

The Loko Scheme Developer's Manual

Göran Weinholt

This manual is for Loko Scheme 0.4.4, an optimizing R6RS Scheme compiler.

Copyright © 2019-2020 Göran Weinholt

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Table of Contents

Preface	1
Purpose, audience and scope	1
Credits	1
The Loko Scheme License	2
1 Introduction	3
1.1 Where Loko fits in	3
2 Using Loko	5
2.1 Building Loko	5
2.2 Running	5
2.2.1 Running under Linux	5
2.2.2 Running under KVM (QEMU)	6
2.2.3 Running on bare metal	6
2.2.3.1 Network booting	6
2.2.4 Running in Docker	7
2.3 Compiling a program	7
3 Library reference	9
3.1 Standard libraries	9
3.1.1 SRFI implementations	9
3.2 Base library	9
3.3 Fibers	13
3.4 Unsafe procedures	15
3.5 Target libraries	17
4 Repair instructions	19
4.1 Tools support	19
4.1.1 Disassembly	19
4.1.2 Debugging	19
4.1.3 Debug logs	19
4.1.4 Profiling	19
4.1.5 Memory checking	20
4.1.6 Fuzzing	20
5 Other resources	21
5.1 Loko Scheme website	21
5.2 Online communities	21
5.3 Issue tracker	21
5.4 Package repositories	21

6	Loko internals	23
6.1	Concurrency	23
6.1.1	Why fibers are not preemptible	24
6.1.2	Loko processes	24
6.2	Drivers	25
6.2.1	Driver abstractions	25
6.2.2	Hardware access	25
6.2.3	Future directions for drivers	26
6.3	Interrupt handling	27
6.3.1	Historical background	27
6.3.2	An experimental approach to IRQs	28
6.3.3	Loko's use of traps	29
6.3.4	Interrupts on bare hardware AMD64	29
6.3.4.1	Traps	29
6.3.4.2	IRQs	29
6.3.4.3	Observed bad IRQ behavior	30
6.3.5	Interrupts on Linux AMD64	31
	Appendix A GNU Free Documentation License ..	33
	Index	41

Preface

I've been writing Scheme code for some time now. While doing so I was also working on Scheme implementations. This is the first one that seems to have come out okay.

Göran Weinholt, 2019

Purpose, audience and scope

This manual has two major parts: usage and internals. The first part is intended to let the developer start using Loko to write useful software. The second part goes into details on how Loko works and why things are the way they are.

Some knowledge of Scheme is assumed for the usage part. The reader who has no prior knowledge of any Lisp dialect will initially find it difficult to parse the language.

This is not a complete description of the Scheme language. The reader who wants a more detailed description of Scheme may want to read *The Scheme Programming Language* (<https://www.scheme.com/tspl4/>) (TSPL) by R. Kent Dybvig. The language described in that book is the same language that you can use in Loko Scheme.

This manual is also not a replacement for comments and descriptions in the code. Loko Scheme is not meant to be a closed box; you are supposed to open it up and look at the parts. Maybe even fix some to suit your situation. Loko Scheme needs its own source code to compile your programs, so every installation should come with source code.

Credits

Many have brought ideas, techniques and instructions to fruition that later went into making Loko Scheme. It would not be what it is without their contributions to science.

The syntax-case implementation is from *r6rs-libraries* by Abdulaziz Ghuloum and R. Kent Dybvig, with bug fixes and improvements from Llewellyn Pritchard.

The high-level optimizer *cp0* is based on the chapter *Fast and Effective Procedure Integration* from *Extending the Scope of Syntactic Abstraction* by Oscar Waddell (Ph.D. thesis).

The low-level optimizer is based on concepts taught in a course given by David Whalley in 2011.

The register allocator, except for the bugs, is from *Register Allocation via Graph Coloring* by Preston Briggs (Ph.D. thesis).

The letrec handling is from *Fixing Letrec (reloaded)* by Abdulaziz Ghuloum and R. Kent Dybvig.

The bignum algorithms are based on algorithms from *BigNum Math* by Tom St Denis.

The `list?` procedure uses Olin Shiver's version of Robert W. Floyd's cycle-finding algorithm.

The `equal?` procedure is from the paper *Efficient Nondestructive Equality Checking for Trees and Graphs* by Michael D. Adams and R. Kent Dybvig.

Some intricate parts of the records implementation are from the reference implementation of SRFI-76 by Michael Sperber.

The list sorting code is from SLIB, was written Richard A. O'Keefe and is based on Prolog code by David H. D. Warren.

The dynamic-wind code is from SLIB and was written by Aubrey Jaffer.

The division magic, and many other wonderful hacks, is from the excellent book *Hacker's Delight* by Henry S. Warren, Jr., with foreword by one Guy L. Steele, Jr.!

The fibers library is loosely based on *Parallel Concurrent ML* by John Reppy, Claudio V. Russo and Yingqi Xiao. The API is based on Guile fibers by Andy Wingo and the implementation is closely related to his blog post *a new concurrent ml*.

Thanks also to Abdulaziz Ghuloum for *An Incremental Approach to Compiler Construction*, which helped me consolidate the Scheme compiler experience I had already accumulated through experimentation.

The Loko Scheme License

Loko Scheme is copyrighted software. The default legal state of software is that no rights are granted. However, Loko Scheme is licensed under a free software license. This license grants many permissions, but they are conditional on following the terms of the license.

- Loko Scheme as a whole, including the compiler, the runtime, the drivers, etc., is published under the terms of the GNU Affero General Public License version 3 or later (AGPL-3.0-or-later). See the file `COPYING` in the source code tree.

Any binaries produced by the compiler are also under this license. This means that all binaries created by the Loko Scheme compiler must use a compatible license due to the restrictions in AGPL-3.0-or-later.

If this is a problem for your use of Loko Scheme, then please see issue #2 in the bug tracker.

- This manual is published under the GNU Free Documentation License (see Appendix A [GNU Free Documentation License], page 33).
- The files under the `samples` directory are published under the terms of the MIT license. See the file `samples/LICENSE.txt` in the source code tree.
- The files under the `srfi` directory are also published under the terms of the MIT license. See the file `srfi/LICENSE.txt` in the source code tree. This license is conventional for SRFI implementations.

Most files carry their license information in brief using an SPDX-License-Identifier.

1 Introduction

1.1 Where Loko fits in

Scheme has many implementations. Every known way to implement a programming language has probably been tried with Scheme. There are Scheme implementations for basically all operating systems and all types of machines. Some even say there are more implementations than applications. And Loko Scheme is one of those implementations.

Every Scheme implementation has something that makes it unique. This is what is peculiar about Loko:

- Loko Scheme runs on bare metal.
- Loko Scheme runs and self-compiles without relying on any other language implementation.
- Loko Scheme's runtime is concurrent throughout.
- Loko Scheme is incompatible with C and does not use any libc.

Due to some of the above, Loko is not suitable for every use case. There are plenty of other Scheme implementations available if Loko can't work for your application.

2 Using Loko

2.1 Building Loko

Download a release tarball from <https://scheme.fail> or clone the git repository: `git clone https://scheme.fail/git/loko.git/`. Release tarballs, git commits and tags are signed with the OpenPGP key 0xE33E61A2E9B8C3A2.

Loko Scheme needs an R6RS Scheme implementation for bootstrapping. It can currently be bootstrapped with Chez Scheme. It also depends on some packages from Akku. Follow these steps to compile Loko:

- Install Chez Scheme (<https://cisco.github.io/ChezScheme/>), version 9.5 or later, as a bootstrap compiler. You can skip this step if you have a pre-built `loko` binary (see below).
- Install the package manager Akku.scm (<https://akkuscm.org>), version 1.0.0 or later.
- Run `make` to compile Loko.
- Install with `make install`.
- (Optional) Build and install the manual with `make install-info`.

Once Loko is bootstrapped you can optionally build it with itself using ‘`make selfcompile`’.

You can get the pre-compiled binary from the GitLab CI system. Go to <https://gitlab.com/weinholt/loko/-/tags>, click the download drop-down on the right, and select the “build” artifact under “Download artifacts” (not one of the links on the top of the drop-down).

The binaries tend toward being reproducible. There are some minor differences between the results from Chez Scheme and Loko Scheme due to different gensym implementations. This is fixable.

2.2 Running

Loko Scheme runs either under an existing operating system kernel (currently only Linux), as a hardware virtual machine or directly on bare hardware. The same binary can be loaded on all three environments.

2.2.1 Running under Linux

The `loko` binary is a statically linked ELF binary that Linux can load directly. It needs to be marked executable (`chmod 755 loko`).

Loko doesn’t come with a built-in line editor. It is convenient to use `rlwrap` when running Loko: `rlwrap loko`. The `rlwrap` program adds `readline` on top of any program, providing line editing and history.

Loko uses the environment variable `LOKO_LIBRARY_PATH` to find libraries. This is a colon-separated list of directories. It uses the file extensions `.loko.sls` and `.sls`. This is automatically handled by the package manager Akku.

R6RS top-level programs can be run from the command line with `loko --program program.sps`.

The `loko` binary is also meant to be installed under the name `scheme-script`. If it is invoked with this name it will load a Scheme script, as described in the non-normative R6RS appendix. It is often used like this:

```
#!/usr/bin/env scheme-script
(import (rnrs))
(display "Hello, world!\n")
```

By marking such a file executable the system will hand it over to `scheme-script`, which will then run it as a Scheme top-level program. But the name `scheme-script` is usually handled by the alternatives system, so it could be another Scheme that runs the script.

Such scripts can also be compiled to static binaries that can be run directly. See Section 2.3 [Compilation], page 7.

2.2.2 Running under KVM (QEMU)

To get a repl on the serial port:

```
qemu-system-x86_64 -enable-kvm -kernel loko -m 1024 -serial stdio
```

There is no echo or line editing, but it works alright as an inferior Scheme for Emacs. You can also try `rlwrap -a`.

If you create a script with this command then you can easily run it as an "Inferior Scheme" in e.g. Emacs. There are some additional options you can try:

- Add files to `/boot` using `-initrd filename1,filename2,etc.`
- Set environment variables with e.g. `-append LOKO_LIBRARY_PATH=/boot.`
- Pass command line arguments in `-append` by adding them after `--`, e.g. `-append 'VAR=abc -- --program foo.sps'`.

See the `samples` directory in the source distribution for more examples.

2.2.3 Running on bare metal

Loko on bare metal does not yet come with an adequate user interface. There is rudimentary log output to the text console during boot. This can be redirected, see Section 4.1.3 [Debug logs], page 19.

The first user process will be attached to the COM1 serial port (115200, 8n1). This is adequate for development until there is networking support. The `loko` program's first process is a repl, but if you compile a program then it will be your top-level program.

The `loko` binary should work with any Multiboot boot loader, such as GRUB 2. See the menu entry examples below.

2.2.3.1 Network booting

Network booting is possible if your hardware supports it. It has been tested with GRUB 2 and should also be possible with PXELINUX.

You will need to add an entry in the network's DHCP server and you need a computer where you can install a TFTP server such as `tftpd-hpa`.

- After installing the TFTP server, create a network directory with GRUB:

```
grub-mknetdir --net-directory /tftpboot
```

The TFTP server might be serving up another directory such as `/srv/tftp`.

- Create a configuration in `/tftboot/boot/grub/grub.cfg` with an entry for Loko, such as this:

```
menuentry "Loko Scheme" {
    multiboot /loko loko
}
```

You can also include files that will be available in `/boot`:

```
menuentry "Loko Scheme with foo library" {
    multiboot /loko loko LOKO_LIBRARY_PATH=/boot
    module /foo.sls foo.sls
}
```

Or start a program in the interpreter:

```
menuentry "Hello world" {
    multiboot /loko loko -- --program /boot/hello.sps
    module /hello.sps hello.sps
}
```

- Copy the Loko binary to the `/tftboot` directory.
- Add an entry in the DHCP server. It can look like this if you're using ISC `dhcpd`:

```
host darkstar {
    hardware ethernet 00:11:22:33:44:55;
    filename "boot/grub/i386-pc/core.0";
    next-server 192.168.0.2;
}
```

The hardware address must be changed to match the interface used for network booting (the machine that will run Loko). The server address is the address of the TFTP server.

2.2.4 Running in Docker

Loko can be run in a Docker container. There is sometimes no need to provide any other files in the container; Loko Scheme is self-sufficient.

These Docker images are already available:

- `'weinholt/loko:base'` – Loko Scheme only, a base image.
- `'weinholt/loko:latest'` – Loko Scheme with Debian GNU/Linux stable.
- `'akkuscm/akku:loko'` – Loko Scheme with Debian GNU/Linux and the package manager Akku.

2.3 Compiling a program

Loko can compile R6RS top-level programs:

```
loko --compile hello.sps --output hello
```

The above will compile the R6RS top-level program `hello.sps` and create the binary `hello`, which will run on Linux and bare metal.

Libraries are looked up from the `LOKO_LIBRARY_PATH` environment variable (which is automatically set by the package manager Akku). The use of `eval` is disabled by default to speed up builds, but can be enabled with `-feval`:

```
loko -feval --compile hello.sps --output hello
```

The supported targets can be changed with `-ftarget=TARGET`. The default target is `pc+linux`. Other targets are `linux` for Linux only, `netbsd` for NetBSD only and `pc` for only bare metal.

An additional weird target is provided. It is called `polyglot` and creates binaries that run on Linux, NetBSD and bare metal. This is even less useful than it seems.

You can also omit the normal Scheme libraries. If you use `-ffreestanding` then only the assembler based runtime is added on top of the libraries that your program uses. This is useful mostly when you're working on the compiler. The source distribution has an example where this is used, `samples/hello/just-hello.sps`.

You can tune the parameters for the source-level optimizer (cp0). The argument `-fcp0-size-limit=N` sets the size limit and `-fcp0-effort-limit=N` sets the effort limit.

The command line is very inflexible, so try to stick to the fixed argument order for now.

Loko integrates its run-time into the resulting binary and Loko's source code needs to be available for compilation to succeed. The location is decided by `PREFIX` when compiling Loko, but can be overridden using the `LOKO_SOURCE` environment variable.

3 Library reference

3.1 Standard libraries

The standard R6RS Scheme (<https://r6rs.org>) libraries are provided. Please see the documents on the website for the official versions. Unofficial versions of R6RS updated with errata (<https://weinholt.se/scheme/r6rs/>) are also available online.

3.1.1 SRFI implementations

These following SRFIs are provided with Loko.

- `(srfi :98 os-environment-variables)` – read-only access to environment variables. Documented at <https://srfi.schemers.org/srfi-98/srfi-98.html>.
- `(srfi :170 posix)` – access to common POSIX operations, currently only available on Linux. This is up to date with draft #7. Documented at <https://srfi.schemers.org/srfi-170/srfi-170.html>.

See the package `chez-srfi` (<https://akkuscm.org/packages/chez-srfi/>) for many other SRFIs.

The package `loko-srfi` (<https://akkuscm.org/packages/loko-srfi/>) also provides SRFIs. This package is for SRFIs that create too much entanglement through their dependencies to be suitable for either `chez-srfi` or Loko Scheme proper. At the time of writing it provides SRFI 106 (basic sockets) for Linux.

3.2 Base library

The `(loko)` library is automatically loaded into the repl. It provides all exports from the R6RS libraries, SRFI 98, and the exports shown below. It is intended as a convenient starting point for the repl user and a place to put features that are expected from a Scheme implementation, but that do not belong to any other library.

`interaction-environment` [Procedure]

`load filename` [Procedure]

`include filename` [Procedure]

Include code into the program from *filename*, as if it had been written in the place of the `include` form.

`void` [Procedure]

Returns a void object.

The value that came from nowhere. Even though it has never been standardized, it is customary for Scheme implementations to use `void` for `(if #f #f)`, `(values)` and the value of `set!` and other mutating operations.

In Loko Scheme, void objects carry a copy of the program counter.

```
(list (void))
⇒ (#<void #x21731F>)
```

library-directories [Parameter]

This parameter is a list of strings that name directories to check when importing libraries.

Default: (".")

library-extensions [Parameter]

This parameter is a list of strings with file extensions to use when importing libraries.

Default: (".loko.sls" ".sls" ".ss" ".scm")

installed-libraries [Procedure]

For use in the repl. Returns a list of libraries.

uninstall-library *name* [Procedure]

For use in the repl. Uninstalls the *name* library.

expand *expr* [Procedure]

Expands the expression, returning core forms. The format of the returned forms should not be relied on.

expand/optimize *expr* [Procedure]

Expands and optimizes the expression, returning core forms. The format of the returned forms should not be relied on.

cp0-size-limit [Parameter]

Limits how much the source-level optimizer cp0 will allow the code to grow.

Default: 16

cp0-effort-limit [Parameter]

Limits the effort spent by the source-level optimizer cp0.

Default: 50

disassemble *procedure* [Procedure]

Print the disassembly of *procedure*. It is annotated with labels for local jump destinations and some simple code equivalents.

```
> (disassemble car)
Disassembly for #<procedure car loko/libs/pairs.loko.sls:3224>
```

```
entry:
  206E00 83F8F8      (cmp eax #xFFFFFFFF8)
  206E03 0F8505000000 (jnz L0)
; (set! rax (car rdi))
  206E09 488B47FE      (mov rax (mem64+ rdi #x-2))
  206E0D C3            (ret)
L0:
  206E0E E96DA2FFFF    (jmp (+ rip #x-5D93))
```

machine-type [Procedure]

The machine type that Loko is running on. This is a vector where the first element is the CPU type `amd64` and the second is the system environment (`linux`, `netbsd` or `pc`).

`time expr` [Syntax]
Run the procedure *thunk* once with no arguments and print some numbers of memory allocation and elapsed time.

`time-it what thunk` [Procedure]
This is the procedural version of `time`.

`time-it* what iterations thunk` [Procedure]
Run *thunk* repeatedly *iterations* times and print some bogus statistics. The aim is that this procedure should be the best way to do micro benchmarks.

```
> (time-it* "fx+" 10000000 (lambda () (fx+ x 1)))
Timing fx+ to find the minimum cycle time:
New minimum is 1819 cycles with 10000000 iterations to go.
...
New minimum is 234 cycles with 6257346 iterations to go.
```

```
The cycle count varied between 234 and 83160784
(Arithmetic mean)       $\mu$  = 248.75
(Standard deviation)    $\sigma$  = 24.33
(Population variance)  $\sigma^2$  = 592.08
                    min x_i =  $\mu - .61\sigma$ 
Used 9736890 samples (263110 outliers discarded).
```

```
234
> (time-it* "+" 10000000 (lambda () (+ x 1)))
Timing + to find the minimum cycle time:
New minimum is 1751 cycles with 10000000 iterations to go.
...
New minimum is 240 cycles with 9968540 iterations to go.
```

```
The cycle count varied between 240 and 84141254
(Arithmetic mean)       $\mu$  = 252.96
(Standard deviation)    $\sigma$  = 30.46
(Population variance)  $\sigma^2$  = 927.82
                    min x_i =  $\mu - .43\sigma$ 
Used 9979862 samples (20138 outliers discarded).
```

```
240
```

Note that `cp0` will optimize the *thunk* before it runs, so you may end up benchmarking something other than what you thought. Check with `expand/optimize`. If the code is entered in the REPL then you also measure the overhead of `eval`.

Modern computers are notoriously difficult to get any consistent results from. An improvement in cycles could be because the code slightly moved in memory. See Producing Wrong Data Without Doing Anything Obviously Wrong (<https://john.cs.olemiss.edu/~hcc/researchMethods/notes/localcopy/mytkowicz-wrong-data.pdf>) (2009, Mytkowicz, et al). A more lively view of the problem is the presentation Performance Matters (<https://youtu.be/r-TLSBdHe1A>) (2019, Emery Berger at Strange Loop).

open-output-string [Procedure]
 Make a new string output port that accumulates characters in memory. The accumulated string can be extracted with `get-output-string`.

get-output-string *string-output-port* [Procedure]
 Extract the accumulated string in `string-output-port` and reset it. Returns the string.

port-file-descriptor *port* [Procedure]
 Get the file descriptor associated with *port*. Returns `#f` if there is no associated file descriptor.

port-file-descriptor-set! *port fd* [Procedure]
 Set the file descriptor associated with *port* to *fd*.
 This procedure is primarily intended to allow custom ports to have file descriptors. It is unspecified whether changing a port's file descriptor affects the file descriptor used for subsequent operations on the port.

gensym [Procedure]
 Generate an uninterned symbol. These are symbols which are not `eq?` to any other symbol.

make-parameter *default-value* [*fender*] [Procedure]
 Create a new parameter object. Parameters are typically used to implement dynamically scoped variables together with `parameterize`. A parameter's current value can be queried by calling it with no arguments and its value can be modified by calling it with one argument, the new value.
 The optional *fender* procedure is applied to the value whenever the parameter is modified. The return value of *fender* is used in place of the new value. A typical use of this procedure is to do some type checks on the new value.

parameterize ((*name value*) ...) *body*... [Syntax]
 Parameterize rebinds the parameter *name* to *value* for the dynamic extent of *body*. This means that while *body* is running, *name* will be set to *value*. The value is possibly filtered by a fender procedure.
 Whenever the program leaves the body, either by a normal return or a non-local exit (such as in a `guard` expression or by calling a continuation created by `call/cc`), the value is reset to the value it has outside of the body. If control reenters body, as in a call to a continuation created inside the body, the parameter will return to the value established by `parameterize`.

Although it has the same name, this syntax is a faster variant that is not fully compatible with SRFI-39. This variant is very common in Scheme implementations and matches the one used in e.g. Chez Scheme.

loko-version [Procedure]
 The version number of the Loko Scheme runtime. This is a SemVer number and may include build information in the future.

putenv *name value* [Procedure]

Set the environment variable *name* to *value*. The *name* is always string.

If the *value* is a string then the variable is set to that value.

If the *value* is **#f** then the variable is removed.

This updates the environment used by SRFI 98 and can also be expected to be visible to any child processes started after the call.

The strings must not contain any **#\nul** characters. The *name* must not contain a **#\=** character. These limitations are not checked.

The names and values are always transcoded to/from UTF-8 in the POSIX interfaces.

Please beware that other Scheme implementations commonly leak memory through this procedure.

collections [Procedure]

The number of garbage collections.

3.3 Fibers

The (**loko system fibers**) library exports procedures for fibers. Fibers are a form of lightweight concurrency based on Concurrent ML. For an overview, see Section 6.1 [Concurrency], page 23.

spawn-fiber *thunk* [Procedure]

Create a new fiber that will start running the procedure *thunk*, which takes no arguments.

make-channel [Procedure]

Create a new channel. Channels are places where two fibers can rendezvous to exchange a message. There is no buffering in a channel.

channel? *obj* [Procedure]

True if *obj* is a channel.

put-message *channel obj* [Procedure]

Put the message *obj* on the channel *channel*. Blocks until another fiber picks up the message. Returns unspecified values.

get-message *channel* [Procedure]

Get a message from the channel *channel*. Blocks until another fiber has arrived with a message. Returns the message.

sleep *time* [Procedure]

Block the fiber for *time* seconds.

put-operation *channel obj* [Procedure]

Returns an operation object that represents putting the message *obj* on the channel *channel*.

get-operation *channel* [Procedure]

Returns an operation object that represents getting a message from the channel *channel*.

wrap-operation *op f* [Procedure]
 Returns an operation object that is the same as the operation *op*, except that the values are wrapped by the procedure *f*.

sleep-operation *seconds* [Procedure]
 Returns an operation object that represents waiting until *seconds* have passed from the time of the call to this procedure.

timer-operation *a* [Procedure]
 Return an operation object that represents waiting until absolute time *a* (in internal time units).

choice-operation *op ...* [Procedure]
 Returns an operation object that represents a choice between the given operations *op ...*. If multiple operations can be performed then one is selected non-deterministically.

It is not an error to call this procedure with no arguments. It is in fact a useful construction when gathering operations.

If **wrap-operation** is used on a choice operation then every operation will be wrapped.

perform-operation *op* [Procedure]
 Perform the operation *op*, possibly blocking the fiber until the operation is ready.

With **choice-operation** and **perform-operation** it's possible to write code that waits for one of several operations. This can be something simple like waiting for a message with a timeout:

```
(perform-operation (get-operation ch) (sleep-operation 1))
```

This example will wait for a message on the channel *ch* for up to one second. In order to distinguish between a message and a timeout, **wrap-operation** is used:

```
(perform-operation
 (choice-operation
  (wrap-operation (get-operation ch) (lambda (x) (cons 'msg x)))
  (wrap-operation (sleep-operation 1) (lambda _ 'timeout))))
```

This code will either return `(msg . x)` where *x* is the received message; but if more than one second passes without a message it returns `timeout`.

The object returned from **choice-operation** can be returned from a procedure, stored in a data structure, sent over a channel, etc.

make-cvar [Procedure]
 Make a new *condition variable* (in Concurrent ML's terminology). These allow a program to wait for a condition to be signalled. See the procedures below.

cvar? *obj* [Procedure]
 True if *obj* is a condition variable.

signal-cvar! *cvar* [Procedure]
 Signal the condition variable *cvar*, unblocking any fibers that are waiting for it.

- wait *cvar*** [Procedure]
 Wait for the condition variable *cvar* to be signalled, blocking until it is.
- wait-operation *cvar*** [Procedure]
 Return an operation that represents waiting for the condition variable *cvar* to be signalled.
- yield-current-task** [Procedure]
 Yield the current task and and let another fiber run. This is generally not needed in I/O-bound programs, but is provided to let CPU-bound programs cooperate and voluntarily let other fibers run.
- exit-current-task** [Procedure]
 Stops the running fiber.
- run-fibers *init-thunk*** [Procedure]
 Provided for compatibility with Guile. It runs the procedure *init-thunk* in the fibers scheduler. This procedure can return earlier in Loko than in does in Guile. Guile provides it because fibers are not an integrated feature in its runtime, so it needs an entry point for when to start and stop the fibers facility.

3.4 Unsafe procedures

The (`loko system unsafe`) library provides raw access to kernel services, linear memory and I/O bus registers.

- syscall *n arg ...*** [Procedure]
 Calls the kernel's system call number *n* with the arguments *arg ...*. Returns a fixnum.
 Example fork on Linux amd64:
- ```
(when (zero? (syscall 57)) ; __NR_fork
 (display "child process\n")
 (exit)) ; child become a zombie
```

Scheme programs should generally not be written directly with syscalls any less than C programs would do the same. There are usually interactions with the standard library that should be considered, such as flushing of ports to prevent duplicated output.

- bytevector-address *bytevector*** [Procedure]  
 Get the linear address of the first byte of *bytevector*.  
 The first byte is guaranteed to have an alignment of eight bytes.  
 A moving garbage collector is used for bytevectors created with `make-bytevector`. There is no way to ensure that they do not move during GC, so their addresses should not be used to perform bus-mastering DMA.  
 Returns a fixnum.

- get-mem-u8 *addr*** [Procedure]  
**get-mem-u16 *addr*** [Procedure]  
**get-mem-u32 *addr*** [Procedure]

`get-mem-s61 addr` [Procedure]

Read a u8, u16, u32 or s61, respectively, from linear address *addr* and return it as a fixnum.

- If *addr* is unaligned then an exception is raised.
- On targets where fixnums are not wide enough to hold the result, a bignum will be used instead, but the bus access will be the correct width.
- Calls to these procedures will not be optimized away or reordered by the compiler.
- Each call corresponds to one memory read instruction of the matching size. For the signed 61-bit procedure this only applies to targets that can issue 64-bit memory bus operations.
- These calls may be issued out of order by the processor or go to the cache. This is usually only a problem when the memory area is backed by RAM. On AMD64 it is safe to read memory-mapped hardware registers with these procedures thanks to MTRR. Other situations may require locking instructions or memory barriers.

`put-mem-u8 addr n` [Procedure]

`put-mem-u16 addr n` [Procedure]

`put-mem-u32 addr n` [Procedure]

`put-mem-s61 addr n` [Procedure]

Write *n* as a u8, u16, u32 or s61, respectively, to linear address *addr*.

- If *addr* is unaligned then an exception is raised.
- Calls to these procedures will not be optimized away or reordered by the compiler.
- Each call corresponds to one memory write instruction of the matching size. For the signed 61-bit procedure this only applies to targets that can issue 64-bit memory bus operations.
- These calls may be issued out of order or be merged by the processor. This is usually only a problem when the memory area is backed by RAM. On AMD64 it is safe to read memory-mapped hardware registers with these procedures thanks to MTRR. Other situations may require locking instructions, manual flushes or memory barriers.

Returns unspecified values.

`get-i/o-u8 busaddr` [Procedure]

`get-i/o-u16 busaddr` [Procedure]

`get-i/o-u32 busaddr` [Procedure]

Read a u8, u16 or u32, respectively, from I/O bus address *busaddr* and return it as a fixnum.

The `get-i/o-u32` procedure may return a bignum on targets where  $(\leq (\text{fixnum-width}) 32)$ , but the bus access will be 32-bit.

`put-i/o-u8 busaddr n` [Procedure]

`put-i/o-u16 busaddr n` [Procedure]

`put-i/o-u32 busaddr n` [Procedure]

Write *n* as a u8, u16 or u32, respectively, to I/O bus address *busaddr*.

Returns unspecified values.

`get-i/o-u8-n! busaddr addr n` [Procedure]

`get-i/o-u16-n! busaddr addr n` [Procedure]

`get-i/o-u32-n! busaddr addr n` [Procedure]

Read *n* units of u8, u16 or u32, respectively, from I/O bus address *busaddr* and write them to memory starting at linear address *addr*.

Returns unspecified values.

### 3.5 Target libraries

The following Linux-specific libraries are provided. Their usage mirrors 1:1 the functionality in the Linux manpages, section 2.

- `(loko arch amd64 linux-numbers)` – constants used in the Linux syscall interface (UAPI).
- `(loko arch amd64 linux-syscalls)` – thin wrappers around Linux syscalls.



## 4 Repair instructions

### 4.1 Tools support

This section describes some tools can be used together with Loko Scheme when things go wrong.

#### 4.1.1 Disassembly

Loko can be disassembled using any regular disassembler that supports ELF and amd64, such as the one in GNU binutils and GNU gdb. There is also a built-in disassembler for procedures, see Section 3.2 [Base library], page 9.

#### 4.1.2 Debugging

Loko can be debugged with GNU gdb (<https://www.gnu.org/software/gdb/>) with the help of `loko-gdb.py`. Stack traces and pretty printing should work. Line information is currently missing.

Scheme objects can get very big. Limit the output with e.g. `set print elements 5`.

Traps from the processor are translated into conditions with a `&program-counter` condition, e.g. this error from trying to evaluate `(#f)`:

```
The condition has 5 components:
 1. &assertion &violation &serious
 2. &who: apply
 3. &message: "Tried to call a non-procedural object"
 4. &irritants: (#f)
 5. &program-counter: #x309795
End of condition components.
```

You can look up the trapping instruction with a disassembler.

#### 4.1.3 Debug logs

The PC port of Loko can have debug logging redirected from the VGA console with the `CONSOLE` environment variable:

- `CONSOLE=vga` prints to the VGA text mode console.
- `CONSOLE=com1` prints to COM1 (BIOS's default baud rate).
- `CONSOLE=debug` prints to QEMU's debug console. This is enabled with e.g. `-debugcon vc`.

#### 4.1.4 Profiling

Loko on Linux can be profiled with `perf` ([https://perf.wiki.kernel.org/index.php/Main\\_Page](https://perf.wiki.kernel.org/index.php/Main_Page)).

Some micro benchmarks can be done from inside Loko, see Section 3.2 [Base library], page 9.

### 4.1.5 Memory checking

It's possible to run Loko in Valgrind (<http://valgrind.org/>). Valgrind does not support alignment checking, so Loko will print a warning about that. Loko also does not use the “red zone”, so Valgrind will think that a lot of what Loko is doing uses uninitialized memory. Don't believe it.

### 4.1.6 Fuzzing

Loko can be run with AFL (<http://lcamtuf.coredump.cx/afl/>). This requires an instrumented binary, which has some overhead. The support needs to be enabled in `(loko arch amd64 codegen)` first. Change `use-branch-instrumentation` to `#t` and recompile.



## 5 Other resources

### 5.1 Loko Scheme website

The Loko Scheme (<https://scheme.fail>) website has official release tarballs and an online copy of the manual.

### 5.2 Online communities

There are two chat channels available on IRC. They are on the IRC network Freenode and they are `#scheme` and `#loko`. The `#scheme` channel is for the whole Scheme community. You will need an IRC client to join these channels, but Freenode also offers a web-based client.

If you're more of a forum or mailing list user, then there aren't any direct options like this yet. But we have Usenet, which has filled this need even since before there was an Internet. The community-wide Scheme group is called `'comp.lang.scheme'`. It is accessible via any Usenet provider, even free ones like Eternal September (<http://www.eternal-september.org/>). It is also accessible via one of the big search engine companies.

There is also a community-wide Scheme Wiki at <http://community.schemewiki.org>.

In my opinion it would be better for the Scheme community if it was not as fragmented as it is. For that reason there will not be any Loko-specific mailing lists.

### 5.3 Issue tracker

There is an issue tracker for Loko Scheme at [gitlab.com](https://gitlab.com/weinholt/loko/issues). The URL is <https://gitlab.com/weinholt/loko/issues>. You can submit issues via email by sending them to the `'bugs'` alias at the Loko Scheme website's domain.

### 5.4 Package repositories

There are repositories of packages online where you can find many useful libraries that work with Loko:

- Akku.scm (<https://akkuscm.org/>) – an R6RS/R7RS package manager. Akku also translates R7RS libraries to R6RS to try and make them work with R6RS implementations.
- Snow (<http://snow-fort.org/>) – an R7RS package manager. These packages are also included in Akku's repository.



## 6 Loko internals

### 6.1 Concurrency

The concurrency in Loko exists at two different levels:

- Loko processes, which are preemptible processes running at the same privilege level and in the same page table as the Loko runtime.
- Loko fibers, which are lightweight processes that run inside a Loko process. They are an implementation of Concurrent ML (<https://people.cs.uchicago.edu/~jhr/papers/cml.html>).

When Loko starts it immediately sets up a Loko process for the scheduler (also called pid 0). The boot loader has already created a heap and stack for it. The scheduler is responsible for starting pid 1, handling interrupts, managing preemption, message passing between processes and other maintenance tasks. The first scheduler that starts is also responsible for booting all other processors in the system.

Currently all other processors boot up but stop after initialization. Some work needs to be done to allocate new scheduler processes for them.

At least one more type of process is planned: old heavy-weight processes with their own page tables. These will be where we put all the FORTRAN programs. For now, this feature is missing and all code on the system is Scheme code running in a Loko process.

Fibers are a lightweight concurrency system based on Concurrent ML. The implementation is heavily inspired by a Andy Wingo’s articles Growing Fibers (<https://wingolog.org/archives/2017/06/27/growing-fibers>). and A New Concurrent ML (<https://wingolog.org/archives/2017/06/29/a-new-concurrent-ml>). Concurrent ML forms a fundamental principle for concurrency that can be used to implement higher-level concurrency like Go channels or Erlang processes.

For details on how to use fibers in your program, see Section 3.3 [Fibers], page 13. You can also consult the Guile fibers manual (<https://github.com/wingo/fibers/wiki/Manual>) to some extent; it has a lot of background information.

The implementation in Loko is different from Guile fibers in these ways:

- Loko fibers need the import (`loko system fibers`) rather than (`fibers`).
- Loko fibers are based on pure `call/cc` to switch between fibers and the fiber scheduler. This is mostly because the Loko runtime doesn’t have delimited continuations, but it made it easier to quickly get the correct semantics for parameters.
- However, basing the implementation on `call/cc` brings a space leak. It means that there is some extra overhead from lugging around the unused parts of the continuation and the dynamic winders. In some programs, this state can potentially grow without bound. This could be solved without implementing delimited continuations, but extra support from the runtime is needed either way.
- Loko fibers are not preemptively scheduled. This is because of several reasons, explained in the next section. This means that you shouldn’t run long-running computations in a fiber that shares a Loko process with other fibers that need to be responsive (unless you explicitly yield the fiber every now and then).

- Loko fibers in a process do not run in parallel; they are not shared between processors. The way memory allocation is done in Loko means that two processors can't manage the same heap.

*XXX: The implementation described above has a broken `dynamic-wind`. See issue #26 in the bug tracker.*

The API is compatible with Guile fibers, so the concurrency parts of a program written for Guile fibers should work with Loko fibers. The largest exception is Guile's (`fibers internals`) library, which manages fiber schedulers. Loko only has one of those per Loko process. Another difference is that I/O is non-blocking by default on Loko.

A large part of what makes fibers attractive is that code can be written as if it were non-concurrent. You're free to read and write to pipes and network streams without explicitly dealing with polling for when data is available or when file descriptors are ready for writing. One of the sample programs is a tiny web-server that spawns a fiber per connected client.

Loko on Linux takes care to set `O_NONBLOCK` on file descriptors and suspends the current fiber when syscalls return `EAGAIN` (also called `EWOULDBLOCK`). Loko uses `epoll` to find out when the file descriptor will be ready.

But Linux does not implement `EAGAIN` for regular files, so reading from a file can block. When this happens it prevents other fibers from running. Standard I/O is non-blocking, but other programs that use the same terminal can either get confused by that and/or turn it off. Even memory accesses can be blocking on Linux because pages can be swapped out to disk. (Turning off swap is not a good idea). The binary for your program was `mmap'd` by Linux when it got started and unmodified pages can be read back from disk, so Linux can freely evict the pages from memory. So having anything like a guaranteed responsive program on Linux is challenging. If anything goes wrong with the disk, all your processes can end up in uninterruptible sleep.

Loko on bare hardware has no operations that block other fibers from running.

### 6.1.1 Why fibers are not preemptible

Here's the excuse. Loko processes can temporarily use registers in such a way that they contain arbitrary bit patterns. If such a register were to be saved to a continuation object, the garbage collector would choke on it. Other fibers in the same Loko process use the same heap, so a fiber can't simply be suspended and left alone as Loko processes are when they're preempted.

One common solution to this problem is that the compiler inserts counters at various points in the code. These counters are incremented at safe points in the code and are used to implement software-based timer interrupts. This solution brings with it some overhead and needs special care to not ruin the performance of tight loops. It may be done later unless another solution is found.

### 6.1.2 Loko processes

The use case for processes is pretty slim at this time, but they are the only way to get preemptive concurrency.

Loko on Linux currently has a very rudimentary scheduler that can't handle more than one process.

## 6.2 Drivers

A driver is a piece of code that allows for abstract access to a hardware device. In Loko Scheme they are collected in the `drivers` directory.

### 6.2.1 Driver abstractions

A large part of creating drivers is to adapt the hardware to the abstractions used in the rest of the system. Sometimes this means that the full capabilities of the hardware are hidden. A serial port supported by a UART may appear as a Scheme input port and an output port, which will allow you to hook it up to any Scheme code that works with ports. But a UART can do much more than a Scheme port can. An output port can't usually control baud rates or send breaks.

Scheme lacks abstractions for hardware. In the RnRS standards, files and ports are the only things that work as abstractions for hardware. This is not so bad, because Unix systems have shown that a lot of hardware can be represented as special files in `/dev` and the file descriptors you get when opening the special files.

However, these file descriptors are just handles that let user space communicate with the driver. Sometimes the communication is through an abstraction layer and sometimes it goes almost directly to the driver. User space needs to use special syscalls (e.g. `ioctl`) to actually do anything interesting that is not a read/write operation.

Drivers for Loko should not be locked in to any particular way of designing the user space interface. Decisions like that are taken on a different level. This will let developers experiment with different user space designs. Designs like Barrelfish (<http://www.barrelfish.org>), where hardware virtualized devices are handed out directly to user space, should be possible to express.

Drivers in Loko should be written as libraries that provide convenient APIs. These APIs should provide basic functionality in a way that is common between similar hardware of the same type. The interfaces will need to be developed over time.

When it is necessary to have concurrency (the most common case for modern devices), drivers should use channels to communicate with the rest of the system. Concurrent drivers can start as many fibers as they like. The messages sent on these channels should preferably be simple objects like vectors, symbols, pairs and fixnums. The messages become part of the driver API.

### 6.2.2 Hardware access

Drivers need access to their devices. The way this is done depends on what type of bus the device is attached to.

Modern PCs have busses that can be probed, like PCI and USB. This process provides enough information that you can easily detect the type of devices that are on the bus and how to access them. Hardware tends to appear like a tree-like structure, so it is natural that a bus driver will pass along a reference to the bus down in the call stack.

Older devices on the PC do not appear on the PCI bus and should be detected and started by just knowing that they ought to be there because it's a PC. Most ARM systems do not come with a PCI bus and have all their devices on addresses that need to be known ahead of time, but are different between platforms. A popular solution is to use DeviceTree to encode this information and that is certainly something that should be explored for Loko.

The major hardware interaction points are:

1. Scanning and configuring the bus; detecting new and removed devices.
2. Setting up access to the device.
3. Interfacing with the device through its registers, channels, etc.
4. Allowing the device to write to system memory.
5. Waiting on interrupts from the device.

The way these things are done depends on the bus. Further documentation is needed. For now, please consult the source code or ask.

### 6.2.3 Future directions for drivers

There is an interesting thing that can be done with drivers for PCI devices when `eval` uses online compilation. PCI devices can appear anywhere in memory and sometimes even anywhere in I/O space. Register access can look like this:

```
(define (driver-pci-uhci dev controller)
 ;; The UHCI registers are mapped to the location in BAR4
 (let ((bar (vector-ref (pcidev-BARs dev) 4)))
 ;; Disable keyboard and mouse legacy support
 (pci-put-u16 dev #xC0 #x0000)
 (driver-uhci (if (pcibar-i/o? bar) 'i/o 'mem)
 (pcibar-base bar)
 (pcibar-size bar)
 (pcidev-irq dev)
 controller)))

(define (driver-uhci reg-type reg-base reg-size irq controller)
 ;; Access to the device registers (independent of i/o vs mem)
 (define (reg-u8-ref offset)
 (assert (fx<? -1 offset reg-size))
 (case reg-type
 ((i/o) (get-i/o-u8 (fx+ reg-base offset)))
 ((mem) (get-mem-u8 (fx+ reg-base offset)))
 (else (assert #f))))
 ...)
```

It would be interesting if `driver-pci-uhci` used `eval` to compile a specialized version of the driver where `reg-u8-ref` (etc.) had been inlined by `cp0`. After compilation, each `reg-u8-ref` call would be a single instruction. Specializing, compiling and starting the driver can be as simple as this:

```
(let ((driver-uhci
 (eval '(lambda (controller)
 (driver-source-uhci ,reg-type ,reg-base
 ,reg-size ,irq
 controller))
 (apply environment driver-environment-uhci))))
 (driver-uhci controller))
```

In principle this kind of code would work even today, but the driver would be slowed down because `eval` is slow.

The same principle can be applied to embedded systems that use DeviceTree. If a static DeviceTree is used then this could even be done as part of the build process. If a dynamic DeviceTree is used (to allow the same kernel to run on different ARM platforms) then boot time may become an issue. But then drivers could be designed to initially use a non-specialized driver, call `eval` asynchronously, and tell the running driver to switch to the specialized driver when `eval` has returned.

## 6.3 Interrupt handling

Loko in essence treats device drivers as communicating sequential processes. IRQs from devices are treated as primitive messages from the device to the driver. Requests to configure the device or transfer data between the device and the rest of the system is also done with messages. Device drivers are not fundamentally different from other Scheme programs.

### 6.3.1 Historical background

This section uses the term *interrupt* to denote an interruption in the processor's normal execution sequence; *IRQ* to denote an interrupt from a hardware device and *trap* to denote an interrupt from the processor itself. There are more types of interrupts, but they are not discussed here.

IRQs are handled differently from how they are handled in normal kernels. An anecdote from The Rise of Worse is Better (<https://www.dreamsongs.com/RiseOfWorseIsBetter.html>) by Richard P. Gabriel is relevant here:

Two famous people, one from MIT and another from Berkeley (but working on Unix) once met to discuss operating system issues. The person from MIT was knowledgeable about ITS (the MIT AI Lab operating system) and had been reading the Unix sources. He was interested in how Unix solved the PC loser-ing problem. The PC loser-ing problem occurs when a user program invokes a system routine to perform a lengthy operation that might have significant state, such as IO buffers. If an interrupt occurs during the operation, the state of the user program must be saved. Because the invocation of the system routine is usually a single instruction, the PC of the user program does not adequately capture the state of the process. The system routine must either back out or press forward. The right thing is to back out and restore the user program PC to the instruction that invoked the system routine so that resumption of the user program after the interrupt, for example, re-enters the system routine. It is called PC loser-ing because the PC is being coerced into loser mode, where loser is the affectionate name for user at MIT.

The MIT guy did not see any code that handled this case and asked the New Jersey guy how the problem was handled. The New Jersey guy said that the Unix folks were aware of the problem, but the solution was for the system routine to always finish, but sometimes an error code would be returned that signaled that the system routine had failed to complete its action. A correct user program, then, had to check the error code to determine whether to simply

try the system routine again. The MIT guy did not like this solution because it was not the right thing.

From the New Jersey guy we get the `errno` value `-EINTR`, so he was clearly the more successful one. The MIT guy's approach would have been to carefully design syscalls so that they keep their full state in a structure or in the arguments to the syscall, which is kind of like doing a very manual and tedious `call/cc` when an interrupt arrives. And nobody has time for that.

But when the New Jersey guy feels like even `-EINTR` is too difficult to manage, we get *uninterruptible sleep*. This is a situation where programs aren't running, but they can't be killed either. This happens when a program is blocked in a syscall and is holding an unknown number of locks and other state in the kernel. Maybe it was reading from a file, but something got wedged and stopped responding. Instead of an error code, the process is in uninterruptible sleep. This slowly creeps from program to program as other parties try to communicate with the frozen process.

So Loko takes a different approach to all this.

### 6.3.2 An experimental approach to IRQs

IRQs are generated by hardware devices when they want the attention of the processor. An example is a UART (serial port controller) that has just received a byte. It will generate an IRQ to get the driver to read the byte from its buffer.

The normal idea of how to handle an IRQ is basically as follows: install the interrupt service routine (ISR) that comes as part of the driver, reset the device to its normal state, then enable the IRQ in the interrupt controller. An ISR is a special piece of code that must be prepared to run at potentially any time (except when interrupts are disabled). Usually there is a priority order on IRQs so that they can in turn be interrupted by higher-priority IRQs. Either way, when an IRQ arrives the driver quickly services the hardware. In the UART example it would read a byte from the UART and place it in a software controlled buffer.

In this way of doing things, there are unfortunately severe restrictions on what can be done in an ISR. An ISR runs in *interrupt context* where many usual kernel services are simply unavailable. Anything that would block the program is usually unavailable and access to memory is limited. Arranging things so that arbitrary Scheme code can run in interrupt context is difficult.

For this reason, Loko does not run driver code in ISRs. The ISRs are instead minimal pieces of code that cooperate with the process scheduler to make the driver's process runnable. An IRQ is sent as a message to the process and it handles it at its leisure.

This approach is somewhat experimental, and may have some problems with latency in legacy hardware such as UARTs, but modern devices do bus mastering DMA and are generally not sensitive to interrupt servicing latency. Bus mastering DMA means that the device has access to the system's memory. Generally such a device cooperates with the driver to maintain queues in system memory that describes data transfers.

There are pros and cons to Loko's approach. The cons are that IRQs are handled with some latency, but this is usually not a problem for modern devices. The pros are that it makes driver code *much* easier to write and maintain. Since it eliminates many of the usual difficulties with writing drivers, it may even mean that most competent Schemers with



access to the hardware programming manual can write device drivers. (Some difficulties still remain with manual memory management, but they are not that hard to deal with).

Even if latency is a potential problem, Loko's approach should be good for throughput. A driver can easily decide that data is coming in at such a high rate that it doesn't need to use interrupts, and switch over to periodic polling instead. Linux uses this technique in its networking stack under the cryptic name "New API" (NAPI).

Another benefit of Loko's approach is that the dilemma in the "worse is better" story is resolved. User programs never need to be given an equivalent of `-EINTR` and the programmer does not need to manually keep track of where they are in the handling of a system call.

### 6.3.3 Loko's use of traps

Loko offloads as much error checking as possible on built-in mechanisms in the hardware. Instead of using explicit type checks, it lets the hardware do the type checks (where possible).

Programming errors like `(/ 1 0)` are often signalled by the hardware in most programming environments. Loko takes this further and extends it to errors like `(car #f)`, `(vector-ref "foo" 0)` and `(1 + 2)`. Loko uses the processor's alignment checking feature to trap wrong uses of pairs, procedures, strings, vectors and bytevectors. You can read more about this in *Faster Dynamic Type Checks* (<https://weinholt.se/articles/alignment-check/>).

### 6.3.4 Interrupts on bare hardware AMD64

The interrupt handlers are in `(loko arch amd64 pc-interrupts)` and are written in assembly.

There are two fundamentally different types of interrupts that both go under the name `interrupt` and that use similar mechanisms, but have very different sources.

#### 6.3.4.1 Traps

Traps are interrupts that are triggered by some classes of errors in the running program. For these interrupts the interrupt handlers cause the program to invoke an error handler written in Scheme. No care is taken to preserve the program's current stack frame or current register values, which means that some useful debugging information is lost. While that is unfortunate, and should be fixed, it is still semantically correct since these errors are all categorized as `&serious`.

Traps are handled by entries in the interrupt descriptor table (IDT). The IDT controls whether the processor should change privilege levels, which address it should jump to, which stack it should use and whether interrupts should be automatically disabled or not.

The processor pushes some of the program state on the stack and gives control to the handler. For traps, the handlers identify the cause of the trap and decide which library function should take control over the program. It then executes a tail-call to that function. These functions live in `(loko arch amd64 lib-traps)` and are responsible for calling the error handler, which is written in Scheme.

#### 6.3.4.2 IRQs

The IRQ handling in Loko relies on the processor's *interrupt stack table* (IST) and the interrupt controller's *specific end of interrupt* (SEOI) and *special mask mode*. On AMD64

systems the SEOI feature is available in the legacy PIC and the APIC (except in early versions). Interrupt priorities (nesting) are not meant to be used.

Interrupt masking is a very important part of Loko's IRQ handling. There are three places where interrupts can be masked: the device itself, the interrupt controller and the processor. The driver configures the device to generate interrupts in a way that suites the driver. The interrupt controller is responsible for delivering interrupts to the processor and can be asked by the processor to mask an interrupt. The interrupt controller also keeps track of which interrupts are being serviced and temporarily masks them until they are acknowledged. Finally, the processor can also mask interrupts with the 'RFLAGS.IF' bit in the flags register. When 'RFLAGS.IF' is clear, no interrupts are delivered.

Loko's IRQ handling system relies on masking interrupts with the 'RFLAGS.IF' bit and delaying acknowledgement. 'RFLAGS.IF' is used to mask interrupts while the scheduler is running. When inside the scheduler, interrupts can only happen in one very controlled circumstance: in the `sys_hlt` syscall. This is used when no processes are runnable and it lets the processor save energy. The `sys_hlt` syscall returns when an IRQ has been delivered to the processor. The scheduler then finds the driver process that has the IRQ registered, enqueues it as a message and makes the process runnable.

Normal processes always run with interrupts unmasked, except maybe for some brief and tricky moments during task switching. When an interrupt arrives, the processor uses the IST to switch to a different stack. This is necessary to avoid messing up the process's Scheme stack. All registers are saved in the process control block so that the process can be resumed later. The IRQ handler resumes the scheduler and lets it know what happened.

When the process resumes it will have an IRQ number in its message queue. It is up to the process when it wants to dequeue this message, but usually a driver process will be waiting for an IRQ to arrive. Whenever a process calls the scheduler it can also receive a pending message. The two cases are then that a process either becomes runnable due to an IRQ and that a process is already runnable and will see the IRQ later.

Processes that are notified about IRQs will handle them, doing whatever task is needed to service the IRQ in the hardware, and afterwards tell the scheduler to acknowledge the IRQ. The scheduler then sends an instruction to the interrupt controller to acknowledge that specific interrupt. At this point the device may again choose to generate an interrupt when the conditions are right.

There are a few classes of programming errors when it comes to IRQ code in drivers. The driver may decide that it has handled an IRQ and sends an acknowledgement. However, if it missed an interrupt reason, the device may generate the same IRQ again, immediately after acknowledgement. This slows down the system with unnecessary work in the driver.

The opposite can also happen if a driver does not properly handle all the interrupt reasons. The symptom can be that an IRQ arrives, is seemingly handled by the driver, but then no more IRQs ever arrive and the device seems to have frozen. Which class of error is more likely depends on if the device uses edge-triggered or level-triggered interrupts. A way to check if this has happened is to add a timeout when waiting for interrupts.

### 6.3.4.3 Observed bad IRQ behavior

Loko's handling of IRQs is experimental, as mentioned. These problems have been observed so far:

- There are sometimes bytes lost on the PC UARTs.
- Wrong hardware implementations of the i8042 controller (the PS/2 bus chip) seem to assume that IRQ 1 always interrupts the IRQ 12 handler. This results in keyboard data being read as mouse data.
- The PIT's IRQ 0 does not reach the CPU if the scheduler is running. The PIT is not used normally, so this is not a problem.

The UART and i8042 may need special interrupt handlers. But they are legacy devices and much amd64 hardware doesn't even have them.

### 6.3.5 Interrupts on Linux AMD64

The signal handlers are in `(loko arch amd64 linux-start)`.

Signals under Linux are similar to interrupts. Just like with interrupts, some signals are traps ('BUS', 'SEGV', 'FPE', 'ILL', etc). Linux places the processor's trap number in the sigcontext of the signal handler, which makes it easy to handle them identically to the way traps are handled on bare hardware.

Other signals are from external sources and the program is innocent, so the current stack frame must be preserved. On bare hardware this is accomplished by the IST and the equivalent mechanism on Linux is called `sigaltstack(2)`.

The normal userspace ABI, documented in System V Application Binary Interface AMD64 Architecture Processor Supplement (<https://github.com/hjl-tools/x86-psABI/wiki/X86-psABI>), is not followed. Normally interrupts would be delivered on the same stack as the program is currently using, but that would mess up Scheme code due to the way Loko manages stack frames. The "red zone" concept does not exist in Loko.



# Appendix A GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.  
<https://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

## 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

### 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

### 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any,

- be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
  - C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
  - D. Preserve all the copyright notices of the Document.
  - E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
  - F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
  - G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
  - H. Include an unaltered copy of this License.
  - I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
  - J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
  - K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
  - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
  - M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
  - N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
  - O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their



titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements.”

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/licenses/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

## 11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

## ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C) year your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts. A copy of the license is included in the section entitled ‘‘GNU
Free Documentation License’’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

# Index

## A

AGPL ..... 2

## B

bytevector-address ..... 15

## C

channel? ..... 13  
 choice-operation ..... 14  
 collections ..... 13  
 cvar? ..... 14

## D

disassemble ..... 10

## E

exit-current-task ..... 15  
 expand ..... 10  
 expand/optimize ..... 10

## G

gensym ..... 12  
 get-i/o-u16 ..... 16  
 get-i/o-u16-n! ..... 17  
 get-i/o-u32 ..... 16  
 get-i/o-u32-n! ..... 17  
 get-i/o-u8 ..... 16  
 get-i/o-u8-n! ..... 17  
 get-mem-s61 ..... 15  
 get-mem-u16 ..... 15  
 get-mem-u32 ..... 15  
 get-mem-u8 ..... 15  
 get-message ..... 13  
 get-operation ..... 13  
 get-output-string ..... 12

## I

include ..... 9  
 installed-libraries ..... 10  
 interaction-environment ..... 9

## L

license ..... 2  
 load ..... 9  
 loko-version ..... 12

## M

machine-type ..... 10  
 make-channel ..... 13  
 make-cvar ..... 14  
 make-parameter ..... 12

## O

open-output-string ..... 12

## P

parameterize ..... 12  
 perform-operation ..... 14  
 port-file-descriptor ..... 12  
 port-file-descriptor-set! ..... 12  
 put-i/o-u16 ..... 16  
 put-i/o-u32 ..... 16  
 put-i/o-u8 ..... 16  
 put-mem-s61 ..... 16  
 put-mem-u16 ..... 16  
 put-mem-u32 ..... 16  
 put-mem-u8 ..... 16  
 put-message ..... 13  
 put-operation ..... 13  
 putenv ..... 13

## R

run-fibers ..... 15

## S

signal-cvar! ..... 14  
 sleep ..... 13  
 sleep-operation ..... 14  
 spawn-fiber ..... 13  
 syscall ..... 15

## T

time ..... 11  
 time-it ..... 11  
 time-it\* ..... 11  
 timer-operation ..... 14

## U

uninstall-library ..... 10

## V

void ..... 9

**W**

|                                   |    |
|-----------------------------------|----|
| <code>wait</code> .....           | 15 |
| <code>wait-operation</code> ..... | 15 |
| <code>wrap-operation</code> ..... | 14 |

**Y**

|                                       |    |
|---------------------------------------|----|
| <code>yield-current-task</code> ..... | 15 |
|---------------------------------------|----|