the ‘Siblings’ field of the current predicate. The ‘Access’ fields show how many times the current predicate accesses each of the callees.

The predicates have a menu that allows changing the view of the detail window to the given caller or callee, showing the documentation (if it is a built-in) and/or jumping to the source.

The statistics shown in the report-field of figure 4.1 show the following information:

- *samples*
Number of times the call-tree was sampled for collecting time statistics. On most hardware the resolution of SIGPROF is 1/100 second. This number must be sufficiently large to get reliable timing figures. The Time menu allows viewing time as samples, relative time or absolute time.
- *sec*
Total user CPU time with the profiler active.
- *predicates*
Total count of predicates that have been called at least one time during the profile.
- *nodes*
Number of nodes in the call-tree.

- *distortion*

How much of the time is spend building the call-tree as a percentage of the total execution time. Timing samples while the profiler is building the call-tree are not added to the call-tree.

4.41.3 Information gathering

While the program executes under the profiler, the system builds a *dynamic* call-tree. It does this using three hooks from the kernel: one that starts a new goal (*profCall*), one that tells the system which goal is resumed after an *exit* (*profExit*) and one that tells the system which goal is resumed after a *fail* (i.e. which goal is used to *retry* (*profRedo*)). The *profCall*() function finds or creates the subnode for the argument predicate below the current node, increments the call-count of this link and returns the subnode which is recorded in the Prolog stack-frame. Choice-points are marked with the current profiling node. *profExit*() and *profRedo*() pass the profiling node where execution resumes.

Just using the above algorithm would create a much too big tree due to recursion. For this reason the system performs detection of recursion. In the simplest case, recursive procedures increment the 'recursive' count on the current node. Mutual recursion however is not easily detected. For example, `call/1` can call a predicate that uses `call/1` itself. This can be viewed as a recursive invocation, but this is generally not desirable. Recursion is currently assumed if the same predicate *with the same parent* appears higher in the call-graph. Early experience with a some arbitrary non-trivial programs are promising.

The last part of the profiler collects statistics on the CPU-time used in each node. On systems providing `setitimer()` with `SIGPROF`, it 'ticks' the current node of the call-tree each time the timer fires. On Windows a MM-timer in a separate thread checks 100 times per second how much time is spent in the profiled thread and adds this to the current node. See section 4.41.3 for details.

Profiling in the Windows Implementation

Profiling in the Windows version is similar but, especially on Windows 95/98/ME one should be aware of the implementation.³⁶ Windows does not provide timers that fire asynchronously, frequent and proportional to the CPU time used by the process. Windows does provide multi-media timers that can run at high frequency. Such timers however run in a separate thread of execution and they are fired on the wall-clock rather than the amount of CPU time used. The profiler installs such a timer running, for saving CPU time, rather inaccurately at about 100 Hz. Each time it is fired, it determines the milliseconds CPU time used by Prolog since the last time it was fired. If this value is non-zero, active predicates are incremented with this value.

On Windows 95/98/ME (DOS-based Windows), there is no possibility to get the CPU-time used by a thread or process. **Therefore, on these systems profiling results does not count CPU-time, but elapsed time.** For sensible results on these systems, ensure the system has no other active tasks and be aware that I/O operations include the time Prolog is blocked waiting for data.

4.42 Memory Management

Note: `limit_stack/2` and `trim_stacks/0` have no effect on machines that do not offer dynamic stack expansion. On these machines these predicates simply succeed to improve portability.

³⁶We hereby acknowledge Lionel Fourquaux, who suggested the design described here after a newsgroup enquiry.

garbage_collect

Invoke the global- and trail stack garbage collector. Normally the garbage collector is invoked automatically if necessary. Explicit invocation might be useful to reduce the need for garbage collections in time critical segments of the code. After the garbage collection `trim_stacks/0` is invoked to release the collected memory resources.

garbage_collect_atoms

Reclaim unused atoms. Normally invoked after `agc_margin` (a prolog flag) atoms have been created. On multi-threaded versions the actual collection is delayed until there are no threads performing normal garbage collection. In this case `garbage_collect_atoms/0` returns immediately. Note this implies there is no guarantee it will *ever* happen as there may always be threads performing garbage collection.

limit_stack(+Key, +Kbytes)

Limit one of the stack areas to the specified value. *Key* is one of `local`, `global` or `trail`. The limit is an integer, expressing the desired stack limit in K bytes. If the desired limit is smaller than the currently used value, the limit is set to the nearest legal value above the currently used value. If the desired value is larger than the maximum, the maximum is taken. Finally, if the desired value is either 0 or the atom `unlimited` the limit is set to its maximum. The maximum and initial limit is determined by the command line options `-L`, `-G` and `-T`.

trim_stacks

Release stack memory resources that are not in use at this moment, returning them to the operating system. Trim stack is a relatively cheap call. It can be used to release memory resources in a backtracking loop, where the iterations require typically seconds of execution time and very different, potentially large, amounts of stack space. Such a loop should be written as follows:

```
loop :-
    generator,
        trim_stacks,
        potentially_expensive_operation,
    stop_condition, !.
```

The prolog top level loop is written this way, reclaiming memory resources after every user query.

stack_parameter(+Stack, +Key, -Old, +New)

Query/set a parameter for the runtime stacks. *Stack* is one of `local`, `global`, `trail` or `argument`. The table below describes the *Key/Value* pairs. *Old* is first unified with the current value.

<code>limit</code>	Maximum size of the stack in bytes
<code>min.free</code>	Minimum free space at entry of foreign predicate

This predicate is currently only available on versions that use the stack-shifter to enlarge the runtime stacks when necessary. It's definition is subject to change.

4.43 Windows DDE interface

The predicates in this section deal with MS-Windows ‘Dynamic Data Exchange’ or DDE protocol.³⁷ A Windows DDE conversation is a form of interprocess communication based on sending reserved window-events between the communicating processes.

See also section 7.4 for loading Windows DLL’s into SWI-Prolog.

4.43.1 DDE client interface

The DDE client interface allows Prolog to talk to DDE server programs. We will demonstrate the use of the DDE interface using the Windows PROGMAN (Program Manager) application:

```
1 ?- open_dde_conversation(progman, progman, C).

C = 0
2 ?- dde_request(0, groups, X)

--> Unifies X with description of groups

3 ?- dde_execute(0, '[CreateGroup("DDE Demo")]').

Yes

4 ?- close_dde_conversation(0).

Yes
```

For details on interacting with `progman`, use the SDK online manual section on the Shell DDE interface. See also the Prolog `library(progman)`, which may be used to write simple Windows setup scripts in Prolog.

open_dde_conversation(+Service, +Topic, -Handle)

Open a conversation with a server supporting the given service name and topic (atoms). If successful, *Handle* may be used to send transactions to the server. If no willing server is found this predicate fails silently.

close_dde_conversation(+Handle)

Close the conversation associated with *Handle*. All opened conversations should be closed when they’re no longer needed, although the system will close any that remain open on process termination.

dde_request(+Handle, +Item, -Value)

Request a value from the server. *Item* is an atom that identifies the requested data, and *Value* will be a string (CF_TEXT data in DDE parlance) representing that data, if the request is successful. If unsuccessful, *Value* will be unified with a term of form `error(⟨Reason⟩)`, identifying the problem. This call uses SWI-Prolog string objects to return the value rather than atoms to reduce the load on the atom-space. See section 4.23 for a discussion on this data type.

³⁷This interface is contributed by Don Dwiggins.

dde_execute(+Handle, +Command)

Request the DDE server to execute the given command-string. Succeeds if the command could be executed and fails with error message otherwise.

dde_poke(+Handle, +Item, +Command)

Issue a POKE command to the server on the specified *Item*. Command is passed as data of type CF_TEXT.

4.43.2 DDE server mode

The (autoload) `library(dde)` defines primitives to realise simple DDE server applications in SWI-Prolog. These features are provided as of version 2.0.6 and should be regarded prototypes. The C-part of the DDE server can handle some more primitives, so if you need features not provided by this interface, please study `library(dde)`.

dde_register_service(+Template, +Goal)

Register a server to handle DDE request or DDE execute requests from other applications. To register a service for a DDE request, *Template* is of the form:

+Service(+Topic, +Item, +Value)

Service is the name of the DDE service provided (like `progman` in the client example above). *Topic* is either an atom, indicating *Goal* only handles requests on this topic or a variable that also appears in *Goal*. *Item* and *Value* are variables that also appear in *Goal*. *Item* represents the request data as a Prolog atom.³⁸

The example below registers the Prolog `current_prolog_flag/2` predicate to be accessible from other applications. The request may be given from the same Prolog as well as from another application.

```
?- dde_register_service(prolog(current_prolog_flag, F, V),
                       current_prolog_flag(F, V)).

?- open_dde_conversation(prolog, current_prolog_flag, Handle),
   dde_request(Handle, home, Home),
   close_dde_conversation(Handle).

Home = '/usr/local/lib/pl-2.0.6/'
```

Handling DDE `execute` requests is very similar. In this case the template is of the form:

+Service(+Topic, +Item)

Passing a *Value* argument is not needed as `execute` requests either succeed or fail. If *Goal* fails, a 'not processed' is passed back to the caller of the DDE request.

³⁸Upto version 3.4.5 this was a list of character codes. As recent versions have atom garbage collection there is no need for this anymore.

dde_unregister_service(+Service)

Stop responding to *Service*. If Prolog is halted, it will automatically call this on all open services.

dde_current_service(-Service, -Topic)

Find currently registered services and the topics served on them.

dde_current_connection(-Service, -Topic)

Find currently open conversations.

4.44 Miscellaneous

dwim_match(+Atom1, +Atom2)

Succeeds if *Atom1* matches *Atom2* in ‘Do What I Mean’ sense. Both *Atom1* and *Atom2* may also be integers or floats. The two atoms match if:

- They are identical
- They differ by one character (spy ≡ spu)
- One character is inserted/deleted (debug ≡ deug)
- Two characters are transposed (trace ≡ tarce)
- ‘Sub-words’ are glued differently (existsfile ≡ existsFile ≡ exists_file)
- Two adjacent sub words are transposed (existsFile ≡ fileExists)

dwim_match(+Atom1, +Atom2, -Difference)

Equivalent to `dwim_match/2`, but unifies *Difference* with an atom identifying the the difference between *Atom1* and *Atom2*. The return values are (in the same order as above): `equal`, `mismatched_char`, `inserted_char`, `transposed_char`, `separated` and `transposed_word`.

wildcard_match(+Pattern, +String)

Succeeds if *String* matches the wildcard pattern *Pattern*. *Pattern* is very similar the the Unix csh pattern matcher. The patterns are given below:

- ? Matches one arbitrary character.
- * Matches any number of arbitrary characters.
- [. . .] Matches one of the characters specified between the brackets.
 ⟨char1⟩-⟨char2⟩ indicates a range.
- { . . . } Matches any of the patterns of the comma separated list between the braces.

Example:

```
?- wildcard_match('[a-z]*.{pro,pl}[%~]', 'a_hello.pl%').
```

Yes

gensym(+Base, -Unique)

Generate a unique atom from base *Base* and unify it with *Unique*. *Base* should be an atom. The first call will return $\langle base \rangle 1$, the next $\langle base \rangle 2$, etc. Note that this is no warrant that the atom is unique in the system.³⁹

sleep(+Time)

Suspend execution *Time* seconds. *Time* is either a floating point number or an integer. Granularity is dependent on the system's timer granularity. A negative time causes the timer to return immediately. On most non-realtime operating systems we can only ensure execution is suspended for **at least** *Time* seconds.

On Unix systems the `sleep/1` predicate is realised—in order of preference—by `nanosleep()`, `usleep()`, `select()` if the time is below 1 minute or `sleep()`. On Windows systems `Sleep()` is used.

³⁹BUG: I plan to supply a real `gensym/2` which does give this warrant for future versions.

5

Using Modules

5.1 Why Using Modules?

In traditional Prolog systems the predicate space was flat. This approach is not very suitable for the development of large applications, certainly not if these applications are developed by more than one programmer. In many cases, the definition of a Prolog predicate requires sub-predicates that are intended only to complete the definition of the main predicate. With a flat and global predicate space these support predicates will be visible from the entire program.

For this reason, it is desirable that each source module has its own predicate space. A module consists of a declaration for its name, its *public predicates* and the predicates themselves. This approach allows the programmer to use short (local) names for support predicates without worrying about name conflicts with the support predicates of other modules. The module declaration also makes explicit which predicates are meant for public usage and which for private purposes. Finally, using the module information, cross reference programs can indicate possible problems much better.

5.2 Name-based versus Predicate-based Modules

Two approaches to realize a module system are commonly used in Prolog and other languages. The first one is the *name based* module system. In these systems, each atom read is tagged (normally prefixed) with the module name, with the exception of those atoms that are defined *public*. In the second approach, each module actually implements its own predicate space.

A critical problem with using modules in Prolog is introduced by the meta-predicates that transform between Prolog data and Prolog predicates. Consider the case where we write:

```
:- module(extend, [add_extension/3]).

add_extension(Extension, Plain, Extended) :-
    maplist(extend_atom(Extension), Plain, Extended).

extend_atom(Extension, Plain, Extended) :-
    atom_concat(Plain, Extension, Extended).
```

In this case we would like `maplist` to call `extend_atom/3` in the module `extend`. A name based module system will do this correctly. It will tag the atom `extend_atom` with the module and `maplist` will use this to construct the tagged term `extend_atom/3`. A name based module however, will not only tag the atoms that will eventually be used to refer to a predicate, but **all** atoms that are not declared `public`. So, with a name based module system also data is local to the module. This introduces another serious problem:

```
:- module(action, [action/3]).

action(Object, sleep, Arg) :- ....
action(Object, awake, Arg) :- ....

:- module(process, [awake_process/2]).

awake_process(Process, Arg) :-
    action(Process, awake, Arg).
```

This code uses a simple object-oriented implementation technique where atoms are used as method selectors. Using a name based module system, this code will not work, unless we declare the selectors public atoms in all modules that use them. Predicate based module systems do not require particular precautions for handling this case.

It appears we have to choose either to have local data, or to have trouble with meta-predicates. Probably it is best to choose for the predicate based approach as novice users will not often write generic meta-predicates that have to be used across multiple modules, but are likely to write programs that pass data around across modules. Experienced Prolog programmers should be able to deal with the complexities of meta-predicates in a predicate based module system.

5.3 Defining a Module

Modules normally are created by loading a *module file*. A module file is a file holding a `module/2` directive as its first term. The `module/2` directive declares the name and the public (i.e., externally visible) predicates of the module. The rest of the file is loaded into the module. Below is an example of a module file, defining `reverse/2`.

```
:- module(reverse, [reverse/2]).

reverse(List1, List2) :-
    rev(List1, [], List2).

rev([], List, List).
rev([Head|List1], List2, List3) :-
    rev(List1, [Head|List2], List3).
```

5.4 Importing Predicates into a Module

As explained before, in the predicate based approach adapted by SWI-Prolog, each module has its own predicate space. In SWI-Prolog, a module initially is completely empty. Predicates can be added to a module by loading a module file as demonstrated in the previous section, using `assert` or by *importing* them from another module.

Two mechanisms for importing predicates explicitly from another module exist. The `use_module/[1,2]` predicates load a module file and import (part of the) public predicates of the file. The `import/1` predicate imports any predicate from any module.

use_module(+File)

Load the file(s) specified with *File* just like `ensure_loaded/1`. The files should all be module files. All exported predicates from the loaded files are imported into the context module. The difference between this predicate and `ensure_loaded/1` becomes apparent if the file is already loaded into another module. In this case `ensure_loaded/1` does nothing; `use_module` will import all public predicates of the module into the current context module.

use_module(+File, +ImportList)

Load the file specified with *File* (only one file is accepted). *File* should be a module file. *ImportList* is a list of name/arity pairs specifying the predicates that should be imported from the loaded module. If a predicate is specified that is not exported from the loaded module a warning will be printed. The predicate will nevertheless be imported to simplify debugging.

import(+Head)

Import predicate *Head* into the current context module. *Head* should specify the source module using the `<module>:<term>` construct. Note that predicates are normally imported using one of the directives `use_module/[1,2]`. `import/1` is meant for handling imports into dynamically created modules.

It would be rather inconvenient to have to import each predicate referred to by the module, including the system predicates. For this reason each module is assigned a *default module*. All predicates in the default module are available without extra declarations. Their definition however can be overruled in the local module. This schedule is implemented by the exception handling mechanism of SWI-Prolog: if an undefined predicate exception is raised for a predicate in some module, the exception handler first tries to import the predicate from one of the module's *import modules*. On success, normal execution is resumed.

5.4.1 Reserved Modules

SWI-Prolog contains two special modules. The first one is the module `system`. This module contains all built-in predicates described in this manual. Module `system` has no default module assigned to it. The second special module is the module `user`. This module forms the initial working space of the user. Initially it is empty. The import module of module `user` is `system`, making all built-in predicate definitions available as defaults. Built-in predicates thus can be overruled by defining them in module `user` before they are used.

All other modules import from the module `user`. This implies they can use all predicates imported into `user` without explicitly importing them.

5.5 Using the Module System

The current structure of the module system has been designed with some specific organisations for large programs in mind. Many large programs define a basic library layer on top of which the actual program itself is defined. The module `user`, acting as the default module for all other modules of the program can be used to distribute these definitions over all program module without introducing the need to import this common layer each time explicitly. It can also be used to redefine built-in predicates if this is required to maintain compatibility to some other Prolog implementation. Typically, the loadfile of a large application looks like this:

```

:- use_module(compatibility).    % load XYZ prolog compatibility

:- use_module(
    [ error
      , goodies
      , debug
      , virtual_machine
      , ...
    ])
    % load generic parts
    % errors and warnings
    % general goodies (library extensions)
    % application specific debugging
    % virtual machine of application
    % more generic stuff

:- ensure_loaded(
    [ ...
    ])
    % the application itself

```

The ‘use_module’ declarations will import the public predicates from the generic modules into the user module. The ‘ensure_loaded’ directive loads the modules that constitute the actual application. It is assumed these modules import predicates from each other using `use_module/[1,2]` as far as necessary.

In combination with the object-oriented schema described below it is possible to define a neat modular architecture. The generic code defines general utilities and the message passing predicates (`invoke/3` in the example below). The application modules define classes that communicate using the message passing predicates.

5.5.1 Object Oriented Programming

Another typical way to use the module system is for defining classes within an object oriented paradigm. The class structure and the methods of a class can be defined in a module and the explicit module-boundary overruling describes in section 5.6.2 can be used by the message passing code to invoke the behaviour. An outline of this mechanism is given below.

```

%      Define class point

:- module(point, []).          % class point, no exports

%      name          type,          default access
%                  value

variable(x,          integer,       0,      both).
variable(y,          integer,       0,      both).

%      method name  predicate name arguments

behaviour(mirror,   mirror,        []).

mirror(P) :-
    fetch(P, x, X),

```

```

fetch(P, y, Y),
store(P, y, X),
store(P, x, Y).

```

The predicates `fetch/3` and `store/3` are predicates that change instance variables of instances. The figure below indicates how message passing can easily be implemented:

```

%      invoke(+Instance, +Selector, ?ArgumentList)
%      send a message to an instance

invoke(I, S, Args) :-
    class_of_instance(I, Class),
    Class:behaviour(S, P, ArgCheck), !,
    convert_arguments(ArgCheck, Args, ConvArgs),
    Goal =.. [P|ConvArgs],
    Class:Goal.

```

The construct `<Module>:<Goal>` explicitly calls `Goal` in module `Module`. It is discussed in more detail in section 5.6.

5.6 Meta-Predicates in Modules

As indicated in the introduction, the problem with a predicate based module system lies in the difficulty to find the correct predicate from a Prolog term. The predicate ‘`solution(Solution)`’ can exist in more than one module, but ‘`assert(solution(4))`’ in some module is supposed to refer to the correct version of `solution/1`.

Various approaches are possible to solve this problem. One is to add an extra argument to all predicates (e.g. ‘`assert(Module, Term)`’). Another is to tag the term somehow to indicate which module is desired (e.g. ‘`assert(Module:Term)`’). Both approaches are not very attractive as they make the user responsible for choosing the correct module, inviting unclear programming by asserting in other modules. The predicate `assert/1` is supposed to assert in the module it is called from and should do so without being told explicitly. For this reason, the notion *context module* has been introduced.

5.6.1 Definition and Context Module

Each predicate of the program is assigned a module, called its *definition module*. The definition module of a predicate is always the module in which the predicate was originally defined. Each active goal in the Prolog system has a *context module* assigned to it.

The context module is used to find predicates from a Prolog term. By default, this module is the definition module of the predicate running the goal. For meta-predicates however, this is the context module of the goal that invoked them. We call this *module-transparent* in SWI-Prolog. In the ‘using `maplist`’ example above, the predicate `maplist/3` is declared `module_transparent`. This implies the context module remains `extend`, the context module of `add_extension/3`. This way `maplist/3` can decide to call `extend_atom` in module `extend` rather than in its own definition module.

All built-in predicates that refer to predicates via a Prolog term are declared `module_transparent`. Below is the code defining `maplist`.

```

:- module(maplist, maplist/3).

:- module_transparent maplist/3.

%      maplist(+Goal, +List1, ?List2)
%      True if Goal can successfully be applied to all succes-
sive pairs
%      of elements of List1 and List2.

maplist(_, [], []).
maplist(Goal, [Elem1|Tail1], [Elem2|Tail2]) :-
    apply(Goal, [Elem1, Elem2]),
    maplist(Goal, Tail1, Tail2).

```

5.6.2 Overruling Module Boundaries

The mechanism above is sufficient to create an acceptable module system. There are however cases in which we would like to be able to overrule this schema and explicitly call a predicate in some module or assert explicitly in some module. The first is useful to invoke goals in some module from the user's toplevel or to implement a object-oriented system (see above). The latter is useful to create and modify *dynamic modules* (see section 5.7).

For this purpose, the reserved term `:/2` has been introduced. All built-in predicates that transform a term into a predicate reference will check whether this term is of the form '`⟨Module⟩:⟨Term⟩`'. If so, the predicate is searched for in *Module* instead of the goal's context module. The `:` operator may be nested, in which case the inner-most module is used.

The special calling construct `⟨Module⟩:⟨Goal⟩` pretends *Goal* is called from *Module* instead of the context module. Examples:

```

?- assert(world:done).    % asserts done/0 into module world
?- world:assert(done).   % the same
?- world:done.           % calls done/0 in module world

```

5.7 Dynamic Modules

So far, we discussed modules that were created by loading a module-file. These modules have been introduced on facilitate the development of large applications. The modules are fully defined at load-time of the application and normally will not change during execution. Having the notion of a set of predicates as a self-contained world can be attractive for other purposes as well. For example, assume an application that can reason about multiple worlds. It is attractive to store the data of a particular world in a module, so we extract information from a world simply by invoking goals in this world.

Dynamic modules can easily be created. Any built-in predicate that tries to locate a predicate in a specific module will create this module as a side-effect if it did not yet exist. Example:

```

?- assert(world_a:consistent),
   world_a:unknown(_, fail).

```

These calls create a module called ‘world_a’ and make the call ‘world_a:consistent’ succeed. Undefined predicates will not start the tracer or autoloader for this module (see `unknown/2`).

Import and export from dynamically created world is arranged via the predicates `import/1` and `export/1`:

```
?- world_b:export(solve(_,_)).           % exports solve/2 from world_b
?- world_c:import(world_b:solve(_,_)). % and import it to world_c
```

5.8 Module Handling Predicates

This section gives the predicate definitions for the remaining built-in predicates that handle modules.

`:- module(+Module, +PublicList)`

This directive can only be used as the first term of a source file. It declares the file to be a *module file*, defining *Module* and exporting the predicates of *PublicList*. *PublicList* is a list of name/arity pairs.

`module_transparent +Preds`

Preds is a comma separated list of name/arity pairs (like `dynamic/1`). Each goal associated with a transparent declared predicate will inherit the *context module* from its parent goal.

`meta_predicate +Heads`

This predicate is defined in `quintus` and provides a partial emulation of the Quintus predicate. See section [5.9.1](#) for details.

`current_module(-Module)`

Generates all currently known modules.

`current_module(?Module, ?File)`

Is true if *File* is the file from which *Module* was loaded. *File* is the internal canonical filename. See also `source_file/[1,2]`.

`context_module(-Module)`

Unify *Module* with the context module of the current goal. `context_module/1` itself is transparent.

`export(+Head)`

Add a predicate to the public list of the context module. This implies the predicate will be imported into another module if this module is imported with `use_module/[1,2]`. Note that predicates are normally exported using the directive `module/2`. `export/1` is meant to handle export from dynamically created modules.

`export_list(+Module, ?Exports)`

Unifies *Exports* with a list of terms. Each term has the name and arity of a public predicate of *Module*. The order of the terms in *Exports* is not defined. See also `predicate_property/2`.

import_module(+Module, -Import)

True if *Import* is defined as an import module for *Module*. All normal modules only import from *user*, which imports from *system*. The predicates `add_import_module/3` and `delete_import_module/2` can be used to manipulate the import list.

add_import_module(+Module, +Import, +StartOrEnd)

If *Import* is not already an import module for *Module*, add it to this list at the *start* or *end* depending on *StartOrEnd*. See also `import_module/2` and `delete_import_module/2`.

delete_import_module(+Module, +Import)

Delete *Import* from the list of import modules for *Module*. Fails silently if *Import* is not in the list.

default_module(+Module, -Default)

Successively unifies *Default* with the module names from which a call in *Module* attempts to use the definition. For the module *user*, this will generate *user* and *system*. For any other module, this will generate the module itself, followed by *user* and *system*.

Backward compatibility. New code should use `import_module/2`.

module(+Module)

The call `module(Module)` may be used to switch the default working module for the interactive toplevel (see `prolog/0`). This may be used to when debugging a module. The example below lists the clauses of `file_of_label/2` in the module `tex`.

```
1 ?- module(tex).
Yes
tex: 2 ?- listing(file_of_label/2).
...
```

5.9 Compatibility of the Module System

The principles behind the module system of SWI-Prolog differ in a number of aspects from the Quintus Prolog module system.

- The SWI-Prolog module system allows the user to redefine system predicates.
- All predicates that are available in the `system` and `user` modules are visible in all other modules as well.
- Quintus has the `'meta_predicate/1'` declaration were SWI-Prolog has the `module_transparent/1` declaration.

The `meta_predicate/1` declaration causes the compiler to tag arguments that pass module sensitive information with the module using the `:/2` operator. This approach has some disadvantages:

- Changing a `meta_predicate` declaration implies all predicates **calling** the predicate need to be reloaded. This can cause serious consistency problems.

```

:- module(findall, [findall/3]).

:- dynamic
    solution/1.

:- module_transparent
    findall/3,
    store/2.

findall(Var, Goal, Bag) :-
    assert(findall:solution('$mark')),
    store(Var, Goal),
    collect(Bag).

store(Var, Goal) :-
    Goal,                                % refers to context module of
                                        % caller of findall/3
    assert(findall:solution(Var)),
    fail.
store(_, _).

collect(Bag) :-
    ...,

```

Figure 5.1: findall/3 using modules

- It does not help for dynamically defined predicates calling module sensitive predicates.
- It slows down the compiler (at least in the SWI-Prolog architecture).
- At least within the SWI-Prolog architecture the run-time overhead is larger than the overhead introduced by the transparent mechanism.

Unfortunately the transparent predicate approach also has some disadvantages. If a predicate A passes module sensitive information to a predicate B , passing the same information to a module sensitive system predicate both A and B should be declared transparent. Using the Quintus approach only A needs to be treated special (i.e., declared with `meta_predicate/1`)¹. A second problem arises if the body of a transparent predicate uses module sensitive predicates for which it wants to refer to its own module. Suppose we want to define `findall/3` using `assert/1` and `retract/1`². The example in figure 5.1 gives the solution.

5.9.1 Emulating `meta_predicate/1`

The Quintus `meta_predicate/1` directive can in many cases be replaced by the transparent declaration. Below is the definition of `meta_predicate/1` as available from `quintus`.

¹Although this would make it impossible to call B directly.

²The system version uses `recordz/2` and `recorded/3`.

```
:- op(1150, fx, (meta_predicate)).

meta_predicate((Head, More)) :- !,
    meta_predicatel(Head),
    meta_predicate(More).
meta_predicate(Head) :-
    meta_predicatel(Head).

meta_predicatel(Head) :-
    Head =.. [Name|Arguments],
    member(Arg, Arguments),
    module_expansion_argument(Arg), !,
    functor(Head, Name, Arity),
    module_transparent(Name/Arity).
meta_predicatel(_) .                % just a mode declaration

module_expansion_argument(:).
module_expansion_argument(N) :- integer(N).
```

The discussion above about the problems with the transparent mechanism show the two cases in which this simple transformation does not work.

6

Multi-threaded applications

SWI-Prolog multithreading is based on standard C-language multithreading support. It is not like *ParLog* or other parallel implementations of the Prolog language. Prolog threads have their own stacks and only share the Prolog *heap*: predicates, records, flags and other global non-backtrackable data. SWI-Prolog thread support is designed with the following goals in mind.

- *Multi-threaded server applications*
Today's computing services often focus on (internet) server applications. Such applications often have need for communication between services and/or fast non-blocking service to multiple concurrent clients. The shared heap provides fast communication and thread creation is relatively cheap.¹
- *Interactive applications*
Interactive applications often need to perform extensive computation. If such computations are executed in a new thread, the main thread can process events and allow the user to cancel the ongoing computation. User interfaces can also use multiple threads, each thread dealing with input from a distinct group of windows. See also section 6.7.
- *Natural integration with foreign code*
Each Prolog thread runs in a native thread of the operating system, automatically making them cooperate with *MT-safe* foreign-code. In addition, any foreign thread can create its own Prolog engine for dealing with calling Prolog from C-code.

SWI-Prolog multi-threading is based on the POSIX thread standard [Butenhof, 1997] used on most popular systems except for MS-Windows. On Windows it uses the pthread-win32 emulation of POSIX threads mixed with the Windows native API for smoother and faster operation.

6.1 Creating and destroying Prolog threads

thread.create(:Goal, -Id, +Options)

Create a new Prolog thread (and underlying C-thread) and start it by executing *Goal*. If the thread is created successfully, the thread-identifier of the created thread is unified to *Id*. *Options* is a list of options. Currently defined options are:

local(*K-Bytes*)

Set the limit to which the local stack of this thread may grow. If omitted, the limit of the calling thread is used. See also the `-L` commandline option.

¹On an dual AMD-Athlon 1600, SWI-Prolog 5.1.0 creates and joins 4,957 threads per second elapsed time.

global(*K-Bytes*)

Set the limit to which the global stack of this thread may grow. If omitted, the limit of the calling thread is used. See also the `-G` commandline option.

trail(*K-Bytes*)

Set the limit to which the trail stack of this thread may grow. If omitted, the limit of the calling thread is used. See also the `-T` commandline option.

argument(*K-Bytes*)

Set the limit to which the argument stack of this thread may grow. If omitted, the limit of the calling thread is used. See also the `-A` commandline option.

stack(*K-Bytes*)

Set the limit to which the system stack of this thread may grow. The default, minimum and maximum values are system-dependant.

alias(*AliasName*)

Associate an 'alias-name' with the thread. This named may be used to refer to the thread and remains valid until the thread is joined (see `thread_join/2`).

detached(*Bool*)

If `false` (default), the thread can be waited for using `thread_join/2`. `thread_join/2` must be called on this thread to reclaim the all resources associated to the thread. If `true`, the system will reclaim all associated resources automatically after the thread finishes. Please note that thread identifiers are freed for reuse after a detached thread finishes or a normal thread has been joined. See also `thread_join/2` and `thread_detach/1`.

The *Goal* argument is *copied* to the new Prolog engine. This implies further instantiation of this term in either thread does not have consequences for the other thread: Prolog threads do not share data from their stacks.

thread_self(*-Id*)

Get the Prolog thread identifier of the running thread. If the thread has an alias, the alias-name is returned.

thread_join(*+Id, -Status*)

Wait for the termination of thread with given *Id*. Then unify the result-status of the thread with *Status*. After this call, *Id* becomes invalid and all resources associated with the thread are reclaimed. Note that threads with the attribute `detached(true)` cannot be joined. See also `current_thread/2`.

A thread that has been completed without `thread_join/2` being called on it is partly reclaimed: the Prolog stacks are released and the C-thread is destroyed. A small data-structure representing the exit-status of the thread is retained until `thread_join/2` is called on the thread. Defined values for *Status* are:

true

The goal has been proven successfully.

false

The goal has failed.

exception(*Term*)

The thread is terminated on an exception. See `print_message/2` to turn system exceptions into readable messages.

exited(*Term*)

The thread is terminated on `thread_exit/1` using the argument *Term*.

thread_detach(+*Id*)

Switch thread into detached-state (see `detached(Bool)` option at `thread_create/3`) at runtime. *Id* is the identifier of the thread placed in detached state. This may be the result of `PL_thread_self/1`.

One of the possible applications is to simplify debugging. Threads that are created as *detached* leave no traces if they crash. For not-detached threads the status can be inspected using `current_thread/2`. Threads nobody is waiting for may be created normally and detach themselves just before completion. This way they leave no traces on normal completion and their reason for failure can be inspected.

thread_exit(+*Term*)

Terminates the thread immediately, leaving `exited(Term)` as result-state for `thread_join/2`. If the thread has the attribute `detached(true)` it terminates, but its exit status cannot be retrieved using `thread_join/2` making the value of *Term* irrelevant. The Prolog stacks and C-thread are reclaimed.

thread_at_exit(:*Goal*)

Run *Goal* just before releasing the thread resources. This is to be compared to `at_halt/1`, but only for the current thread. These hooks are ran regardless of why the execution of the thread has been completed. As these hooks are run, the return-code is already available through `current_thread/2` using the result of `thread_self/1` as thread-identifier.

thread_setconcurrency(-*Old*, +*New*)

Determine the concurrency of the process, which is defined as the maximum number of concurrently active threads. 'Active' here means they are using CPU time. This option is provided if the thread-implementation provides `pthread_setconcurrency()`. Solaris is a typical example of this family. On other systems this predicate unifies *Old* to 0 (zero) and succeeds silently.

6.2 Monitoring threads

Normal multi-threaded applications should not need these the predicates from this section because almost any usage of these predicates is unsafe. For example checking the existence of a thread before signalling it is of no use as it may vanish between the two calls. Catching exceptions using `catch/3` is the only safe way to deal with thread-existence errors.

These predicates are provided for diagnosis and monitoring tasks. See also section 6.5, describing more high-level primitives.

current_thread(?*Id*, ?*Status*)

Enumerates identifiers and status of all currently known threads. Calling `current_thread/2` does not influence any thread. See also `thread_join/2`. For threads that have an alias-name, this name is returned in *Id* instead of the numerical thread identifier. *Status* is one of:

running

The thread is running. This is the initial status of a thread. Please note that threads waiting for something are considered running too.

false

The *Goal* of the thread has been completed and failed.

true

The *Goal* of the thread has been completed and succeeded.

exited(*Term*)

The *Goal* of the thread has been terminated using `thread_exit/1` with *Term* as argument. If the underlying native thread has exited (using `pthread_exit()`) *Term* is unbound.

exception(*Term*)

The *Goal* of the thread has been terminated due to an uncaught exception (see `throw/1` and `catch/3`).

thread_statistics(+*Id*, +*Key*, -*Value*)

Obtains statistical information on thread *Id* as `statistics/2` does in single-threaded applications. This call returns all keys of `statistics/2`, although only information statistics about the stacks and CPU time yield different values for each thread.²

mutex_statistics

Print usage statistics on internal mutexes and mutexes associated with dynamic predicates. For each mutex two numbers are printed: the number of times the mutex was acquired and the number of *collisions*: the number times the calling thread has to wait for the mutex. The collision-count is not available on Windows as this would break portability to Windows-95/98/ME or significantly harm performance. Generally collision count is close to zero on single-CPU hardware.

6.3 Thread communication

6.3.1 Message queues

Prolog threads can exchange data using dynamic predicates, database records, and other globally shared data. These provide no suitable means to wait for data or a condition as they can only be checked in an expensive polling loop. *Message queues* provide a means for threads to wait for data or conditions without using the CPU.

Each thread has a message-queue attached to it that is identified by the thread. Additional queues are created using `message_queue_create/2`.

thread_send_message(+*QueueOrThreadId*, +*Term*)

Place *Term* in the given queue or default queue of the indicated thread (which can even be the message queue of itself (see `thread_self/1`)). Any term can be placed in a message queue, but note that the term is copied to the receiving thread and variable-bindings are thus lost. This call returns immediately.

²Getting the CPU-time of a different thread is not supported on all platforms. For Microsoft, it does not work in 95/98/ME. For POSIX systems it requires `times()` to return values specific for the calling thread.

If more than one thread is waiting for messages on the given queue and at least one of these is waiting with a partially instantiated *Term*, the waiting threads are *all* sent a wakeup signal, starting a rush for the available messages in the queue. This behaviour can seriously harm performance with many threads waiting on the same queue as all-but-the-winner perform a useless scan of the queue. If there is only one waiting thread or all waiting threads wait with an unbound variable an arbitrary thread is restarted to scan the queue.³

thread_get_message(?Term)

Examines the thread message-queue and if necessary blocks execution until a term that unifies to *Term* arrives in the queue. After a term from the queue has been unified to *Term*, the term is deleted from the queue and this predicate returns.

Please note that not-unifying messages remain in the queue. After the following has been executed, thread 1 has the term `b(gnu)` in its queue and continues execution using `A` is `gnat`.

```
<thread 1>
thread_get_message(a(A)),

<thread 2>
thread_send_message(b(gnu)),
thread_send_message(a(gnat)),
```

See also `thread_peek_message/1`.

thread_peek_message(?Term)

Examines the thread message-queue and compares the queued terms with *Term* until one unifies or the end of the queue has been reached. In the first case the call succeeds (possibly instantiating *Term*). If no term from the queue unifies this call fails.

message_queue_create(?Queue)

If *Queue* is an atom, create a named queue. To avoid ambiguity of `thread_send_message/2`, the name of a queue may not be in use as a thread-name. If *Queue* is unbound an anonymous queue is created and *Queue* is unified to its identifier.

message_queue_destroy(+Queue)

Destroy a message queue created with `message_queue_create/1`. It is *not* allowed to destroy the queue of a thread. Neither is it allowed to destroy a queue other threads are waiting for or, for anonymous message queues, may try to wait for later.⁴

thread_get_message(+Queue, ?Term)

As `thread_get_message/1`, operating on a given queue. It is allowed (but not advised) to get messages from the queue of other threads.

thread_get_message(+Queue, ?Term)

As `thread_get_message/1`, operating on a given queue. It is allowed to peek into another

³See the documentation for the POSIX thread functions `pthread_cond_signal()` v.s. `pthread_cond_broadcastt()` for background information.

⁴BUG: None of these constraints are properly enforced by the system in the current implementation. It is therefore advised not to delete queues unless you are absolutely sure it is safe.

thread's message queue, an operation that can be used to check whether a thread has swallowed a message sent to it.

Explicit message queues are designed with the *worker-pool* model in mind, where multiple threads wait on a single queue and pick up the first goal to execute. Below is a simple implementation where the workers execute arbitrary Prolog goals. Note that this example provides no means to tell when all work is done. This must be realised using additional synchronisation.

```
%      create_workers(+Id, +N)
%
%      Create a pool with given Id and number of workers.

create_workers(Id, N) :-
    message_queue_create(Id),
    forall(between(1, N, _),
           thread_create(do_work(Id), _, [])).

do_work(Id) :-
    repeat,
        thread_get_message(Id, Goal),
        (   catch(Goal, E, print_message(error, E))
        -> true
        ;   print_message(error, goal_failed(Goal, worker(Id)))
        ),
    fail.

%      work(+Id, +Goal)
%
%      Post work to be done by the pool

work(Id, Goal) :-
    thread_send_message(Id, Goal).
```

6.3.2 Signalling threads

These predicates provide a mechanism to make another thread execute some goal as an *interrupt*. Signalling threads is safe as these interrupts are only checked at safe points in the virtual machine. Nevertheless, signalling in multi-threaded environments should be handled with care as the receiving thread may hold a *mutex* (see `with_mutex`). Signalling probably only makes sense to start debugging threads and to cancel no-longer-needed threads with `throw/1`, where the receiving thread should be designed carefully do handle exceptions at any point.

thread_signal(+ThreadId, :Goal)

Make thread *ThreadId* execute *Goal* at the first opportunity. In the current implementation, this implies at the first pass through the *Call-port*. The predicate `thread_signal/2` itself places *Goal* into the signalled-thread's signal queue and returns immediately.

Signals (interrupts) do not cooperate well with the world of multi-threading, mainly because the status of mutexes cannot be guaranteed easily. At the call-port, the Prolog virtual machine holds no locks and therefore the asynchronous execution is safe.

Goal can be any valid Prolog goal, including `throw/1` to make the receiving thread generate an exception and `trace/0` to start tracing the receiving thread.

In the Windows version, the receiving thread immediately executes the signal if it reaches a Windows `GetMessage()` call, which generally happens of the thread is waiting for (user-)input.

6.3.3 Threads and dynamic predicates

Besides queues (section 6.3.1) threads can share and exchange data using dynamic predicates. The multi-threaded version knows about two types of dynamic predicates. By default, a predicate declared *dynamic* (see `dynamic/1`) is shared by all threads. Each thread may assert, retract and run the dynamic predicate. Synchronisation inside Prolog guarantees the consistency of the predicate. Updates are *logical*: visible clauses are not affected by `assert/retract` after a query started on the predicate. In many cases primitive from section 6.4 should be used to ensure application invariants on the predicate are maintained.

Besides shared predicates, dynamic predicates can be declared with the `thread_local/1` directive. Such predicates share their attributes, but the clause-list is different in each thread.

thread_local *+Functor/+Arity, ...*

This directive is related to the `dynamic/1` directive. It tells the system that the predicate may be modified using `assert/1`, `retract/1`, etc. during execution of the program. Unlike normal shared dynamic data however each thread has its own clause-list for the predicate. As a thread starts, this clause list is empty. If there are still clauses as the thread terminates these are automatically reclaimed by the system (see also `volatile/1`). The `thread_local` property implies the properties `dynamic` and `volatile`.

Thread-local dynamic predicates are intended for maintaining thread-specific state or intermediate results of a computation.

It is not recommended to put clauses for a thread-local predicate into a file as in the example below as the clause is only visible from the thread that loaded the source-file. All other threads start with an empty clause-list.

```
:- thread_local
   foo/1.
```

```
foo(gnat).
```

DISCLAIMER Whether or not this declaration is appropriate in the sense of the proper mechanism to reach the goal is still debated. If you have strong feeling in favour or against, please share them in the SWI-Prolog mailing list.

6.4 Thread synchronisation

All internal Prolog operations are thread-safe. This implies two Prolog threads can operate on the same dynamic predicate without corrupting the consistency of the predicate. This section deals with

user-level *mutexes* (called *monitors* in ADA or *critical-sections* by Microsoft). A mutex is a **MUTual EXclusive** device, which implies at most one thread can *hold* a mutex.

Mutexes are used to realise related updates to the Prolog database. With ‘related’, we refer to the situation where a ‘transaction’ implies two or more changes to the Prolog database. For example, we have a predicate `address/2`, representing the address of a person and we want to change the address by retracting the old and asserting the new address. Between these two operations the database is invalid: this person has either no address or two addresses, depending on the assert/retract order.

Here is how to realise a correct update:

```
:- initialization
    mutex_create(addressbook).

change_address(Id, Address) :-
    mutex_lock(addressbook),
    retractall(address(Id, _)),
    asserta(address(Id, Address)),
    mutex_unlock(addressbook).
```

mutex_create(?MutexId)

Create a mutex. If *MutexId* is an atom, a *named* mutex is created. If it is a variable, an anonymous mutex reference is returned. There is no limit to the number of mutexes that can be created.

mutex_destroy(+MutexId)

Destroy a mutex. After this call, *MutexId* becomes invalid and further references yield an `existence_error` exception.

with_mutex(+MutexId, :Goal)

Execute *Goal* while holding *MutexId*. If *Goal* leaves choicepoints, these are destroyed (as in `once/1`). The mutex is unlocked regardless of whether *Goal* succeeds, fails or raises an exception. An exception thrown by *Goal* is re-thrown after the mutex has been successfully unlocked. See also `mutex_create/2` and `call_cleanup/3`.

Although described in the thread-section, this predicate is also available in the single-threaded version, where it behaves simply as `once/1`.

mutex_lock(+MutexId)

Lock the mutex. Prolog mutexes are *recursive* mutexes: they can be locked multiple times by the same thread. Only after unlocking it as many times as it is locked, the mutex becomes available for locking by other threads. If another thread has locked the mutex the calling thread is suspended until the mutex is unlocked.

If *MutexId* is an atom, and there is no current mutex with that name, the mutex is created automatically using `mutex_create/1`. This implies named mutexes need not be declared explicitly.

Please note that locking and unlocking mutexes should be paired carefully. Especially make sure to unlock mutexes even if the protected code fails or raises an exception. For most common cases use `with_mutex/2`, which provides a safer way for handling prolog-level mutexes. The predicate `call_cleanup/[2-3]` is another way to guarantee that the mutex is unlocked while retaining non-determinism.

mutex_trylock(+MutexId)

As `mutex_lock/1`, but if the mutex is held by another thread, this predicate fails immediately.

mutex_unlock(+MutexId)

Unlock the mutex. This can only be called if the mutex is held by the calling thread. If this is not the case, a `permission_error` exception is raised.

mutex_unlock_all

Unlock all mutexes held by the current thread. This call is especially useful to handle thread-termination using `abort/0` or exceptions. See also `thread_signal/2`.

current_mutex(?MutexId, ?ThreadId, ?Count)

Enumerates all existing mutexes. If the mutex is held by some thread, *ThreadId* is unified with the identifier of the holding thread and *Count* with the recursive count of the mutex. Otherwise, *ThreadId* is `[]` and *Count* is 0.

6.5 Thread-support library(threadutil)

This library defines a couple of useful predicates for demonstrating and debugging multi-threaded applications. This library is certainly not complete.

threads

Lists all current threads and their status. In addition, all ‘zombie’ threads (finished threads that are not detached, nor waited for) are joined to reclaim their resources.

interactor

Create a new console and run the Prolog toplevel in this new console. See also `attach_console/0`. In the Windows version a new interactor can also be created from the Run/New thread menu.

attach_console

If the current thread has no console attached yet, attach one and redirect the user streams (input, output, and error) to the new console window. On Unix systems the console is an `xterm` application. On Windows systems this requires the GUI version `plwin.exe` rather than the console based `plcon.exe`.

This predicate has a couple of useful applications. One is to separate (debugging) I/O of different threads. Another is to start debugging a thread that is running in the background. If thread 10 is running, the following sequence starts the tracer on this thread:

```
?- thread_signal(10, (attach_console, trace)).
```

6.6 Multi-threaded mixed C and Prolog applications

All foreign-code linked to the multi-threading version of SWI-Prolog should be thread-safe (*reentrant*) or guarded in Prolog using `with_mutex/2` from simultaneous access from multiple Prolog

threads. If you want to write mixed multi-threaded C and Prolog application you should first familiarise yourself with writing multi-threaded applications in C (C++).

If you are using SWI-Prolog as an embedded engine in a multi-threaded application you can access the Prolog engine from multiple threads by creating an *engine* in each thread from which you call Prolog. Without creating an engine, a thread can only use functions that do *not* use the `term_t` type (for example `PL_new_atom()`).

The system supports two models. Section 6.6.1 describes the original one-to-one mapping. In this schema a native thread attaches a Prolog thread if it needs to call Prolog and detaches it when finished, as opposed to the model from section 6.6.2 where threads temporarily use a Prolog engine.

Please note that the interface below will only work if threading in your application is based on the same thread-library as used to compile SWI-Prolog.

6.6.1 A Prolog thread for each native thread (one-to-one)

In the one-to-one model, the thread that called `PL_initialise()` has a Prolog engine attached. If another C-thread in the system wishes to call Prolog it must first attach an engine using `PL_thread_attach_engine()` and call `PL_thread_destroy_engine()` after all Prolog work is finished. This model is especially suitable with long running threads that need to do Prolog work regularly. See section 6.6.2 for the alternative many-to-many model.

`int PL_thread_self()`

Returns the integer Prolog identifier of the engine or -1 if the calling thread has no Prolog engine. This function is also provided in the single-threaded version of SWI-Prolog, where it returns -2.

`int PL_thread_attach_engine(const PL_thread_attr_t *attr)`

Creates a new Prolog engine in the calling thread. If the calling thread already has an engine the reference count of the engine is incremented. The *attr* argument can be `NULL` to create a thread with default attributes. Otherwise it is a pointer to a structure with the definition below. For any field with value '0', the default is used. The `cancel` field may be filled with a pointer to a function that is called when `PL_cleanup()` terminates the running Prolog engines. If this function is not present or returns `FALSE` `pthread_cancel()` is used.

```
typedef struct
{ unsigned long    local_size;    /* Stack sizes (K-bytes) */
  unsigned long    global_size;
  unsigned long    trail_size;
  unsigned long    argument_size;
  char *           alias;        /* alias name */
  int              (*cancel)(int thread);
} PL_thread_attr_t;
```

The structure may be destroyed after `PL_thread_attach_engine()` has returned. On success it returns the Prolog identifier for the thread (as returned by `PL_thread_self()`). If an error occurs, -1 is returned. If this Prolog is not compiled for multi-threading, -2 is returned.

int **PL_thread_destroy_engine()**

Destroy the Prolog engine in the calling thread. Only takes effect if `PL_thread_destroy_engine()` is called as many times as `PL_thread_attach_engine()` in this thread. Returns `TRUE` on success and `FALSE` if the calling thread has no engine or this Prolog does not support threads.

Please note that construction and destruction of engines are relatively expensive operations. Only destroy an engine if performance is not critical and memory is a critical resource.

int **PL_thread_at_exit**(void (*function)(void *), void *closure, int global)

Register a handle to be called as the Prolog engine is destroyed. The handler function is called with one `void *` argument holding *closure*. If *global* is `TRUE`, the handler is installed *for all threads*. Globally installed handlers are executed after the thread-local handlers. If the handler is installed local for the current thread only (*global* == `FALSE`) it is stored in the same FIFO queue as used by `thread_at_exit/1`.

6.6.2 Pooling Prolog engines (many-to-many)

In this model Prolog engines live as entities that are independent from threads. If a thread needs to call Prolog it takes one of the engines from the pool and returns the engine when done. This model is suitable in the following identified cases:

- *Compatibility with the single-threaded version*
In the single-threaded version, foreign threads must serialise access the the one and only thread engine. Functions from this section allow sharing one engine among multiple threads.
- *Many native threads with infrequent Prolog work*
Prolog threads are expensive in terms of memory and time to create and destroy them. Systems that use a large number of threads that only infrequently need to call Prolog are better take an engine from a pool and return it there.
- *Prolog status must be handed to another thread*
This situation has been identified by Uwe Lesta when creating a .NET interface for SWI-Prolog. .NET distributes work for active internet connection over a pool of threads. If a Prolog engine contains state for a connection, it must be possible to detach the engine from a thread and re-attach it to another thread handling the same connection.

PL_engine_t **PL_create_engine**(PL_thread_attr_t *attributes)

Create a new Prolog engine. *attributes* is described with `PL_thread_attach_engine()`. Any thread can make this call after `PL_initialise()` returned success. The returned engine is not attached to any thread and lives until `PL_destroy_engine()` is used on the returned handle.

In the single-threaded version this call always returns `NULL`, indicating failure.

int **PL_destroy_engine**(PL_engine_t e)

Destroy the given engine. Destroying an engine is only allowed if the engine is not attached to any thread or attached to the calling thread. On success this function returns `TRUE`, on failure the return value is `FALSE`.

```
int PL_set_engine(PL_engine_t engine, PL_engine_t *old)
```

Make the calling thread ready to use *engine*. If *old* is non-NULL the current engine associated with the calling thread is stored at the given location. If *engine* equals `PL_ENGINE_MAIN` the initial engine is attached to the calling thread. If *engine* is `PL_ENGINE_CURRENT` the engine is not changed. This can be used to query the current engine. This call returns `PL_ENGINE_SET` if the engine was switched successfully, `PL_ENGINE_INVALID` if *engine* is not a valid engine handle and `PL_ENGINE_INUSE` if the engine is currently in use by another thread.

Engines can be changed at any time. For example, it is allowed to select an engine to initiate a Prolog goal, detach it and at a later moment execute the goal from another thread. Note however that the `term_t`, `qid_t` and `fid_t` types are interpreted relative to the engine for which they are created. Behaviour when passing one of these types from one engine to another is undefined.

In the single-threaded version this call only succeeds if *engine* refers to the main engine.

Engines in single-threaded SWI-Prolog

In theory it is possible to port the API of section 6.6.2 to the single-threaded version of SWI-Prolog. This allows C-programs to control multiple Prolog engines concurrently. This has not yet been realised.

6.7 Multithreading and the XPCE graphics system

GUI applications written in XPCE can benefit from the multi-threaded version of XPCE/SWI-Prolog if they need to do expensive computations that block to UI in the single-threaded version.

Due to various technical problems on both Windows and Unix/X11 threading is best exploited by handing long computations to their own thread.

The XPCE message passing system is guarded with a single *mutex*, which synchronises both access from Prolog and activation through the GUI. In MS-Windows, GUI events are processed by the thread that created the window in which the event occurred, whereas in Unix/X11 they are processed by the thread that dispatches messages.

Some tentative work is underway to improve the integration between XPCE and multi-threaded SWI-Prolog.

```
pce_dispatch(+Options)
```

Create a Prolog thread with the alias-name `pce` for XPCE event-handling. In the X11 version this call creates a thread that executes the X11 event-dispatch loop. In MS-Windows it creates a thread that executes a windows event-dispatch loop. The XPCE event-handling thread has the alias `pce`. *Options* specifies the thread-attributes as `thread.create/3`.

```
pce_call(:Goal)
```

Post *Goal* to the `pce` thread, executing it synchronous with the thread's event-loop. The `pce.call/1` predicate returns immediately without waiting. Note that *Goal* is *copied* to the `pce` thread.

For further information about XPCE in threaded applications, please visit <http://gollem.swi.psy.uva.nl/twiki/pl/bin/view/Development/MultiThreadsXPCE>

6.8 Threaded Prolog: Status and plans

With SWI-Prolog 5.1.0 Prolog threads have become portable and the performance gap between the single and multi-threaded version has become acceptable.⁵ Most of the basic machinery and API have been tested on some large applications and appear to be stable. Especially as we do not anticipate problems as long as multi-threading is not used or threads do not interact much we believe it is time to move towards making the multi-threaded version the default distribution.

⁵Dual AMD 1600+ running SuSE Linux 7.3: 18%, Single P-III 733 running SuSE Linux 8.0: 12%, Single P-III 450 running NT-4/TSE: 1%.

7

Foreign Language Interface

SWI-Prolog offers a powerful interface to C [[Kernighan & Ritchie, 1978](#)]. The main design objectives of the foreign language interface are flexibility and performance. A foreign predicate is a C-function that has the same number of arguments as the predicate represented. C-functions are provided to analyse the passed terms, convert them to basic C-types as well as to instantiate arguments using unification. Non-deterministic foreign predicates are supported, providing the foreign function with a handle to control backtracking.

C can call Prolog predicates, providing both a query interface and an interface to extract multiple solutions from a non-deterministic Prolog predicate. There is no limit to the nesting of Prolog calling C, calling Prolog, etc. It is also possible to write the 'main' in C and use Prolog as an embedded logical engine.

7.1 Overview of the Interface

A special include file called `SWI-Prolog.h` should be included with each C-source file that is to be loaded via the foreign interface. The installation process installs this file in the directory `include` in the SWI-Prolog home directory (`?- current_prolog_flag(home, Home).`). This C-header file defines various data types, macros and functions that can be used to communicate with SWI-Prolog. Functions and macros can be divided into the following categories:

- Analysing Prolog terms
- Constructing new terms
- Unifying terms
- Returning control information to Prolog
- Registering foreign predicates with Prolog
- Calling Prolog from C
- Recorded database interactions
- Global actions on Prolog (halt, break, abort, etc.)

7.2 Linking Foreign Modules

Foreign modules may be linked to Prolog in two ways. Using *static linking*, the extensions, a (short) file defining `main()` which attaches the extensions calls Prolog and the SWI-Prolog kernel distributed as a C-library are linked together to form a new executable. Using *dynamic linking*, the extensions

are linked to a shared library (.so file on most Unix systems) or dynamic-link library (.DLL file on Microsoft platforms) and loaded into the the running Prolog process.¹

7.2.1 What linking is provided?

The *static linking* schema can be used on all versions of SWI-Prolog. Whether or not dynamic linking is supported can be deduced from the prolog-flag `open_shared_object` (see `current_prolog_flag/2`). If this prolog-flag yields true, `open_shared_object/2` and related predicates are defined. See section 7.4 for a suitable high-level interface to these predicates.

7.2.2 What kind of loading should I be using?

All described approaches have their advantages and disadvantages. Static linking is portable and allows for debugging on all platforms. It is relatively cumbersome and the libraries you need to pass to the linker may vary from system to system, though the utility program `plld` described in section 7.7 often hides these problems from the user.

Loading shared objects (DLL files on Windows) provides sharing and protection and is generally the best choice. If a saved-state is created using `qsave_program/[1,2]`, an `initialization/1` directive may be used to load the appropriate library at startup.

Note that the definition of the foreign predicates is the same, regardless of the linking type used.

7.3 Dynamic Linking of shared libraries

The interface defined in this section allows the user to load shared libraries (.so files on most Unix systems, .dll files on Windows). This interface is portable to Windows as well as to Unix machines providing `dlopen(2)` (Solaris, Linux, FreeBSD, Irix and many more) or `shl_open(2)` (HP/UX). It is advised to use the predicates from section 7.4 in your application.

`open_shared_object(+File, -Handle)`

File is the name of a .so file (see your C programmers documentation on how to create a .so file). This file is attached to the current process and *Handle* is unified with a handle to the shared object. Equivalent to `open_shared_object(File, [global], Handle)`. See also `load_foreign_library/[1,2]`.

On errors, an exception `shared_object(Action, Message)` is raised. *Message* is the return value from `dLError()`.

`open_shared_object(+File, +Options, -Handle)`

As `open_shared_object/2`, but allows for additional flags to be passed. *Options* is a list of atoms. `now` implies the symbols are resolved immediately rather than lazy (default). `global` implies symbols of the loaded object are visible while loading other shared objects (by default they are local). Note that these flags may not be supported by your operating system. Check the documentation of `dlopen()` or equivalent on your operating system. Unsupported flags are silently ignored.

¹The system also contains code to load .o files directly for some operating systems, notably Unix systems using the BSD a.out executable format. As the number of Unix platforms supporting this gets quickly smaller and this interface is difficult to port and slow, it is no longer described in this manual. The best alternative would be to use the `dld` package on machines do not have shared libraries

close_shared_object(+Handle)

Detach the shared object identified by *Handle*.

call_shared_object_function(+Handle, +Function)

Call the named function in the loaded shared library. The function is called without arguments and the return-value is ignored. Normally this function installs foreign language predicates using calls to `PL_register_foreign()`.

7.4 Using the library shlib for .DLL and .SO files

This section discusses the functionality of the (autoload) library `shlib.pl`, providing an interface to shared libraries. This library can only be used if the prolog-flag `open_shared_object` is enabled.

load_foreign_library(+Lib, +Entry)

Search for the given foreign library and link it to the current SWI-Prolog instance. The library may be specified with or without the extension. First, `absolute_file_name/3` is used to locate the file. If this succeeds, the full path is passed to the low-level function to open the library. Otherwise, the plain library name is passed, exploiting the operating-system defined search mechanism for the shared library. The `file_search_path/2` alias mechanism defines the alias `foreign`, which refers to the directories `<plhome>/lib/<arch>` and `<plhome>/lib`, in this order.

If the library can be loaded, the function called *Entry* will be called without arguments. The return value of the function is ignored.

The *Entry* function will normally call `PL_register_foreign()` to declare functions in the library as foreign predicates.

load_foreign_library(+Lib)

Equivalent to `load_foreign_library/2`. For the entry-point, this function first identifies the ‘base-name’ of the library, which is defined to be the file-name with path nor extension. It will then try the entry-point `install-<base>`. On failure it will try to function `install()`. Otherwise no install function will be called.

unload_foreign_library(+Lib)

If the foreign library defines the function `uninstall-<base>()` or `uninstall()`, this function will be called without arguments and its return value is ignored. Next, `abolish/2` is used to remove all known foreign predicates defined in the library. Finally the library itself is detached from the process.

current_foreign_library(-Lib, -Predicates)

Query the currently loaded foreign libraries and their predicates. *Predicates* is a list with elements of the form `Module:Head`, indicating the predicates installed with `PL_register_foreign()` when the entry-point of the library was called.

Figure 7.1 connects a Windows message-box using a foreign function. This example was tested using Windows NT and Microsoft Visual C++ 2.0.


```

#include <windows.h>
#include <SWI-Prolog.h>

static foreign_t
pl_say_hello(term_t to)
{ char *a;

  if ( PL_get_atom_chars(to, &a) )
    { MessageBox(NULL, a, "DLL test", MB_OK|MB_TASKMODAL);

      PL_succeed;
    }

  PL_fail;
}

install_t
install()
{ PL_register_foreign("say_hello", 1, pl_say_hello, 0);
}

```

Figure 7.1: MessageBox() example in Windows NT

7.4.1 Static Linking

Below is an outline of the files structure required for statically linking SWI-Prolog with foreign extensions. `\ldots/pl` refers to the SWI-Prolog home directory (see `current_prolog_flag/2`). `<arch>` refers to the architecture identifier that may be obtained using `current_prolog_flag/2`.

<code>.../pl/runtime/<arch>/libpl.a</code>	SWI-Library
<code>\ldots/pl/include/SWI-Prolog.h</code>	Include file
<code>\ldots/pl/include/SWI-Stream.h</code>	Stream I/O include file
<code>\ldots/pl/include/SWI-Exports</code>	Export declarations (AIX only)
<code>\ldots/pl/include/stub.c</code>	Extension stub

The definition of the foreign predicates is the same as for dynamic linking. Unlike with dynamic linking however, there is no initialisation function. Instead, the file `\ldots/pl/include/stub.c` may be copied to your project and modified to define the foreign extensions. Below is `stub.c`, modified to link the lowercase example described later in this chapter:

```

#include <stdio.h>
#include <SWI-Prolog.h>

extern foreign_t pl_lowercase(term, term);

PL_extension predicates[] =
{

```

```

/*{ "name",      arity,  function,      PL_FA_<flags> },*/

  { "lowercase", 2      pl_lowercase,  0 },
  { NULL,        0,     NULL,          0 }      /* terminat-
ing line */
};

int
main(int argc, char **argv)
{ PL_register_extensions(predicates);

  if ( !PL_initialise(argc, argv) )
    PL_halt(1);

  PL_install_readline();          /* delete if not re-
quired */

  PL_halt(PL_toplevel() ? 0 : 1);
}

```

Now, a new executable may be created by compiling this file and linking it to `libpl.a` from the runtime directory and the libraries required by both the extensions and the SWI-Prolog kernel. This may be done by hand, or using the `plld` utility described in `secrplld`. If the linking is preformed ‘by hand’, the commandline option `-dump-runtime-variables` (see section 2.4) can be used to obtain the required paths, libraries and linking options to link the new executable.

7.5 Interface Data types

7.5.1 Type `term_t`: a reference to a Prolog term

The principal data-type is `term_t`. Type `term_t` is what Quintus calls `QP_term_ref`. This name indicates better what the type represents: it is a *handle* for a term rather than the term itself. Terms can only be represented and manipulated using this type, as this is the only safe way to ensure the Prolog kernel is aware of all terms referenced by foreign code and thus allows the kernel to perform garbage-collection and/or stack-shifts while foreign code is active, for example during a callback from C.

A term reference is a C unsigned long, representing the offset of a variable on the Prolog environment-stack. A foreign function is passed term references for the predicate-arguments, one for each argument. If references for intermediate results are needed, such references may be created using `PL_new_term_ref()` or `PL_new_term_refs()`. These references normally live till the foreign function returns control back to Prolog. Their scope can be explicitly limited using `PL_open_foreign_frame()` and `PL_close_foreign_frame()/PL_discard_foreign_frame()`.

A `term_t` always refers to a valid Prolog term (variable, atom, integer, float or compound term). A term lives either until backtracking takes us back to a point before the term was created, the garbage

collector has collected the term or the term was created after a `PL_open_foreign_frame()` and `PL_discard_foreign_frame()` has been called.

The foreign-interface functions can either *read*, *unify* or *write* to term-references. In this document we use the following notation for arguments of type `term_t`:

```
term_t +t   Accessed in read-mode. The '+' indicates the argument is
            'input'.
term_t -t   Accessed in write-mode.
term_t ?t   Accessed in unify-mode.
```

Term references are obtained in any of the following ways.

- *Passed as argument*
The C-functions implementing foreign predicates are passed their arguments as term-references. These references may be read or unified. Writing to these variables causes undefined behaviour.
- *Created by `PL_new_term_ref()`*
A term created by `PL_new_term_ref()` is normally used to build temporary terms or be written by one of the interface functions. For example, `PL_get_arg()` writes a reference to the term-argument in its last argument.
- *Created by `PL_new_term_refs(int n)`*
This function returns a set of term refs with the same characteristics as `PL_new_term_ref()`. See `PL_open_query()`.
- *Created by `PL_copy_term_ref(term_t t)`*
Creates a new term-reference to the same term as the argument. The term may be written to. See figure 7.3.

Term-references can safely be copied to other C-variables of type `term_t`, but all copies will always refer to the same term.

`term_t PL_new_term_ref()`

Return a fresh reference to a term. The reference is allocated on the *local* stack. Allocating a term-reference may trigger a stack-shift on machines that cannot use sparse-memory management for allocation the Prolog stacks. The returned reference describes a variable.

`term_t PL_new_term_refs(int n)`

Return *n* new term references. The first term-reference is returned. The others are *t + 1*, *t + 2*, etc. There are two reasons for using this function. `PL_open_query()` expects the arguments as a set of consecutive term references and *very* time-critical code requiring a number of term-references can be written as:

```
plmypredicate(term_t a0, term_t a1)
{ term_t t0 = PL_new_term_refs(2);
  term_t t1 = t0+1;

  ...
}
```

`term_t` **PL_copy_term_ref**(*term_t from*)

Create a new term reference and make it point initially to the same term as *from*. This function is commonly used to copy a predicate argument to a term reference that may be written.

`void` **PL_reset_term_refs**(*term_t after*)

Destroy all term references that have been created after *after*, including *after* itself. Any reference to the invalidated term references after this call results in undefined behaviour.

Note that returning from the foreign context to Prolog will reclaim all references used in the foreign context. This call is only necessary if references are created inside a loop that never exits back to Prolog. See also `PL_open_foreign_frame()`, `PL_close_foreign_frame()` and `PL_discard_foreign_frame()`.

Interaction with the garbage collector and stack-shifter

Prolog implements two mechanisms for avoiding stack overflow: garbage collection and stack expansion. On machines that allow for it, Prolog will use virtual memory management to detect stack overflow and expand the runtime stacks. On other machines Prolog will reallocate the stacks and update all pointers to them. To do so, Prolog needs to know which data is referenced by C-code. As all Prolog data known by C is referenced through term references (`term_t`), Prolog has all information necessary to perform its memory management without special precautions from the C-programmer.

7.5.2 Other foreign interface types

atom_t An atom in Prolog's internal representation. Atoms are pointers to an opaque structure. They are a unique representation for represented text, which implies that atom *A* represents the same text as atom *B* if-and-only-if *A* and *B* are the same pointer.

Atoms are the central representation for textual constants in Prolog. The transformation of C a character string to an atom implies a hash-table lookup. If the same atom is needed often, it is advised to store its reference in a global variable to avoid repeated lookup.

functor_t A functor is the internal representation of a name/arity pair. They are used to find the name and arity of a compound term as well as to construct new compound terms. Like atoms they live for the whole Prolog session and are unique.

predicate_t Handle to a Prolog predicate. Predicate handles live forever (although they can lose their definition).

qid_t Query Identifier. Used by `PL_open_query()`/`PL_next_solution()`/`PL_close_query()` to handle backtracking from C.

fid_t Frame Identifier. Used by `PL_open_foreign_frame()`/`PL_close_foreign_frame()`.

module_t A module is a unique handle to a Prolog module. Modules are used only to call predicates in a specific module.

foreign_t Return type for a C-function implementing a Prolog predicate.

control_t Passed as additional argument to non-deterministic foreign functions. See `PL_retry*()` and `PL_foreign_context*()`.

install_t Type for the `install()` and `uninstall()` functions of shared or dynamic link libraries. See `secrefshlib`.

7.6 The Foreign Include File

7.6.1 Argument Passing and Control

If Prolog encounters a foreign predicate at run time it will call a function specified in the predicate definition of the foreign predicate. The arguments $1, \dots, \langle \text{arity} \rangle$ pass the Prolog arguments to the goal as Prolog terms. Foreign functions should be declared of type `foreign_t`. Deterministic foreign functions have two alternatives to return control back to Prolog:

(return) foreign_t **PL_succeed()**

Succeed deterministically. `PL_succeed` is defined as `return TRUE`.

(return) foreign_t **PL_fail()**

Fail and start Prolog backtracking. `PL_fail` is defined as `return FALSE`.

Non-deterministic Foreign Predicates

By default foreign predicates are deterministic. Using the `PL_FA_NONDETERMINISTIC` attribute (see `PL_register_foreign()`) it is possible to register a predicate as a non-deterministic predicate. Writing non-deterministic foreign predicates is slightly more complicated as the foreign function needs context information for generating the next solution. Note that the same foreign function should be prepared to be simultaneously active in more than one goal. Suppose the `natural_number_below_n/2` is a non-deterministic foreign predicate, backtracking over all natural numbers lower than the first argument. Now consider the following predicate:

```
quotient_below_n(Q, N) :-
    natural_number_below_n(N, N1),
    natural_number_below_n(N, N2),
    Q ::= N1 / N2, !.
```

In this predicate the function `natural_number_below_n/2` simultaneously generates solutions for both its invocations.

Non-deterministic foreign functions should be prepared to handle three different calls from Prolog:

- *Initial call* (`PL_FIRST_CALL`)
Prolog has just created a frame for the foreign function and asks it to produce the first answer.
- *Redo call* (`PL_REDO`)
The previous invocation of the foreign function associated with the current goal indicated it was possible to backtrack. The foreign function should produce the next solution.
- *Terminate call* (`PL_CUTTED`)
The choice point left by the foreign function has been destroyed by a cut. The foreign function is given the opportunity to clean the environment.

Both the context information and the type of call is provided by an argument of type `control_t` appended to the argument list for deterministic foreign functions. The macro `PL_foreign_control()` extracts the type of call from the control argument. The foreign function can pass a context handle using the `PL_retry*` macros and extract the handle from the extra argument using the `PL_foreign_context*` macro.

(return) foreign_t **PL_retry(long)**

The foreign function succeeds while leaving a choice point. On backtracking over this goal the foreign function will be called again, but the control argument now indicates it is a ‘Redo’ call and the macro `PL_foreign_context()` returns the handle passed via `PL_retry()`. This handle is a 30 bits signed value (two bits are used for status indication). Defined as `return _PL_retry(n)`. See also `PL_succeed()`.

(return) foreign_t **PL_retry_address(void *)**

As `PL_retry()`, but ensures an address as returned by `malloc()` is correctly recovered by `PL_foreign_context_address()`. Defined as `return _PL_retry_address(n)`. See also `PL_succeed()`.

int **PL_foreign_control(control_t)**

Extracts the type of call from the control argument. The return values are described above. Note that the function should be prepared to handle the `PL_CUTTED` case and should be aware that the other arguments are not valid in this case.

long **PL_foreign_context(control_t)**

Extracts the context from the context argument. In the call type is `PL_FIRST_CALL` the context value is 0L. Otherwise it is the value returned by the last `PL_retry()` associated with this goal (both if the call type is `PL_REDO` as `PL_CUTTED`).

*void ** **PL_foreign_context_address(control_t)**

Extracts an address as passed in by `PL_retry_address()`.

Note: If a non-deterministic foreign function returns using `PL_succeed` or `PL_fail`, Prolog assumes the foreign function has cleaned its environment. **No** call with control argument `PL_CUTTED` will follow.

The code of figure 7.2 shows a skeleton for a non-deterministic foreign predicate definition.

7.6.2 Atoms and functors

The following functions provide for communication using atoms and functors.

atom_t **PL_new_atom(const char *)**

Return an atom handle for the given C-string. This function always succeeds. The returned handle is valid as long as the atom is referenced (see section 7.6.2).

*const char** **PL_atom_chars(atom_t atom)**

Return a C-string for the text represented by the given atom. The returned text will not be changed by Prolog. It is not allowed to modify the contents, not even ‘temporary’ as the string may reside in read-only memory. The returned string becomes invalid if the atom is garbage-collected (see section 7.6.2). Foreign functions that require the text from an atom passed in a `term_t` normally use `PL_get_atom_chars()` or `PL_get_atom_nchars()`.

```
typedef struct                                /* define a context structure */
{ ...
} context;

foreign_t
my_function(term_t a0, term_t a1, control_t handle)
{ struct context * ctxt;

  switch( PL_foreign_control(handle) )
  { case PL_FIRST_CALL:
      ctxt = malloc(sizeof(struct context));
      ...
      PL_retry_address(ctxt);
    case PL_REDO:
      ctxt = PL_foreign_context_address(handle);
      ...
      PL_retry_address(ctxt);
    case PL_CUTTED:
      ctxt = PL_foreign_context_address(handle);
      ...
      free(ctxt);
      PL_succeed;
  }
}
```

Figure 7.2: Skeleton for non-deterministic foreign functions

`functor_t` **PL_new_functor**(*atom_t name*, *int arity*)

Returns a *functor identifier*, a handle for the name/arity pair. The returned handle is valid for the entire Prolog session.

`atom_t` **PL_functor_name**(*functor_t f*)

Return an atom representing the name of the given functor.

`int` **PL_functor_arity**(*functor_t f*)

Return the arity of the given functor.

Atoms and atom-garbage collection

With the introduction of atom-garbage collection in version 3.3.0, atoms no longer live as long as the process. Instead, their lifetime is guaranteed only as long as they are referenced. In the single-threaded version, atom garbage collections are only invoked at the *call-port*. In the multi-threaded version (see section 6, they appear asynchronously, except for the invoking thread.

For dealing with atom garbage collection, two additional functions are provided:

`void` **PL_register_atom**(*atom_t atom*)

Increment the reference count of the atom by one. `PL_new_atom()` performs this automatically, returning an atom with a reference count of at least one.²

`void` **PL_unregister_atom**(*atom_t atom*)

Decrement the reference count of the atom. If the reference-count drops below zero, an assertion error is raised.

Please note that the following two calls are different with respect to atom garbage collection:

```
PL_unify_atom_chars(t, "text");
PL_unify_atom(t, PL_new_atom("text"));
```

The latter increments the reference count of the atom `text`, which effectively ensures the atom will never be collected. It is advised to use the `*_chars()` or `*_nchars()` functions whenever applicable.

7.6.3 Analysing Terms via the Foreign Interface

Each argument of a foreign function (except for the control argument) is of type `term_t`, an opaque handle to a Prolog term. Three groups of functions are available for the analysis of terms. The first just validates the type, like the Prolog predicates `var/1`, `atom/1`, etc and are called `PL_is_*`. The second group attempts to translate the argument into a C primitive type. These predicates take a `term_t` and a pointer to the appropriate C-type and return `TRUE` or `FALSE` depending on successful or unsuccessful translation. If the translation fails, the pointed-to data is never modified.

²Otherwise asynchronous atom garbage collection might destroy the atom before it is used.

Testing the type of a term

int **PL_term_type**(*term_t*)

Obtain the type of a term, which should be a term returned by one of the other interface predicates or passed as an argument. The function returns the type of the Prolog term. The type identifiers are listed below. Note that the extraction functions `PL_get_t*`() also validate the type and thus the two sections below are equivalent.

```

if ( PL_is_atom(t) )
{ char *s;

  PL_get_atom_chars(t, &s);
  ...;
}

```

or

```

char *s;
if ( PL_get_atom_chars(t, &s) )
{ ...;
}

```

PL_VARIABLE	An unbound variable. The value of term as such is a unique identifier for the variable.
PL_ATOM	A Prolog atom.
PL_STRING	A Prolog string.
PL_INTEGER	A Prolog integer.
PL_FLOAT	A Prolog floating point number.
PL_TERM	A compound term. Note that a list is a compound term <code>./2</code> .

The functions `PL_is_⟨type⟩` are an alternative to `PL_term_type()`. The test `PL_is_variable(term)` is equivalent to `PL_term_type(term) == PL_VARIABLE`, but the first is considerably faster. On the other hand, using a switch over `PL_term_type()` is faster and more readable than using an if-then-else using the functions below. All these functions return either TRUE or FALSE.

int **PL_is_variable**(*term_t*)

Returns non-zero if *term* is a variable.

int **PL_is_atom**(*term_t*)

Returns non-zero if *term* is an atom.

int **PL_is_string**(*term_t*)

Returns non-zero if *term* is a string.

int **PL_is_integer**(*term_t*)

Returns non-zero if *term* is an integer.

- `int PL_is_float(term_t)`
Returns non-zero if *term* is a float.
- `int PL_is_compound(term_t)`
Returns non-zero if *term* is a compound term.
- `int PL_is_functor(term_t, functor_t)`
Returns non-zero if *term* is compound and its functor is *functor*. This test is equivalent to `PL_get_functor()`, followed by testing the functor, but easier to write and faster.
- `int PL_is_list(term_t)`
Returns non-zero if *term* is a compound term with functor `./2` or the atom `[]`.
- `int PL_is_atomic(term_t)`
Returns non-zero if *term* is atomic (not variable or compound).
- `int PL_is_number(term_t)`
Returns non-zero if *term* is an integer or float.

Reading data from a term

The functions `PL_get_*()` read information from a Prolog term. Most of them take two arguments. The first is the input term and the second is a pointer to the output value or a term-reference.

- `int PL_get_atom(term_t +t, atom_t *a)`
If *t* is an atom, store the unique atom identifier over *a*. See also `PL_atom_chars()` and `PL_new_atom()`. If there is no need to access the data (characters) of an atom, it is advised to manipulate atoms using their handle. As the atom is referenced by *t*, it will live at least as long as *t* does. If longer live-time is required, the atom should be locked using `PL_register_atom()`.
- `int PL_get_atom_chars(term_t +t, char **s)`
If *t* is an atom, store a pointer to a 0-terminated C-string in *s*. It is explicitly **not** allowed to modify the contents of this string. Some built-in atoms may have the string allocated in read-only memory, so ‘temporary manipulation’ can cause an error.
- `int PL_get_string_chars(term_t +t, char **s, int *len)`
If *t* is a string object, store a pointer to a 0-terminated C-string in *s* and the length of the string in *len*. Note that this pointer is invalidated by backtracking, garbage-collection and stack-shifts, so generally the only save operations are to pass it immediately to a C-function that doesn’t involve Prolog.
- `int PL_get_chars(term_t +t, char **s, unsigned flags)`
Convert the argument term *t* to a 0-terminated C-string. *flags* is a bitwise disjunction from two groups of constants. The first specifies which term-types should converted and the second how the argument is stored. Below is a specification of these constants. `BUF_RING` implies, if the data is not static (as from an atom), the data is copied to the next buffer from a ring of 16 buffers. This is a convenient way of converting multiple arguments passed to a foreign predicate to C-strings. If `BUF_MALLOC` is used, the data must be freed using `PL_free()` when not needed any longer.

CVT_ATOM	Convert if term is an atom
CVT_STRING	Convert if term is a string
CVT_LIST	Convert if term is a list of integers between 1 and 255
CVT_INTEGER	Convert if term is an integer (using %d)
CVT_FLOAT	Convert if term is a float (using %f)
CVT_NUMBER	Convert if term is a integer or float
CVT_ATOMIC	Convert if term is atomic
CVT_VARIABLE	Convert variable to print-name
CVT_WRITE	Convert any term that is not converted by any of the other flags using write/1. If no BUF_* is provided, BUF_RING is implied.
CVT_ALL	Convert if term is any of the above, except for CVT_VARIABLE and CVT_WRITE
BUF_DISCARDABLE	Data must copied immediately
BUF_RING	Data is stored in a ring of buffers
BUF_MALLOC	Data is copied to a new buffer returned by PL_malloc(3). When no longer needed the user must call PL_free() on the data.

int **PL_get_list_chars**(+term_t l, char **, unsigned flags)

Same as PL_get_chars(l, s, CVT_LIST—flags), provided flags contains no of the CVT_* flags.

int **PL_get_integer**(+term_t t, int *i)

If t is a Prolog integer, assign its value over i. On 32-bit machines, this is the same as PL_get_long(), but avoids a warning from the compiler. See also PL_get_long().

int **PL_get_long**(term_t +t, long *i)

If t is a Prolog integer, assign its value over i. Note that Prolog integers have limited value-range. If t is a floating point number that can be represented as a long, this function succeeds as well.

int **PL_get_bool**(term_t +t, int *val)

If t has the value true or false, set val to the C constant TRUE or FALSE and return success. otherwise return failure.

int **PL_get_pointer**(term_t +t, void **ptr)

In the current system, pointers are represented by Prolog integers, but need some manipulation to make sure they do not get truncated due to the limited Prolog integer range. PL_put_pointer()/PL_get_pointer() guarantees pointers in the range of malloc() are handled without truncating.

int **PL_get_float**(term_t +t, double *f)

If t is a float or integer, its value is assigned over f.

int **PL_get_functor**(term_t +t, functor_t *f)

If t is compound or an atom, the Prolog representation of the name-arity pair will be assigned over f. See also PL_get_name.arity() and PL.is_functor().

int **PL_get_name_arity**(*term_t t*, *atom_t *name*, *int *arity*)

If *t* is compound or an atom, the functor-name will be assigned over *name* and the arity over *arity*. See also `PL_get_functor()` and `PL_is_functor()`.

int **PL_get_module**(*term_t t*, *module_t *module*)

If *t* is an atom, the system will lookup or create the corresponding module and assign an opaque pointer to it over *module*.

int **PL_get_arg**(*int index*, *term_t t*, *term_t -a*)

If *t* is compound and *index* is between 1 and arity (including), assign *a* with a term-reference to the argument.

int **PL_get_arg**(*int index*, *term_t t*, *term_t -a*)

Same as `PL_get_arg()`, but no checking is performed, nor whether *t* is actually a term, nor whether *index* is a valid argument-index.

Exchanging text using length and string

All internal text-representation of SWI-Prolog is represented using `char *` plus length and allow for *0-bytes* in them. The foreign library supports this by implementing a `*_nchars()` function for each applicable `*_chars()` function. Below we briefly present the signatures of these functions. For full documentation consult the `*_chars()` function.

int **PL_get_atom_nchars**(*term_t t*, *unsigned int *len*, *char **s*)

See `PL_get_atom_chars()`.

int **PL_get_list_nchars**(*term_t t*, *unsigned int *len*, *char **s*)

See `PL_get_list_chars()`.

int **PL_get_nchars**(*term_t t*, *unsigned int *len*, *char **s*, *unsigned int flags*)

See `PL_get_chars()`.

int **PL_put_atom_nchars**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_put_atom_chars()`.

int **PL_put_string_nchars**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_put_string_chars()`.

int **PL_put_list_ncodes**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_put_list_codes()`.

int **PL_put_list_nchars**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_put_list_chars()`.

int **PL_unify_atom_nchars**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_unify_atom_chars()`.

int **PL_unify_string_nchars**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_unify_string_chars()`.

int **PL_unify_list_ncodes**(*term_t t*, *unsigned int len*, *const char *s*)

See `PL_unify_codes()`.

```
int PL_unify_list_nchars(term_t t, unsigned int len, const char *s)
    See PL_unify_list_chars().
```

In addition, the following functions are available for creating and inspecting atoms:

```
atom_t PL_new_atom_nchars(unsigned int len, const char *s)
    Create a new atom as PL_new_atom(), but from length and characters.
```

```
const char * PL_atom_nchars(atom_t a, unsigned int *len)
    Extract text and length of an atom.
```

Reading a list

The functions from this section are intended to read a Prolog list from C. Suppose we expect a list of atoms, the following code will print the atoms, each on a line:

```
foreign_t
pl_write_atoms(term_t l)
{ term_t head = PL_new_term_ref();      /* variable for the ele-
ments */
  term_t list = PL_copy_term_ref(l);    /* copy as we need to write */

  while( PL_get_list(list, head, list) )
  { char *s;

    if ( PL_get_atom_chars(head, &s) )
      Sprintf("%s\n", s);
    else
      PL_fail;
  }

  return PL_get_nil(list);              /* test end for [] */
}
```

```
int PL_get_list(term_t +l, term_t -h, term_t -t)
    If l is a list and not [] assign a term-reference to the head to h and to the tail to t.
```

```
int PL_get_head(term_t +l, term_t -h)
    If l is a list and not [] assign a term-reference to the head to h.
```

```
int PL_get_tail(term_t +l, term_t -t)
    If l is a list and not [] assign a term-reference to the tail to t.
```

```
int PL_get_nil(term_t +l)
    Succeeds if l represents the atom [].
```

An example: defining `write/1` in C

Figure 7.3 shows a simplified definition of `write/1` to illustrate the described functions. This simplified version does not deal with operators. It is called `display/1`, because it mimics closely the behaviour of this Edinburgh predicate.

```

foreign_t
pl_display(term_t t)
{ functor_t functor;
  int arity, len, n;
  char *s;

  switch( PL_term_type(t) )
  { case PL_VARIABLE:
    case PL_ATOM:
    case PL_INTEGER:
    case PL_FLOAT:
      PL_get_chars(t, &s, CVT_ALL);
      Sprintf("%s", s);
      break;
    case PL_STRING:
      PL_get_string_chars(t, &s, &len);
      Sprintf("\"%s\"", s);
      break;
    case PL_TERM:
      { term_t a = PL_new_term_ref();

        PL_get_name_arity(t, &name, &arity);
        Sprintf("%s(", PL_atom_chars(name));
        for(n=1; n<=arity; n++)
          { PL_get_arg(n, t, a);
            if ( n > 1 )
              Sprintf(", ");
            pl_display(a);
          }
        Sprintf(")");
        break;
      }
    default:
      PL_fail; /* should not happen */
  }

  PL_succeed;
}

```

Figure 7.3: A Foreign definition of display/1

7.6.4 Constructing Terms

Terms can be constructed using functions from the `PL_put_*()` and `PL_cons_*()` families. This approach builds the term ‘inside-out’, starting at the leaves and subsequently creating compound terms. Alternatively, terms may be created ‘top-down’, first creating a compound holding only variables and subsequently unifying the arguments. This section discusses functions for the first approach. This approach is generally used for creating arguments for `PL_call()` and `PL_open_query`.

`void PL_put_variable(term_t -t)`

Put a fresh variable in the term. The new variable lives on the global stack. Note that the initial variable lives on the local stack and is lost after a write to the term-reference. After using this function, the variable will continue to live.

`void PL_put_atom(term_t -t, atom_t a)`

Put an atom in the term reference from a handle. See also `PL_new_atom()` and `PL_atom_chars()`.

`void PL_put_atom_chars(term_t -t, const char *chars)`

Put an atom in the term-reference constructed from the 0-terminated string. The string itself will never be references by Prolog after this function.

`void PL_put_string_chars(term_t -t, const char *chars)`

Put a zero-terminated string in the term-reference. The data will be copied. See also `PL_put_string_nchars()`.

`void PL_put_string_nchars(term_t -t, unsigned int len, const char *chars)`

Put a string, represented by a length/start pointer pair in the term-reference. The data will be copied. This interface can deal with 0-bytes in the string. See also section [7.6.17](#).

`void PL_put_list_chars(term_t -t, const char *chars)`

Put a list of ASCII values in the term-reference.

`void PL_put_integer(term_t -t, long i)`

Put a Prolog integer in the term reference.

`void PL_put_pointer(term_t -t, void *ptr)`

Put a Prolog integer in the term-reference. Provided `ptr` is in the ‘malloc()-area’, `PL_get_pointer()` will get the pointer back.

`void PL_put_float(term_t -t, double f)`

Put a floating-point value in the term-reference.

`void PL_put_functor(term_t -t, functor_t functor)`

Create a new compound term from `functor` and bind `t` to this term. All arguments of the term will be variables. To create a term with instantiated arguments, either instantiate the arguments using the `PL_unify_*()` functions or use `PL_cons_functor()`.

`void PL_put_list(term_t -l)`

Same as `PL_put_functor(l, PL_new_functor(PL_new_atom("."), 2))`.

```
void PL_put_nil(term_t l)
    Same as PL_put_atom_chars( "[ ]").
```

```
void PL_put_term(term_t t1, term_t t2)
    Make t1 point to the same term as t2.
```

```
void PL_cons_functor(term_t h, functor_t f, ... )
    Create a term, whose arguments are filled from variable argument list holding the same number of term_t objects as the arity of the functor. To create the term animal(gnu, 50), use:
```

```
{ term_t a1 = PL_new_term_ref();
  term_t a2 = PL_new_term_ref();
  term_t t  = PL_new_term_ref();
  functor_t animal2;

  /* animal2 is a constant that may be bound to a global
     variable and re-used
  */
  animal2 = PL_new_functor(PL_new_atom("animal"), 2);

  PL_put_atom_chars(a1, "gnu");
  PL_put_integer(a2, 50);
  PL_cons_functor(t, animal2, a1, a2);
}
```

After this sequence, the term-references *a1* and *a2* may be used for other purposes.

```
void PL_cons_functor_v(term_t h, functor_t f, term_t a0)
    Creates a compound term like PL_cons_functor(), but a0 is an array of term references as returned by PL_new_term_refs(). The length of this array should match the number of arguments required by the functor.
```

```
void PL_cons_list(term_t l, term_t h, term_t t)
    Create a list (cons-) cell in l from the head and tail. The code below creates a list of atoms from a char **. The list is built tail-to-head. The PL_unify_*() functions can be used to build a list head-to-tail.
```

```
void
put_list(term_t l, int n, char **words)
{ term_t a = PL_new_term_ref();

  PL_put_nil(l);
  while( --n >= 0 )
  { PL_put_atom_chars(a, words[n]);
    PL_cons_list(l, a, l);
  }
}
```

Note that *l* can be redefined within a PL_cons_list call as shown here because operationally its old value is consumed before its new value is set.

7.6.5 Unifying data

The functions of this sections *unify* terms with other terms or translated C-data structures. Except for `PL_unify()`, the functions of this section are specific to SWI-Prolog. They have been introduced to make translation of old code easier, but also because they provide for a faster mechanism for returning data to Prolog that requires less term-references. Consider the case where we want a foreign function to return the host name of the machine Prolog is running on. Using the `PL_get_*()` and `PL_put_*()` functions, the code becomes:

```
foreign_t
pl_hostname(term_t name)
{ char buf[100];

  if ( gethostname(buf, sizeof(buf)) )
  { term_t tmp = PL_new_term_ref();

    PL_put_atom_chars(tmp, buf);
    return PL_unify(name, tmp);
  }

  PL_fail;
}
```

Using `PL_unify_atom_chars()`, this becomes:

```
foreign_t
pl_hostname(term_t name)
{ char buf[100];

  if ( gethostname(buf, sizeof(buf)) )
    return PL_unify_atom_chars(name, buf);

  PL_fail;
}
```

int PL_unify(*term_t ?t1, term_t ?t2*)
Unify two Prolog terms and return non-zero on success.

int PL_unify_atom(*term_t ?t, atom_t a*)
Unify *t* with the atom *a* and return non-zero on success.

int PL_unify_atom_chars(*term_t ?t, const char *chars*)
Unify *t* with an atom created from *chars* and return non-zero on success.

int PL_unify_list_chars(*term_t ?t, const char *chars*)
Unify *t* with a list of ASCII characters constructed from *chars*.

void PL_unify_string_chars(*term_t ?t, const char *chars*)
Unify *t* with a Prolog string object created from the zero-terminated string *chars*. The data will be copied. See also `PL_unify_string_nchars()`.

`void PL_unify_string_nchars(term_t ?t, unsigned int len, const char *chars)`
 Unify *t* with a Prolog string object created from the string created from the *len/chars* pair. The data will be copied. This interface can deal with 0-bytes in the string. See also section 7.6.17.

`int PL_unify_integer(term_t ?t, long n)`
 Unify *t* with a Prolog integer from *n*.

`int PL_unify_float(term_t ?t, double f)`
 Unify *t* with a Prolog float from *f*.

`int PL_unify_pointer(term_t ?t, void *ptr)`
 Unify *t* with a Prolog integer describing the pointer. See also `PL_put_pointer()` and `PL_get_pointer()`.

`int PL_unify_functor(term_t ?t, functor_t f)`
 If *t* is a compound term with the given functor, just succeed. If it is unbound, create a term and bind the variable, else fails. Not that this function does not create a term if the argument is already instantiated.

`int PL_unify_list(term_t ?l, term_t -h, term_t -t)`
 Unify *l* with a list-cell (`./2`). If successful, write a reference to the head of the list to *h* and a reference to the tail of the list in *t*. This reference may be used for subsequent calls to this function. Suppose we want to return a list of atoms from a `char **`. We could use the example described by `PL_put_list()`, followed by a call to `PL_unify()`, or we can use the code below. If the predicate argument is unbound, the difference is minimal (the code based on `PL_put_list()` is probably slightly faster). If the argument is bound, the code below may fail before reaching the end of the word-list, but even if the unification succeeds, this code avoids a duplicate (garbage) list and a deep unification.

```
foreign_t
pl_get_environ(term_t env)
{ term_t l = PL_copy_term_ref(env);
  term_t a = PL_new_term_ref();
  extern char **environ;
  char **e;

  for(e = environ; *e; e++)
  { if ( !PL_unify_list(l, a, l) ||
        !PL_unify_atom_chars(a, *e) )
      PL_fail;
  }

  return PL_unify_nil(l);
}
```

`int PL_unify_nil(term_t ?l)`
 Unify *l* with the atom `[]`.

`int PL_unify_arg(int index, term_t ?t, term_t ?a)`
 Unifies the *index-th* argument (1-based) of *t* with *a*.

`int PL_unify_term(term_t ?t, ...)`
 Unify *t* with a (normally) compound term. The remaining arguments is a sequence of a type identifier, followed by the required arguments. This predicate is an extension to the Quintus and SICStus foreign interface from which the SWI-Prolog foreign interface has been derived, but has proved to be a powerful and comfortable way to create compound terms from C. Due to the vararg packing/unpacking and the required type-switching this interface is slightly slower than using the primitives. Please note that some bad C-compilers have fairly low limits on the number of arguments that may be passed to a function.

Special attention is required when passing numbers. C ‘promotes’ any integral smaller than `int` to `int`. I.e. the types `char`, `short` and `int` are all passed as `int`. In addition, on most 32-bit platforms `int` and `long` are the same. Upto version 4.0.5, only `PL_INTEGER` could be specified which was taken from the stack as `long`. Such code fails when passing small integral types on machines where `int` is smaller than `long`. It is advised to use `PL_SHORT`, `PL_INT` or `PL_LONG` as appropriate. Similar, C compilers promote `float` to `double` and therefore `PL_FLOAT` and `PL_DOUBLE` are synonyms.

The type identifiers are:

`PL_VARIABLE` *none*

No op. Used in arguments of `PL_FUNCTOR`.

`PL_BOOL` *int*

Unify the argument with `true` or `false`.

`PL_ATOM` *atom t*

Unify the argument with an atom, as in `PL_unify_atom()`.

`PL_CHARS` *const char **

Unify the argument with an atom created from a 0-terminated string.

`PL_NCHARS` *unsigned int, const char **

Unify the argument with an atom created from the given number of characters starting at the given address. See also `PL_unify_atom_nchars()`;

`PL_SHORT` *short*

Unify the argument with an integer, as in `PL_unify_integer()`. As `short` is promoted to `int`, `PL_SHORT` is a synonym for `PL_INT`.

`PL_INT` *int*

Unify the argument with an integer, as in `PL_unify_integer()`.

`PL_LONG` *long*

Unify the argument with an integer, as in `PL_unify_integer()`.

`PL_INTEGER` *long*

Unify the argument with an integer, as in `PL_unify_integer()`.

`PL_DOUBLE` *double*

Unify the argument with a float, as in `PL_unify_float()`. Note that, as the argument is passed using the C vararg conventions, a float must be casted to a double explicitly.

PL_FLOAT *double*

Unify the argument with a float, as in `PL_unify_float()`.

PL_POINTER *void **

Unify the argument with a pointer, as in `PL_unify_pointer()`.

PL_STRING *const char **

Unify the argument with a string object, as in `PL_unify_string_chars()`.

PL_TERM *term t*

Unify a subterm. Note this may be the return value of a `PL_new_term_ref()` call to get access to a variable.

PL_CHARS *const char **

Unify the argument with an atom, constructed from the C `char *`, as in `PL_unify_atom_chars()`.

PL_FUNCTOR *functor t, ...*

Unify the argument with a compound term. This specification should be followed by exactly as many specifications as the number of arguments of the compound term.

PL_FUNCTOR_CHARS *const char *name, int arity, ...*

Create a functor from the given name and arity and then behave as `PL_FUNCTOR`.

PL_LIST *int length, ...*

Create a list of the indicated length. The following arguments contain the elements of the list.

For example, to unify an argument with the term `language(dutch)`, the following skeleton may be used:

```
static functor_t FUNCTOR_language1;

static void
init_constants()
{ FUNCTOR_language1 = PL_new_functor(PL_new_atom("language"), 1);
}

foreign_t
pl_get_lang(term_t r)
{ return PL_unify_term(r,
                      PL_FUNCTOR, FUNCTOR_language1,
                      PL_CHARS, "dutch");
}

install_t
install()
{ PL_register_foreign("get_lang", 1, pl_get_lang, 0);
  init_constants();
}
```

int **PL_chars_to_term**(const char *chars, term_t t)

Parse the string *chars* and put the resulting Prolog term into *t*. *chars* may or may not be closed using a Prolog full-stop (i.e., a dot followed by a blank). Returns FALSE if a syntax error was encountered and TRUE after successful completion. In addition to returning FALSE, the exception-term is returned in *t* on a syntax error. See also `term_to_atom/2`.

The following example build a goal-term from a string and calls it.

```
int
call_chars(const char *goal)
{ fid_t fid = PL_open_foreign_frame();
  term_t g = PL_new_term_ref();
  BOOL rval;

  if ( PL_string_to_term(goal, g) )
    rval = PL_call(goal, NULL);
  else
    rval = FALSE;

  PL_discard_foreign_frame(fid);
  return rval;
}

...
call_chars("consult(load)");
...
```

char * **PL_quote**(int chr, const char *string)

Return a quoted version of *string*. If *chr* is `'\'`, the result is a quoted atom. If *chr* is `'"`, the result is a string. The result string is stored in the same ring of buffers as described with the `BUF_RING` argument of `PL_get_chars()`;

In the current implementation, the string is surrounded by *chr* and any occurrence of *chr* is doubled. In the future the behaviour will depend on the `character_escape` prolog-flag. See `current_prolog_flag/2`.

7.6.6 Calling Prolog from C

The Prolog engine can be called from C. There are two interfaces for this. For the first, a term is created that could be used as an argument to `call/1` and next `PL_call()` is used to call Prolog. This system is simple, but does not allow to inspect the different answers to a non-deterministic goal and is relatively slow as the runtime system needs to find the predicate. The other interface is based on `PL_open_query()`, `PL_next_solution()` and `PL_cut_query()` or `PL_close_query()`. This mechanism is more powerful, but also more complicated to use.

Predicate references

This section discusses the functions used to communicate about predicates. Though a Prolog predicate may defined or not, redefined, etc., a Prolog predicate has a handle that is not destroyed, nor moved.

This handle is known by the type `predicate_t`.

`predicate_t PL_pred(funcctor_t f, module_t m)`

Return a handle to a predicate for the specified name/arity in the given module. This function always succeeds, creating a handle for an undefined predicate if no handle was available.

`predicate_t PL_predicate(const char *name, int arity, const char* module)`

Same as `PL_pred()`, but provides a more convenient interface to the C-programmer.

`void PL_predicate_info(predicate_t p, atom_t *n, int *a, module_t *m)`

Return information on the predicate `p`. The name is stored over `n`, the arity over `a`, while `m` receives the definition module. Note that the latter need not be the same as specified with `PL_predicate()`. If the predicate was imported into the module given to `PL_predicate()`, this function will return the module where the predicate was defined.

Initiating a query from C

This section discusses the functions for creating and manipulating queries from C. Note that a foreign context can have at most one active query. This implies it is allowed to make strictly nested calls between C and Prolog (Prolog calls C, calls Prolog, calls C, etc., but it is **not** allowed to open multiple queries and start generating solutions for each of them by calling `PL_next_solution()`. Be sure to call `PL_cut_query()` or `PL_close_query()` on any query you opened before opening the next or returning control back to Prolog.

`qid_t PL_open_query(module_t ctx, int flags, predicate_t p, term_t +t0)`

Opens a query and returns an identifier for it. This function always succeeds, regardless whether the predicate is defined or not. `ctx` is the *context module* of the goal. When `NULL`, the context module of the calling context will be used, or `user` if there is no calling context (as may happen in embedded systems). Note that the context module only matters for *module_transparent* predicates. See `context_module/1` and `module_transparent/1`. The `p` argument specifies the predicate, and should be the result of a call to `PL_pred()` or `PL_predicate()`. Note that it is allowed to store this handle as global data and reuse it for future queries. The term-reference `t0` is the first of a vector of term-references as returned by `PL_new_term_refs(n)`. The `flags` arguments provides some additional options concerning debugging and exception handling. It is a bitwise or of the following values:

`PL_Q_NORMAL`

Normal operation. The debugger inherits its settings from the environment. If an exception occurs that is not handled in Prolog, a message is printed and the tracer is started to debug the error.³

`PL_Q_NODEBUG`

Switch off the debugger while executing the goal. This option is used by many calls to hook-predicates to avoid tracing the hooks. An example is `print/1` calling `portray/1` from foreign code.

³Do not pass the integer 0 for normal operation, as this is interpreted as `PL_Q_NODEBUG` for backward compatibility reasons.

PL_Q_CATCH_EXCEPTION

If an exception is raised while executing the goal, do not report it, but make it available for `PL_exception()`.

PL_Q_PASS_EXCEPTION

As `PL_Q_CATCH_EXCEPTION`, but do not invalidate the exception-term while calling `PL_close_query()`. This option is experimental.

The example below opens a query to the predicate `is_a/2` to find the ancestor of for some name.

```
char *
ancestor(const char *me)
{ term_t a0 = PL_new_term_refs(2);
  static predicate_t p;

  if ( !p )
    p = PL_predicate("is_a", 2, "database");

  PL_put_atom_chars(a0, me);
  PL_open_query(NULL, PL_Q_NORMAL, p, a0);
  ...
}
```

int PL_next_solution(*qid_t qid*)

Generate the first (next) solution for the given query. The return value is `TRUE` if a solution was found, or `FALSE` to indicate the query could not be proven. This function may be called repeatedly until it fails to generate all solutions to the query.

void PL_cut_query(*qid*)

Discards the query, but does not delete any of the data created by the query. It just invalidate *qid*, allowing for a new call to `PL_open_query()` in this context.

void PL_close_query(*qid*)

As `PL_cut_query()`, but all data and bindings created by the query are destroyed.

int PL_call_predicate(*module_t m, int flags, predicate_t pred, term_t +t0*)

Shorthand for `PL_open_query()`, `PL_next_solution()`, `PL_cut_query()`, generating a single solution. The arguments are the same as for `PL_open_query()`, the return value is the same as `PL_next_solution()`.

int PL_call(*term_t, module_t*)

Call term just like the Prolog predicate `once/1`. *Term* is called in the specified module, or in the context module if `module_t = NULL`. Returns `TRUE` if the call succeeds, `FALSE` otherwise. Figure 7.4 shows an example to obtain the number of defined atoms. All checks are omitted to improve readability.

7.6.7 Discarding Data

The Prolog data created and term-references needed to setup the call and/or analyse the result can in most cases be discarded right after the call. `PL_close_query()` allows for destructing the data, while leaving the term-references. The calls below may be used to destroy term-references and data. See figure 7.4 for an example.

`fid_t PL_open_foreign_frame()`

Created a foreign frame, holding a mark that allows the system to undo bindings and destroy data created after it as well as providing the environment for creating term-references. This function is called by the kernel before calling a foreign predicate.

`void PL_close_foreign_frame(fid_t id)`

Discard all term-references created after the frame was opened. All other Prolog data is retained. This function is called by the kernel whenever a foreign function returns control back to Prolog.

`void PL_discard_foreign_frame(fid_t id)`

Same as `PL_close_foreign_frame()`, but also undo all bindings made since the open and destroy all Prolog data.

`void PL_rewind_foreign_frame(fid_t id)`

Undo all bindings and discard all term-references created since the frame was created, but does not pop the frame. I.e. the same frame can be rewinded multiple times, and must eventually be closed or discarded.

It is obligatory to call either of the two closing functions to discard a foreign frame. Foreign frames may be nested.

7.6.8 Foreign Code and Modules

Modules are identified via a unique handle. The following functions are available to query and manipulate modules.

`module_t PL_context()`

Return the module identifier of the context module of the currently active foreign predicate.

`int PL_strip_module(term_t +raw, module_t *m, term_t -plain)`

Utility function. If *raw* is a term, possibly holding the module construct $\langle module \rangle : \langle rest \rangle$ this function will make *plain* a reference to $\langle rest \rangle$ and fill *module ** with $\langle module \rangle$. For further nested module constructs the inner most module is returned via *module **. If *raw* is not a module construct *arg* will simply be put in *plain*. If *module ** is NULL it will be set to the context module. Otherwise it will be left untouched. The following example shows how to obtain the plain term and module if the default module is the user module:

```
{ module m = PL_new_module(PL_new_atom("user"));
  term_t plain = PL_new_term_ref();

  PL_strip_module(term, &m, plain);
  ...
```



```

int
count_atoms()
{ fid_t fid = PL_open_foreign_frame();
  term_t goal = PL_new_term_ref();
  term_t a1 = PL_new_term_ref();
  term_t a2 = PL_new_term_ref();
  functor_t s2 = PL_new_functor(PL_new_atom("statistics"), 2);
  int atoms;

  PL_put_atom_chars(a1, "atoms");
  PL_cons_functor(goal, s2, a1, a2);
  PL_call(goal, NULL); /* call it in current module */

  PL_get_integer(a2, &atoms);
  PL_discard_foreign_frame(fid);

  return atoms;
}

```

Figure 7.4: Calling Prolog

atom_t **PL_module_name**(*module_t*)

Return the name of *module* as an atom.

module_t **PL_new_module**(*atom_t name*)

Find an existing or create a new module with name specified by the atom *name*.

7.6.9 Prolog exceptions in foreign code

This section discusses `PL_exception()`, `PL_throw()` and `PL_raise_exception()`, the interface functions to detect and generate Prolog exceptions from C-code. `PL_throw()` and `PL_raise_exception()` from the C-interface to raise an exception from foreign code. `PL_throw()` exploits the C-function `longjmp()` to return immediately to the innermost `PL_next_solution()`. `PL_raise_exception()` registers the exception term and returns `FALSE`. If a foreign predicate returns `FALSE`, while an exception-term is registered a Prolog exception will be raised by the virtual machine.

Calling these functions outside the context of a function implementing a foreign predicate results in undefined behaviour.

`PL_exception()` may be used after a call to `PL_next_solution()` fails, and returns a term reference to an exception term if an exception was raised, and 0 otherwise.

If a C-function, implementing a predicate calls Prolog and detects an exception using `PL_exception()`, it can handle this exception, or return with the exception. Some caution is required though. It is **not** allowed to call `PL_close_query()` or `PL_discard_foreign_frame()` afterwards, as this will invalidate the exception term. Below is the code that calls a Prolog defined arithmetic function (see `arithmetic_function/1`).

If `PL_next_solution()` succeeds, the result is analysed and translated to a number, after

which the query is closed and all Prolog data created after `PL_open_foreign_frame()` is destroyed. On the other hand, if `PL_next_solution()` fails and if an exception was raised, just pass it. Otherwise generate an exception (`PL_error()` is an internal call for building the standard error terms and calling `PL_raise_exception()`). After this, the Prolog environment should be discarded using `PL_cut_query()` and `PL_close_foreign_frame()` to avoid invalidating the exception term.

```
static int
prologFunction(ArithFunction f, term_t av, Number r)
{ int arity = f->proc->definition->functor->arity;
  fid_t fid = PL_open_foreign_frame();
  qid_t qid;
  int rval;

  qid = PL_open_query(NULL, PL_Q_NORMAL, f->proc, av);

  if ( PL_next_solution(qid) )
  { rval = valueExpression(av+arity-1, r);
    PL_close_query(qid);
    PL_discard_foreign_frame(fid);
  } else
  { term_t except;

    if ( (except = PL_exception(qid)) )
    { rval = PL_throw(except);          /* pass exception */
    } else
    { char *name = stringAtom(f->proc->definition->functor->name);

      /* generate exception */
      rval = PL_error(name, arity-1, NULL, ERR_FAILED, f->proc);
    }

    PL_cut_query(qid);                  /* donot destroy data */
    PL_close_foreign_frame(fid);        /* same */
  }

  return rval;
}
```

int PL_raise_exception(*term_t exception*)

Generate an exception (as `throw/1`) and return `FALSE`. Below is an example returning an exception from foreign predicate:

```
foreign_t
pl_hello(term_t to)
{ char *s;
```

```

if ( PL_get_atom_chars(to, &s) )
{ Sprintf("Hello \"%s\"\n", s);

  PL_succeed;
} else
{ term_t except = PL_new_term_ref();

  PL_unify_term(except,
                PL_FUNCTOR_CHARS, "type_error", 2,
                PL_CHARS, "atom",
                PL_TERM, to);

  return PL_raise_exception(except);
}
}

```

int **PL_throw**(*term_t exception*)

Similar to `PL_raise_exception()`, but returns using the C `longjmp()` function to the innermost `PL_next_solution()`.

term_t **PL_exception**(*qid_t qid*)

If `PL_next_solution()` fails, this can be due to normal failure of the Prolog call, or because an exception was raised using `throw/1`. This function returns a handle to the exception term if an exception was raised, or 0 if the Prolog goal simply failed.⁴

7.6.10 Catching Signals (Software Interrupts)

SWI-Prolog offers both a C and Prolog interface to deal with software interrupts (signals). The Prolog mapping is defined in section 4.10. This subsection deals with handling signals from C.

If a signal is not used by Prolog and the handler does not call Prolog in any way, the native signal interface routines may be used.

Some versions of SWI-Prolog, notably running on popular Unix platforms, handle `SIG_SEGV` for guarding the Prolog stacks. If the application wishes to handle this signal too, it should use `PL_signal()` to install its handler after initialising Prolog. SWI-Prolog will pass `SIG_SEGV` to the user code if it detected the signal is not related to a Prolog stack overflow.

Any handler that wishes to call one of the Prolog interface functions should call `PL_signal()` for its installation.

void (*)() **PL_signal**(*sig, func*)

This function is equivalent to the BSD-Unix `signal()` function, regardless of the platform used.

The signal handler is blocked while the signal routine is active, and automatically reactivated after the handler returns.

⁴This interface differs in two ways from Quintus. The calling predicates simply signal failure if an exception was raised, and a term referenced is returned, rather passed and filled with the error term. Exceptions can only be handled using the `PL_next_solution()` interface, as a handle to the query is required

After a signal handler is registered using this function, the native signal interface redirects the signal to a generic signal handler inside SWI-Prolog. This generic handler validates the environment, creates a suitable environment for calling the interface functions described in this chapter and finally calls the registered user-handler.

By default, signals are handled asynchronously (i.e. at the time they arrive). It is inherently dangerous to call extensive code fragments, and especially exception related code from asynchronous handlers. The interface allows for *synchronous* handling of signals. In this case the native OS handler just schedules the signal using `PL_raise()`, which is checked by `PL_handle_signals()` at the call- and redo-port. This behaviour is realised by or-ing *sig* with the constant `PL_SIGSYNC`.⁵

Signal handling routines may raise exceptions using `PL_raise_exception()`. The use of `PL_throw()` is not safe. If a synchronous handler raises an exception, the exception is delayed to the next call to `PL_handle_signals()`;

`int PL_handle_signals(void)`

Handle any signals pending from `PL_raise()`. `PL_handle_signals()` is called at each pass through the call- and redo-port at a safe point. Exceptions raised by the handler using `PL_raise_exception()` are properly passed to the environment.

The user may call this function inside long-running foreign functions to handle scheduled interrupts. This routine returns the number of signals handled. If a handler raises an exception, the return value is -1 and the calling routine should return with `FALSE` as soon as possible.

7.6.11 Miscellaneous

Term Comparison

`int PL_compare(term_t t1, term_t t2)`

Compares two terms using the standard order of terms and returns -1, 0 or 1. See also `compare/3`.

`int PL_same_compound(term_t t1, term_t t2)`

Yields `TRUE` if *t1* and *t2* refer to physically the same compound term and `FALSE` otherwise.

Recorded database

In some applications it is useful to store and retrieve Prolog terms from C-code. For example, the XPCE graphical environment does this for storing arbitrary Prolog data as slot-data of XPCE objects.

Please note that the returned handles have no meaning at the Prolog level and the recorded terms are not visible from Prolog. The functions `PL_recorded()` and `PL_erase()` are the only functions that can operate on the stored term.

Two groups of functions are provided. The first group (`PL_record()` and friends) store Prolog terms on the Prolog heap for retrieval during the same session. These functions are also used by `recorda/3` and friends. The recorded database may be used to communicate Prolog terms between threads.

⁵A better default would be to use synchronous handling, but this interface preserves backward compatibility.

`record_t` **PL_record**(*term_t* +*t*)

Record the term *t* into the Prolog database as `recorda/3` and return an opaque handle to the term. The returned handle remains valid until `PL_erase()` is called on it. `PL_recorded()` is used to copy recorded terms back to the Prolog stack.

`void` **PL_recorded**(*record_t* *record*, *term_t* -*t*)

Copy a recorded term back to the Prolog stack. The same record may be used to copy multiple instances at any time to the Prolog stack. See also `PL_record()` and `PL_erase()`.

`void` **PL_erase**(*record_t* *record*)

Remove the recorded term from the Prolog database, reclaiming all associated memory resources.

The second group (headed by `PL_record_external()`) provides the same functionality, but the returned data has properties that enable storing the data on an external device. It has been designed to make it possible to store Prolog terms fast and compact in an external database. Here are the main features:

- *Independent of session*
Records can be communicated to another Prolog session and made visible using `PL_recorded_external()`.
- *Binary*
The representation is binary for maximum performance. The returned data may contain 0-bytes.
- *Byte-order independent*
The representation can be transferred between machines with different byte-order.
- *No alignment restrictions*
There are no memory alignment restrictions and copies of the record can thus be moved freely. For example, it is possible to use this representation to exchange terms using shared memory between different Prolog processes.
- *Compact*
It is assumed that a smaller memory footprint will eventually outperform slightly faster representations.
- *Stable*
The format is designed for future enhancements without breaking compatibility with older records.

`char *` **PL_record_external**(*term_t* +*t*, *unsigned int* **len*)

Record the term *t* into the Prolog database as `recorda/3` and return an opaque handle to the term. The returned handle remains valid until `PL_erase()` is called on it.

It is allowed to copy the data and use `PL_recorded_external()` on the copy. The user is responsible for the memory management of the copy. After copying, the original may be discarded using `PL_erase_external()`.

`PL_recorded_external()` is used to copy such recorded terms back to the Prolog stack.

int **PL_recorded_external**(*const char *record, term_t -t*)

Copy a recorded term back to the Prolog stack. The same record may be used to copy multiple instances at any time to the Prolog stack. See also `PL_record_external()` and `PL_erase_external()`.

int **PL_erase_external**(*char *record*)

Remove the recorded term from the Prolog database, reclaiming all associated memory resources.

Getting file names

The function `PL_get_file_name()` provides access to Prolog filenames and its file-search mechanism described with `absolute_file_name/3`. Its existence is motivated to realise a uniform interface to deal with file-properties, search, naming conventions etc. from foreign code.

int **PL_get_file_name**(*term_t spec, char **name, int flags*)

Translate a Prolog term into a file name. The name is stored in the static buffer ring described with `PL_get_chars()` option `BUF_RING`. *flags* is a bit-mask controlling the conversion process. Options are:

`PL_FILE_ABSOLUTE`

Return an absolute path to the requested file.

`PL_FILE_OSPATH`

Return a the name using the hosting OS conventions. On MS-Windows, `\` is used to separate directories rather than the canonical `/`.

`PL_FILE_SEARCH`

Invoke `absolute_file_name/3`. This implies rules from `file_search_path/2` are used.

`PL_FILE_EXIST`

Demand the path to refer to an existing entity.

`PL_FILE_READ`

Demand read-access on the result.

`PL_FILE_WRITE`

Demand write-access on the result.

`PL_FILE_EXECUTE`

Demand execute-access on the result.

`PL_FILE_NOERRORS`

Do not raise any exceptions.

7.6.12 Errors and warnings

`PL_warning()` prints a standard Prolog warning message to the standard error (`user_error`) stream. Please note that new code should consider using `PL_raise_exception()` to raise a Prolog exception. See also section 4.9.

PL_ACTION_TRACE	Start Prolog tracer (<code>trace/0</code>). Requires no arguments.
PL_ACTION_DEBUG	Switch on Prolog debug mode (<code>debug/0</code>). Requires no arguments.
PL_ACTION_BACKTRACE	Print backtrace on current output stream. The argument (an int) is the number of frames printed.
PL_ACTION_HALT	Halt Prolog execution. This action should be called rather than Unix <code>exit()</code> to give Prolog the opportunity to clean up. This call does not return. The argument (an int) is the exit code. See <code>halt/1</code> .
PL_ACTION_ABORT	Generate a Prolog abort (<code>abort/0</code>). This call does not return. Requires no arguments.
PL_ACTION_BREAK	Create a standard Prolog break environment (<code>break/0</code>). Returns after the user types the end-of-file character. Requires no arguments.
PL_ACTION_GUIAPP	Win32: Used to indicate the kernel that the application is a GUI application if the argument is not 0 and a console application if the argument is 0. If a fatal error occurs, the system uses a windows messagebox to report this on a GUI application and simply prints the error and exits otherwise.
PL_ACTION_WRITE	Write the argument, a <code>char *</code> to the current output stream.
PL_ACTION_FLUSH	Flush the current output stream. Requires no arguments.
PL_ACTION_ATTACH_CONSOLE	Attach a console to a thread if it does not have one. See <code>attach_console/0</code> .

Table 7.1: `PL_action()` options

`int PL_warning(format, a1, ...)`

Print an error message starting with `'[WARNING: '`, followed by the output from `format`, followed by a `']` and a newline. Then start the tracer. `format` and the arguments are the same as for `printf(2)`. Always returns `FALSE`.

7.6.13 Environment Control from Foreign Code

`int PL_action(int, ...)`

Perform some action on the Prolog system. `int` describes the action, Remaining arguments depend on the requested action. The actions are listed in table 7.1.

7.6.14 Querying Prolog

`C_type PL_query(int)`

Obtain status information on the Prolog system. The actual argument type depends on the information required. `int` describes what information is wanted. The options are given in table 7.2.

PL_QUERY_ARGC	Return an integer holding the number of arguments given to Prolog from Unix.
PL_QUERY_ARGV	Return a char ** holding the argument vector given to Prolog from Unix.
PL_QUERY_SYMBOLFILE	Return a char * holding the current symbol file of the running process.
PL_MAX_INTEGER	Return a long, representing the maximal integer value represented by a Prolog integer.
PL_MIN_INTEGER	Return a long, representing the minimal integer value.
PL_QUERY_VERSION	Return a long, representing the version as $10,000 \times M + 100 \times m + p$, where M is the major, m the minor version number and p the patch-level. For example, 20717 means 2.7.17.

Table 7.2: PL_query() options

7.6.15 Registering Foreign Predicates

int **PL_register_foreign**(const char *name, int arity, foreign_t (*function)(), int flags)

Register a C-function to implement a Prolog predicate. After this call returns successfully a predicate with name *name* (a char *) and arity *arity* (a C int) is created. As a special case, *name* may consist of a sequence of alpha-numerical characters followed by the colon (:). In this case the name until the colon is taken to be the destination module and the rest of the name the predicate name.

When called in Prolog, Prolog will call *function*. *flags* forms bitwise or'ed list of options for the installation. These are:

PL_FA_NOTRACE	Predicate cannot be seen in the tracer
PL_FA_TRANSPARENT	Predicate is module transparent
PL_FA_NONDETERMINISTIC	Predicate is non-deterministic. See also PL_retry().
PL_FA_VARARGS	Use alternative calling convention.

void **PL_load_extensions**(PL_extension *e)

Register foreign predicates from a table of structures. This is an alternative to multiple calls to PL_register_foreign() and simplifies code that wishes to use PL_register_extensions() as an alternative. The type PL_extension is defined as:

```
typedef struct _PL_extension
{ char      *predicate_name; /* Name of the predicate */
  short     arity;          /* Arity of the predicate */
  pl_function_t function;   /* Implementing functions */
  short     flags;         /* Or of PL_FA_... */
} PL_extension;
```

void **PL_register_extensions**(PL_extension *e)

The function PL_register_extensions() behaves as PL_load_extensions(), but is the only PL_* function that may be called **before** PL_initialise(). The predicates are

registered **into the module** user after registration of the SWI-Prolog builtin foreign predicates and before loading the initial saved state. This implies that `initialization/1` directives can refer to them.

Here is an example of its usage:

```
static PL_extension predicates[] = {
  { "foo",      1,      pl_foo, 0 },
  { "bar",      2,      pl_bar, PL_FA_NONDETERMINISTIC },
  { NULL,       0,      NULL,   0 }
};

main(int argc, char **argv)
{ PL_register_extensions(predicates);

  if ( !PL_initialise(argc, argv) )
    PL_halt(1);

  ...
}
```

7.6.16 Foreign Code Hooks

For various specific applications some hooks are provided.

`PL_dispatch_hook_t` **PL_dispatch_hook**(*PL_dispatch_hook_t*)

If this hook is not NULL, this function is called when reading from the terminal. It is supposed to dispatch events when SWI-Prolog is connected to a window environment. It can return two values: `PL_DISPATCH_INPUT` indicates Prolog input is available on file descriptor 0 or `PL_DISPATCH_TIMEOUT` to indicate a timeout. The old hook is returned. The type `PL_dispatch_hook_t` is defined as:

```
typedef int (*PL_dispatch_hook_t)(void);
```

`void` **PL_abort_hook**(*PL_abort_hook_t*)

Install a hook when `abort/0` is executed. SWI-Prolog `abort/0` is implemented using `setjmp()/longjmp()` construct. The hooks are executed in the reverse order of their registration after the `longjmp()` took place and before the Prolog toplevel is reinvoked. The type `PL_abort_hook_t` is defined as:

```
typedef void (*PL_abort_hook_t)(void);
```

`int` **PL_abort_unhook**(*PL_abort_hook_t*)

Remove a hook installed with `PL_abort_hook()`. Returns `FALSE` if no such hook is found, `TRUE` otherwise.

`void PL_on_halt(void (*)(int, void *), void *closure)`

Register the function f to be called if SWI-Prolog is halted. The function is called with two arguments: the exit code of the process (0 if this cannot be determined on your operating system) and the *closure* argument passed to the `PL_on_halt()` call. See also `at_halt/1`.

`PL_agc_hook_t PL_agc_hook(PL_agc_hook_t new)`

Register a hook with the atom-garbage collector (see `garbage_collect_atoms/0` that is called on any atom that is reclaimed. The old hook is returned. If no hook is currently defined, `NULL` is returned. The argument of the called hook is the atom that is to be garbage collected. The return value is an `int`. If the return value is zero, the atom is **not** reclaimed. The hook may invoke any Prolog predicate.

The example below defines a foreign library for printing the garbage collected atoms for debugging purposes.

```
#include <SWI-Stream.h>
#include <SWI-Prolog.h>

static int
atom_hook(atom_t a)
{ Sdprintf("AGC: deleting %s\n", PL_atom_chars(a));

  return TRUE;
}

static PL_agc_hook_t old;

install_t
install()
{ old = PL_agc_hook(atom_hook);
}

install_t
uninstall()
{ PL_agc_hook(old);
}
```

7.6.17 Storing foreign data

This section provides some hints for handling foreign data in Prolog. With foreign data, we refer to data that is used by foreign language predicates and needs to be passed around in Prolog. Excluding combinations, there are three principal options for storing such data

- *Natural Prolog data*
E.i. using the representation one would choose if there was no foreign interface required.
- *Opaque packed Prolog data*
Data can also be represented in a foreign structure and stored on the Prolog stacks using

`PL_put_string_nchars()` and retrieved using `PL_get_string_chars()`. It is generally good practice to wrap the string in a compound term with arity 1, so Prolog can identify the type. `portray/1` rules may be used to streamline printing such terms during development.

- *Natural foreign data, passing a pointer*

An alternative is to pass a pointer to the foreign data. Again, this functor may be wrapped in a compound term.

The choice may be guided using the following distinctions

- *Is the data opaque to Prolog*

With ‘opaque’ data, we refer to data handled in foreign functions, passed around in Prolog, but of which Prolog never examines the contents of the data itself. If the data is opaque to Prolog, the chosen representation does not depend on simple analysis by Prolog, and the selection will be driven solely by simplicity of the interface and performance (both in time and space).

- *How big is the data*

Is efficient encoding required? For example, a boolean array may be expressed as a compound term, holding integers each of which contains a number of bits, or as a list of `true` and `false`.

- *What is the nature of the data*

For examples in C, constants are often expressed using ‘enum’ or #define’d integer values. If Prolog needs to handle this data, atoms are a more logical choice. Whether or not this mapping is used depends on whether Prolog needs to interpret the data, how important debugging is and how important performance is.

- *What is the lifetime of the data*

We can distinguish three cases.

1. The lifetime is dictated by the accessibility of the data on the Prolog stacks. There is no way by which the foreign code when the data becomes ‘garbage’, and the data thus needs to be represented on the Prolog stacks using Prolog data-types. (2),
2. The data lives on the ‘heap’ and is explicitly allocated and deallocated. In this case, representing the data using native foreign representation and passing a pointer to it is a sensible choice.
3. The data lives as during the lifetime of a foreign predicate. If the predicate is deterministic, foreign automatic variables are suitable. If the predicate is non-deterministic, the data may be allocated using `malloc()` and a pointer may be passed. See section 7.6.1.

Examples for storing foreign data

In this section, we will outline some examples, covering typical cases. In the first example, we will deal with extending Prolog’s data representation with integer-sets, represented as bit-vectors. In the second example, we look at handling a ‘netmask’. Finally, we discuss the outline of the DDE interface.

Integer sets with not-to-far-apart upper- and lower-bounds can be represented using bit-vectors. Common set operations, such as union, intersection, etc. are reduced to simple and’ing and or’ing the bitvectors. This can be done in Prolog, using a compound term holding integer arguments. Especially

if the integers are kept below the maximum tagged integer value (see `current_prolog_flag/2`), this representation is fairly space-efficient (wasting 1 word for the functor and 7 bits per integer for the tags). Arithmetic can all be performed in Prolog too.

For really demanding applications, foreign representation will perform better, especially time-wise. Bit-vectors are naturally expressed using string objects. If the string is wrapped in `bitvector/1`, lower-bound of the vector is 0, and the upperbound is not defined, an implementation for getting and putting the sets as well as the union predicate for it is below.

```
#include <SWI-Prolog.h>

#define max(a, b) ((a) > (b) ? (a) : (b))
#define min(a, b) ((a) < (b) ? (a) : (b))

static functor_t FUNCTOR_bitvector1;

static int
get_bitvector(term_t in, int *len, unsigned char **data)
{ if ( PL_is_functor(in, FUNCTOR_bitvector1) )
  { term_t a = PL_new_term_ref();

    PL_get_arg(1, in, a);
    return PL_get_string(a, (char **)data, len);
  }

  PL_fail;
}

static int
unify_bitvector(term_t out, int len, const unsigned char *data)
{ if ( PL_unify_functor(out, FUNCTOR_bitvector1) )
  { term_t a = PL_new_term_ref();

    PL_get_arg(1, out, a);

    return PL_unify_string_nchars(a, len, (const char *)data);
  }

  PL_fail;
}

static foreign_t
pl_bitvector_union(term_t t1, term_t t2, term_t u)
{ unsigned char *s1, *s2;
  int l1, l2;

  if ( get_bitvector(t1, &l1, &s1) &&
        get_bitvector(t2, &l2, &s2) )
```

```

{ int l = max(l1, l2);
  unsigned char *s3 = alloca(l);

  if ( s3 )
  { int n;
    int ml = min(l1, l2);

    for(n=0; n<ml; n++)
      s3[n] = s1[n] | s2[n];
    for( ; n < l1; n++)
      s3[n] = s1[n];
    for( ; n < l2; n++)
      s3[n] = s2[n];

    return unify_bitvector(u, l, s3);
  }

  return PL_warning("Not enough memory");
}

PL_fail;
}

install_t
install()
{ PL_register_foreign("bitvector_union", 3, pl_bitvector_union, 0);

  FUNCTOR_bitvector1 = PL_new_functor(PL_new_atom("bitvector"), 1);
}

```

Netmask's are used with TCP/IP configuration. Suppose we have an application dealing with reasoning about a network configuration. Such an application requires communicating netmask structures from the operating system, reasoning about them and possibly communicate them to the user. A netmask consists of 4 bitmasks between 0 and 255. C-application normally see them as an 4-byte wide unsigned integer. SWI-Prolog cannot do that, as integers are always signed.

We could use the string approach outlined above, but this makes it hard to handle these terms in Prolog. A better choice is a compound term `netmask/4`, holding the 4 submasks as integer arguments.

As the implementation is trivial, we will omit this here.

The DDE interface (see section 4.43) represents another common usage of the foreign interface: providing communication to new operating system features. The DDE interface requires knowledge about active DDE server and client channels. These channels contains various foreign data-types. Such an interface is normally achieved using an open/close protocol that creates and destroys a *handle*. The handle is a reference to a foreign data-structure containing the relevant information.

There are a couple of possibilities for representing the handle. The choice depends on responsibilities and debugging facilities. The simplest approach is to use `PL_unify_pointer()` and `PL_get_pointer()`. This approach is fast and easy, but has the drawbacks of (untyped) pointers: there is no reliable way to detect the validity of the pointer, not to verify it is pointing to a structure of the desired type. The pointer may be wrapped into a compound term with arity 1 (i.e., `dde_channel(<Pointer>)`), making the type-problem less serious.

Alternatively (used in the DDE interface), the interface code can maintain a (preferably variable length) array of pointers and return the index in this array. This provides better protection. Especially for debugging purposes, wrapping the handle in a compound is a good suggestion.

7.6.18 Embedding SWI-Prolog in other applications

With embedded Prolog we refer to the situation where the ‘main’ program is not the Prolog application. Prolog is sometimes embedded in C, C++, Java or other languages to provide logic based services in a larger application. Embedding loads the Prolog engine as a library to the external language. Prolog itself only provides for embedding in the C-language (compatible to C++). Embedding in Java is achieved using JPL using a C-glu between the Java and Prolog C-interfaces.

The most simple embedded program is below. The interface function `PL_initialise()` **must** be called before any of the other SWI-Prolog foreign language functions described in this chapter, except for `PL_initialise_hook()`, `PL_new_atom()` and `PL_register_foreign()`. `PL_initialise()` interprets all the command-line arguments, except for the `-t toplevel` flag that is interpreted by `PL_toplevel()`.

```
int
main(int argc, char **argv)
{
#ifdef READLINE /* Remove if you don't want readline */
    PL_initialise_hook(install_readline);
#endif

    if ( !PL_initialise(argc, argv) )
        PL_halt(1);

    PL_halt(PL_toplevel() ? 0 : 1);
}
```

```
int PL_initialise(int argc, char **argv)
```

Initialises the SWI-Prolog heap and stacks, restores the Prolog state, loads the system and personal initialisation files, runs the `at_initialization/1` hooks and finally runs the `-g goal` hook.

Special consideration is required for `argv[0]`. On **Unix**, this argument passes the part of the commandline that is used to locate the executable. Prolog uses this to find the file holding the running executable. The **Windows** version uses this to find a *module* of the running executable. If the specified module cannot be found, it tries the module `libpl.dll`, containing the Prolog runtime kernel. In all these cases, the resulting file is used for two purposes

- See whether a Prolog saved-state is appended to the file. If this is the case, this state will be loaded instead of the default `boot.prc` file from the SWI-Prolog home directory. See also `qsave_program/[1,2]` and section 7.7.
- Find the Prolog home directory. This process is described in detail in section 7.8.

`PL_initialise()` returns 1 if all initialisation succeeded and 0 otherwise.⁶

In most cases, `argc` and `argv` will be passed from the main program. It is allowed to create your own argument vector, provided `argv[0]` is constructed according to the rules above. For example:

```
int
main(int argc, char **argv)
{ char *av[10];
  int ac = 0;

  av[ac++] = argv[0];
  av[ac++] = "-x";
  av[ac++] = "mystate";
  av[ac] = NULL;

  if ( !PL_initialise(ac, av) )
    PL_halt(1);
  ...
}
```

Please note that the passed argument vector may be referred from Prolog at any time and should therefore be valid as long as the Prolog engine is used.

A good setup in Windows is to add SWI-Prolog's `bin` directory to your `PATH` and either pass a module holding a saved-state, or `"libpl.dll"` as `argv[0]`. If the Prolog state is attached to a DLL (see the `-dll` option of `plld`, pass the name of this DLL.

`int PL_is_initialised(int *argc, char ***argv)`

Test whether the Prolog engine is already initialised. Returns `FALSE` if Prolog is not initialised and `TRUE` otherwise. If the engine is initialised and `argc` is not `NULL`, the argument count used with `PL_initialise()` is stored in `argc`. Same for the argument vector `argv`.

`void PL_install_readline()`

Installs the GNU-readline line-editor. Embedded applications that do not use the Prolog toplevel should normally delete this line, shrinking the Prolog kernel significantly. Note that the Windows version does not use GNU readline.

`int PL_toplevel()`

Runs the goal of the `-t toplevel` switch (default `prolog/0`) and returns 1 if successful, 0 otherwise.

⁶BUG: Various fatal errors may cause `PL_initialise` to call `PL_halt(1)`, preventing it from returning at all.

`void PL_cleanup(int status)`

This function performs the reverse of `PL_initialise()`. It runs the `PL_on_halt()` and `at_halt/1` handlers, closes all streams (except for the ‘standard I/O’ streams which are flushed only), deallocates all memory and restores all signal handlers. The *status* argument is passed to the various termination hooks and indicates the *exit-status*.

This function allows deleting and restarting the Prolog system in the same process. Use it with care, as `PL_initialise()` is a costly function. Unix users should consider using `exec()` (available as part of the `clib` package.).

`int PL_halt(int status)`

Cleanup the Prolog environment using `PL_cleanup()` and calls `exit()` with the status argument. As `PL_cleanup()` can only be called from the main thread, this function returns `FALSE` when called from another thread as the main one.⁷

Threading, Signals and embedded Prolog

This section applies to Unix-based environments that have signals or multi-threading. The Windows version is compiled for multi-threading and Windows lacks proper signals.

We can distinguish two classes of embedded executables. There are small C/C++-programs that act as an interfacing layer around Prolog. Most of these programs can be replaced using the normal Prolog executable extended with a dynamically loaded foreign extension and in most cases this is the preferred route. In other cases, Prolog is embedded in a complex application that—like Prolog—wants to control the process environment. A good example is Java. Embedding Prolog is generally the only way to get these environments together in one process image. Java applications however are by nature multi-threaded and appear to do signal-handling (software interrupts).

To make Prolog operate smoothly in such environments it must be told not to alter the process environment. This is partly done at build-time and partly execution time. At build-time we must specify the use of software stack-overflow rather than the default hardware checks. This is done using

```
sh configure --disable-segv-handling
```

The resulting Prolog executable is about 10% slower than the normal executable, but behaves much more reliable in complicated embedded situations. In addition, as the process no longer handles segmentation violations, debugging foreign code linked to it is much easier.

At runtime, it is advised to pass the flag `-nosignals`, which inhibits all default signal handling. This has a few consequences though:

- It is no longer possible to break into the tracer using an interrupt signal (Control-C).
- `SIGPIPE` is normally set to be ignored. Prolog uses return-codes to diagnose broken pipes. Depending on the situation one should take appropriate action if Prolog streams are connected to pipes.
- Fatal errors normally cause Prolog to call `PL_cleanup()` and `exit()`. It is advised to call `PL_cleanup()` as part of the exit-procedure of your application.

⁷BUG: Eventually it may become possible to call `PL_halt()` from any thread.

7.7 Linking embedded applications using plld

The utility program `plld` (Win32: `plld.exe`) may be used to link a combination of C-files and Prolog files into a stand-alone executable. `plld` automates most of what is described in the previous sections.

In the normal usage, a copy is made of the default embedding template `\ldots\pl\include\stub.c`. The `main()` routine is modified to suit your application. `PLinitialise()` **must** be passed the program-name (`argv[0]`) (Win32: the executing program can be obtained using `GetModuleFileName()`). The other elements of the command-line may be modified. Next, `plld` is typically invoked as:

```
plld -o output stubfile.c [other-c-or-o-files] [plfiles]
```

`plld` will first split the options into various groups for both the C-compiler and the Prolog compiler. Next, it will add various default options to the C-compiler and call it to create an executable holding the user's C-code and the Prolog kernel. Then, it will call the SWI-Prolog compiler to create a saved state from the provided Prolog files and finally, it will attach this saved state to the created emulator to create the requested executable.

Below, it is described how the options are split and which additional options are passed.

-help

Print brief synopsis.

-pl *prolog*

Select the prolog to use. This prolog is used for two purposes: get the home-directory as well as the compiler/linker options and create a saved state of the Prolog code.

-ld *linker*

Linker used to link the raw executable. Default is to use the C-compiler (Win32: `link.exe`).

-cc *C-compiler*

Compiler for `.c` files found on the commandline. Default is the compiler used to build SWI-Prolog (see `current_prolog_flag/2`) (Win32: `cl.exe`).

-c++ *C++-compiler*

Compiler for C++ sources (extensions `.cpp`, `.cxx`, `.cc` or `.C`) files found on the commandline. Default is `c++` or `g++` if the C-compiler is `gcc`) (Win32: `cl.exe`).

-nostate

Just relink the kernel, do not add any Prolog code to the new kernel. This is used to create a new kernel holding additional foreign predicates on machines that do not support the shared-library (DLL) interface, or if building the state cannot be handled by the default procedure used by `plld`. In the latter case the state is created separately and appended to the kernel using `cat <kernel> <state> > <out>` (Win32: `copy /b <kernel>+<state> <out>`)

-shared

Link C, C++ or object files into a shared object (DLL) that can be loaded by the `load_foreign_library/1` predicate. If used with `-c` it sets the proper options to compile a C or C++ file ready for linking into a shared object

-dll

Windows only. Embed SWI-Prolog into a DLL rather than an executable.

-c

Compile C or C++ source-files into object files. This turns `plld` into a replacement for the C or C++ compiler where proper options such as the location of the include directory are passed automatically to the compiler.

-E

Invoke the C preprocessor. Used to make `plld` a replacement for the C or C++ compiler.

-pl-options ,...

Additional options passed to Prolog when creating the saved state. The first character immediately following `pl-options` is used as separator and translated to spaces when the argument is built. Example: `-pl-options , -F , xpce` passed `-F xpce` as additional flags to Prolog.

-ld-options ,...

Passes options to the linker, similar to `-pl-options`.

-cc-options ,...

Passes options to the C/C++ compiler, similar to `-pl-options`.

-v

Select verbose operation, showing the various programs and their options.

-o *outfile*

Reserved to specify the final output file.

-library

Specifies a library for the C-compiler. By default, `-lpl` (Win32: `libpl.lib`) and the libraries needed by the Prolog kernel are given.

-Llibrary-directory

Specifies a library directory for the C-compiler. By default the directory containing the Prolog C-library for the current architecture is passed.

-g | -Iinclude-directory | -Ddefinition

These options are passed to the C-compiler. By default, the include directory containing `SWI-Prolog.h` is passed. `plld` adds two additional `* -Ddef` flags:

-D__SWI_PROLOG__

Indicates the code is to be connected to SWI-Prolog.

-D__SWI_EMBEDDED__

Indicates the creation of an embedded program.

***.o | *.c | *.C | *.cxx | *.cpp**

Passed as input files to the C-compiler

***.pl | *.qlf**

Passed as input files to the Prolog compiler to create the saved-state.

I.e. all other options. These are passed as linker options to the C-compiler.

7.7.1 A simple example

The following is a very simple example going through all the steps outlined above. It provides an arithmetic expression evaluator. We will call the application `calc` and define it in the files `calc.c` and `calc.pl`. The Prolog file is simple:

```
calc(Atom) :-
    term_to_atom(Expr, Atom),
    A is Expr,
    write(A),
    nl.
```

The C-part of the application parses the command-line options, initialises the Prolog engine, locates the `calc/1` predicate and calls it. The coder is in figure 7.5.

The application is now created using the following command-line:

```
% plld -o calc calc.c calc.pl
```

The following indicates the usage of the application:

```
% calc pi/2
1.5708
```

7.8 The Prolog ‘home’ directory

Executables embedding SWI-Prolog should be able to find the ‘home’ directory of the development environment unless a self-contained saved-state has been added to the executable (see `qsave_program/[1,2]` and section 7.7).

If Prolog starts up, it will try to locate the development environment. To do so, it will try the following steps until one succeeds.

1. If the environment variable `SWI_HOME_DIR` is defined and points to an existing directory, use this.
2. If the environment variable `SWIPL` is defined and points to an existing directory, use this.
3. Locate the primary executable or (Windows only) a component (*module*) thereof and check whether the parent directory of the directory holding this file contains the file `swipl`. If so, this file contains the (relative) path to the home directory. If this directory exists, use this. This is the normal mechanism used by the binary distribution.
4. If the precompiled path exists, use it. This is only useful for a source installation.

If all fails and there is no state attached to the executable or provided Windows module (see `PL_initialise()`), SWI-Prolog gives up. If a state is attached, the current working directory is used.

The `file_search_path/2` alias `swi` is set to point to the home directory located.

```

#include <stdio.h>
#include <SWI-Prolog.h>

#define MAXLINE 1024

int
main(int argc, char **argv)
{ char expression[MAXLINE];
  char *e = expression;
  char *program = argv[0];
  char *plav[2];
  int n;

  /* combine all the arguments in a single string */

  for(n=1; n<argc; n++)
  { if ( n != 1 )
    *e++ = ' ';
    strcpy(e, argv[n]);
    e += strlen(e);
  }

  /* make the argument vector for Prolog */

  plav[0] = program;
  plav[1] = NULL;

  /* initialise Prolog */

  if ( !PL_initialise(1, plav) )
    PL_halt(1);

  /* Lookup calc/1 and make the arguments and call */

  { predicate_t pred = PL_predicate("calc", 1, "user");
    term_t h0 = PL_new_term_refs(1);
    int rval;

    PL_put_atom_chars(h0, expression);
    rval = PL_call_predicate(NULL, PL_Q_NORMAL, pred, h0);

    PL_halt(rval ? 0 : 1);
  }

  return 0;
}

```

Figure 7.5: C-source for the calc application

7.9 Example of Using the Foreign Interface

Below is an example showing all stages of the declaration of a foreign predicate that transforms atoms possibly holding uppercase letters into an atom only holding lower case letters. Figure 7.6 shows the C-source file, figure 7.7 illustrates compiling and loading of foreign code.

```
/* Include file depends on local installation */
#include <SWI-Prolog.h>
#include <stdlib.h>
#include <ctype.h>

foreign_t
pl_lowercase(term_t u, term_t l)
{ char *copy;
  char *s, *q;
  int rval;

  if ( !PL_get_atom_chars(u, &s) )
    return PL_warning("lowercase/2: instantiation fault");
  copy = malloc(strlen(s)+1);

  for( q=copy; *s; q++, s++)
    *q = (isupper(*s) ? tolower(*s) : *s);
  *q = '\0';

  rval = PL_unify_atom_chars(l, copy);
  free(copy);

  return rval;
}

install_t
install()
{ PL_register_foreign("lowercase", 2, pl_lowercase, 0);
}
```

Figure 7.6: Lowercase source file

```
% gcc -I/usr/local/lib/pl-\plversion/include -fpic -c lowercase.c
% gcc -shared -o lowercase.so lowercase.o
% pl
Welcome to SWI-Prolog (Version \plversion)
Copyright (c) 1993-1996 University of Amsterdam. All rights reserved.

For help, use ?- help(Topic). or ?- apropos(Word).

1 ?- load_foreign_library(lowercase).

Yes
2 ?- lowercase('Hello World!', L).

L = 'hello world!'

Yes
```

Figure 7.7: Compiling the C-source and loading the object file

7.10 Notes on Using Foreign Code

7.10.1 Memory Allocation

SWI-Prolog's heap memory allocation is based on the `malloc(3)` library routines. The stacks are allocated using `mmap()` on most Unix machines and using `VirtualAlloc()` on windows. SWI-Prolog provides the functions below as a wrapper around `malloc()`. Allocation errors in these functions trap SWI-Prolog's fatal-error handler, in which case `PL_malloc()` or `PL_realloc()` do not return.

Portable applications must use `PL_free()` to release strings returned by `PL_get_chars()` using the `BUF_MALLOC` argument. Portable applications may use both `PL_malloc()` and friends or `malloc()` and friends but should not mix these two sets of functions on the same memory.⁸

```
void * PL_malloc(size_t bytes)
```

Allocate *bytes* of memory. On failure SWI-Prolog's fatal error handler is called and `PL_malloc()` does not return. Memory allocated using these functions must use `PL_realloc()` and `PL_free()` rather than `realloc()` and `free()`.

```
void * PL_realloc(void *mem, size_t size)
```

Change the size of the allocated chunk, possibly moving it. The *mem* argument must be obtained from a previous `PL_malloc()` or `PL_realloc()` call.

```
void PL_free(void *mem)
```

Release memory. The *mem* argument must be obtained from a previous `PL_malloc()` or `PL_realloc()` call.

7.10.2 Compatibility between Prolog versions

Great care is taken to ensure binary compatibility of foreign extensions between different Prolog versions. Only much less frequently used stream interface has been responsible for binary incompatibilities.

Source-code that relies on new features of the foreign interface can use the macro `PLVERSION` to find the version of `SWI-Prolog.h` and `PL_query()` using the option `PL_QUERY_VERSION` to find the version of the attached Prolog system. Both follow the same numbering schema explained with `PL_query()`.

7.10.3 Debugging Foreign Code

Statically linked foreign code or embedded systems can be debugged normally. Most modern environments provide debugging tools for dynamically loaded shared objects or dynamic load libraries. The following example traces the code of `lowercase` using `gdb(1)` in a Unix environment.

```
% gcc -I/usr/local/lib/pl-2.2.0/include -fpic -c -g lowercase.c
% gcc -shared -o lowercase.so lowercase.o
% gdb pl
(gdb) r
Welcome to SWI-Prolog (Version \plversion)
Copyright (c) 1993-1996 University of Amsterdam. All rights reserved.
```

⁸These functions were introduced in SWI-Prolog 5.0.9 to realise guaranteed portability. Foreign code that must be compatible with older versions can check the `PLVERSION` macro.

For help, use `?- help(Topic).` or `?- apropos(Word).`

```
?- load_foreign_library(lowercase).
<type Control-C>
(gdb) shared                % loads symbols for shared objects
(gdb) break pl_lowercase
(gdb) continue
?- lowercase('HELLO', X).
```

7.10.4 Name Conflicts in C modules

In the current version of the system all public C functions of SWI-Prolog are in the symbol table. This can lead to name clashes with foreign code. Someday I should write a program to strip all these symbols from the symbol table (why does Unix not have that?). For now I can only suggest to give your function another name. You can do this using the C preprocessor. If—for example—your foreign package uses a function `warning()`, which happens to exist in SWI-Prolog as well, the following macro should fix the problem.

```
#define warning warning_
```

Note that shared libraries do not have this problem as the shared library loader will only look for symbols in the main executable for symbols that are not defined in the library itself.

7.10.5 Compatibility of the Foreign Interface

The term-reference mechanism was first used by Quintus Prolog version 3. SICStus Prolog version 3 is strongly based on the Quintus interface. The described SWI-Prolog interface is similar to using the Quintus or SICStus interfaces, defining all foreign-predicate arguments of type `+term`. SWI-Prolog explicitly uses type `functor_t`, while Quintus and SICStus uses `<name>` and `<arity>`. As the names of the functions differ from Prolog to Prolog, a simple macro layer dealing with the names can also deal with this detail. For example:

```
#define QP_put_functor(t, n, a) PL_put_functor(t, PL_new_functor(n, a))
```

The `PL_unify_*()` functions are lacking from the Quintus and SICStus interface. They can easily be emulated or the `put/unify` approach should be used to write compatible code.

The `PL_open_foreign_frame()/PL_close_foreign_frame()` combination is lacking from both other Prologs. SICStus has `PL_new_term_refs(0)`, followed by `PL_reset_term_refs()` that allows for discarding term references.

The Prolog interface for the graphical user interface package XPCP shares about 90% of the code using a simple macro layer to deal with different naming and calling conventions of the interfaces.

Generating Runtime Applications

8

This chapter describes the features of SWI-Prolog for delivering applications that can run without the development version of the system installed.

A SWI-Prolog runtime executable is a file consisting of two parts. The first part is the *emulator*, which is machine dependent. The second part is the *resource archive*, which contains the compiled program in a machine-independent format, startup options and possibly user-defined *resources*, see `resource/3` and `open_resource/3`.

These two parts can be connected in various different ways. The most common way for distributed runtime applications is to *concatenate* the two parts. This can be achieved using external commands (Unix: `cat`, Windows: `copy`), or using the `stand_alone` option to `qsave_program/2`. The second option is to attach a startup script in front of the resource that starts the emulator with the proper options. This is the default under Unix. Finally, an emulator can be told to use a specified resource file using the `-x` commandline switch.

qsave_program(+File, +ListOfOptions)

Saves the current state of the program to the file *File*. The result is a resource archive containing a saved-state that expresses all Prolog data from the running program and all user-defined resources. Depending on the `stand_alone` option, the resource is headed by the emulator, a Unix shell-script or nothing.

ListOfOptions is a list of `<Key> = <Value>` or `<Key>(<Value>)` pairs. The available keys are described in table 8.1.

Before writing the data to file, `qsave_program/2` will run `autoload/0` to all required autoloading the system can discover. See `autoload/0`.

Provided the application does not require any of the Prolog libraries to be loaded at runtime, the only file from the SWI-Prolog development environment required is the emulator itself. The emulator may be built in two flavours. The default is the *development emulator*. The *runtime emulator* is similar, but lacks the tracer.

If the option `stand_alone(on)` is present, the emulator is the first part of the state. If the emulator is started it will test whether a boot-file (state) is attached to the emulator itself and load this state. Provided the application has all libraries loaded, the resulting executable is completely independent of the runtime environment or location where it was build.

See also section 6.

qsave_program(+File)

Equivalent to `qsave_program(File, [])`.

autoload

Check the current Prolog program for predicates that are referred to, are undefined and have a definition in the Prolog library. Load the appropriate libraries.

Key	Option	Type	Description
local	-L	K-bytes	Size (Limit) of local stack
global	-G	K-bytes	Size (Limit) of global stack
trail	-T	K-bytes	Size (Limit) of trail stack
argument	-A	K-bytes	Size (Limit) of argument stack
goal	-g	atom	Initialisation goal
toplevel	-t	atom	Prolog toplevel goal
init_file	-f	atom	Personal initialisation file
class		atom	If runtime, only read resources from the state (default). If kernel, lock all predicates as system predicates. If development, save the predicates in their current state and keep reading resources from their source (if present). See also <code>resource/3</code> .
autoload		bool	If true, run <code>autoload/0</code> first
map		file	File to write info on dump
op		save/standard	Save operator declarations?
stand_alone		bool	Include the emulator in the state
emulator		file	Emulator attached to the (stand-alone) executable. Default is the running emulator.

Table 8.1: $\langle Key \rangle = \langle Value \rangle$ pairs for `qsave_program/2`

This predicate is used by `qsave_program/[1,2]` to ensure the saved state will not depend on one of the libraries. The predicate `autoload/0` will find all **direct** references to predicates. It does not find predicates referenced via meta-predicates. The predicate `log/2` is defined in the library(`quintus`) to provide a `quintus` compatible means to compute the natural logarithm of a number. The following program will behave correctly if its state is executed in an environment where the library(`quintus`) is not available:

```
logtable(From, To) :-
    From > To, !.
logtable(From, To) :-
    log(From, Value),
    format('~d~t~8|~2f~n', [From, Value]),
    F is From + 1,
    logtable(F, To).
```

However, the following implementation refers to `log/2` through the meta-predicate `maplist/3`. `Autoload` will not be able to find the reference. This problem may be fixed either by loading the module `library(quintus)` explicitly or use `require/1` to tell the system that the predicate `log/2` is required by this module.

```
logtable(From, To) :-
    findall(X, between(From, To, X), Xlist),
```

```

maplist(log, Xlist, SineList),
write_table(Xlist, SineList).

write_table([], []).
write_table([I|IT], [V|VT]) :-
    format('~d~t~8|~2f~n', [I, V]),
    write_table(IT, VT).

```

volatile *+Name/Arity, ...*

Declare that the clauses of specified predicates should **not** be saved to the program. The volatile declaration is normally used to avoid that the clauses of dynamic predicates that represent data for the current session is saved in the state file.

8.1 Limitations of qsave_program

There are three areas that require special attention when using `qsave_program/[1,2]`.

- If the program is an embedded Prolog application or uses the foreign language interface, care has to be taken to restore the appropriate foreign context. See section 8.2 for details.
- If the program uses directives (`:- goal.` lines) that perform other actions than setting predicate attributes (dynamic, volatile, etc.) or loading files (consult, etc.), the directive may need to be prefixed with `initialization/1`.
- Database references as returned by `clause/3`, `recorded/3`, etc. are not preserved and may thus not be part of the database when saved.

8.2 Runtimes and Foreign Code

Some applications may need to use the foreign language interface. Object code is by definition machine-dependent and thus cannot be part of the saved program file.

To complicate the matter even further there are various ways of loading foreign code:

- *Using the library(shlib) predicates*
This is the preferred way of dealing with foreign code. It loads quickly and ensures an acceptable level of independence between the versions of the emulator and the foreign code loaded. It works on Unix machines supporting shared libraries and library functions to load them. Most modern Unixes, as well as Win32 (Windows 95/NT) satisfy this constraint.
- *Static linking*
This mechanism works on all machines, but generally requires the same C-compiler and linker to be used for the external code as is used to build SWI-Prolog itself.

To make a runtime executable that can run on multiple platforms one must make runtime checks to find the correct way of linking. Suppose we have a source-file `myextension` defining the installation function `install()`.

If this file is compiled into a shared library, `load_foreign_library/1` will load this library and call the installation function to initialise the foreign code. If it is loaded as a static extension, define `install()` as the predicate `install/0`:

```

static foreign_t
pl_install()
{ install();

  PL_succeed;
}

PL_extension PL_extensions [] =
{
/*{ "name",      arity,  function,      PL_FA_<flags> },*/

  { "install",  0,      pl_install,    0 },
  { NULL,      0,      NULL,          0 }      /* terminat-
ing line */
};

```

Now, use the following Prolog code to load the foreign library:

```

load_foreign_extensions :-
    current_predicate(install, install), !, % static loaded
    install.
load_foreign_extensions :- % shared library
    load_foreign_library(foreign(myextension)).

:- initialization load_foreign_extensions.

```

The path alias `foreign` is defined by `file_search_path/2`. By default it searches the directories `<home>/lib/<arch>` and `<home>/lib`. The application can specify additional rules for `file_search_path/2`.

8.3 Using program resources

A *resource* is very similar to a file. Resources however can be represented in two different formats: on files, as well as part of the resource *archive* of a saved-state (see `qsave_program/2`).

A resource has a *name* and a *class*. The *source* data of the resource is a file. Resources are declared by declaring the predicate `resource/3`. They are accessed using the predicate `open_resource/3`.

Before going into details, let us start with an example. Short texts can easily be expressed in Prolog sourcecode, but long texts are cumbersome. Assume our application defines a command ‘help’ that prints a helptext to the screen. We put the content of the helptext into a file called `help.txt`. The following code implements our help command such that `help.txt` is incorporated into the runtime executable.

```

resource(help, text, 'help.txt').

help :-
    open_resource(help, text, In),

```

```

        copy_stream(In, user_output),
        close(In).

copy_stream(In, Out) :-
    get0(In, C),
    copy_stream(C, In, Out).

copy_stream(-1, _, _) :- !.
copy_stream(C, In, Out) :-
    put(Out, C),
    get0(In, C2),
    copy_stream(C2, In, Out).

```

The predicate `help/0` opens the resource as a Prolog stream. If we are executing this from the development environment, this will actually return a stream to the `gelp.txt` itself. When executed from the saved-state, the stream will actually be a stream opened on the program resource file, taking care of the offset and length of the resource.

8.3.1 Predicates Definitions

resource(+Name, +Class, +FileSpec)

This predicate is defined as a dynamic predicate in the module `user`. Clauses for it may be defined in any module, including the `user` module. *Name* is the name of the resource (an atom). A resource name may contain any character, except for `$` and `:`, which are reserved for internal usage by the resource library. *Class* describes the what kind of object is stored in the resource. In the current implementation, it is just an atom. *FileSpec* is a file specification that may exploit `file_search_path/2` (see `absolute_file_name/2`).

Normally, resources are defined as unit clauses (facts), but the definition of this predicate also allows for rules. For proper generation of the saved state, it must be possible to enumerate the available resources by calling this predicate with all its arguments unbound.

Dynamic rules are useful to turn all files in a certain directory into resources, without specifying a resources for each file. For example, assume the `file_search_path/2 icons` refers to the resource directory containing icon-files. The following definition makes all these images available as resources:

```

resource(Name, image, icons(XpmName)) :-
    atom(Name), !,
    file_name_extension(Name, xpm, XpmName).
resource(Name, image, XpmFile) :-
    var(Name),
    absolute_file_name(icons(.), [type(directory)], Dir)
    concat(Dir, '/*.xpm', Pattern),
    expand_file_name(Pattern, XpmFiles),
    member(XpmFile, XpmFiles).

```

open_resource(+Name, ?Class, -Stream)

Opens the resource specified by *Name* and *Class*. If the latter is a variable, it will be unified to the class of the first resource found that has the specified *Name*. If successful, *Stream* becomes a handle to a binary input stream, providing access to the content of the resource.

The predicate `open_resource/3` first checks `resource/3`. When successful it will open the returned resource source-file. Otherwise it will look in the programs resource database. When creating a saved-state, the system normally saves the resource contents into the resource archive, but does not save the resource clauses.

This way, the development environment uses the files (and modifications to the `resource/3` declarations and/or files containing resource info thus immediately affect the running environment, while the runtime system quickly accesses the system resources.

8.3.2 The plrc program

The utility program `plrc` can be used to examine and manipulate the contents of a SWI-Prolog resource file. The options are inspired by the Unix `ar` program. The basic command is:

```
% plrc option resource-file member ...
```

The options are described below.

l

List contents of the archive.

x

Extract named (or all) members of the archive into the current directory.

a

Add files to the archive. If the archive already contains a member with the same name, the contents is replaced. Anywhere in the sequence of members, the options `--class=class` and `--encoding=encoding` may appear. They affect the class and encoding of subsequent files. The initial class is `data` and encoding `none`.

d

Delete named members from the archive.

This command is also described in the `pl(1)` Unix manual page.

8.4 Finding Application files

If your application uses files that are not part of the saved program such as database files, configuration files, etc., the runtime version has to be able to locate these files. The `file_search_path/2` mechanism in combination with the `-palias` command-line argument is the preferred way to locate runtime files. The first step is to define an alias for the toplevel directory of your application. We will call this directory `gnatdir` in our examples.

A good place for storing data associated with SWI-Prolog runtime systems is below the emulator's home-directory. `swi` is a predefined alias for this directory. The following is a useful default definition for the search path.

```
user:file_search_path(gnatdir, swi(gnat)).
```

The application should locate all files using `absolute_file_name`. Suppose `gnatdir` contains a file `config.pl` to define local configuration. Then use the code below to load this file:

```
configure_gnat :-
    ( absolute_file_name(gnatdir('config.pl'), ConfigFile)
      -> consult(ConfigFile)
      ; format(user_error, 'gnat: Cannot locate config.pl~n'),
        halt(1)
      ).
```

8.4.1 Passing a path to the application

Suppose the system administrator has installed the SWI-Prolog runtime environment in `/usr/local/lib/rt-pl-3.2.0`. A user wants to install `gnat`, but `gnat` will look for its configuration in `/usr/local/lib/rt-pl-3.2.0/gnat` where the user cannot write.

The user decides to install the `gnat` runtime files in `/users/bob/lib/gnat`. For one-time usage, the user may decide to start `gnat` using the command:

```
% gnat -p gnatdir=/users/bob/lib/gnat
```

8.5 The Runtime Environment

8.5.1 The Runtime Emulator

The sources may be used to build two versions of the emulator. By default, the *development emulator* is built. This emulator contains all features for interactive development of Prolog applications. If the system is configured using `--enable-runtime`, `make(1)` will create a *runtime version* of the emulator. This emulator is equivalent to the development version, except for the following features:

- *No input editing*
The GNU library `-lreadline` that provides EMACS compatible editing of input lines will not be linked to the system.
- *No tracer*
The tracer and all its options are removed, making the system a little faster too.
- *No profiler*
`profile/3` and friends are not supported. This saves some space and provides better performance.
- *No interrupt*
Keyboard interrupt (Control-C normally) is not rebound and will normally terminate the application.

- *current_prolog_flag(runtime, true) succeeds*
This may be used to verify your application is running in the runtime environment rather than the development environment.
- *clause/[2,3] do not work on static predicates*
This prolog-flag inhibits listing your program. It is only a very limited protection however.


The following fragment is an example for building the runtime environment in `\env{HOME}/lib/rt-pl-3.2.0`. If possible, the shared-library interface should be configured to ensure it can serve a large number of applications.

```
% cd pl-3.2.0
% mkdir runtime
% cd runtime
% ../src/configure --enable-runtime --prefix=$HOME
% make
% make rt-install
```

The runtime directory contains the components listed below. This directory may be tar'ed and shipped with your application.

README.RT	Info on the runtime environment
bin/ <i><arch></i> /pl	The emulator itself
man/pl.1	Manual page for pl
swipl	pointer to the home directory (.)
lib/	directory for shared libraries
lib/ <i><arch></i> /	machine-specific shared libraries

The SWI-Prolog library



This chapter documents the SWI-Prolog library. As SWI-Prolog provides auto-loading, there is little difference between library predicates and built-in predicates. Part of the library is therefore documented in the rest of the manual. Library predicates differ from built-in predicates in the following ways.

- User-definition of a built-in leads to a permission-error, while using the name of a library predicate is allowed.
- If autoloading is disabled explicitly or because trapping unknown predicates is disabled (see `unknown/2` and `current_prolog_flag/2`), library predicates must be loaded explicitly.
- Using libraries reduce the footprint of applications that don't need them.

The documentation of the library is just started. Material from the standard packages should be moved here, some material from other parts of the manual should be moved too and various libraries are not documented at all.

A.1 lists: List Manipulation

This library provides commonly accepted basic predicates for list manipulation in the Prolog community. Some additional list manipulations are built-in. Their description is in section [4.29](#).

append(?List1, ?List2, ?List3)

Succeeds when *List3* unifies with the concatenation of *List1* and *List2*. The predicate can be used with any instantiation pattern (even three variables).

member(?Elem, ?List)

Succeeds when *Elem* can be unified with one of the members of *List*. The predicate can be used with any instantiation pattern.

nextto(?X, ?Y, ?List)

Succeeds when *Y* follows *X* in *List*.

delete(+List1, ?Elem, ?List2)

Delete all members of *List1* that simultaneously unify with *Elem* and unify the result with *List2*.

select(?Elem, ?List, ?Rest)

Select *Elem* from *List* leaving *Rest*. It behaves as `member/2`, returning the remaining elements in *Rest*. Note that besides selecting elements from a list, it can also be used to insert elements.¹

¹BUG: Upto SWI-Prolog 3.3.10, the definition of this predicate was not according to the de-facto standard. The first two arguments were in the wrong order.

nth0(?Index, ?List, ?Elem)

Succeeds when the *Index*-th element of *List* unifies with *Elem*. Counting starts at 0.

nth1(?Index, ?List, ?Elem)

Succeeds when the *Index*-th element of *List* unifies with *Elem*. Counting starts at 1.

last(?List, ?Elem)

Succeeds if *Elem* unifies with the last element of *List*. If *List* is a proper list `last/2` is deterministic. If *List* has an unbound tail, backtracking will cause *List* to grow.²

reverse(+List1, -List2)

Reverse the order of the elements in *List1* and unify the result with the elements of *List2*.

permutation(?List1, ?List2)

Permutation is true when *List1* is a permutation of *List2*. The implementation can solve for *List2* given *List1* or *List1* given *List2*, or even enumerate *List1* and *List2* together.

flatten(+List1, -List2)

Transform *List1*, possibly holding lists as elements into a ‘flat’ list by replacing each list with its elements (recursively). Unify the resulting flat list with *List2*. Example:

```
?- flatten([a, [b, [c, d], e]], X).
```

```
X = [a, b, c, d, e]
```

sumlist(+List, -Sum)

Unify *Sum* to the result of adding all elements in *List*. *List* must be a proper list holding numbers. See `number/1` and `is/2`. for details on arithmetic.

numlist(+Low, +High, -List)

If *Low* and *High* are integers with $Low \leq High$, unify *List* to a list $[Low, Low + 1, \dots High]$. See also `between/3`.

A.1.1 Set Manipulation

is_set(+Set)

Succeeds if *Set* is a list (see `is_list/1`) without duplicates.

list_to_set(+List, -Set)

Unifies *Set* with a list holding the same elements as *List* in the same order. If *list* contains duplicates, only the first is retained. See also `sort/2`. Example:

```
?- list_to_set([a,b,a], X)
```

```
X = [a,b]
```

²The argument order of this predicate was changed in 5.1.12 for compatibility reasons.

intersection(+Set1, +Set2, -Set3)

Succeeds if *Set3* unifies with the intersection of *Set1* and *Set2*. *Set1* and *Set2* are lists without duplicates. They need not be ordered.

subtract(+Set, +Delete, -Result)

Delete all elements of set 'Delete' from 'Set' and unify the resulting set with 'Result'.

union(+Set1, +Set2, -Set3)

Succeeds if *Set3* unifies with the union of *Set1* and *Set2*. *Set1* and *Set2* are lists without duplicates. They need not be ordered.

subset(+Subset, +Set)

Succeeds if all elements of *Subset* are elements of *Set* as well.

A.2 check: Elementary completeness checks

This library defines the predicate `check/0` and a few friends that allow for a quick-and-dirty cross-referencing.

check

Performs the three checking passes implemented by `list_undefined/0`, `list_autoload/0` and `list_redefined/0`. Please check the definition of these predicates for details.

The typical usage of this predicate is right after loading your program to get a quick overview on the completeness and possible conflicts in your program.

list_undefined

Scans the database for predicates that have no definition. A predicate is considered defined if it has clauses, is declared using `dynamic/1` or `multifile/1`. As a program is compiled calls are translated to predicates. If the called predicate is not yet defined it is created as a predicate without definition. The same happens with runtime generated calls. This predicate lists all such undefined predicates that are referenced and not defined in the library. See also `list_autoload/0`. Below is an example from a real program and an illustration how to edit the referencing predicate using `edit/1`.

```
?- list_undefined.
```

```
Warning: The predicates below are not defined. If these are defined
```

```
Warning: at runtime using assert/1, use :- dynamic Name/Arity.
```

```
Warning:
```

```
Warning: rdf_edit:rdfe_retract/4, which is referenced by
```

```
Warning:      1-st clause of rdf_edit:undo/4
```

```
Warning: rdf_edit:rdfe_retract/3, which is referenced by
```

```
Warning:      1-st clause of rdf_edit:delete_object/1
```

```
Warning:      1-st clause of rdf_edit:delete_subject/1
```

```
Warning:      1-st clause of rdf_edit:delete_predicate/1
```

```
?- edit(rdf_edit:undo/4).
```

list_autoload

Lists all undefined (see `list_undefined/0`) predicates that have a definition in the library along with the file from which they will be autoloaded when accessed. See also `autoload/0`.

list_redefined

Lists predicates that are defined in the global module `user` as well as in a normal module. I.e. predicates for which the local definition overrules the global default definition.

A.3 debug: Some reusable code to help debugging applications

This library provides an structured alternative for putting print-statements into your source-code to trace what is going on. Debug messages are organised in *topics* that can be activated and de-activated without changing the source. In addition, if the application is compiled with the `-O` flag these predicates are removed using `goal_expansion/2`.

Although this library can be used through the normal demand-loading mechanism it is advised to load it explicitly before code using it to profit from goal-expansion, which removes these calls if compiled with optimisation on and records the topics from `debug/3` and `debugging/1` for `list_debug_topics/0`.

debug(+Topic, +Format, +Args)

If *Topic* is a selected debugging topic (see `debug/1`) a message is printed using `print_message/2` with level `informational`. *Format* and *Args* are interpreted by `format/2`. Here is a typical example:

```

    . . . ,
    debug(init, 'Initialised ~w', [Module]),
    . . . ,

```

Topic can be any Prolog term. Compound terms can be used to make categories of topics that can be activated using `debug/1`.

debugging(+Topic)

Succeeds if *Topic* is a selected debugging topic. It is intended to execute arbitrary code depending on the users debug topic selection. The construct `(debugging(Topic) -> Code ; true)` is removed if the code is compiled in optimise mode.

debug(+Topic)

Select all registered topics that unify with *Topic* for debugging. This call is normally used from the toplevel to activate a topic for debugging. Topics are de-activated using `nodebug/1`.

nodebug(+Topic)

Deactivates topics for debugging. See `debug/1` for the arguments.

list_debug_topics

List the current status of registered topics. See also `debugging/0`.

assume(*Goal*)

This predicate is to be compared to the C-library `assert()` function. By inserting this goal you explicitly state you expect *Goal* to succeed at this place. As `assume/1` calls are removed when compiling in optimized mode *Goal* should not have side-effects. Typical examples are type-tests and validating invariants defined by your application.

If *Goal* fails the system prints a message and starts the debugger.

A.4 readutil: Reading lines, streams and files

This library contains primitives to read lines, files, multiple terms, etc.

read_line_to_codes(+Stream, -Codes)

Read the next line of input from *Stream* and unify the result with *Codes* after the line has been read. A line is ended by a newline character or end-of-file. Unlike `read_line_to_codes/3`, this predicate removes trailing newline character.

On end-of-file the atom `end_of_file` is returned. See also `at_end_of_stream/[0,1]`.

read_line_to_codes(+Stream, -Codes, ?Tail)

Difference-list version to read an input line to a list of character codes. Reading stops at the newline or end-of-file character, but unlike `read_line_to_codes/2`, the newline is retained in the output. This predicate is especially useful for reading a block of lines upto some delimiter. The following example reads an HTTP header ended by a blank line:

```
read_header_data(Stream, Header) :-
    read_line_to_codes(Stream, Header, Tail),
    read_header_data(Header, Stream, Tail).

read_header_data("\r\n", _, _) :- !.
read_header_data("\n", _, _) :- !.
read_header_data("", _, _) :- !.
read_header_data(_, Stream, Tail) :-
    read_line_to_codes(Stream, Tail, NewTail),
    read_header_data(Tail, Stream, NewTail).
```

read_stream_to_codes(+Stream, -Codes)

Read all input until end-of-file and unify the result to *Codes*.

read_stream_to_codes(+Stream, -Codes, ?Tail)

Difference-list version of `read_stream_to_codes/2`.

read_file_to_codes(+Spec, -Codes, +Options)

Read a file to a list of character codes. *Spec* is a file-specification for `absolute_file_name/3`. *Codes* is the resulting code-list. *Options* is a list of options for `absolute_file_name/3` and `open/4`. In addition, the option `tail(Tail)` is defined, forming a difference-list.

read_file_to_terms(+Spec, -Terms, +Options)

Read a file to a list of prolog terms (see `read/1`). *Spec* is a file-specification for `absolute_file_name/3`. *Terms* is the resulting list of Prolog terms. *Options* is a list of options for `absolute_file_name/3` and `open/4`. In addition, the option `tail(Tail)` is defined, forming a difference-list.

A.5 netscape: Activating your Web-browser

This library deals with the very system dependent task of opening a web page in a browser. See also `url` and the HTTP package.

www_open_url(+URL)

Open *URL* in an external web-browser. The reason to place this in the library is to centralise the maintenance on this highly platform and browser specific task. It distinguishes between the following cases:

- *MS-Windows*
If it detects MS-Windows it uses `win_shell/2` to open the *URL*. The behaviour and browser started depends on the Window and Windows-shell configuration, but in general it should be the behaviour expected by the user.
- *Other platforms*
On other platforms it tests the environment variable (see `getenv/1`) named `BROWSER` or uses `netscape` if this variable is not set. If the browser is either `mozilla` or `netscape`, `www_open_url/1` first tries to open a new window on a running using the `-remote` option of `netscape`. If this fails or the browser is not `mozilla` or `netscape` the system simply passes the *URL* as first argument to the program.

A.6 registry: Manipulating the Windows registry

The `registry` is only available on the MS-Windows version of SWI-Prolog. It loads the foreign extension `plregistry.dll`, providing the predicates described below. This library only makes the most common operations on the registry available through the Prolog user. The underlying DLL provides a more complete coverage of the Windows registry API. Please consult the sources in `pl/src/win32/foreign/plregistry.c` for further details.

In all these predicates, *Path* refers to a '/' separated path into the registry. This is *not* an atom containing '/'-characters as used for filenames, but a term using the functor `//2`. Windows defines the following roots for the registry: `classes_root`, `current_user`, `local_machine` and `users`

registry_get_key(+Path, -Value)

Get the principal (default) value associated to this key. Fails silently if the key does not exist.

registry_get_key(+Path, +Name, -Value)

Get a named value associated to this key.

registry_set_key(+Path, +Value)

Set the principal (default) value of this key. Creates (a path to) the key if this does not already exist.

registry_set_key(+Path, +Name, +Value)

Associated a named value to this key. Creates (a path to) the key if this does not already exist.

registry_delete_key(+Path)

Delete the indicated key.

shell_register_file_type(+Ext, +Type, +Name, +OpenAction)

Register a file-type. *Ext* is the extension to associate. *Type* is the type name, often something like `prolog.type`. *Name* is the name visible in the Windows file-type browser. Finally, *OpenAction* defines the action to execute when a file with this extension is opened in the Windows explorer.

shell_register_dde(+Type, +Action, +Service, +Topic, +Command, +IfNotRunning)

Associate DDE actions to a type. *Type* is the same type as used for the 2nd argument of `shell_register_file_type/4`, *Action* is the a action to perform, *Service* and *Topic* specify the DDE topic to address and *Command* is the command to execute on this topic. Finally, *IfNotRunning* defines the command to execute if the required DDE server is not present.

shell_register_prolog(+Ext)

Default registration of SWI-Prolog, which is invoked as part of the initialisation process on Windows systems. As the source also explains the above predicates, it is given as an example:

```
shell_register_prolog(Ext) :-
    current_prolog_flag(argv, [Me|_]),
    concat_atom(['"', Me, '" "%1"', OpenCommand),
    shell_register_file_type(Ext, 'prolog.type', 'Pro-
log Source',
                                OpenCommand),
    shell_register_dde('prolog.type', consult,
                    prolog, control, 'con-
sult(''%1'''), Me),
    shell_register_dde('prolog.type', edit,
                    prolog, control, 'edit(''%1'''), Me).
```

A.7 url: Analysing and constructing URL

This library deals with the analysis and construction of a URL, **U**niversal **R**esource **L**ocator. URL is the basis for communicating locations of resources (data) on the web. A URL consists of a protocol identifier (e.g. HTTP, FTP), and a protocol-specific syntax further defining the location. URLs are standardized in RFC-1738.

The implementation in this library covers only a small portion of the defined protocols. Though the initial implementation followed RFC-1738 strictly, the current is more relaxed to deal with frequent violations of the standard encountered in practical use.

This library contains code by Jan Wielemaker who wrote the initial version and Lukas Faulstich who added various extensions.

parse_url(?URL, ?Parts)

Construct or analyse a *URL*. *URL* is an atom holding a URL or a variable. *Parts* is a list of components. Each component is of the format *Name(Value)*. Defined components are:

protocol(Protocol)

The used protocol. This is, after the optional `url:`, an identifier separated from the remainder of the URL using `:`. `parse_url/2` assumes the `http` protocol if no protocol is specified and the URL can be parsed as a valid HTTP url. In addition to the RFC-1738 specified protocols, the `file:` protocol is supported as well.

host(Host)

Host-name or IP-address on which the resource is located. Supported by all network-based protocols.

port(Port)

Integer port-number to access on the *Host*. This only appears if the port is explicitly specified in the URL. Implicit default ports (e.g. 80 for HTTP) do *not* appear in the part-list.

path(Path)

(File-) path addressed by the URL. This is supported for the `ftp`, `http` and `file` protocols. If no path appears, the library generates the path `/`.

search(ListOfNameValue)

Search-specification of HTTP URL. This is the part after the `?`, normally used to transfer data from HTML forms that use the 'GET' protocol. In the URL it consists of a www-form-encoded list of *Name=Value* pairs. This is mapped to a list of Prolog *Name=Value* terms with decoded names and values.

fragment(Fragment)

Fragment specification of HTTP URL. This is the part after the `#` character.

The example below illustrates the all this for an HTTP URL.

```
?- parse_url('http://swi.psy.uva.nl/message.cgi?msg=Hello+World%21#x',
             P).
P = [ protocol(http),
      host('swi.psy.uva.nl'),
      fragment(x),
      search([ msg = 'Hello World!'
              ]),
      path('/message.cgi')
    ].
```

By instantiating the parts-list this predicate can be used to create a URL.

parse_url(?URL, +BaseURL, ?Parts)

Same as `parse_url/2`, but dealing a url that is relative to the given *BaseURL*. This is used to analyse or construct a URI found in the document behind *BaseURL*.

global_url(+URL, +BaseURL, -AbsoluteUrl)

Transform a (possibly) relative URL into a global one.

http_location(?Parts, ?Location)

Similar to `parse_url/2`, but only deals with the location part of an HTTP URL. That is, the path, search and fragment specifiers. In the HTTP protocol, the first line of a message is

Action Location [HTTP /HttpVersion]

Location is either an atom or a code-list.

www_form_encode(?Value, ?WwwFormEncoded)

Translate between a string-literal and the x-www-form-encoded representation used in path and search specifications of the HTTP protocol.

Encoding implies mapping space to `+`, preserving alpha-numercial characters, map newlines to `%0D%0A` and anything else to `%XX`. When decoding, newlines appear as a single newline (10) character.

Hackers corner

B

This appendix describes a number of predicates which enable the Prolog user to inspect the Prolog environment and manipulate (or even redefine) the debugger. They can be used as entry points for experiments with debugging tools for Prolog. The predicates described here should be handled with some care as it is easy to corrupt the consistency of the Prolog system by misusing them.

B.1 Examining the Environment Stack

prolog_current_frame(-Frame)

Unify *Frame* with an integer providing a reference to the parent of the current local stack frame. A pointer to the current local frame cannot be provided as the predicate succeeds deterministically and therefore its frame is destroyed immediately after succeeding.

prolog_frame_attribute(+Frame, +Key, -Value)

Obtain information about the local stack frame *Frame*. *Frame* is a frame reference as obtained through `prolog_current_frame/1`, `prolog_trace_interception/4` or this predicate. The key values are described below.

alternative

Value is unified with an integer reference to the local stack frame in which execution is resumed if the goal associated with *Frame* fails. Fails if the frame has no alternative frame.

has_alternatives

Value is unified with `true` if *Frame* still is a candidate for backtracking. `false` otherwise.

goal

Value is unified with the goal associated with *Frame*. If the definition module of the active predicate is not `user` the goal is represented as $\langle module \rangle : \langle goal \rangle$. Do not instantiate variables in this goal unless you **know** what you are doing! Note that the returned term may contain references to the frame and should be discarded before the frame terminates.¹

parent_goal

If *Value* is instantiated to a callable term, find a frame executing the predicate described by *Value* and unify the arguments of *Value* to the goal arguments associated with the frame. This is intended to check the current execution context. The user must ensure the checked parent goal is not removed from the stack due to last-call optimisation and be aware of the slow operation on deeply nested calls.

¹The returned term is actually an illegal Prolog term that may hold references from the global- to the local stack to preserve the variable names.

clause

Value is unified with a reference to the currently running clause. Fails if the current goal is associated with a foreign (C) defined predicate. See also `nth_clause/3` and `clause_property/2`.

level

Value is unified with the recursion level of *Frame*. The top level frame is at level '0'.

parent

Value is unified with an integer reference to the parent local stack frame of *Frame*. Fails if *Frame* is the top frame.

context_module

Value is unified with the name of the context module of the environment.

top

Value is unified with `true` if *Frame* is the top Prolog goal from a recursive call back from the foreign language. `false` otherwise.

hidden

Value is unified with `true` if the frame is hidden from the user, either because a parent has the `hide-children` attribute (all system predicates), or the system has no `trace-me` attribute.

pc

Value is unified with the program-pointer saved on behalf of the parent-goal if the parent-goal is not owned by a foreign predicate.

argument(*N*)

Value is unified with the *N*-th slot of the frame. Argument 1 is the first argument of the goal. Arguments above the arity refer to local variables. Fails silently if *N* is out of range.

prolog_choice_attribute(+ChoicePoint, +Key, -Value)

Extract attributes of a choice-point. *ChoicePoint* is a reference to a choice-point as passed to `prolog_trace_interception/4` on the 3-th argument. *Key* specifies the requested information:

parent

Requests a reference to the first older choice-point.

frame

Requests a reference to the frame to which the choice-point refers.

type

Requests the type. Defined values are `clause` (the goal has alternative clauses), `foreign` (non-deterministic foreign predicate), `jump` (clause internal choicepoint), `top` (first dummy choice-point), `catch` (`catch/3` to allow for undo), `debug` (help the debugger), or `none` (has been deleted).

This predicate is used for the graphical debugger to show the choicepoint stack.

deterministic(-Boolean)

Unifies its argument with `true` if the clause in which it appears has not created any choice-points since it was started. There are few realistic situations for using this predicate. It is used by the `prolog/0` toplevel to check whether Prolog should prompt the user for alternatives.

B.2 Intercepting the Tracer

prolog_trace_interception(+Port, +Frame, +Choice, -Action)

Dynamic predicate, normally not defined. This predicate is called from the SWI-Prolog debugger just before it would show a port. If this predicate succeeds the debugger assumes the trace action has been taken care of and continues execution as described by *Action*. Otherwise the normal Prolog debugger actions are performed.

Port denotes the reason to activate the tracer ('port' in the 4/5-port, but with some additions:

call

Normal entry through the call-port of the 4-port debugger.

redo

Normal entry through the call-port of the 4-port debugger. The redo port signals resuming a predicate to generate alternative solutions.

unify

The unify-port represents the *neck* instruction, signalling the end of the head-matching process. This port is normally invisible. See `leash/1` and `visible/1`.

exit

The exit-port signals the goal is proved. It is possible for the goal to have alternative. See `prolog_frame_attribute/3` to examine the goal-stack.

fail

The fail-port signals final failure of the goal.

exception(*Except*)

An exception is raised and still pending. This port is activated on each parent frame of the frame generating the exception until the exception is caught or the user restarts normal computation using `retry`. *Except* is the pending exception-term.

break(*PC*)

A break instruction is executed. *PC* is program counter. This port is used by the graphical debugger.

cut_call(*PC*)

A cut is encountered at *PC*. This port is used by the graphical debugger. to visualise the effect of the cut.

cut_exit(*PC*)

A cut has been executed. See `cut_call(PC)` for more information.

Frame is a reference to the current local stack frame, which can be examined using `prolog_frame_attribute/4`. *Choice* is a reference to the last choice-point and can be examined using `prolog_choice_attribute/3`. *Action* should be unified with one of the atoms `continue` (just continue execution), `retry` (retry the current goal) or `fail` (force the current goal to fail). Leaving it a variable is identical to `continue`.

Together with the predicates described in section 4.39 and the other predicates of this chapter this predicate enables the Prolog user to define a complete new debugger in Prolog. Besides this it enables the Prolog programmer monitor the execution of a program. The example below records all goals trapped by the tracer in the database.

```

prolog_trace_interception(Port, Frame, _PC, continue) :-
    prolog_frame_attribute(Frame, goal, Goal),
    prolog_frame_attribute(Frame, level, Level),
    recordz(trace, trace(Port, Level, Goal)).

```

To trace the execution of ‘go’ this way the following query should be given:

```
?- trace, go, notrace.
```

prolog_skip_level(-Old, +New)

Unify *Old* with the old value of ‘skip level’ and then set this level according to *New*. *New* is an integer, or the special atom `very_deep` (meaning don’t skip). The ‘skip level’ is a global variable of the Prolog system that disables the debugger on all recursion levels deeper than the level of the variable. Used to implement the trace options ‘skip’ (sets skip level to the level of the frame) and ‘up’ (sets skip level to the level of the parent frame (i.e., the level of this frame minus 1)).

B.3 Hooks using the exception/3 predicate

This section describes the predicate `exception/3`, which may be defined by the user in the module `user` as a multifile predicate. Unlike the name suggests, this is actually a *hook* predicate. Exceptions are handled by the ISO predicates `catch/3` and `throw/1`. They all frames created after the matching `catch/3` to be discarded immediately.

The predicate `exception/3` is called by the kernel on a couple of events, allowing the user to alter the behaviour on some predefined events.

exception(+Exception, +Context, -Action)

Dynamic predicate, normally not defined. Called by the Prolog system on run-time exceptions. Currently `exception/3` is only used for trapping undefined predicates. Future versions might handle signal handling, floating exceptions and other runtime errors via this mechanism. The values for *Exception* are described below.

undefined_predicate

If *Exception* is `undefined_predicate` *Context* is instantiated to a term *Name/Arity*. *Name* refers to the name and *Arity* to the arity of the undefined predicate. If the definition module of the predicate is not `user`, *Context* will be of the form `<Module>:<Name>/<Arity>`. If the predicate fails Prolog will generate an `existence_error` exception. If the predicate succeeds it should instantiate the last argument either to the atom `fail` to tell Prolog to fail the predicate, the atom `retry` to tell Prolog to retry the predicate or `error` to make the system generate an exception. The action `retry` only makes sense if the exception handler has defined the predicate.

B.4 Hooks for integrating libraries

Some libraries realise an entirely new programming paradigm on top of Prolog. An example is XPCP which adds an object-system to Prolog as well as an extensive set of graphical primitives. SWI-Prolog

provides several hooks to improve the integration of such libraries. See also section 4.4 for editing hooks and section 4.9.3 for hooking into the message system.

prolog_list_goal(:Goal)

Hook, normally not defined. This hook is called by the 'L' command of the tracer in the module `user` to list the currently called predicate. This hook may be defined to list only relevant clauses of the indicated *Goal* and/or show the actual source-code in an editor. See also `portray/1` and `multifile/1`.

prolog:debug_control_hook(:Action)

Hook for the debugger-control predicates that allows the creator of more high-level programming languages to use the common front-end predicates to control de debugger. For example, XPCE uses these hooks to allow for spying methods rather than predicates. *Action* is one of:

spy(Spec)

Hook in `spy/1`. If the hook succeeds `spy/1` takes no further action.

nospy(Spec)

Hook in `nospy/1`. If the hook succeeds `spy/1` takes no further action. If `spy/1` is hooked, it is advised to place a complementary hook for `nospy/1`.

nospyall

Hook in `nospyall/0`. Should remove all spy-points. This hook is called in a failure-driven loop.

debugging

Hook in `debugging/0`. It can be used in two ways. It can report the status of the additional debug-points controlled by the above hooks and fail to let the system report the others or it succeed, overruling the entire behaviour of `debugging/0`.

prolog:help_hook(+Action)

Hook into `help/0` and `help/1`. If the hook succeeds, the built-in actions are not executed. For example, `?- help(picture).` is caught by the XPCE help-hook to give help on the class *picture*. Defined actions are:

help

User entered plain `help/0` to give default help. The default performs `help(help/1)`, giving help on help.

help(What)

Hook in `help/1` on the topic *What*.

apropos(What)

Hook in `apropos/1` on the topic *What*.

B.5 Hooks for loading files

All loading of source-files is achieved by `load_files/2`. The hook `prolog_load_file/2` can be used to load Prolog code from non-files or even load entirely different information, such as foreign files.

user:prolog_load_file(+Spec, +Options)

Load a single object. If this call succeeds, `load_files/2` assumes the action has been taken care of. This hook is only called if *Options* does not contain the `stream(Input)` option.

The `http_load` provides an example, loading Prolog sources directly from an HTTP server.

B.6 Readline Interaction

The following predicates are available if `current_prolog_flag(readline, true)` succeeds. They allow for direct interaction with the GNU readline library. See also `readline(3)`

rl_read_init_file(+File)

Read a readline initialisation file. Readline by default reads `~/ .inputrc`. This predicate may be used to read alternative readline initialisation files.

rl_add_history(+Line)

Add a line to the Control-P/Control-N history system of the readline library.

Glossary of Terms



anonymous [variable]

The variable `_` is called the *anonymous* variable. Multiple occurrences of `_` in a single *term* are not *shared*.

arguments

Arguments are *terms* that appear in a *compound term*. *A1* and *a2* are the first and second argument of the term `myterm(A1, a2)`.

arity

Argument count (is number of arguments) of a *compound term*.

assert

Add a *clause* to a *predicate*. Clauses can be added at either end of the clause-list of a *predicate*. See `assert/1` and `assertz/1`.

atom

Textual constant. Used as name for *compound terms*, to represent constants or text.

backtracking

Searching process used by Prolog. If a predicate offers multiple *clauses* to solve a *goal*, they are tried one-by-one until one *succeeds*. If a subsequent part of the prove is not satisfied with the resulting *variable binding*, it may ask for an alternative *solution* (= *binding* of the *variables*), causing Prolog to reject the previously chosen *clause* and try the next one.

binding [of a variable]

Current value of the *variable*. See also *backtracking* and *query*.

built-in [predicate]

Predicate that is part of the Prolog system. Built in predicates cannot be redefined by the user, unless this is overruled using `redefine_system_predicate/1`.

body

Part of a *clause* behind the *neck* operator (`: -`).

clause

‘Sentence’ of a Prolog program. A *clause* consists of a *head* and *body* separated by the *neck* operator (`: -`) or it is a *fact*. For example:

```
parent(X) :-  
    father(X, _).
```


Expressed “X is a parent if X is a father of someone”. See also *variable* and *predicate*.

compile

Process where a Prolog *program* is translated to a sequence of instructions. See also *interpreted*. SWI-Prolog always compiles your program before executing it.

compound [term]

Also called *structure*. It consists of a name followed by *N arguments*, each of which are *terms*. *N* is called the *arity* of the term.

context module

If a *term* is referring to a *predicate* in a *module*, the *context module* is used to find the target module. The context module of a *goal* is the module in which the *predicate* is defined, unless this *predicate* is *module transparent*, in which case the *context module* is inherited from the parent *goal*. See also `module_transparent/1`.

dynamic [predicate]

A *dynamic* predicate is a predicate to which *clauses* may be *asserted* and from which *clauses* may be *retracted* while the program is running. See also *update view*.

exported [predicate]

A *predicate* is said to be *exported* from a *module* if it appears in the *public list*. This implies that the predicate can be *imported* into another module to make it visible there. See also `use_module/[1, 2]`.

fact

Clause without a *body*. This is called a fact because interpreted as logic, there is no condition to be satisfied. The example below states john is a person.

```
person(john).
```

fail

A *goal* is said to have failed if it could not be *proven*.

float

Computers crippled representation of a real number. Represented as ‘IEEE double’.

foreign

Computer code expressed in other languages than Prolog. SWI-Prolog can only cooperate directly with the C and C++ computer languages.

functor

Combination of name and *arity* of a *compound* term. The term $f_{OO}(a, b, c)$ is said to be a term belonging to the functor $f_{OO}/3$. $f_{OO}/0$ is used to refer to the *atom* f_{OO} .

goal

Question stated to the Prolog engine. A *goal* is either an *atom* or a *compound* term. A *goal* succeeds, in which case the *variables* in the *compound* terms have a *binding* or *fails* if Prolog fails to prove the *goal*.

hashing

Indexing technique used for quick lookup.

head

Part of a *clause* before the *neck* instruction. This is an atom or *compound* term.

imported [predicate]

A *predicate* is said to be *imported* into a *module* if it is defined in another *module* and made available in this *module*. See also chapter 5.

indexing

Indexing is a technique used to quickly select candidate *clauses* of a *predicate* for a specific *goal*. In most Prolog systems, including SWI-Prolog, indexing is done on the first *argument* of the *head*. If this argument is instantiated to an *atom*, *integer*, *float* or *compound* term with *functor*, *hashing* is used quickly select all *clauses* of which the first argument may *unify* with the first argument of the *goal*.

integer

Whole number. On most current machines, SWI-Prolog integers are represented as '32-bit signed values', ranging from -2147483648 to 2147483647. See also `current_prolog_flag/2`.

interpreted

As opposed to *compiled*, interpreted means the Prolog system attempts to prove a *goal* by directly reading the *clauses* rather than executing instructions from an (abstract) instruction set that is not or only indirectly related to Prolog.

meta predicate

A *predicate* that reasons about other *predicates*, either by calling them, (re)defining them or querying *properties*.

module

Collection of predicates. Each module defines a name-space for predicates. *built-in* predicates are accessible from all modules. Predicates can be published (*exported*) and *imported* to make their definition available to other modules.

module transparent [predicate]

A *predicate* that does not change the *context module*. Sometimes also called a *meta predicate*.

multifile [predicate]

Predicate for which the definition is distributed over multiple source-files. See `multi_file/1`.

neck

Operator (`: -`) separating *head* from *body* in a *clause*.

operator

Symbol (*atom*) that may be placed before its *operant* (prefix), after its *operant* (postfix) or between its two *operants* (infix).

In Prolog, the expression `a+b` is exactly the same as the canonical term `+(a,b)`.

operant

Argument of an operator.

precedence

The *priority* of an *operator*. Operator precedence is used to interpret $a+b*c$ as $+(a, *(b, c))$.

predicate

Collection of *clauses* with the same *functor* (name/arity). If a *goal* is proved, the system looks for a *predicate* with the same functor, then used *indexing* to select candidate *clauses* and then tries these *clauses* one-by-one. See also *backtracking*.

priority

In the context of *operators* a synonym for *precedence*.

program

Collection of *predicates*.

property

Attribute of an object. SWI-Prolog defines various **_property* predicates to query the status of predicates, clauses. etc.

prove

Process where Prolog attempts to prove a *query* using the available *predicates*.

public list

List of *predicates* exported from a *module*.

query

See *goal*.

retract

Remove a *clause* from a *predicate*. See also *dynamic*, *update view* and *assert*.

shared

Two *variables* are called *shared* after they are *unified*. This implies if either of them is *bound*, the other is bound to the same value:

```
?- A = B, A = a.
```

```
A = a,
```

```
B = a
```

singleton [variable]

Variable appearing only one time in a *clause*. SWI-Prolog normally warns for this to avoid you making spelling mistakes. If a variable appears on purpose only once in a clause, write it as `_` (see *anonymous*) or make sure the first character is a `_`. See also the `style_check/1` option `singletons`.

solution

Bindings resulting from a successfully *proven goal*.

structure

Synonym for *compound* term.

string

Used for the following representations of text: a packed array (see section 4.23, SWI-Prolog specific), a list of character codes or a list of one-character *atoms*.

succeed

A *goal* is said to have *succeeded* if it has been *proven*.

term

Value in Prolog. A *term* is either a *variable*, *atom*, integer, float or *compound* term. In addition, SWI-Prolog also defines the type *string*

transparent

See *module transparent*.

unify

Prolog process to make two terms equal by assigning variables in one term to values at the corresponding location of the other term. For example:

```
?- foo(a, B) = foo(A, b).
```

```
A = a,
```

```
B = b
```

Unlike assignment (which does not exist in Prolog), unification is not directed.

update view

How Prolog behaves when a *dynamic predicate* is changed while it is running. There are two models. In most older Prolog systems the change becomes immediately visible to the *goal*, in modern systems including SWI-Prolog, the running *goal* is not affected. Only new *goals* ‘see’ the new definition.

variable

A Prolog variable is a value that ‘is not yet bound’. After *binding* a variable, it cannot be modified. *Backtracking* to a point in the execution before the variable was bound will turn it back into a variable:

```
?- A = b, A = c.
```

```
No
```

```
?- (A = b; true; A = c).
```

```
A = b ;
```

```
A = _G283 ;
```

```
A = c ;
```

```
No
```

See also *unify*.

SWI-Prolog License Conditions and Tools



SWI-Prolog licensing aims at a large audience, combining ideas from the Free Software Foundation and the less principal Open Source Initiative. The license aims at:

- Make SWI-Prolog itself and its libraries are ‘As free as possible’.
- Allow for easy integration of contributions. See section [D.2](#).
- Free software can build on SWI-Prolog without limitations.
- Non-free (open or proprietary) software can be produced using SWI-Prolog, although contributed pure GPL-ed components cannot be used.

To achieve this, different parts of the system have different licenses. SWI-Prolog programs consists of a mixture of ‘native’ code (source compiled to machine instructions) and ‘virtual machine’ code (Prolog source compiled to SWI-Prolog virtual machine instructions, covering both compiled SWI-Prolog libraries and your compiled application).

For maximal coherence between free licenses, we start with the two prime licenses from the Free Software Foundation, the GNU General Public License (GPL) and the Lesser GNU General Public License (LGPL), after which we add a proven (used by the GNU-C compiler runtime library as well as the GNU *ClassPath* project) exception to deal with the specific nature of compiled virtual machine code in a saved state.

D.1 The SWI-Prolog kernel and foreign libraries

The SWI-Prolog kernel and our foreign libraries are distributed under the **LGPL**. A Prolog executable consists of the combination of these ‘native’ code components and Prolog virtual machine code. The SWI-Prolog `plrc` utility allows for disassembling and re-assembling these parts, a process satisfying article **6b** of the LGPL.

Under the LGPL SWI-Prolog can be linked to code distributed under arbitrary licenses, provided a number of requirements are fulfilled. The most important requirement is that, if an application relies on a *modified* version of SWI-Prolog, the modified sources must be made available.

D.1.1 The SWI-Prolog Prolog libraries

Lacking a satisfactory technical solution to handle article **6** of the LGPL, this license cannot be used for the Prolog source code that is part of the SWI-Prolog system (both libraries and kernel code). This situation is comparable to `libgcc`, the runtime library used with the GNU C-compiler. Therefore, we use the same proven license terms as this library. The `libgcc` license is the with a special exception. Below we rephrased this exception adjusted to our needs:

As a special exception, if you link this library with other files, compiled with a Free Software compiler, to produce an executable, this library does not by itself cause the resulting executable to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the executable file might be covered by the GNU General Public License.

D.2 Contributing to the SWI-Prolog project

To achieve maximal coherence using SWI-Prolog for Free and Non-Free software we advise the use of the LGPL for contributed foreign code and the use of the GPL with SWI-Prolog exception for Prolog code for contributed modules.

As a rule of thumb it is advised to use the above licenses whenever possible and only use a strict GPL compliant license only if the module contains other code under strict GPL compliant licenses.

D.3 Software support to keep track of license conditions

Given the above, it is possible that SWI-Prolog packages and extensions will rely on the GPL.¹ The predicates below allow for registering license requirements for Prolog files and foreign modules. The predicate `eval_license/0` reports which components from the currently configured system are distributed under copy-left and open source enforcing licenses (the GPL) and therefore must be replaced before distributing linked applications under non-free license conditions.

eval_license

Evaluate the license conditions of all loaded components. If the system contains one or more components that are licenced under GPL-like restrictions the system indicates this program may only be distributed under the GPL license as well as which components prohibit the use of other license conditions.

license(+LicenseId, +Component)

Register the fact that *Component* is distributed under a license identified by *LicenseId*. The most important *LicenseId*'s are:

swipl

Indicates this module is distributed under the GNU General Public License (GPL) with the SWI-Prolog exception:²

As a special exception, if you link this library with other files, compiled with SWI-Prolog, to produce an executable, this library does not by itself cause the resulting executable to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the executable file might be covered by the GNU General Public License.

¹On the Unix version, the default toplevel uses the GNU readline library for command-line editing. This library is distributed under the GPL. In practice this problem is small as most final applications have Prolog embedded, without direct access to the commandline and therefore without need for `libreadline`.

²This exception is a straight re-phrasing of the license used for `libgcc`, the GNU-C runtime library facing similar technical issues.

This should be the default for software contributed to the SWI-Prolog project as it allows the community to prosper both in the free and non-free world. Still, people using SWI-Prolog to create non-free applications must contribute sources to improvements they make to the community.

lgpl

This is the default license for foreign-libraries linked with SWI-Prolog. Use `PL_license()` to register the condition from foreign code.

gpl

Indicates this module is strictly Free Software, which implies it cannot be used together with any module that is incompatible to the GPL. Please only use these conditions when forced by other code used in the component.

license(+LicenseId)

Intended as a directive in Prolog source files. It takes the current filename and calls `license/2`.

void PL_license(const char *LicenseId, const char *Component)

Intended for the `install()` procedure of foreign libraries. This call can be made *before* `PL_initialise()`.

Summary

E

E.1 Predicates

The predicate summary is used by the Prolog predicate `apropos/1` to suggest predicates from a keyword.

<code>!/0</code>	Cut (discard choicepoints)
<code>!/1</code>	Cut block. See <code>block/3</code>
<code>,/2</code>	Conjunction of goals
<code>->/2</code>	If-then-else
<code>*->/2</code>	Soft-cut
<code>./2</code>	Consult. Also list constructor
<code>;/2</code>	Disjunction of goals. Same as <code> /2</code>
<code></2</code>	Arithmetic smaller
<code>=/2</code>	Unification
<code>==/2</code>	“Univ.” Term to list conversion
<code>:=/2</code>	Arithmetic equal
<code>=</2</code>	Arithmetic smaller or equal
<code>==/2</code>	Identical
<code>==@/2</code>	Structural identical
<code>=\=/2</code>	Arithmetic not equal
<code>>/2</code>	Arithmetic larger
<code>>=/2</code>	Arithmetic larger or equal
<code>@</2</code>	Standard order smaller
<code>@=</2</code>	Standard order smaller or equal
<code>@>/2</code>	Standard order larger
<code>@>=/2</code>	Standard order larger or equal
<code>\+/1</code>	Negation by failure. Same as <code>not/1</code>
<code>\=/2</code>	Not unifyable
<code>\==/2</code>	Not identical
<code>\==@/2</code>	Not structural identical
<code>^/2</code>	Existential quantification (<code>bagof/3</code> , <code>setof/3</code>)
<code> /2</code>	Disjunction of goals. Same as <code>;/2</code>
<code>abolish/1</code>	Remove predicate definition from the database
<code>abolish/2</code>	Remove predicate definition from the database
<code>abort/0</code>	Abort execution, return to top level
<code>absolute_file_name/2</code>	Get absolute path name
<code>absolute_file_name/3</code>	Get absolute path name with options
<code>access_file/2</code>	Check access permissions of a file

<code>add_import_module/3</code>	Add module to the auto-import list
<code>append/1</code>	Append to a file
<code>apply/2</code>	Call goal with additional arguments
<code>apropos/1</code>	<code>online_help</code> Search manual
<code>arg/3</code>	Access argument of a term
<code>arithmetic_function/1</code>	Register an evaluable function
<code>assert/1</code>	Add a clause to the database
<code>assert/2</code>	Add a clause to the database, give reference
<code>asserta/1</code>	Add a clause to the database (first)
<code>asserta/2</code>	Add a clause to the database (first)
<code>assertz/1</code>	Add a clause to the database (last)
<code>assertz/2</code>	Add a clause to the database (last)
<code>assume/1</code>	Make assertions about your program
<code>attach_console/0</code>	Attach I/O console to thread
<code>at_end_of_stream/0</code>	Test for end of file on input
<code>at_end_of_stream/1</code>	Test for end of file on stream
<code>at_halt/1</code>	Register goal to run at <code>halt/1</code>
<code>at_initialization/1</code>	Register goal to run at start-up
<code>atom/1</code>	Type check for an atom
<code>atom_chars/2</code>	Convert between atom and list of characters
<code>atom_codes/2</code>	Convert between atom and list of characters codes
<code>atom_length/2</code>	Determine length of an atom
<code>atom_prefix/2</code>	Test for start of atom
<code>atom_number/2</code>	Convert between atom and number
<code>atom_to_term/3</code>	Convert between atom and term
<code>atomic/1</code>	Type check for primitive
<code>autoload/0</code>	Autoload all predicates now
<code>bagof/3</code>	Find all solutions to a goal
<code>between/3</code>	Integer range checking/generating
<code>block/3</code>	Start a block ('catch'/'throw')
<code>break/0</code>	Start interactive toplevel
<code>call/1</code>	Call a goal
<code>call/[2..]</code>	Call with additional arguments
<code>call_cleanup/3</code>	Guard a goal with a cleanup-handler
<code>call_cleanup/2</code>	Guard a goal with a cleanup-handler
<code>call_shared_object_function/2</code>	UNIX: Call C-function in shared (.so) file
<code>call_with_depth_limit/3</code>	Prove goal with bounded depth
<code>callable/1</code>	Test for atom or compound term
<code>catch/3</code>	Call goal, watching for exceptions
<code>char_code/2</code>	Convert between character and character code
<code>char_conversion/2</code>	Provide mapping of input characters
<code>char_type/2</code>	Classify characters
<code>character_count/2</code>	Get character index on a stream
<code>chdir/1</code>	Compatibility: change working directory
<code>checklist/2</code>	Invoke goal on all members of a list
<code>clause/2</code>	Get clauses of a predicate
<code>clause/3</code>	Get clauses of a predicate

clause_property/2	Get properties of a clause
close/1	Close stream
close/2	Close stream (forced)
close_dde_conversation/1	Win32: Close DDE channel
close_shared_object/1	UNIX: Close shared library (.so file)
compare/3	Compare, using a predicate to determine the order
compile_predicates/1	Compile dynamoc code to static
compiling/0	Is this a compilation run?
compound/1	Test for compound term
atom_concat/3	Append two atoms
code_type/2	Classify a character-code
concat_atom/2	Append a list of atoms
concat_atom/3	Append a list of atoms with separator
consult/1	Read (compile) a Prolog source file
context_module/1	Get context module of current goal
convert_time/8	Break time stamp into fields
convert_time/2	Convert time stamp to string
copy_stream_data/2	Copy all data from stream to stream
copy_stream_data/3	Copy n bytes from stream to stream
copy_term/2	Make a copy of a term
current_arithmetic_function/1	Examine evaluable functions
current_atom/1	Examine existing atoms
current_char_conversion/2	Query input character mapping
current_flag/1	Examine existing flags
current_foreign_library/2	shlib Examine loaded shared libraries (.so files)
current_format_predicate/2	Enumerate user-defined format codes
current_functor/2	Examine existing name/arity pairs
current_input/1	Get current input stream
current_key/1	Examine existing database keys
current_module/1	Examine existing modules
current_module/2	Examine existing modules
current_mutex/3	Examine existing mutexes
current_op/3	Examine current operator declarations
current_output/1	Get the current output stream
current_predicate/1	Examine existing predicates (ISO)
current_predicate/2	Examine existing predicates
current_signal/3	Current software signal mapping
current_stream/3	Examine open streams
current_thread/2	Examine Prolog threads
dde_current_connection/2	Win32: Examine open DDE connections
dde_current_service/2	Win32: Examine DDE services provided
dde_execute/2	Win32: Execute command on DDE server
dde_register_service/2	Win32: Become a DDE server
dde_request/3	Win32: Make a DDE request
dde_poke/3	Win32: POKE operation on DDE server
dde_unregister_service/1	Win32: Terminate a DDE service
debug/0	Test for debugging mode

<code>debug/1</code>	Select topic for debugging
<code>debug/3</code>	Print debugging message on topic
<code>debug_control_hook/1</code>	(hook) Extend <code>spy/1</code> , etc.
<code>debugging/0</code>	Show debugger status
<code>debugging/1</code>	Test where we are debugging topic
<code>default_module/2</code>	Get the default modules of a module
<code>delete_directory/1</code>	Remove a folder from the file system
<code>delete_file/1</code>	Remove a file from the file system
<code>delete_import_module/2</code>	Remove module from import list
<code>deterministic/1</code>	Test deterministicy of current clause
<code>discontiguous/1</code>	Indicate distributed definition of a predicate
<code>downcase_atom/2</code>	Convert atom to lower-case
<code>dwim_match/2</code>	Atoms match in “Do What I Mean” sense
<code>dwim_match/3</code>	Atoms match in “Do What I Mean” sense
<code>dwim_predicate/2</code>	Find predicate in “Do What I Mean” sense
<code>dynamic/1</code>	Indicate predicate definition may change
<code>edit/0</code>	Edit current script- or associated file
<code>edit/1</code>	Edit a file, predicate, module (extensible)
<code>ensure_loaded/1</code>	Consult a file if that has not yet been done
<code>erase/1</code>	Erase a database record or clause
<code>eval_license/0</code>	Evaluate licenses of loaded modules
<code>exception/3</code>	(hook) Handle runtime exceptions
<code>exists_directory/1</code>	Check existence of directory
<code>exists_file/1</code>	Check existence of file
<code>exit/2</code>	Exit from named block. See <code>block/3</code>
<code>expand_answer/2</code>	Expand answer of query
<code>expand_file_name/2</code>	Wildcard expansion of file names
<code>expand_file_search_path/2</code>	Wildcard expansion of file paths
<code>expand_goal/2</code>	Compiler: expand goal in clause-body
<code>expand_query/4</code>	Expanded entered query
<code>expand_term/2</code>	Compiler: expand read term into clause(s)
<code>explain/1</code>	<code>explain</code> Explain argument
<code>explain/2</code>	<code>explain</code> 2nd argument is explanation of first
<code>export/1</code>	Export a predicate from a module
<code>export_list/2</code>	List of public predicates of a module
<code>fail/0</code>	Always false
<code>fail/1</code>	Immediately fail named block. See <code>block/3</code>
<code>current_prolog_flag/2</code>	Get system configuration parameters
<code>file_base_name/2</code>	Get file part of path
<code>file_directory_name/2</code>	Get directory part of path
<code>file_name_extension/3</code>	Add, remove or test file extensions
<code>file_search_path/2</code>	Define path-aliases for locating files
<code>fileerrors/2</code>	Do/Don’t warn on file errors
<code>findall/3</code>	Find all solutions to a goal
<code>flag/3</code>	Simple global variable system
<code>float/1</code>	Type check for a floating point number
<code>flush_output/0</code>	Output pending characters on current stream

flush_output/1	Output pending characters on specified stream
forall/2	Prove goal for all solutions of another goal
format/1	Formatted output
format/2	Formatted output with arguments
format/3	Formatted output on a stream
format_predicate/2	Program <code>format/[1,2]</code>
free_variables/2	Find unbound variables in a term
functor/3	Get name and arity of a term or construct a term
garbage_collect/0	Invoke the garbage collector
garbage_collect_atoms/0	Invoke the atom garbage collector
gensym/2	Generate unique atoms from a base
get/1	Read first non-blank character
get/2	Read first non-blank character from a stream
get0/1	Read next character
get0/2	Read next character from a stream
get_byte/1	Read next byte (ISO)
get_byte/2	Read next byte from a stream (ISO)
get_char/1	Read next character as an atom (ISO)
get_char/2	Read next character from a stream (ISO)
get_code/1	Read next character (ISO)
get_code/2	Read next character from a stream (ISO)
get_single_char/1	Read next character from the terminal
get_time/1	Get current time
getenv/2	Get shell environment variable
goal_expansion/2	Hook for macro-expanding goals
ground/1	Verify term holds no unbound variables
guitracer/0	Install hooks for the graphical debugger
halt/0	Exit from Prolog
halt/1	Exit from Prolog with status
hash_term/2	Hash-value of ground term
help/0	Give help on help
help/1	Give help on predicates and show parts of manual
help_hook/1	(hook) User-hook in the help-system
ignore/1	Call the argument, but always succeed
import/1	Import a predicate from a module
import_module/2	Query import modules
include/1	Include a file with declarations
index/1	Change clause indexing
initialization/1	Initialization directive
int_to_atom/2	Convert from integer to atom
int_to_atom/3	Convert from integer to atom (non-decimal)
integer/1	Type check for integer
interactor/0	Start new thread with console and toplevel
is/2	Evaluate arithmetic expression
is_absolute_file_name/1	True if arg defines an absolute path
is_list/1	Type check for a list
keysort/2	Sort, using a key

<code>last/2</code>	Last element of a list
<code>leash/1</code>	Change ports visited by the tracer
<code>length/2</code>	Length of a list
<code>library_directory/1</code>	(hook) Directories holding Prolog libraries
<code>license/1</code>	Define license for current file
<code>license/2</code>	Define license for named module
<code>limit_stack/2</code>	Limit stack expansion
<code>line_count/2</code>	Line number on stream
<code>line_position/2</code>	Character position in line on stream
<code>list_debug_topics/0</code>	List registered topics for debugging
<code>listing/0</code>	List program in current module
<code>listing/1</code>	List predicate
<code>load_files/2</code>	Load source files with options
<code>load_foreign_library/1</code>	<code>shlib</code> Load shared library (.so file)
<code>load_foreign_library/2</code>	<code>shlib</code> Load shared library (.so file)
<code>make/0</code>	Reconsult all changed source files
<code>make_directory/1</code>	Create a folder on the file system
<code>make_fat_filemap/1</code>	Win32: Create file containing non-FAT filenames
<code>make_library_index/1</code>	Create autoload file INDEX.pl
<code>make_library_index/2</code>	Create selective autoload file INDEX.pl
<code>maplist/3</code>	Transform all elements of a list
<code>memberchk/2</code>	Deterministic <code>member/2</code>
<code>merge/3</code>	Merge two sorted lists
<code>merge_set/3</code>	Merge two sorted sets
<code>message_hook/3</code>	Intercept <code>print_message/2</code>
<code>message_queue_create/1</code>	Create queue for thread communication
<code>message_queue_destroy/1</code>	Destroy queue for thread communication
<code>message_to_string/2</code>	Translate message-term to string
<code>meta_predicate/1</code>	Quintus compatibility
<code>module/1</code>	Query/set current type-in module
<code>module/2</code>	Declare a module
<code>module_transparent/1</code>	Indicate module based meta predicate
<code>msort/2</code>	Sort, do not remove duplicates
<code>multifile/1</code>	Indicate distributed definition of predicate
<code>mutex_create/1</code>	Create a thread-synchronisation device
<code>mutex_destroy/1</code>	Destroy a mutex
<code>mutex_lock/1</code>	Become owner of a mutex
<code>mutex_statistics/0</code>	Print statistics on mutex usage
<code>mutex_trylock/1</code>	Become owner of a mutex (non-blocking)
<code>mutex_unlock/1</code>	Release ownership of mutex
<code>mutex_unlock_all/0</code>	Release ownership of all mutexes
<code>name/2</code>	Convert between atom and list of character codes
<code>nl/0</code>	Generate a newline
<code>nl/1</code>	Generate a newline on a stream
<code>nodebug/0</code>	Disable debugging
<code>nodebug/1</code>	Disable debug-topic
<code>noguitracer/0</code>	Disable the graphical debugger

nonvar/1	Type check for bound term
noprofile/1	Hide (meta) predicate for the profiler
noprotocol/0	Disable logging of user interaction
nospy/1	Remove spy point
nospyall/0	Remove all spy points
not/1	Negation by failure (argument not provable). Same as <code>\+/1</code>
notrace/0	Stop tracing
notrace/1	Do not debug argument goal
nth_clause/3	N-th clause of a predicate
number/1	Type check for integer or float
number_chars/2	Convert between number and one-char atoms
number_codes/2	Convert between number and character codes
numbervars/3	Number unbound variables of a term
numbervars/4	Number unbound variables of a term
on_signal/3	Handle a software signal
once/1	Call a goal deterministically
op/3	Declare an operator
open/3	Open a file (creating a stream)
open/4	Open a file (creating a stream)
open_dde_conversation/3	Win32: Open DDE channel
open_null_stream/1	Open a stream to discard output
open_resource/3	Open a program resource as a stream
open_shared_object/2	UNIX: Open shared library (.so file)
open_shared_object/3	UNIX: Open shared library (.so file)
pce_dispatch/1	Run XPCE GUI in separate thread
pce_call/1	Run goal in XPCE GUI thread
peek_byte/1	Read byte without removing
peek_byte/2	Read byte without removing
peek_char/1	Read character without removing
peek_char/2	Read character without removing
peek_code/1	Read character-code without removing
peek_code/2	Read character-code without removing
phrase/2	Activate grammar-rule set
phrase/3	Activate grammar-rule set (returning rest)
please/3	Query/change environment parameters
plus/3	Logical integer addition
portray/1	(hook) Modify behaviour of <code>print/1</code>
portray_clause/1	Pretty print a clause
portray_clause/2	Pretty print a clause to a stream
predicate_property/2	Query predicate attributes
predsort/3	Sort, using a predicate to determine the order
preprocessor/2	Install a preprocessor before the compiler
print/1	Print a term
print/2	Print a term on a stream
print_message/2	Print message from (exception) term
print_message_lines/3	Print message to stream
profile/1	Obtain execution statistics

profile/3	Obtain execution statistics
profile_count/3	Obtain profile results on a predicate
profiler/2	Obtain/change status of the profiler
prolog/0	Run interactive toplevel
prolog_choice_attribute/3	Examine the choice-point stack
prolog_current_frame/1	Reference to goal's environment stack
prolog_edit:locate/2	Locate targets for <code>edit/1</code>
prolog_edit:locate/3	Locate targets for <code>edit/1</code>
prolog_edit:edit_source/1	Call editor for <code>edit/1</code>
prolog_edit:edit_command/2	Specify editor activation
prolog_edit:load/0	Load <code>edit/1</code> extensions
prolog_file_type/2	Define meaning of file extension
prolog_frame_attribute/3	Obtain information on a goal environment
prolog_ide/1	Program access to the development environment
prolog_list_goal/1	(hook) Intercept tracer 'L' command
prolog_load_context/2	Context information for directives
prolog_load_file/2	(hook) Program <code>load_files/2</code>
prolog_skip_level/2	Indicate deepest recursion to trace
prolog_to_os_filename/2	Convert between Prolog and OS filenames
prolog_trace_interception/4	<code>user</code> Intercept the Prolog tracer
prompt/1/1	Change prompt for 1 line
prompt/2	Change the prompt used by <code>read/1</code>
protocol/1	Make a log of the user interaction
protocola/1	Append log of the user interaction to file
protocolling/1	On what file is user interaction logged
put/1	Write a character
put/2	Write a character on a stream
put_byte/1	Write a byte
put_byte/2	Write a byte on a stream
put_char/1	Write a character
put_char/2	Write a character on a stream
put_code/1	Write a character-code
put_code/2	Write a character-code on a stream
qcompile/1	Compile source to Quick Load File
qsave_program/1	Create runtime application
qsave_program/2	Create runtime application
read/1	Read Prolog term
read/2	Read Prolog term from stream
read_clause/1	Read clause
read_clause/2	Read clause from stream
read_history/6	Read using history substitution
read_link/3	Read a symbolic link
read_term/2	Read term with options
read_term/3	Read term with options from stream
recorda/2	Record term in the database (first)
recorda/3	Record term in the database (first)
recorded/2	Obtain term from the database

recorded/3	Obtain term from the database
recordz/2	Record term in the database (last)
recordz/3	Record term in the database (last)
redefine_system_predicate/1	Abolish system definition
reload_library_index/0	Force reloading the autoload index
rename_file/2	Change name of file
repeat/0	Succeed, leaving infinite backtrack points
require/1	This file requires these predicates
reset_profiler/0	Clear statistics obtained by the profiler
resource/3	Declare a program resource
retract/1	Remove clause from the database
retractall/1	Remove unifying clauses from the database
same_file/2	Succeeds if arguments refer to same file
see/1	Change the current input stream
seeing/1	Query the current input stream
seek/4	Modify the current position in a stream
seen/0	Close the current input stream
set_input/1	Set current input stream from a stream
set_output/1	Set current output stream from a stream
set_prolog_flag/2	Define a system feature
set_stream/2	Set stream attribute
set_stream_position/2	Seek stream to position
set_tty/2	Set 'tty' stream
setarg/3	Destructive assignment on term
setenv/2	Set shell environment variable
setof/3	Find all unique solutions to a goal
sformat/2	Format on a string
sformat/3	Format on a string
shell/0	Execute interactive subshell
shell/1	Execute OS command
shell/2	Execute OS command
show_profile/1	Show results of the profiler
show_profile/2	Show results of the profiler
size_file/2	Get size of a file in characters
skip/1	Skip to character in current input
skip/2	Skip to character on stream
rl_add_history/1	Add line to readline(3) history
rl_read_init_file/1	Read readline(3) init file
sleep/1	Suspend execution for specified time
sort/2	Sort elements in a list
source_file/1	Examine currently loaded source files
source_file/2	Obtain source file of predicate
source_location/2	Location of last read term
spy/1	Force tracer on specified predicate
stack_parameter/4	Some systems: Query/Set runtime stack parameter
statistics/0	Show execution statistics
statistics/2	Obtain collected statistics

<code>stream_property/2</code>	Get stream properties
<code>string/1</code>	Type check for string
<code>string_concat/3</code>	<code>atom_concat/3</code> for strings
<code>string_length/2</code>	Determine length of a string
<code>string_to_atom/2</code>	Conversion between string and atom
<code>string_to_list/2</code>	Conversion between string and list of character codes
<code>style_check/1</code>	Change level of warnings
<code>sub_atom/5</code>	Take a substring from an atom
<code>sublist/3</code>	Determine elements that meet condition
<code>sub_string/5</code>	Take a substring from a string
<code>succ/2</code>	Logical integer successor relation
<code>writef/2</code>	Formatted write on a string
<code>writef/3</code>	Formatted write on a string
<code>tab/1</code>	Output number of spaces
<code>tab/2</code>	Output number of spaces on a stream
<code>tell/1</code>	Change current output stream
<code>telling/1</code>	Query current output stream
<code>term_expansion/2</code>	(hook) Convert term before compilation
<code>term_to_atom/2</code>	Convert between term and atom
<code>thread_at_exit/1</code>	Register goal to be called at exit
<code>thread_create/3</code>	Create a new Prolog task
<code>thread_detach/1</code>	Make thread cleanup after completion
<code>thread_exit/1</code>	Terminate Prolog task with value
<code>thread_get_message/1</code>	Wait for message
<code>thread_get_message/2</code>	Wait for message in a queue
<code>thread_join/2</code>	Wait for Prolog task-completion
<code>thread_local/1</code>	Declare thread-specific clauses for a predicate
<code>thread_peek_message/1</code>	Test for message
<code>thread_peek_message/2</code>	Test for message in a queue
<code>thread_self/1</code>	Get identifier of current thread
<code>thread_send_message/2</code>	Send message to another thread
<code>thread_setconcurrency/2</code>	Number of active threads
<code>thread_signal/2</code>	Execute goal in another thread
<code>thread_statistics/3</code>	Get statistics of another thread
<code>threads/0</code>	List running threads
<code>throw/1</code>	Raise an exception (see <code>catch/3</code>)
<code>time/1</code>	Determine time needed to execute goal
<code>time_file/2</code>	Get last modification time of file
<code>tmp_file/2</code>	Create a temporary filename
<code>told/0</code>	Close current output
<code>trace/0</code>	Start the tracer
<code>trace/1</code>	Set trace-point on predicate
<code>trace/2</code>	Set/Clear trace-point on ports
<code>tracing/0</code>	Query status of the tracer
<code>trim_stacks/0</code>	Release unused memory resources
<code>true/0</code>	Succeed
<code>tty_get_capability/3</code>	Get terminal parameter

tty_goto/2	Goto position on screen
tty_put/2	Write control string to terminal
tty_size/2	Get row/column size of the terminal
ttyflush/0	Flush output on terminal
unify_with_occurs_check/2	Logically sound unification
unix/1	OS interaction
unknown/2	Trap undefined predicates
unload_foreign_library/1	shlib Detach shared library (.so file)
unsetenv/1	Delete shell environment variable
upcase_atom/2	Convert atom to upper-case
use_module/1	Import a module
use_module/2	Import predicates from a module
var/1	Type check for unbound variable
visible/1	Ports that are visible in the tracer
volatile/1	Predicates that are not saved
wait_for_input/3	Wait for input with optional timeout
wildcard_match/2	Csh(1) style wildcard match
win_exec/2	Win32: spawn Windows task
win_has_menu/0	Win32: true if console menu is available
win_insert_menu/2	plwin.exe: add menu
win_insert_menu_item/4	plwin.exe: add item to menu
win_shell/2	Win32: open document through Shell
win_registry_get_value/3	Win32: get registry value
win_window_pos/1	Win32: change size and position of window
window_title/2	Win32: change title of window
with_mutex/2	Run goal while holding mutex
working_directory/2	Query/change CWD
write/1	Write term
write/2	Write term to stream
writeln/1	Write term, followed by a newline
write_canonical/1	Write a term with quotes, ignore operators
write_canonical/2	Write a term with quotes, ignore operators on a stream
write_term/2	Write term with options
write_term/3	Write term with options to stream
writeln/1	Formatted write
writeln/2	Formatted write on stream
writelnq/1	Write term, insert quotes
writelnq/2	Write term, insert quotes on stream

E.2 Library predicates

E.2.1 lists

append/3	Concatenate lists
delete/3	Delete all matching members from a list
flatten/2	Transform nested list into flat list
intersection/3	Set intersection
is_set/1	Type check for a set
list_to_set/2	Remove duplicates
member/2	Element is member of a list
nextto/3	Y follows X in List
nth0/3	N-th element of a list (0-based)
nth1/3	N-th element of a list (1-based)
numlist/3	Create list of integers in interval
permutation/2	Test/generate permutations of a list
reverse/2	Inverse the order of the elements in a list
select/3	Select element of a list
subset/2	Check subset relation for unordered sets
subtract/3	Delete elements that do not satisfy condition
sumlist/2	Add all numbers in a list
union/3	Union of two sets

E.2.2 check

check/0	Program completeness and consistency
list_undefined/0	List undefined predicates
list_autoload/0	List predicates that require autoload
list_redefined/0	List locally redefined predicates

E.2.3 readutil

read_line_to_codes/2	Read line from a stream
read_line_to_codes/3	Read line from a stream
read_stream_to_codes/2	Read contents of stream
read_stream_to_codes/3	Read contents of stream
read_file_to_codes/3	Read contents of file
read_file_to_terms/3	Read contents of file to Prolog terms

E.2.4 netscape

www_open_url/1	Open a web-page in a browser
----------------	------------------------------

E.2.5 registry

registry_get_key/2	Get principal value of key
--------------------	----------------------------

registry_get_key/3	Get associated value of key
registry_set_key/2	Set principal value of key
registry_set_key/3	Set associated value of key
registry_delete_key/1	Remove a key
shell_register_file_type/4	Register a file-type
shell_register_dde/6	Register DDE action
shell_register_prolog/1	Register Prolog

E.2.6 url

parse_url/2	Analyse or construct a URL
parse_url/3	Analyse or construct a relative URL
global_url/3	Make relative URL global
http_location/2	Analyse or construct location
www_form_encode/2	Encode or decode form-data

E.3 Arithmetic Functions

<code>*/2</code>	Multiplication
<code>**/2</code>	Power function
<code>+/2</code>	Addition
<code>-/1</code>	Unary minus
<code>-/2</code>	Subtraction
<code>//2</code>	Division
<code>//2</code>	Integer division
<code>/\2</code>	Bitwise and
<code><</2</code>	Bitwise left shift
<code>>>/2</code>	Bitwise right shift
<code>./2</code>	List of one character: character code
<code>\1</code>	Bitwise negation
<code>\/2</code>	Bitwise or
<code>^/2</code>	Power function
<code>abs/1</code>	Absolute value
<code>acos/1</code>	Inverse (arc) cosine
<code>asin/1</code>	Inverse (arc) sine
<code>atan/1</code>	Inverse (arc) tangent
<code>atan/2</code>	Rectangular to polar conversion
<code>ceil/1</code>	Smallest integer larger than arg
<code>ceiling/1</code>	Smallest integer larger than arg
<code>cos/1</code>	Cosine
<code>cputime/0</code>	Get CPU time
<code>e/0</code>	Mathematical constant
<code>exp/1</code>	Exponent (base e)
<code>float/1</code>	Explicitly convert to float
<code>float_fractional_part/1</code>	Fractional part of a float
<code>float_integer_part/1</code>	Integer part of a float
<code>floor/1</code>	Largest integer below argument
<code>integer/1</code>	Round to nearest integer
<code>log/1</code>	Natural logarithm
<code>log10/1</code>	10 base logarithm
<code>max/2</code>	Maximum of two numbers
<code>min/2</code>	Minimum of two numbers
<code>mod/2</code>	Remainder of division
<code>random/1</code>	Generate random number
<code>rem/2</code>	Remainder of division
<code>round/1</code>	Round to nearest integer
<code>truncate/1</code>	Truncate float to integer
<code>pi/0</code>	Mathematical constant
<code>sign/1</code>	Extract sign of value
<code>sin/1</code>	Sine
<code>sqrt/1</code>	Square root
<code>tan/1</code>	Tangent

xor/2

Bitwise exclusive or

E.4 Operators

\$	1	fx	Bind toplevel variable
^	200	xfy	Predicate
^	200	xfy	Arithmetic function
mod	300	xfx	Arithmetic function
*	400	yfx	Arithmetic function
/	400	yfx	Arithmetic function
//	400	yfx	Arithmetic function
<<	400	yfx	Arithmetic function
>>	400	yfx	Arithmetic function
xor	400	yfx	Arithmetic function
+	500	fx	Arithmetic function
-	500	fx	Arithmetic function
?	500	fx	XPCE: obtainer
\	500	fx	Arithmetic function
+	500	yfx	Arithmetic function
-	500	yfx	Arithmetic function
/\	500	yfx	Arithmetic function
\/	500	yfx	Arithmetic function
:	600	xfy	module:term separator
<	700	xfx	Predicate
=	700	xfx	Predicate
= . .	700	xfx	Predicate
:=	700	xfx	Predicate
<	700	xfx	Predicate
==	700	xfx	Predicate
@=	700	xfx	Predicate
=\=	700	xfx	Predicate
>	700	xfx	Predicate
>=	700	xfx	Predicate
@<	700	xfx	Predicate
@=<	700	xfx	Predicate
@>	700	xfx	Predicate
@>=	700	xfx	Predicate
is	700	xfx	Predicate
\=	700	xfx	Predicate
\==	700	xfx	Predicate
@@=	700	xfx	Predicate
not	900	fy	Predicate
\+	900	fy	Predicate
,	1000	xfy	Predicate
->	1050	xfy	Predicate
*->	1050	xfy	Predicate
;	1100	xfy	Predicate
	1100	xfy	Predicate

discontiguous	1150	<i>fx</i>	Predicate
dynamic	1150	<i>fx</i>	Predicate
module_transparent	1150	<i>fx</i>	Predicate
meta_predicate	1150	<i>fx</i>	Head
multifile	1150	<i>fx</i>	Predicate
thread_local	1150	<i>fx</i>	Predicate
volatile	1150	<i>fx</i>	Predicate
initialization	1150	<i>fx</i>	Predicate
: -	1200	<i>fx</i>	Introduces a directive
? -	1200	<i>fx</i>	Introduces a directive
- - >	1200	<i>xfx</i>	DCGrammar: rewrite
: -	1200	<i>xfx</i>	head : - body. separator

Bibliography

- [Anjewierden & Wielemaker, 1989] A. Anjewierden and J. Wielemaker. Extensible objects. ESPRIT Project 1098 Technical Report UvA-C1-TR-006a, University of Amsterdam, March 1989.
- [BIM, 1989] *BIM Prolog release 2.4*. Everberg, Belgium, 1989.
- [Bowen & Byrd, 1983] D. L. Bowen and L. M. Byrd. A portable Prolog compiler. In L. M. Pereira, editor, *Proceedings of the Login Programming Workshop 1983*, Lisbon, Portugal, 1983. Universidade nova de Lisboa.
- [Bratko, 1986] I. Bratko. *Prolog Programming for Artificial Intelligence*. Addison-Wesley, Reading, Massachusetts, 1986.
- [Butenhof, 1997] David R. Butenhof. *Programming with POSIX threads*. Addison-Wesley, Reading, MA, USA, 1997.
- [Clocksin & Melish, 1987] W. F. Clocksin and C. S. Melish. *Programming in Prolog*. Springer-Verlag, New York, Third, Revised and Extended edition, 1987.
- [Deransart *et al.*, 1996] P. Deransart, A. Ed-Dbali, and L. Cervoni. *Prolog: The Standard*. Springer-Verlag, New York, 1996.
- [Graham *et al.*, 1982] Susan L. Graham, Peter B. Kessler, and Marshall K. McKusick. gprof: a call graph execution profiler. In *SIGPLAN Symposium on Compiler Construction*, pages 120–126, 1982.
- [Hodgson, 1998] Jonathan Hodgson. validation suite for conformance with part 1 of the standard, 1998, <http://www.sju.edu/~jhodgson/pub/suite.tar.gz>.
- [Kernighan & Ritchie, 1978] B. W. Kernighan and D. M. Ritchie. *The C Programming Language*. Prentice-Hall, Englewood Cliffs, New Jersey, 1978.
- [O’Keefe, 1990] R. A. O’Keefe. *The Craft of Prolog*. MIT Press, Massachusetts, 1990.
- [Pereira, 1986] F. Pereira. *C-Prolog User’s Manual*, 1986.
- [Qui, 1997] *Quintus Prolog, User Guide and Reference Manual*. Berkhamsted, UK, 1997.
- [Sterling & Shapiro, 1986] L. Sterling and E. Shapiro. *The Art of Prolog*. MIT Press, Cambridge, Massachusetts, 1986.

Index

'MANUAL' library, 23
-disable-segv-handling, 208
-lpl library, 210
-lreadline library, 223
.pl, 46
.pro, 46
=: =/2, 111
/\2, 113
=\2, 111
|/2, 70
, /2, 70
!/0, 69
!/1, 79
/, 46
/ /2, 112
. /2, 113
=/2, 68
==/2, 68
> =/2, 111
> /2, 111
^/2, 114
/ /2, 112
-> /2, 70
=< /2, 111
<< /2, 113
< /2, 111
- /1, 112
- /2, 112
\ =/2, 68
\ /1, 114
\ =/2, 68
\ + /1, 70
\ / /2, 113
+ /2, 112
* * /2, 114
> > /2, 113
i /2, 70
* - > /2, 70
=@ =/2, 68
\ =@ =/2, 68
@ > =/2, 69
@ > /2, 69
* /2, 112
@ = < /2, 69
@ < /2, 69
= . . /2, 101
_PL_get_arg(), 180
\, 46
abolish/1, 13, 80
abolish/2, 80, 168
abolish/[1
2], 35
abort/0, 22, 27, 32, 88, 90, 130, 161, 199, 201
abs/1, 112
absolute_file_name/2, 127
absolute_file_name/3, 127
absolute_file_name/2, 17, 59, 62, 128, 129, 221
absolute_file_name/3, 36, 38, 58, 62, 127, 128,
168, 198, 229, 230
absolute_file_name/[2
3], 37, 61, 128
access_file/2, 126
access_file/2, 34, 128
acos/1, 114
add_import_module/3, 150
add_import_module/3, 150
address/2, 160
Alpha
DEC, 13
append/1, 91, 92
append/3, 104, 225
apply/2, 71
apropos/1, 24, 39, 238, 248
arg/3, 101
arithmetic_function/1, 193
arithmetic_function/1, 115
arithmetic_function/1, 115
asin/1, 114
assert/1, 59, 61, 80–83, 85, 147, 151, 159, 240
assert/2, 81, 82, 86
asserta/1, 22, 61, 81
asserta/2, 81, 82
assertions, 229

- assertz/1, 81, 240
- assertz/2, 81, 82
- assume/1, 229
- at_end_of_stream/0, 96
- at_end_of_stream/1, 97
- at_halt/1, 63
- at_initialization/1, 63
- at_end_of_stream/[0
1], 89, 229
- at_halt/1, 63, 130, 155, 202, 208
- at_initialization/1, 64, 206
- atan/1, 114
- atan/2, 114
- atom/1, 67, 176
- atom_char/2, 103
- atom_chars/2, 15, 103
- atom_codes/2, 103
- atom_concat/3, 104, 257
- atom_length/2, 105
- atom_number/2, 103
- atom_prefix/2, 105
- atom_to_term/3, 104
- atom_chars/2, 34, 58, 95, 103, 107
- atom_codes/2, 15, 34, 58, 103, 107
- atom_concat/3, 104, 108
- atom_length/2, 35, 108
- atom_to_term/3, 97
- atomic/1, 67
- attach_console/0, 161
- attach_console/0, 161, 199
- autoload/0, 61, 217, 218, 228

- backcomp *library*, 14, 15
- bagof/3, 63, 117, 248
- between/3, 110, 226
- block/3, 78, 248, 251
- break/0, 22, 27, 130, 199

- call/1, 14, 63, 67, 71, 73, 132, 136, 137, 189
- call/2, 71
- call/[2-6], 71
- call_cleanup/2, 73
- call_cleanup/3, 72
- call_shared_object_function/2, 168
- call_with_depth_limit/2, 77
- call_with_depth_limit/3, 72
- call_cleanup/2, 73

- call_cleanup/3, 73, 160
- call_cleanup/[2-3], 160
- call_with_depth_limit/3, 72
- callable/1, 67
- catch/3, 13, 14, 73, 74, 78, 94, 99, 136, 155,
156, 235, 237, 257
- ceil/1, 113
- ceiling/1, 113
- char_code/2, 103
- char_conversion/2, 109
- char_type/2, 105
- char_code/2, 58
- char_conversion/2, 33, 109
- char_type/2, 107
- character_count/2, 94
- chat:inv_map_list/5, 136
- chdir/1, 129
- check *library*, 61, 85, 227, 259
- check/0, 61, 227
- check_old_select/0, 16
- checklist/2, 118
- checkselect *library*, 15
- clause/2, 70, 86
- clause/3, 81, 82, 86, 87, 219
- clause/[2
3], 35, 224
- clause_property/2, 87
- clause_property/2, 62, 235
- clib
 - package, 208
- close/1, 87, 89
- close/2, 89
- close_dde_conversation/1, 139
- close_shared_object/1, 168
- code_type/2, 107
- code_type/2, 105, 107
- commandline
 - arguments, 22
- compare/3, 69, 116, 196
- compile_predicates/1, 83
- compile_predicates/1, 83
- compiling/0, 64
- completion
 - TAB, 49
- compound/1, 67
- concat_atom/2, 104
- concat_atom/3, 104

- concat_atom/2, 104
- consult/1, 18, 19, 29, 39, 54, 58, 60, 64, 83, 99, 133
- context_module/1, 149
- context_module/1, 149, 190
- convert_time/2, 125
- convert_time/8, 125
- convert_time/2, 125
- convert_time/8, 125
- convert_time/[2, 8], 127
- copy_stream_data/2, 97
- copy_stream_data/3, 97
- copy_term/2, 102
- copy_term/2, 102
- cos/1, 114
- cputime/0, 114
- ctype *library*, 105
- current_arithmetic_function/1, 115
- current_atom/1, 84
- current_char_conversion/2, 109
- current_flag/1, 84
- current_foreign_library/2, 168
- current_format_predicate/2, 122
- current_functor/2, 84
- current_input/1, 93
- current_key/1, 84
- current_module/1, 149
- current_module/2, 149
- current_mutex/3, 161
- current_op/3, 109
- current_output/1, 93
- current_predicate/1, 85
- current_predicate/2, 84
- current_prolog_flag/2, 32
- current_signal/3, 78
- current_stream/3, 15, 90, 129
- current_thread/2, 155
- current_atom/1, 84
- current_char_conversion/2, 109
- current_input/1, 62, 92
- current_output/1, 92
- current_predicate/1, 85
- current_predicate/2, 85
- current_prolog_flag/2, 15, 21, 22, 32, 39, 41, 42, 58, 59, 68, 74, 75, 97, 98, 100, 130, 133, 140, 167, 169, 189, 204, 209, 225, 242
- current_signal/3, 77
- current_stream/3, 90
- current_thread/2, 154, 155

- DCG, 59, 79
- dde_current_connection/2, 141
- dde_current_service/2, 141
- dde_execute/2, 140
- dde_poke/4, 140
- dde_register_service/2, 140
- dde_request/3, 139
- dde_unregister_service/1, 141
- debug *library*, 228
- debug/0, 26, 29, 74, 132, 199
- debug/1, 228
- debug/3, 228
- debugging
 - exceptions, 74
- debugging/0, 39, 132, 228, 238
- debugging/1, 228
- DEC
 - Alpha, 13
- default_module/2, 150
- delete/3, 225
- delete_directory/1, 129
- delete_file/1, 127
- delete_import_module/2, 150
- delete_import_module/2, 150
- deterministic/1, 235
- Development environment, 46
- discontiguous/1, 83
- display/1, 119, 181, 182
- display/[1, 2], 14
- displayq/1, 120
- displayq/[1, 2], 14
- dld, 167
- downcase_atom/2, 107
- downcase_atom/2, 106, 107
- dup/2, 15
- dup_stream/2, 15
- dup_stream/2, 15
- dwim_match/2, 141
- dwim_match/3, 141

- dwim_predicate/2, 86
- dwim_match/2, 86, 141
- dynamic/1, 36, 80, 83, 85, 149, 159, 227
- e/0, 114
- edit/0, 65
- edit/1, 16, 34, 39, 49, 50, 54, 55, 61, 65, 66, 227, 255
- edit_source/1, 66
- editor *class*, 49, 50
- Emacs, 23
- emacs/[0 1], 50
- emacs/prolog_colour *library*, 52
- emacs/prolog_mode *library*, 52
- emacs/swi_prolog *library*, 16
- ensure_loaded/1, 60
- ensure_loaded/1, 29, 58, 60, 145
- erase/1, 81, 82, 86
- eval_license/0, 246
- eval_license/0, 246
- exception/3, 39, 237
- exceptions
 - debugging, 74
- exists_directory/1, 127
- exists_file/1, 126
- exists_file/1, 34
- exit/2, 78
- exp/1, 112, 114
- expand_answer/2, 130
- expand_file_name/2, 128
- expand_file_search_path/2, 62
- expand_goal/2, 63
- expand_query/4, 130
- expand_term/2, 63
- expand_answer/2, 130
- expand_file_name/2, 34, 124, 127, 128
- expand_goal/2, 36, 63
- expand_term/2, 63, 79
- explain *library*, 251
- explain/1, 24
- explain/2, 24
- export/1, 149
- export_list/2, 149
- fail/0, 69
- fail/1, 79
- feature/2, 15
- file_base_name/2, 127
- file_directory_name/2, 127
- file_name_extension/3, 128
- file_search_path/2, 38, 61
- file_search_path/2, 19, 22, 33, 37, 39–41, 47, 60–62, 168, 198, 211, 220–222
- fileerrors/0, 94
- fileerrors/2, 34, 94
- findall/3, 63, 117, 151
- flag/3, 35, 82, 84
- flatten/2, 226
- float/1, 67, 111, 113
- float_fractional_part/1, 113
- float_integer_part/1, 113
- float_integer_part/1, 113
- floor/1, 113
- flush_output/0, 95
- flush_output/1, 95
- flush_output/0, 95
- flush_output/1, 75
- flush_output/[0 1], 88, 95
- foo/0, 241
- foo/3, 241
- forall/2, 63, 118
- format/1, 75, 120
- format/2, 120, 121, 228
- format/3, 75, 76, 121
- format/[1 2], 35, 97, 118, 252
- format/[2 3], 42
- format_predicate/2, 122
- free_variables/2, 102
- free_variables/2, 99
- FTP, 231
- functor/3, 9, 67, 101
- garbage_collect/0, 138
- garbage_collect_atoms/0, 138
- garbage_collect_atoms/0, 138, 202
- gensym/2, 142
- get/1, 96
- get/2, 96
- get0/1, 88, 96
- get0/2, 96

- get_byte/1, 95
- get_byte/2, 95
- get_char/1, 95
- get_char/2, 96
- get_code/1, 95
- get_code/2, 95
- get_single_char/1, 96
- get_time/1, 125
- get_byte/1, 96
- get_byte/2, 96
- get_byte/[1
2], 58
- get_char/1, 95, 96
- get_char/2, 96
- get_char/[1
2], 58
- get_code/1, 96
- get_code/2, 96
- get_code/[1
2], 58
- get_single_char/1, 22, 37
- get_time/1, 125, 127
- getenv/1, 230
- getenv/2, 124
- global_url/3, 232
- GNU-Emacs, 23
- go/0, 20
- goal_expansion/2, 39, 63
- goal_expansion/2, 59, 63, 228
- Graphics, 10
- ground/1, 67, 83
- GUI, 10
- guitracer/0, 16, 54, 56, 131

- halt/0, 28, 130
- halt/1, 130, 199, 249
- halt/[0
1], 63
- hash_term/2, 83
- hash_term/2, 82
- help/0, 23, 39, 221, 238
- help/1, 23, 24, 38, 39, 238
- helpidx *library*, 23
- hooks, 38
- HTTP, 231
- http_load *library*, 60, 239
- http_location/2, 233

- IDE, 46
- ignore/1, 63, 71, 136
- immediate
 - update view, 82
- import/1, 144, 145, 149
- import_module/2, 150
- import_module/2, 150
- include/1, 58, 60
- include/2, 59
- index/1, 82, 84, 86
- initialization/1, 64, 167, 201, 219
- install/0, 219
- int_to_atom/2, 104
- int_to_atom/3, 104
- integer/1, 67, 113
- interactor/0, 90, 161
- intersection/3, 227
- is/2, 35, 111, 113, 115, 226
- is_absolute_file_name/1, 128
- is_list/1, 115
- is_set/1, 226
- is_list/1, 226

- keysort/2, 116

- last/2, 226
- leash/1, 26, 132, 133, 236
- length/2, 116
- library_directory/1, 38, 61
- library_directory/1, 39, 40, 60
- library_directory/2, 41
- license/1, 247
- license/2, 246, 247
- likes/2, 18
- limit_stack/2, 138
- limit_stack/2, 137
- line_count/2, 94
- line_position/2, 94
- line_count/2, 91, 122
- line_position/2, 91, 122
- list_autoload/0, 228
- list_debug_topics/0, 228
- list_redefined/0, 228
- list_to_set/2, 226
- list_undefined/0, 227
- list_autoload/0, 227
- list_debug_topics/0, 228

- list_redefined/0, 227
- list_undefined/0, 61, 227, 228
- listing/0, 66
- listing/1, 27, 66
- lists *library*, 115, 225, 259
- load_files/2, 59, 255
- load_foreign_library/1, 168
- load_foreign_library/2, 168
- load_files/1, 39
- load_files/2, 39, 59, 60, 238, 239
- load_foreign_library/1, 209, 219
- load_foreign_library/2, 168
- load_foreign_library/[1
2], 61, 167
- log/1, 114
- log10/1, 114
- logical
 - update view, 82
- MacOS X, 10
- main/0, 32
- make/0, 9, 40, 41, 49, 54, 55, 61
- make_directory/1, 129
- make_library_index/1, 40
- make_library_index/2, 40
- make_library_index/1, 40
- make_library_index/2, 40
- make_library_index/[1
2], 41
- manpce/0, 43
- maplist/3, 118, 147, 218
- max/2, 112
- member/2, 27, 73, 86, 116, 128, 225, 253
- memberchk/2, 116
- memory
 - layout, 43
- merge/3, 116
- merge_set/3, 116
- message_hook/3, 38, 76
- message_queue_create/1, 157
- message_queue_destroy/1, 157
- message_to_string/2, 76
- message_hook/3, 14, 75, 76
- message_queue_create/1, 157
- message_queue_create/2, 156
- message_to_string/2, 75, 76
- meta_predicate/1, 149, 151
- meta_predicate/1, 150, 151
- min/2, 112
- mod/2, 112
- module/1, 150
- module/2, 63, 144, 149
- module_transparent/1, 149
- module_transparent/1, 85, 150, 190, 241
- msort/2, 116
- multi_file/1, 242
- multifile/1, 65, 83, 85, 227, 238
- mutex_create/1, 160
- mutex_destroy/1, 160
- mutex_lock/1, 160
- mutex_statistics/0, 156
- mutex_trylock/1, 161
- mutex_unlock/1, 161
- mutex_unlock_all/0, 161
- mutex_create/1, 160
- mutex_create/2, 160
- mutex_lock/1, 161
- name/2, 103, 104
- netmask/4, 205
- netscape *library*, 230, 259
- nextto/3, 225
- nl/0, 94
- nl/1, 94
- nl/[0
1], 119
- nodebug/0, 132
- nodebug/1, 228
- nofileerrors/0, 94
- noguitracer/0, 54, 56, 131
- nonvar/1, 67
- noprofile/1, 136
- noprotocol/0, 131
- nospy/1, 26, 39, 132, 238
- nospyall/0, 39, 132, 238
- not/1, 63, 71, 248
- notrace/0, 131
- notrace/1, 132
- nth0/3, 226
- nth1/3, 226
- nth_clause/3, 86
- nth_clause/3, 87, 235
- number/1, 67, 226
- number_chars/2, 15, 103

- number_codes/2, 103
- number_chars/2, 15, 58, 103
- number_codes/2, 15, 58, 103
- numbervars/3, 102
- numbervars/4, 102
- numlist/3, 226

- on_signal/3, 76
- on_signal/3, 14, 76, 78
- once/1, 63, 71, 72, 77, 132, 133, 160, 191
- online_help *library*, 249
- onsult/1, 60
- op/3, 15, 57, 83, 97, 109
- open/3, 34, 87, 88
- open/4, 12, 57, 87–90, 97, 229, 230
- open_dde_conversation/3, 139
- open_null_stream/1, 89
- open_resource/3, 222
- open_shared_object/2, 167
- open_shared_object/3, 167
- open_resource/3, 14, 217, 220, 222
- open_shared_object/2, 36, 167
- operator
 - and modules, 108

- package
 - clib, 208
- parse_url/2, 232
- parse_url/3, 232
- parse_url/2, 232, 233
- pce_call/1, 164
- pce_dispatch/1, 164
- pce_call/1, 164
- peek_byte/1, 96
- peek_byte/2, 96
- peek_char/1, 96
- peek_char/2, 96
- peek_code/1, 96
- peek_code/2, 96
- peek_byte/[1
 - 2], 58
- peek_char/[1
 - 2], 58
- peek_code/[1
 - 2], 58
- permutation/2, 226
- phrase/2, 79

- phrase/3, 79, 80
- pi/0, 114
- PL_abort_hook(), 201
- PL_abort_unhook(), 201
- PL_action(), 199
- PL_agc_hook(), 202
- PL_atom_chars(), 174
- PL_atom_nchars(), 181
- PL_call(), 191
- PL_call_predicate(), 191
- PL_chars_to_term(), 189
- PL_cleanup(), 208
- PL_close_foreign_frame(), 192
- PL_close_query(), 191
- PL_compare(), 196
- PL_cons_functor(), 184
- PL_cons_functor_v(), 184
- PL_cons_list(), 184
- PL_context(), 192
- PL_copy_term_ref(), 172
- PL_create_engine(), 163
- PL_cut_query(), 191
- PL_destroy_engine(), 163
- PL_discard_foreign_frame(), 192
- PL_dispatch_hook(), 201
- PL_erase(), 197
- PL_erase_external(), 198
- PL_exception(), 195
- PL_fail(), 173
- PL_foreign_context(), 174
- PL_foreign_context_address(), 174
- PL_foreign_control(), 174
- PL_free(), 215
- PL_functor_arity(), 176
- PL_functor_name(), 176
- PL_get_arg(), 180
- PL_get_atom(), 178
- PL_get_atom_chars(), 178
- PL_get_atom_nchars(), 180
- PL_get_bool(), 179
- PL_get_chars(), 178
- PL_get_file_name(), 198
- PL_get_float(), 179
- PL_get_functor(), 179
- PL_get_head(), 181
- PL_get_integer(), 179
- PL_get_list(), 181

- PL_get_list_chars(), 179
- PL_get_list_nchars(), 180
- PL_get_long(), 179
- PL_get_module(), 180
- PL_get_name_arity(), 180
- PL_get_nchars(), 180
- PL_get_nil(), 181
- PL_get_pointer(), 179
- PL_get_string_chars(), 178
- PL_get_tail(), 181
- PL_halt(), 208
- PL_handle_signals(), 196
- PL_initialise(), 206
- PL_install_readline(), 207
- PL_is_atom(), 177
- PL_is_atomic(), 178
- PL_is_compound(), 178
- PL_is_float(), 178
- PL_is_functor(), 178
- PL_is_initialised(), 207
- PL_is_integer(), 177
- PL_is_list(), 178
- PL_is_number(), 178
- PL_is_string(), 177
- PL_is_variable(), 177
- PL_license(), 247
- PL_load_extensions(), 200
- PL_malloc(), 215
- PL_module_name(), 193
- PL_new_atom(), 174
- PL_new_atom_nchars(), 181
- PL_new_functor(), 176
- PL_new_module(), 193
- PL_new_term_ref(), 171
- PL_new_term_refs(), 171
- PL_next_solution(), 191
- PL_on_halt(), 202
- PL_open_foreign_frame(), 192
- PL_open_query(), 190
- PL_pred(), 190
- PL_predicate(), 190
- PL_predicate_info(), 190
- PL_put_atom(), 183
- PL_put_atom_chars(), 183
- PL_put_atom_nchars(), 180
- PL_put_float(), 183
- PL_put_functor(), 183
- PL_put_integer(), 183
- PL_put_list(), 183
- PL_put_list_chars(), 183
- PL_put_list_nchars(), 180
- PL_put_list_ncodes(), 180
- PL_put_nil(), 184
- PL_put_pointer(), 183
- PL_put_string_chars(), 183
- PL_put_string_nchars(), 180, 183
- PL_put_term(), 184
- PL_put_variable(), 183
- PL_query(), 199
- PL_quote(), 189
- PL_raise_exception(), 194
- PL_realloc(), 215
- PL_record(), 197
- PL_record_external(), 197
- PL_recorded(), 197
- PL_recorded_external(), 198
- PL_register_atom(), 176
- PL_register_extensions(), 200
- PL_register_foreign(), 200
- PL_reset_term_refs(), 172
- PL_retry(), 174
- PL_retry_address(), 174
- PL_rewind_foreign_frame(), 192
- PL_same_compound(), 196
- PL_set_engine(), 164
- PL_signal(), 195
- PL_strip_module(), 192
- PL_succeed(), 173
- PL_term_type(), 177
- PL_thread_at_exit(), 163
- PL_thread_attach_engine(), 162
- PL_thread_destroy_engine(), 163
- PL_thread_self(), 162
- PL_throw(), 195
- PL_toplevel(), 207
- PL_unify(), 185
- PL_unify_arg(), 187
- PL_unify_atom(), 185
- PL_unify_atom_chars(), 185
- PL_unify_atom_nchars(), 180
- PL_unify_float(), 186
- PL_unify_functor(), 186
- PL_unify_integer(), 186
- PL_unify_list(), 186

- PL_unify_list_chars(), 185
- PL_unify_list_nchars(), 181
- PL_unify_list_ncodes(), 180
- PL_unify_nil(), 186
- PL_unify_pointer(), 186
- PL_unify_string_chars(), 185
- PL_unify_string_nchars(), 180, 186
- PL_unify_term(), 187
- PL_unregister_atom(), 176
- PL_warning(), 199
- plus/3, 71, 110
- PLVERSION, 215
- portray/1, 28, 38, 48, 97–99, 190, 203, 238
- portray_clause/1, 66
- portray_clause/2, 66
- portray_clause/1, 66
- portray_clause/2, 66
- predicate_property/2, 85
- predicate_property/2, 149
- predsort/3, 116
- preprocessor/2, 64
- print/1, 98, 99, 119, 121, 190, 254
- print/2, 98
- print/[1
2], 97
- print_message/2, 75, 253
- print_message_lines/3, 75
- print_message/2, 14, 37, 38, 59, 63, 75, 76, 99,
155, 228
- print_message_lines/3, 14, 75, 76
- profile file, 19
- profile/1, 135
- profile/3, 135, 223
- profiler/2, 135
- prolog/0, 22, 34, 130, 150, 207, 235
- prolog:debug_control_hook/1, 39, 238
- prolog:help_hook/1, 39, 238
- prolog:show_profile_hook/2, 135
- prolog_choice_attribute/3, 235
- prolog_current_frame/1, 234
- prolog_edit:edit_command/2, 39, 66
- prolog_edit:edit_source/1, 39, 65
- prolog_edit:load/0, 66
- prolog_edit:locate/2, 65
- prolog_edit:locate/3, 39, 65
- prolog_file_type/2, 62
- prolog_frame_attribute/3, 234
- prolog_ide class, 55
- prolog_ide/1, 55
- prolog_list_goal/1, 39, 238
- prolog_load_context/2, 62
- prolog_load_file/2, 39
- prolog_skip_level/2, 237
- prolog_to_os_filename/2, 129
- prolog_trace_interception/4, 39, 236
- prolog_choice_attribute/3, 236
- prolog_current_frame/1, 234
- prolog_edit:edit_source/1, 49, 55, 65
- prolog_edit:locate/3, 65
- prolog_file_type/2, 58, 62, 128
- prolog_frame_attribute/3, 87, 236
- prolog_frame_attribute/4, 236
- prolog_ide/1, 55
- prolog_load_context/2, 62
- prolog_load_file/2, 60, 238
- prolog_to_os_filename/2, 47, 128
- prolog_trace_interception/4, 54, 131, 234, 235
- prompt/2, 100, 101
- prompt1/1, 101
- proper_list/1, 115
- protocol/1, 131
- protocola/1, 131
- protocolling/1, 131
- prove/3, 49
- put/1, 94, 95
- put/2, 94, 95
- put_byte/1, 95
- put_byte/2, 95
- put_char/1, 95
- put_char/2, 95
- put_code/1, 95
- put_code/2, 95
- put_byte/[1
2], 58
- put_char/[1
2], 58
- put_code/[1
2], 58
- qcompile/1, 59, 60, 64
- qsave_program/1, 217
- qsave_program/2, 217, 218
- qsave_program/2, 13, 31, 34, 48, 55, 217, 220
- qsave_program/[1

- 2], 13, 14, 22, 31, 36, 64, 85, 167, 207, 211, 218, 219
- quiet, 21, 75
- quintus *library*, 15, 149, 151
- random/1, 113
- read/1, 32, 33, 35, 43, 88, 95, 98, 99, 101, 108, 133, 230, 255
- read/2, 93, 99
- read_clause/1, 99
- read_clause/2, 99
- read_file_to_codes/3, 229
- read_file_to_terms/3, 230
- read_history/6, 100
- read_line_to_codes/2, 229
- read_line_to_codes/3, 229
- read_link/3, 129
- read_stream_to_codes/2, 229
- read_stream_to_codes/3, 229
- read_term/2, 99
- read_term/3, 100
- read_clause/1, 99, 133
- read_history/6, 100
- read_line_to_codes/2, 229
- read_line_to_codes/3, 229
- read_stream_to_codes/2, 229
- read_term/2, 25, 32, 99, 100, 104
- read_term/3, 100, 109, 130
- read_term/[2
3], 99
- readutil *library*, 48, 229, 259
- reconsult/1, 59
- recorda/2, 81
- recorda/3, 81, 82, 84, 196, 197
- recorded/2, 82
- recorded/3, 81, 82, 151, 219
- recordz/2, 81, 151
- recordz/3, 81
- redefine_system_predicate/1, 81
- redefine_system_predicate/1, 12, 240
- registry, 43
- registry *library*, 230, 259
- registry_delete_key/1, 231
- registry_get_key/2, 230
- registry_get_key/3, 230
- registry_set_key/2, 230
- registry_set_key/3, 231
- reload_library_index/0, 41
- reload_library_index/0, 41
- rem/2, 112
- rename_file/2, 127
- repeat/0, 69, 72
- require/1, 60, 218
- reset_profiler/0, 135
- reset_profiler/0, 135
- resource/3, 14, 36, 39, 217, 218, 220–222
- retract/1, 59, 61, 80–83, 85, 151, 159
- retractall/1, 80, 81
- reverse/2, 144, 226
- RFC-1738, 231
- rl_add_history/1, 239
- rl_read_init_file/1, 239
- rlimit *library*, 16
- round/1, 113
- same_file/2, 127
- see/1, 14, 87, 91, 92
- seeing/1, 91, 92, 129
- seek/4, 89, 90
- seen/0, 92
- select/3, 15, 225
- set_feature/2, 15
- set_input/1, 93
- set_output/1, 93
- set_prolog_flag/2, 38
- set_stream/2, 90
- set_stream_position/2, 90
- set_tty/2, 123
- set_feature/2, 15
- set_input/1, 90, 92
- set_output/1, 92
- set_prolog_flag/2, 15, 24, 32, 110
- set_stream/2, 88, 93
- set_stream_position/2, 89, 90
- setarg/3, 101, 102
- setenv/2, 124
- setof/3, 63, 117, 248
- sformat/2, 122
- sformat/3, 97, 121
- shell/0, 123, 124
- shell/1, 47, 66, 123, 124
- shell/2, 123
- shell/[0-2], 124
- shell/[1

2], 123
 shell_register_dde/6, 231
 shell_register_file_type/4, 231
 shell_register_prolog/1, 231
 shell_register_file_type/4, 231
 shlib *library*, 250, 253, 258
 show_profile/1, 135
 show_profile/2, 135
 show_profile/1, 135
 sign/1, 112
 silent, 75
 sin/1, 112, 114
 size_file/2, 127
 skip/1, 96
 skip/2, 96
 sleep/1, 142
 socket *library*, 93
 Solaris, 155
 sort/2, 116, 117, 226
 source_file/1, 62
 source_file/2, 62
 source_location/2, 62
 source_file/2, 64, 86
 source_file/[1
 2], 149
 source_location/2, 62
 spy/1, 26, 33, 39, 54, 56, 132, 238, 251
 sqrt/1, 114
 stack
 memory management, 43
 stack_parameter/4, 138
 startup file, 19
 statistics *library*, 135
 statistics/0, 133
 statistics/2, 114, 133, 134, 156
 stream_property/2, 89
 stream_property/2, 90, 100
 string/1, 67, 121
 string_concat/3, 108
 string_length/2, 108
 string_to_atom/2, 108
 string_to_list/2, 108
 string_concat/3, 104
 style_check/1, 133
 style_check/1, 43, 45, 83, 243
 sub_atom/5, 105
 sub_string/5, 108

sub_atom/5, 108
 sublist/3, 118
 subset/2, 227
 subtract/3, 227
 succ/2, 110
 sumlist/2, 226
 swi/pce_profile *library*, 135
 swi_edit *library*, 66
 swi_help *library*, 23
 swritef/2, 120
 swritef/3, 119

TAB

 completion, 49
 tab/1, 95
 tab/2, 95
 tan/1, 114
 tell/1, 14, 87, 91, 92
 telling/1, 91, 92, 129
 term_expansion/2, 39, 63
 term_to_atom/2, 104
 term_expansion/2, 39, 59, 63, 64, 130
 term_to_atom/2, 97, 189
 thread_at_exit/1, 155
 thread_create/3, 153
 thread_detach/1, 155
 thread_exit/1, 155
 thread_get_message/1, 157
 thread_get_message/2, 157
 thread_join/2, 154
 thread_local/1, 159
 thread_peek_message/1, 157
 thread_self/1, 154
 thread_send_message/2, 156
 thread_setconcurrency/2, 155
 thread_signal/2, 158
 thread_statistics/3, 156
 thread_at_exit/1, 163
 thread_create/3, 155, 164
 thread_detach/1, 154
 thread_exit/1, 155, 156
 thread_get_message/1, 157
 thread_join/2, 154, 155
 thread_local/1, 83, 85, 159
 thread_peek_message/1, 157
 thread_self/1, 37, 155, 156
 thread_send_message/2, 157

- thread_signal/2, 158, 161
- threads/0, 161
- throw/1, 13, 26, 73, 78, 130, 156, 158, 159, 194, 195, 237
- time *library*, 77
- time/1, 114, 133, 135
- time_file/2, 127
- time_file/2, 125
- tmp_file/2, 129
- told/0, 92
- trace/0, 26, 54, 56, 131, 132, 159, 199
- trace/1, 33, 132
- trace/2, 132
- tracing/0, 131
- trim_stacks/0, 138
- trim_stacks/0, 34, 137, 138
- true/0, 36, 69, 72
- truncate/1, 113
- tty_get_capability/3, 122
- tty_goto/2, 122
- tty_put/2, 123
- tty_size/2, 123
- tty_get_capability/2, 123
- tty_get_capability/3, 123
- tty_goto/2, 123
- tty_put/2, 123
- tty_size/2, 123
- ttyflush/0, 95, 119

- unify_with_occurs_check/2, 68
- union/3, 227
- Unix, 10
- unix, 37
- unix/1, 15, 124
- unknown/2, 40, 133, 149, 225
- unload_foreign_library/1, 168
- unsetenv/1, 124
- upcase_atom/2, 107
- upcase_atom/2, 106
- update view, 82
- URL, 124
- url *library*, 230, 231, 260
- use_module/1, 145
- use_module/2, 145
- use_module/1, 39, 47
- use_module/2, 40
- use_module/[1
 - 2], 29, 54, 58–60, 144–146, 149, 241
- user *library*, 255
- user profile file, 19
- user:prolog_load_file/2, 239
- utf-8, 57

- var/1, 12, 67, 176
- verbose, 21
- visible/1, 133, 236
- volatile/1, 85, 159, 219

- wait_for_input/3, 93
- wait_for_input/3, 91, 93
- wildcard_match/2, 141
- win_exec/2, 123
- win_has_menu/0, 126
- win_insert_menu/2, 126
- win_insert_menu_item/4, 126
- win_registry_get_value/3, 124
- win_shell/2, 124
- win_window_pos/1, 126
- win_exec/2, 123
- win_insert_menu/2, 126
- win_insert_menu_item/4, 126
- win_shell/2, 123, 230
- Window interface, 10
- window_title/2, 125
- Windows, 10
- windows, 38
- with_mutex/2, 160
- with_mutex/2, 160, 161
- working_directory/2, 129
- working_directory/2, 124, 125, 128, 129
- write/1, 34, 98, 104, 105, 119, 121, 179, 181
- write/2, 98
- write_canonical/1, 98
- write_canonical/2, 98
- write_term/2, 97
- write_term/3, 98
- write_canonical/[1
 - 2], 14
- write_term/2, 26, 34, 37, 68, 97, 98, 108, 119, 121
- write_term/3, 38, 102
- write_term/[2
 - 3], 13, 35
- writeln/1, 118

writef/2, 17, 42, 97, 119
writef/[1
2], 118
writeln/1, 118
writeq/1, 98, 119, 121
writeq/2, 98
www_form_encode/2, 233
www_open_url/1, 230
www_open_url/1, 230

X-Windows, 10
X11, 10
xor/2, 113
XPCE, 10