SWI-Prolog C-library

Jan Wielemaker
HCS,
University of Amsterdam
The Netherlands
E-mail: `wielemak@science.uva.nl`

August 29, 2006

Abstract

This document describes commonly used foreign language extensions to SWI-Prolog distributed as a package known under the name *clib*. The package defines a number of Prolog libraries with accompagnying foreign libraries.

On Windows systems, the `unix` and `crypt` libraries can only be used if the whole SWI-Prolog suite is compiled using Cywin. The other libraries have been ported to native Windows.

Contents

1 Introduction

Many useful facilities offered by one or more of the operating systems supported by SWI-Prolog are not supported by the SWI-Prolog kernel distribution. Including these would enlarge the *footprint* and complicate portability matters while supporting only a limited part of the user-community.

This document describes `unix` to deal with the Unix process API, `socket` to deal with inet-domain stream-sockets, `cgi` to deal with getting CGI form-data if SWI-Prolog is used as a CGI scripting language and `crypt` to provide access to Unix password encryption.

2 Unix Process manipulation library

The `unix` library provides the commonly used Unix primitives to deal with process management. These primitives are useful for many tasks, including server management, parallel computation, exploiting and controlling other processes, etc.

The predicates are modelled closely after their native Unix counterparts. Higher-level primitives, especially to make this library portable to non-Unix systems are desirable. Using these primitives and considering that process manipulation is not a very time-critical operation we anticipate these libraries to be developed in Prolog.

fork(*-Pid***)**

Clone the current process into two branches. In the child, *Pid* is unified to `child`. In the original process, *Pid* is unified to the process identifier of the created child. Both parent and child are fully functional Prolog processes running the same program. The processes share open I/O streams that refer to Unix native streams, such as files, sockets and pipes. Data is not shared, though on most Unix systems data is initially shared and duplicated only if one of the programs attempts to modify the data.

Unix `fork()` is the only way to create new processes and `fork/2` is a simple direct interface to it.

exec(*+Command(...Args...)***)**

Replace the running program by starting *Command* using the given commandline arguments. Each command-line argument must be atomic and is converted to a string before passed to the Unix call `execvp()`.

Unix `exec()` is the only way to start an executable file executing. It is commonly used together with `fork/1`. For example to start `netscape` on an URL in the background, do:

```
run_netscape(URL) :-
        (     fork(child),
              exec(netscape(URL))
        ;     true
        ).
```

Using this code, netscape remains part of the process-group of the invoking Prolog process and Prolog does not wait for netscape to terminate. The predicate `wait/2` allows waiting for a child, while `detach_IO/0` disconnects the child as a deamon process.

wait(*-Pid, -Status***)**

Wait for a child to change status. Then report the child that changed status as well as the reason. *Status* is unified with `exited`(*ExitCode*) if the child with pid *Pid* was terminated by calling `exit()` (Prolog `halt/[0,1]`). *ExitCode* is the return=status. *Status* is unified with `signaled`(*Signal*) if the child died due to a software interrupt (see `kill/2`). *Signal* contains the signal number. Finally, if the process suspended execution due to a signal, *Status* is unified with `stopped`(*Signal*).

kill(*+Pid, +Signal***)**

Deliver a software interrupt to the process with identifier *Pid* using software-interrupt number *Signal*. See also `on_signal/2`. The meaning of the signal numbers can be found in the Unix manual.[1].

pipe(*-InSream, -OutStream***)**

Create a communication-pipe. This is normally used to make a child communicate to its parent. After `pipe/2`, the process is cloned and, depending on the desired direction, both processes close the end of the pipe they do not use. Then they use the remaining stream to communicate. Here is a simple example:

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

fork_demo(Result) :-
        pipe(Read, Write),
        fork(Pid),
        (   Pid == child
        ->  close(Read),
            format(Write, '˜q.˜n',
                    [hello(world)]),
            flush_output(Write),
            halt
        ;   close(Write),
            read(Read, Result),
            close(Read)
        ).
```

dup(*+FromStream, +ToStream***)**

Interface to Unix dup2(), copying the underlying filedescriptor and thus making both streams point to the same underlying object. This is normally used together with `fork/1` and `pipe/2` to talk to an external program that is designed to communicate using standard I/O.

Both *FromStream* and *ToStream* either refer to a Prolog stream or an integer descriptor number to refer directly to OS descriptors. See also `demo/pipe.pl` in the source-distribution of this package.

detach_IO

This predicate is intended to create Unix deamon-processes. It preforms two actions. First of

---

[1]kill/2 should support interrupt-names as well

all, the I/O streams `user_input`, `user_output` and `user_error` are closed and rebound to a Prolog stream that returns end-of-file on any attempt to read and starts writing to a file named `/tmp/pl-out.pid` (where ⟨*pid*⟩ is the process-id of the calling Prolog) on any attempt to write. This file is opened only if there is data available. This is intended for debugging purposes.[2] Finally, the process is detached from the current process-group and its controlling terminal.

3 File manipulation library

The `files` library provides additional operations on files from SWI-Prolog. It is currently very incomplete.

set_time_file(+*File, -OldTimes, +NewTimes***)**
> Query and set POSIX time attributes of a file. Both *OldTimes* and *NewTimes* are lists of option-terms. Times are represented in SWI-Prolog's standard floating point numbers. New times may be specified as `now` to indicate the current time. Defined options are:
>
> **access(***Time***)**
> > Describes the time of last access of the file. This value can be read and written.
>
> **modified(***Time***)**
> > Describes the time the contents of the file was last modified. This value can be read and written.
>
> **changed(***Time***)**
> > Describes the time the file-structure itself was changed by adding (link()) or removing (unlink()) names.
>
> Here are some example queries. The first retrieves the access-time, while the second sets the last-modified time to the current time.
>
> ```
> ?- set_time_file(foo, [acess(Access)], []).
> ?- set_time_file(foo, [], [modified(now)]).
> ```

4 Socket library

The `socket` library provides TCP and UDP inet-domain sockets from SWI-Prolog, both client and server-side communication. The interface of this library is very close to the Unix socket interface, also supported by the MS-Windows *winsock* API. SWI-Prolog applications that wish to communicate with multiple sources have three options:

1. Use I/O multiplexing based on `wait_for_input/3`. On Windows systems this can only be used for sockets, not for general (device-) file handles.

2. Use multiple threads, handling either a single blocking socket or a pool using I/O multiplexing as above.

---

[2]More subtle handling of I/O, especially for debugging is required: communicate with the syslog deamon and optionally start a debugging dialog on a newly created (X-)terminal should be considered.

3. Using XPCE's class *socket* which synchronises socket events in the GUI event-loop.

tcp_socket(-*SocketId*)
> Creates an `INET`-domain stream-socket and unifies an identifier to it with *SocketId*. On MS-Windows, if the socket library is not yet initialised, this will also initialise the library.

tcp_close_socket(+*SocketId*)
> Closes the indicated socket, making *SocketId* invalid. Normally, sockets are closed by closing both stream handles returned by `open_socket/3`. There are two cases where `tcp_close_socket/1` is used because there are no stream-handles:
>
> - After `tcp_accept/3`, the server does a `fork/1` to handle the client in a sub-process. In this case the accepted socket is not longer needed from the main server and must be discarded using `tcp_close_socket/1`.
> - If, after discovering the connecting client with `tcp_accept/3`, the server does not want to accept the connection, it should discard the accepted socket immediately using `tcp_close_socket/1`.

tcp_open_socket(+*SocketId, -InStream, -OutStream*)
> Open two SWI-Prolog I/O-streams, one to deal with input from the socket and one with output to the socket. If `tcp_bind/2` has been called on the socket. *OutSream* is useless and will not be created. After closing both *InStream* and *OutSream*, the socket itself is discarded.

tcp_bind(+*Socket, ?Port*)
> Bind the socket to *Port* on the current machine. This operation, together with `tcp_listen/2` and `tcp_accept/3` implement the *server*-side of the socket interface. If *Port* is unbound, the system picks an arbitrary free port and unifies *Port* with the selected port number.

tcp_listen(+*Socket, +Backlog*)
> Tells, after `tcp_bind/2`, the socket to listen for incoming requests for connections. *Backlog* indicates how many pending connection requests are allowed. Pending requests are requests that are not yet acknowledged using `tcp_accept/3`. If the indicated number is exceeded, the requesting client will be signalled that the service is currently not available. A suggested default value is 5.

tcp_accept(+*Socket, -Slave, -Peer*)
> This predicate waits on a server socket for a connection request by a client. On success, it creates a new socket for the client and binds the identifier to *Slave*. *Peer* is bound to the IP-address of the client.

tcp_connect(+*Socket, +Host:+Port*)
> Client-interface to connect a socket to a given *Port* on a given *Host*. After successful completion, `tcp_open_socket/3` can be used to create I/O-Streams to the remote socket.

tcp_setopt(+*Socket, +Option*)
> Set options on the socket. Defined options are:
>
> **reuseaddr**
>> Allow servers to reuse a port without the system being completely sure the port is no longer in use.

broadcast

UDP sockets only: broadcast the package to all addresses matching the address. The address is normally the address of the local subnet (i.e. 192.168.1.255). See `udp_send/4`.

dispatch(*Bool***)**

In GUI environments (using XPCE or the Windows plwin.exe executable) this flags defines whether or not any events are dispatched on behalf of the user interface. Default is `true`. Only very specific situations require setting this to `false`.

tcp_fcntl(+*Stream, +Action, ?Argument***)**

Interface to the Unix `fcntl()` call. Currently only suitable to deal switch stream to non-blocking mode using:

```
...
tcp_fcntl(Stream, setfl. nonblock),
...
```

As of SWI-Prolog 3.2.4, handling of non-blocking stream is supported. An attempt to read from a non-blocking stream returns -1 (or `end_of_file` for `read/1`), but `at_end_of_stream/1` fails. On actual end-of-input, `at_end_of_stream/1` succeeds.

tcp_host_to_address(*?HostName, ?Address***)**

Translate between a machines host-name and it's (IP-)address. If *HostName* is an atom, it is resolved using `gethostbyname()` and the IP-number is unified to *Address* using a term of the format ip(*Byte1, Byte2, Byte3, Byte4*). Otherwise, if *Address* is bound to a `ip/4` term, it is resolved by `gethostbyaddr()` and the canonical hostname is unified with *HostName*.

gethostname(-*Hostname***)**

Return the official fully qualified name of this host. This is achieved by calling gethostname() followed by gethostbyname() and return the official name of the host (`h_name`) of the structure returned by the latter function.

4.1 Server applications

The typical sequence for generating a server application is defined below:

```
create_server(Port) :-
        tcp_socket(Socket),
        tcp_bind(Socket, Port),
        tcp_listen(Socket, 5),
        tcp_open_socket(Socket, AcceptFd, _),
        <dispatch>
```

There are various options for ⟨*dispatch*⟩. One is to keep track of active clients and server-sockets using `wait_for_input/3`. If input arrives at a server socket, use `tcp_accept/3` and add the new connection to the active clients. Otherwise deal with the input from the client. Another is to use (Unix) `fork/1` to deal with the client in a separate process.

Using `fork/1`, ⟨*dispatch*⟩ may be implemented as:

```
dispatch(AcceptFd) :-
        tcp_accept(AcceptFd, Socket, _Peer),
        fork(Pid)
        (   Pid == child
        ->  tcp_open_socket(Socket, In, Out),
            handle_service(In, Out),
            close(In),
            close(Out),
            halt
        ;   tcp_close_socket(Socket)
        ),
        dispatch(AcceptFd).
```

4.2 Client applications

The skeleton for client-communication is given below.

```
create_client(Host, Port) :-
        tcp_socket(Socket),
        tcp_connect(Socket, Host:Port),
        tcp_open_socket(Socket, ReadFd, WriteFd),
        <handle I/O using the two streams>
        close(ReadFd),
        close(WriteFd).
```

To deal with timeouts and multiple connections, wait_for_input/3 and/or non-blocking streams (see tcp_fcntl/3) can be used.

4.3 The stream_pool library

The streampool library dispatches input from multiple streams based on wait_for_input/3. It is part of the clib package as it is used most of the time together with the socket library. On non-Unix systems it often can only be used with socket streams.

   With SWI-Prolog 5.1.x, multi-threading often provides a good alternative to using this library. In this schema one thread watches the listening socket waiting for connections and either creates a thread per connection or processes the accepted connections with a pool of *worker threads*. The library http/thread_httpd provides an example realising a mult-threaded HTTP server.

add_stream_to_pool(+*Stream, :Goal*)

   Add *Stream*, which must be an input stream and —on non-unix systems— connected to a socket to the pool. If input is available on *Stream*, *Goal* is called.

delete_stream_from_pool(+*Stream*)

   Delete the given stream from the pool. Succeeds, even if *Stream* is no member of the pool. If *Stream* is unbound the entire pool is emtied but unlike close_stream_pool/0 the streams are not closed.

close_stream_pool

Empty the pool, closing all streams that are part of it.

dispatch_stream_pool(+*TimeOut*)

Wait for maximum of *TimeOut* for input on any of the streams in the pool. If there is input, call the *Goal* associated with add_stream_to_pool/2. If *Goal* fails or raises an exception a message is printed. *TimeOut* is described with wait_for_input/3.

If *Goal* is called, there is *some* input on the associated stream. *Goal* must be careful not to block as this will block the entire pool.[3]

stream_pool_main_loop

Calls dispatch_stream_pool/1 in a loop until the pool is empty.

Below is a very simple example that reads the first line of input and echos it back.

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

server(Port) :-
        tcp_socket(Socket),
        tcp_bind(Socket, Port),
        tcp_listen(Socket, 5),
        tcp_open_socket(Socket, In, _Out),
        add_stream_to_pool(In, accept(Socket)),
        stream_pool_main_loop.

accept(Socket) :-
        tcp_accept(Socket, Slave, Peer),
        tcp_open_socket(Slave, In, Out),
        add_stream_to_pool(In, client(In, Out, Peer)).

client(In, Out, _Peer) :-
        read_line_to_codes(In, Command),
        close(In),
        format(Out, 'Please to meet you: ~s~n', [Command]),
        close(Out),
        delete_stream_from_pool(In).
```

4.4   UDP protocol support

The current library provides limited support for UDP packets. The UDP protocol is a *connection-less* and *unreliable* datagram based protocol. That means that messages sent may or may not arrive at the client side and may arrive in a different order as they are sent. UDP messages are often used for streaming media or for service discovery using the broadcasting mechanism.

udp_socket(-*Socket*)

Similar to tcp_socket/1, but create a socket using the SOCK_DGRAM protocol, ready for UDP connections.

---

[3]This is hard to achieve at the moment as none of the Prolog read-commands provide for a timeout.

udp_receive(+*Socket, -Data, -From, +Options*)

Wait for and return the next datagram. The data is returned as a Prolog string object (see `string_to_list/2`). *From* is a term of the format ip(*A,B,C,D*):*Port* indicating the sender of the message. *Options* is currently unused. *Socket* can be waited for using `wait_for_input/3`.

udp_send(+*Socket, +Data, +To, +Options*)

Send a UDP message. Data is a string, atom or code-list providing the data. *To* is an address of the form *Host:Port* where Host is either the hostname or a term `ip/4`. *Options* is currently unused.

A broadcast is achieved by using `tcp_setopt`(*Socket, broadcast*) prior to sending the datagram and using the local network broadcast address as a `ip/4` term.

5 CGI Support library

This is currently a very simple library, providing support for obtaining the form-data for a CGI script:

cgi_get_form(-*Form*)

Decodes standard input and the environment variables to obtain a list of arguments passed to the CGI script. This predicate both deals with the CGI **GET** method as well as the **POST** method. If the data cannot be obtained, an `existence_error` exception is raised.

Below is a very simple CGI script that prints the passed parameters. To test it, compile this program using the command below, copy it to your cgi-bin directory (or make it otherwise known as a CGI-script) and make the query `http://myhost.mydomain/cgi-bin/cgidemo?hello=world`

```
% pl -o cgidemo --goal=main --toplevel=halt -c cgidemo.pl

:- use_module(library(cgi)).

main :-
        cgi_get_form(Arguments),
        format('Content-type: text/html~n~n', []),
        format('<HTML>~n', []),
        format('<HEAD>~n', []),
        format('<TITLE>Simple SWI-Prolog CGI script</TITLE>~n', []),
        format('</HEAD>~n~n', []),
        format('<BODY>~n', []),
        format('<P>', []),
        print_args(Arguments),
        format('</BODY>~n</HTML>~n', []).

print_args([]).
print_args([A0|T]) :-
        A0 =.. [Name, Value],
        format('<B>~w</B>=<EM>~w</EM><BR>~n', [Name, Value]),
        print_args(T).
```

5.1 Some considerations

Printing an HTML document using `format/2` is not really a neat way of producing HTML. A high-level alternative is provided by `http/html_write` from the XPCE package.

6 MIME decoding library

MIME (Multipurpose Internet Mail Extensions) is a format for serializing multiple typed data objects. It was designed for E-mail, but it is also used for other applications such packaging multiple values using the HTTP POST request on web-servers. Double Precision, Inc. has produced the C-libraries rfc822 (mail) and rfc2045 (MIME) for decoding and manipulating MIME messages. The `mime` library is a Prolog wrapper around the rfc2045 library for decoding MIME messages.

The general name 'mime' is used for this library as it is anticipated to add MIME-creation functionality to this library.

Currently the mime library defines one predicate:

mime_parse(*Data, Parsed*)

> Parse *Data* and unify the result to *Parsed*. *Data* is one of:

> **stream**(*Stream*)
>> Parse the data from *Stream* upto the end-of-file.

> **stream**(*Stream, Length*)
>> Parse a maximum of *Length* characters from *Stream* or upto the end-of-file, whichever comes first.

> ***Text***
>> Atoms, strings, code- and character lists are treated as valid sources of data.

> *Parsed* is a tree structure of mime(*Attributes, Data, PartList*) terms. Currently either *Data* is the empty atom or *PartList* is an empty list.[4] *Data* is an atom holding the message data. The library automatically decodes *base64* and *quoted-printable* messages. See also the `transfer_encoding` attribute below.

> *PartList* is a list of `mime/3` terms. *Attributes* is a list holding a subset of the following arguments. For details please consult the RFC2045 document.

> **type**(*Atom*)
>> Denotes the Content-Type, how the *Data* should be interpreted.

> **transfer_encoding**(*Atom*)
>> How the *Data* was encoded. This is not very interesting as the library decodes the content of the message.

> **character_set**(*Atom*)
>> The character set used for text data. Note that SWI-Prolog's capabilities for character-set handling are limited.

> **language**(*Atom*)
>> Language in which the text-data is written.

---

[4]It is unclear to me whether a MIME note can contain a mixture of content and parts, but I believe the answer is 'no'.

id(*Atom*)
>   Identifier of the message-part.

description(*Atom*)
>   Descrptive text for the *Data*.

disposition(*Atom*)
>   Where the data comes from. The current library only deals with 'inline' data.

name(*Atom*)
>   Name of the part.

filename(*Atom*)
>   Name of the file the data should be stored in.

7   Unix password encryption library

The `crypt` library defines `crypt/2` for encrypting and testing Unix passwords:

crypt(*+Plain, ?Encrypted*)
>   This predicate can be used in three modes. If *Encrypted* is unbound, it will be unified to a string (list of character-codes) holding a random encryption of *Plain*. If *Encrypted* is bound to a list holding 2 characters and an unbound tail, these two character are used for the *salt* of the encryption. Finally, if *Encrypted* is instantiated to an encrypted password the predicate succeeds iff *Encrypted* is a valid encryption of *Plain*.
>
>   *Plain* is either an atom, SWI-Prolog string, list of characters or list of character-codes. It is not advised to use atoms, as this implies the password will be available from the Prolog heap as a defined atom.

8   Memory files

The `memfile` provides an alternative to temporary files, intended for temporary buffering of data. Memory files in general are faster than temporary files and do not suffer from security riscs or naming conflicts associated with temporary-file management. They do assume proper memory management by the hosting OS and cannot be used to pass data to external processes using a file-name.

There is no limit to the number of memory streams, nor the size of them. However, memory-streams cannot have multiple streams at the same time (i.e. cannot be opened for reading and writing at the same time).

These predicates are first of all intended for building higher-level primitives. See also `sformat/3`, `atom_to_term/3`, `term_to_atom/2` and the XPCE primitive `pce_open/3`.

new_memory_file(*-Handle*)
>   Create a new memory file and return a unique opaque handle to it.

free_memory_file(*+Handle*)
>   Discard the memory file and its contents. If the file is open it is first closed.

open_memory_file(*+Handle, +Mode, -Stream*)
>   Open the memory-file. *Mode* is currently one of `read` or `write`. The resulting *Stream* must be closed using `close/1`.

size memory file(+*Handle, -Bytes*)

    Return the content-length of the memory-file it *Bytes*. The file should be closed and contain data.

atom to memory file(+*Atom, -Handle*)

    Turn an atom into a read-only memory-file containing the (shared) characters of the atom. Opening this memory-file in mode `write` yields a permission error.

memory file to atom(+*Handle, -Atom*)

    Return the content of the memory-file in *Atom*.

memory file to codes(+*Handle, -Codes*)

    Return the content of the memory-file as a list of character-codes in *Codes*.

9 Time and alarm library

The `time` provides timing and alarm functions.

alarm(+*Time, :Callable, -Id, +Option*)

    Schedule *Callable* to be called *Time* seconds from now. *Time* is a number (integer or float). *Callable* is called on the next pass through a call- or redo-port of the Prolog engine, or a call to the PL_handle_signals() routine from SWI-Prolog. *Id* is unified with a reference to the timer.

    The resolution of the alarm depends on the underlying implementation. On Unix systems it is based on setitimer(), on Windows on timeSetEvent() using a resolution specified at 50 milliseconds. Long-running foreign predicates that do not call PL_handle_signals() may further delay the alarm.

    Options is a list of `Name`(*Value*) terms. Defined options are:

    remove(*Bool*)

        If `true` (default `false`), the timer is removed automatically. Otherwise it must be delated explicitly using `remove_alarm/1`.

alarm(+*Time, :Callable, -Id*)

    Same as `alarm`(*Time, Callable, Id, []*).

remove alarm(+*Id*)

    Remove an alarm. If it is not yet fired, it will not be fired any more.

current alarm(*?At, ?:Callable, ?Id, ?Status*)

    Enumerate the not-yet-removed alarms. *Status* is one of `done` if the alarm has been called, `next` if it is the next to be fired and *scheduled* otherwise.

call with time limit(+*Time, :Goal*)

    True if Goal completes within Time. *Goal* is executed as in `once/1`. If *Goal* doesn't complete within *Time* seconds, exit using the exception `time_limit_exceeded`. See `catch/3`.

    Please note that this predicate uses `alarm/4` and therefore is *not* capable to break out of long running goals such as `sleep/1`, blocking I/O or other long-running (foreign) predicates. Blocking I/O can be handled using the timeout option of `read_term/3`.

10 Limiting process resources

The `rlimit` library provides an interface to the POSIX getrlimit()/setrlimit() API that control the maximum resource-usage of a process or group of processes. This call is especially useful for server such as CGI scripts and inetd-controlled servers to avoid an uncontrolled script claiming too much resources.

rlimit(+*Resource, -Old, +New***)**

> Query and/or set the limit for *Resource*. Time-values are in seconds and size-values are counted in bytes. The following values are supported by this library. Please note that not all resources may be available and accessible on all platforms. This predicate can throw a variety of exceptions. In portable code this should be guarded with `catch/3`. The defined resources are:

`cpu`	CPU time in seconds
`fsize`	Maximum filesize
`data`	max data size
`stack`	max stack size
`core`	max core file size
`rss`	max resident set size
`nproc`	max number of processes
`nofile`	max number of open files
`memlock`	max locked-in-memory address

> When the process hits a limit POSIX systems normally send the process a signal that terminates it. These signals may be caught using SWI-Prolog's `on_signal/3` primitive. The code below illustrates this behaviour. Please note that asynchronous signal handling is dangerous, especially when using threads. 100% fail-safe operation cannot be guaranteed, but this procedure will inform the user properly 'most of the time'.

```
rlimit_demo :-
        rlimit(cpu, _, 2),
        on_signal(xcpu, _, cpu_exceeded),
        ( repeat, fail ).

cpu_exceeded(_Sig) :-
        format(user_error, 'CPU time exceeded~n', []),
        halt(1).
```

11 Installation

11.1 Unix systems

Installation on Unix system uses the commonly found *configure*, *make* and *make install* sequence. SWI-Prolog should be installed before building this package. If SWI-Prolog is not installed as `pl`, the environment variable `PL` must be set to the name of the SWI-Prolog executable. Installation is now accomplished using:

```
% ./configure
% make
% make install
```

This installs the foreign libraries in `$PLBASE/lib/$PLARCH` and the Prolog library files in `$PLBASE/library`, where `$PLBASE` refers to the SWI-Prolog 'home-directory'.