

Ikarus Scheme User's Guide

(Preliminary Document)

Version 0.0.3

Abdulaziz Ghuloum
February 2, 2008

Ikarus Scheme User's Guide

Copyright © 2007,2008, Abdulaziz Ghuloum

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License version 3 as published by the Free Software Foundation.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <http://www.gnu.org/licenses/>.

Contents

Contents	iii
1 Getting Started	1
1.1 Introduction	1
1.2 Technology overview	2
1.3 System requirements	2
1.3.1 Hardware	2
1.3.2 Operating systems	3
1.3.3 Additional software	3
1.4 Installation	4
1.4.1 Installation details	4
1.4.2 Uninstalling Ikarus	6
1.5 Command-line switches	7
1.6 Using scheme-script	8
1.7 Mapping library names to file names	10
1.8 Writing cross-implementation libraries	11
2 R⁶RS Crash Course	13
2.1 Writing a simple script	14
2.2 Writing simple libraries	15
2.3 R ⁶ RS record types	17
2.3.1 Defining new record types	17
2.3.2 Extending existing record types	18
2.3.3 Specifying custom constructors	19
2.3.4 Custom constructors for derived record types	20
2.4 Exception handling	21
3 The (ikarus) library	25
3.1 Parameters	28

3.2	Local library imports	31
3.3	Local modules	32
3.4	Gensyms	33
3.5	Printing	37
3.6	Tracing	44
3.7	Timing	47
4	Contributed Libraries	49
4.1	(srfi *)	52
4.2	(math number-theory)	54
4.3	(r6rs-clos)	55
4.4	SRFI-41: (streams)	56
5	Missing Features	57
5.1	List of missing R ⁶ RS procedures	58
	Bibliography	59
	Index	60

Chapter 1

Getting Started

1.1 Introduction

Ikarus Scheme is an implementation of the Scheme programming language. The preliminary release of Ikarus implements the majority of the features found in the current standard, the Revised⁶ report on the algorithmic language Scheme[6] including full R⁶RS library and script syntax, syntax-case, unicode strings, bytevectors, user-defined record types, exception handling, conditions, and enumerations. More than 94% of the R⁶RS procedures and keywords are currently implemented and subsequent releases will proceed towards bringing Ikarus to full R⁶RS conformance.

The main purpose behind releasing Ikarus early is to give Scheme programmers the opportunity to experiment with the various new features that were newly introduced in R⁶RS. The most important of such features is the ability to structure large programs into libraries; where each library extends the language through procedural and syntactic abstractions. Many useful libraries can be written using the currently supported set of R⁶RS features including text processing tools, symbolic logic systems, interpreters and compilers, and many mathematical and scientific packages. It is my hope that this release will encourage the Scheme community to write and to share their most useful R⁶RS libraries.

1.2 Technology overview

Ikarus Scheme provides the programmer with many advantages:

Optimizing code generator: The compiler's backend employs state of the art technologies in code generation that produce fast efficient machine code. When developing computationally intensive programs, one is not constrained by using a slow interpreter.

Fast incremental compilation: Every library and script is quickly compiled to native machine code. When developing large software, one is not constrained by how slow the batch compiler runs.

Robust and fine-tuned standard libraries: The standard libraries are written such that they perform as much error checking as required to provide a safe and fast runtime environment.

Multi-generational garbage collector: The BiBOP[3] based garbage collector used in Ikarus allows the runtime system to expand its memory footprint as needed. The entire 32-bit virtual address space could be used and unneeded memory is released back to the operating system.

Supports many operating systems: Ikarus runs on the most popular and widely used operating systems for servers and personal computers. The supported systems include Mac OS X, GNU/Linux, FreeBSD, NetBSD, and Microsoft Windows.

1.3 System requirements

1.3.1 Hardware

Ikarus Scheme runs on the IA-32 (x86) architecture supporting SSE2 extensions. This includes the Athlon 64, Sempron 64, and Turion 64 processors from AMD and the Pentium 4, Xeon, Celeron, Pentium M, Core, and Core2 processors from Intel. The system does not run on Intel Pentium III or earlier processors.

The Ikarus compiler generates SSE2 instructions to handle Scheme's IEEE floating

point representation (*flonums*) for inexact numbers.

1.3.2 Operating systems

Ikarus is tested under the following operating systems:

- Mac OS X version 10.4 and 10.5.
- Linux 2.6.18 (Debian, Fedora, Gentoo, and Ubuntu).
- FreeBSD version 6.2.
- NetBSD version 3.1.
- Microsoft Windows XP (using Cygwin 1.5.24).

1.3.3 Additional software

- **GMP:** Ikarus uses the GNU Multiple Precision Arithmetic Library (GMP) for some bignum arithmetic operations. To build Ikarus from scratch, GMP version 4.2 or better must be installed along with the required header files. Pre-built GMP packages are available for most operating systems. Alternatively, GMP can be downloaded from <http://gmplib.org/>.

Note: Ikarus runs in 32-bit mode only. To run it in 64-bit environments, you will have to obtain the 32-bit version of GMP, or compile it yourself after adding `ABI=32` to its configuration options.

- **GCC:** The GNU C Compiler is required to build the Ikarus executable (e.g. the garbage collector, loader, and OS-related runtime). GCC versions 4.1 and 4.2 were successfully used to build Ikarus.
- **Autoconf and Automake:** The GNU Autoconf (version 2.61) and GNU Automake (version 1.10) tools are required if one wishes to modify the Ikarus source base. They are not required to build the official release of Ikarus.

- **X_YL^AT_EX**: The X_YL^AT_EX typesetting system is required for building the documentation. X_YL^AT_EX (and X_YT_EX) is an implementation of the L^AT_EX (and T_EX) typesetting system. X_YL^AT_EX can be obtained from <http://scripts.sil.org/xetex> and is included with T_EX-Live¹ and Mac-T_EX² distributions.

1.4 Installation

If you are familiar with installing Unix software on your system, then all you need to know is that Ikarus uses the standard installation method found in most other Unix software. Simply run the following commands from the shell:

```
$ tar -zxf ikarus-n.n.n.tar.gz
$ cd ikarus-n.n.n
$ ./configure [--prefix=path] [CFLAGS=-I/dir] [LDFLAGS=-L/dir]
$ make
$ make install
$
```

The rest of this section describes the build process in more details. It is targeted to users who are unfamiliar with steps mentioned above.

1.4.1 Installation details

1. Download the Ikarus source distribution. The source is distributed as a gzip-compressed tar file (`ikarus-n.n.n.tar.gz` where `n.n.n` is a 3-digit number indicating the current revision). The latest revision can be downloaded from the following URL:

<http://www.cs.indiana.edu/~aghuloum/ikarus/>

2. Unpack the source distribution package. From your shell command, type:

```
$ tar -zxf ikarus-n.n.n.tar.gz
$
```

¹<http://tug.org/texlive/>

²<http://tug.org/mactex/>

This creates the base directory `ikarus-n.n.n`.

3. Configure the build system by running the configure script located in the base directory. To do this, type the following commands:

```
$ cd ikarus-n.n.n
$ ./configure
checking build system type... i386-apple-darwin8.10.1
checking host system type... i386-apple-darwin8.10.1
...
configure: creating ./config.status
config.status: creating Makefile
config.status: creating src/Makefile
config.status: creating scheme/Makefile
config.status: creating doc/Makefile
config.status: executing depfiles commands
$
```

This configures the system to be built then installed in the system-wide location (binaries are installed in `/usr/local/bin`) . If you wish to install it in another location (e.g. in your home directory), you can supply a `--prefix` location to the configure script as follows:

```
$ ./configure --prefix=/path/to/installation/location
```

The configure script will fail if it cannot locate the location where GMP is installed. If running configure fails to locate GMP, you should supply the location in which the GMP header file, `gmp.h`, and the GMP library file, `libgmp.so`, are installed. This is done by supplying the two paths in the `CFLAGS` and `LDFLAGS` arguments:

```
$ ./configure CFLAGS=-I/path/to/include LDFLAGS=-L/path/to/lib
```

4. Build the system by running:

```
$ make
```

This performs two tasks. First, it builds the `ikarus` executable from the C files located in the `src` directory. It then uses the `ikarus` executable and the pre-built `ikarus.boot.orig` boot file to rebuild the Scheme boot image file `ikarus.boot` from the Scheme sources located in the `scheme` directory.

5. Install Ikarus by typing:

```
$ make install
```

If you are installing Ikarus in a system-wide location, you might need to have administrator privileges (use the `sudo` or `su` commands).

6. Test that Ikarus runs from the command line.

```
$ ikarus
Ikarus Scheme version 0.0.3
Copyright (c) 2006-2008 Abdulaziz Ghuloum

>
```

If you get the prompt, then Ikarus was successfully installed on your system. You may need to update the `PATH` variable in your environment to contain the directory in which the `ikarus` executable was installed.

Do not delete the `ikarus-n.n.n` directory from which you configured, built, and installed Ikarus. It will be needed if you decide at a later time to uninstall Ikarus.

1.4.2 Uninstalling Ikarus

To uninstall Ikarus, use the following steps:

```
$ cd path/to/ikarus-n.n.n
$ make uninstall
$
```

1.5 Command-line switches

The `ikarus` executable recognizes a few command-line switches that influence how Ikarus starts.

- `ikarus -h`

The presence of the `-h` flag causes `ikarus` to display a help message then exits. The help message summarizes the command-line switches. No further action is performed.

- `ikarus -b path/to/boot/file.boot`

The `-b` flag (which requires an extra argument) directs `ikarus` to use the specified boot file as the initial system boot file. The boot file is a binary file that contains all the code and data of the Scheme system. In the absence of `-b` flag, the executable will use the default boot file. Running `ikarus -h` shows the location where the default boot file was installed.

The rest of the command-line arguments are recognized by the standard Scheme run time system. They are processed after the boot file is loaded.

- `ikarus --r6rs-script script-file-name [arguments ...]`

The `--r6rs-script` argument instructs Ikarus that the supplied file is an R⁶RS script. See Section 2.1 for a short introduction to writing R⁶RS scripts. The script file name and any additional optional arguments can be obtained by calling the command-line procedure.

```
$ cat test.ss
(import (rnrs))
(write (command-line))
(newline)

$ ikarus --r6rs-script test.ss hi there
("test.ss" "hi" "there")
$
```

- `ikarus files ... [-- arguments ...]`

The lack of an `--r6rs-script` argument causes Ikarus to start in interactive mode. Each of the files is first loaded, in the interaction environment. The interaction environment initially contains all the bindings exported from the `(ikarus)` library (see Chapter 3). The optional arguments following the `--` marker can be obtained by calling the command-line procedure. In interactive mode, the first element of the returned list will be the string `"*interactive*`", corresponding to the script name in R⁶RS-script mode.

Note: The interactive mode is intended for quickly experimenting with the built-in features. It is intended neither for developing applications nor for writing any substantial pieces of code. The main reason for this is that the interaction between R⁶RS libraries and the interactive environment is not well understood. We hope to achieve better interaction between the two subsystems in the future.

1.6 Using scheme-script

Scheme scripts can be executed using the `ikarus --r6rs-script script-name` command as described in the previous section. For convenience, Ikarus follows the R⁶RS recommendations and installs a wrapper program called `scheme-script`. Typically, a script you write would start with a `#!` line that directs your operating system to the interpreter used to evaluate the script file. The following example shows a very simple script that uses the `scheme-script` command.

```
#!/usr/bin/env scheme-script

(import (rnrs))
(display "Hello World\n")
```

If the above script was placed in a file called `hello-world`, then one can make it executable using the `chmod` Unix command.

```
$ cat hello-world
#!/usr/bin/env scheme-script

(import (rnrs))
(display "Hello World\n")

$ chmod 755 hello-world
$ ./hello-world
Hello World
$
```

Under Mac OS X, if a script name ends with the `.command` extension, then it can be executed from the Finder by double-clicking on it. This brings up a terminal window in which the script is executed. The `.command` extension can be hidden from the *Get Info* item from the Finder's File menu.

1.7 Mapping library names to file names

The name of an R⁶RS library consists of a non-empty list of identifiers (symbols), followed by an optional version number. All of the standard R⁶RS libraries are built into Ikarus, thus importing any one of them does not require any special action other than listing the library name in the `import` part of a library or a script. The same holds for the `(ikarus)` library (chapter 3, page 25).

When writing a new library to a file, Ikarus uses a simple mechanism to map library names to file names. A library name is converted to a file path by joining the library identifiers with a path separator, e.g. `"/"`.

Library Name	⇒	File name
<code>(foo)</code>	⇒	<code>foo</code>
<code>(foo bar)</code>	⇒	<code>foo/bar</code>
<code>(foo bar baz)</code>	⇒	<code>foo/bar/baz</code>

Having mapped a library name to a file path, Ikarus attempts to locate that file in one of several locations. The locations attempted depend on two settings: the search path and the file prefix set (e.g., `.sls`, `.ss`, `.scm`, etc.). First, Ikarus attempts to locate the file in the current working directory from which Ikarus was invoked. In the current working directory, Ikarus enumerates all file prefixes first before searching other locations. If the file is not found in the current directory, Ikarus tries to find it in the Ikarus library directory. The Ikarus library directory is determined when Ikarus is installed (based on the `--prefix` argument that was passed to the `configure` script). If Ikarus fails to locate the library file, it raises an exception and exits. See Chapter 4 for more details about the library locations.

Tip: Use simple library names for the libraries that you define. Library names that contain non-printable characters, complex punctuations, or unicode may pose a challenge for some operating systems. If Ikarus cannot find a library, it will raise an error listing the locations in which it looked, helping you move the library file to a place where Ikarus can find it.

1.8 Writing cross-implementation libraries

When searching for a library, Ikarus appends a prefix (e.g., `.ss`) to the appropriate file name (e.g., `foo/bar`). The initial set of file extensions are: `.ikarus.sls`, `.ikarus.ss`, `.ikarus.scm`, `.sls`, `.ss`, and `.scm`.

The list of file extensions are searched sequentially. As a consequence, files ending with the `.ikarus.*` prefixes are given precedence over files that have generic Scheme extensions. The rationale for this behavior is to facilitate writing cross-implementation libraries: ones that take advantage of implementation-specific features, while at the same time provide a fail-safe alternative for other R⁶RS implementations.

Consider for example a program which would like to use the `pretty-print` procedure to format some code, and suppose further that pretty printing is just a nice add-on (e.g., using `write` suffices, but pretty-printing is *just prettier*) Ikarus exports a good pretty-printing facility in its `(ikarus)` library. However, since `pretty-print` is not a standard procedure, a program that uses it would be rendered unportable to other R⁶RS Scheme implementations.

The programmer can put the `.ikarus.*` extensions to use in this situation. First, the programmer writes two versions of a `(pretty-printing)` library: one for use by Ikarus, and one portable for other implementations.

```
(library (pretty-printing) ;;; this is pretty-printing.ikarus.ss
  (export pretty-print)    ;;; can only be used by Ikarus
  (import (only (ikarus) pretty-print)))
```

```
(library (pretty-printing) ;;; this is pretty-printing.sls
  (export pretty-print)    ;;; *portable* though not very pretty.
  (import (rnrs))          ;;; for any other implementation
  (define (pretty-print x port)
    (write x port)
    (newline port)))
```

Chapter 2

R⁶RS Crash Course

The major difference between R⁵RS and R⁶RS is the way in which programs are loaded and evaluated.

In R⁵RS, Scheme implementations typically start as an interactive session (often referred to as the REPL, or read-eval-print-loop). Inside the interactive session, the user enters definitions and expressions one at a time using the keyboard. Files, which also contain definitions and expressions, can be loaded and reloaded by calling the load procedure. The environment in which the interactive session starts often contains implementation-specific bindings that are not found in R⁵RS and users may redefine any of the initial bindings. The semantics of loading a file depends on the state of the environment at the time the file contents are evaluated.

R⁶RS differs from R⁵RS in that it specifies how *whole programs*, or scripts, are compiled and evaluated. An R⁶RS script is closed in the sense that all the identifiers found in the body of the script must either be defined in the script or imported from a library. R⁶RS also specifies how *libraries* can be defined and used. While files in R⁵RS are typically *loaded* imperatively into the top-level environments, R⁶RS libraries are *imported* declaratively in scripts and in other R⁶RS libraries.

2.1 Writing a simple script

An R⁶RS script is a set of definitions and expressions preceded by an `import` form. The `import` form specifies the language (i.e. the variable and keyword bindings) in which the library body is written. A very simple example of an R⁶RS script is listed below.

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

The first line imports the `(rnrs)` library. All the bindings exported from the `(rnrs)` library are made available to be used within the body of the script. The exports of the `(rnrs)` library include variables (e.g. `cons`, `car`, `display`, etc.) and keywords (e.g. `define`, `lambda`, `quote`, etc.). The second line displays the string `Hello World!` followed by a new line character.

In addition to expressions, such as the call to `display` in the previous example, a script may define some variables. The script below defines the variable `greeting` and calls the procedure bound to it.

```
#!/usr/bin/env scheme-script
(import (rnrs))

(define greeting
  (lambda ()
    (display "Hello World!\n")))

(greeting)
```

Additional keywords may be defined within a script. In the example below, we define the `(do-times n exprs ...)` macro that evaluates the expressions `exprs` `n` times. Running the script displays `Hello World` 3 times.

```
#!/usr/bin/env scheme-script
(import (rnrs))

(define greeting
  (lambda ()
    (display "Hello World!\n")))

(define-syntax do-times
  (syntax-rules ()
    [(_ n exprs ...)
     (let f ([i n])
       (unless (zero? i)
         exprs ...
         (f (- i 1))))]))

(do-times 3 (greeting))
```

2.2 Writing simple libraries

A script is intended to be a small piece of the program—useful abstractions belong to libraries. The `do-times` macro that was defined in the previous section may be useful in places other than printing greeting messages. So, we can create a small library, `(iterations)` that contains common iteration forms.

An R⁶RS library form is made of four essential parts: (1) the library name, (2) the set of identifiers that the library exports, (3) the set of libraries that the library imports, and (4) the body of the library.

The library name can be any non-empty list of identifiers. R⁶RS-defined libraries includes `(rnrs)`, `(rnrs unicode)`, `(rnrs bytevectors)`, and so on.

The library exports are a set of identifiers that are made available to importing libraries. Every exported identifier must be bound: it may either be defined in the library or imported using the `import` form. Library exports include variables, keywords, record names, and condition names.

Library imports are similar to script imports: they specify the set of libraries whose exports are made visible within the body of the library.

The body of a library contains definitions (variable, keyword, record, condition, etc.) followed by an optional set of expressions. The expressions are evaluated for side effect when needed.

The (iteration) library may be written as follows:

```
(library (iteration)
  (export do-times)
  (import (rnrs))

  (define-syntax do-times
    (syntax-rules ()
      [(_ n exprs ...)
       (let f ([i n])
         (unless (zero? i)
           exprs ...)
         (f (- i 1))))]))
```

To use the (iteration) library in our script, we add the name of the library to the script's import form. This makes all of (iteration)'s exported identifiers, e.g. do-times, visible in the body of the script.

```
#!/usr/bin/env scheme-script
(import (rnrs) (iteration))

(define greeting
  (lambda ()
    (display "Hello World!\n")))

(do-times 3 (greeting))
```

2.3 R⁶RS record types

R⁶RS provides ways for users to define new types, called record types. A record is a fixed-size data structure with a unique type (called a record type). A record may have any finite number of fields that hold arbitrary values. This section briefly describes what we expect to be the most commonly used features of the record system. Full details are in the R⁶RS Standard Libraries document[7].

2.3.1 Defining new record types

To define a new record type, use the `define-record-type` form. For example, suppose we want to define a new record type for describing points, where a point is a data structure that has two fields to hold the point's x and y coordinates. The following definition achieves just that:

```
(define-record-type point
  (fields x y))
```

The above use of `define-record-type` defines the following procedures automatically for you:

- The constructor `make-point` that takes two arguments, x and y and returns a new record whose type is `point`.
- The predicate `point?` that takes an arbitrary value and returns `#t` if that value is a point, `#f` otherwise.
- The accessors `point-x` and `point-y` that, given a record of type `point`, return the value stored in the x and y fields.

Both the x and y fields of the `point` record type are *immutable*, meaning that once a record is created with specific x and y values, they cannot be changed later. If you want the fields to be *mutable*, then you need to specify that explicitly as in the following example.

```
(define-record-type point
  (fields (mutable x) (mutable y)))
```

This definition gives us, in addition to the constructor, predicate, and accessors, two additional procedures:

- The mutators `point-x-set!` and `point-y-set!` that, given a record of type `point`, and a new value, sets the value stored in the `x` field or `y` field to the new value.

Note: Records in Ikarus have a printable representation in order to enable debugging programs that use records. Records are printed in the `#[type-name field-values ...]` notation. For example, `(write (make-point 1 2))` produces `#[point 1 2]`.

2.3.2 Extending existing record types

A record type may be extended by defining new variants of a record with additional fields. In our running example, suppose we want to define a `colored-point` record type that, in addition to being a `point`, it has an additional field: a *color*. A simple way of achieving that is by using the following record definition:

```
(define-record-type cpoint
  (parent point)
  (fields color))
```

Here, the definition of `cpoint` gives us:

- A constructor `make-cpoint` that takes three arguments (`x`, `y`, and `color` in that order) and returns a `cpoint` record.

- A predicate `cpoint?` that takes a single argument and determines whether the argument is a `cpoint` record.
- An accessor `cpoint-color` that returns the value of the `color` field of a `cpoint` object.

All procedures that are applicable to records of type `point` (`point?`, `point-x`, `point-y`) are also applicable to records of type `cpoint` since a `cpoint` is also a `point`.

2.3.3 Specifying custom constructors

The record type definitions explained so far use the default constructor that takes as many arguments as there are fields and returns a new record type with the values of the fields initialized to the arguments' values. It is sometimes necessary or convenient to provide a constructor that performs more than the default constructor. For example, we can modify the definition of our `point` record so that the constructor takes either no arguments, in which case it would return a point located at the origin, or two arguments specifying the x and y coordinates. We use the `protocol` keyword for specifying such constructor as in the following example:

```
(define-record-type point
  (fields x y)
  (protocol
    (lambda (new)
      (case-lambda
        [(x y) (new x y)]
        [()   (new 0 0)]))))
```

The `protocol` here is a procedure that takes a constructor procedure `new` (`new` takes as many arguments as there are fields.) and returns the desired custom constructor that we want (The actual constructor will be the value of the `case-lambda` expression in the example above). Now the constructor `make-point` would either take two arguments which constructs a point record as before, or no arguments, in which case `(new 0 0)` is called to construct a point at the origin.

Another reason why one might want to use custom constructors is to precompute the initial values of some fields based on the values of other fields. An example of this case is adding a distance field to the record type which is computed as $d = \sqrt{x^2 + y^2}$. The protocol in this case may be defined as:

```
(define-record-type point
  (fields x y distance)
  (protocol
    (lambda (new)
      (lambda (x y)
        (new x y (sqrt (+ (expt x 2) (expt y 2))))))))
```

Note that derived record types need not be modified when additional fields are added to the parent record type. For example, our `cpoint` record type still works unmodified even after we added the new distance field to the parent. Calling `(point-distance (make-cpoint 3 4 #xFF0000))` returns `5.0` as expected.

2.3.4 Custom constructors for derived record types

Just like how base record types (e.g. `point` in the running example) may have a custom constructor, derived record types can also have custom constructors that do other actions. Suppose that you want to construct `cpoint` records using an optional color that, if not supplied, defaults to the value `0`. To do so, we supply a protocol argument to `define-record-type`. The only difference here is that the procedure `new` is a *curried* constructor. It first takes as many arguments as the constructor of the parent record type, and returns a procedure that takes the initial values of the new fields.

In our example, the constructor for the `point` record type takes two arguments. `cpoint` extends `point` with one new field. Therefore, `new` in the definition below first takes the arguments for `point`'s constructor, then takes the initial color value. The definition below shows how the custom constructor may be defined.

```
(define-record-type cpoint
  (parent point)
  (fields color)
  (protocol
    (lambda (new)
      (case-lambda
        [(x y c) ((new x y) c)]
        [(x y)  ((new x y) 0)]))))
```

2.4 Exception handling

The procedure `with-exception-handler` allows the programmer to specify how to handle exceptional situations. It takes two procedures as arguments:

- An exception handler which is a procedure that takes a single argument, the object that was raised.
- A body thunk which is a procedure with no arguments whose body is evaluated with the exception handler installed.

In addition to installing exception handlers, R⁶RS provides two ways of raising exceptions: `raise` and `raise-continuable`. We describe the `raise-continuable` procedure first since it's the simpler of the two. For the code below, assume that `print` is defined as:

```
(define (print who obj)
  (display who)
  (display ": ")
  (display obj)
  (newline))
```

The first example, below, shows how a simple exception handler is installed. Here, the exception handler prints the object it receives and returns the symbol `there`.

The call to `raise-continuable` calls the exception handler, passing it the symbol here. When the handler returns, the returned value becomes the value of the call to `raise-continuable`.

```
(with-exception-handler
  (lambda (obj)
    (print "handling" obj)
    'there)
  (lambda ()
    (print "returned" (raise-continuable 'here))))
```

Exceptional handlers may nest, and in that case, if an exception is raised while evaluating an inner handler, the outer handler is called as the following example illustrates:

```
(with-exception-handler
  (lambda (obj)
    (print "outer" obj)
    'outer)
  (lambda ()
    (with-exception-handler
      (lambda (obj)
        (print "inner" obj)
        (raise-continuable 'there))
      (lambda ()
        (print "returned" (raise-continuable 'here))))))
```

In short, `with-exception-handler` binds an exception handler within the dynamic context of evaluating the thunk, and `raise-continuable` calls it.

The procedure `raise` is similar to `raise-continuable` except that if the handler returns, a new exception is raised, calling the next handler in sequence until the list of handlers is exhausted.

```

(call/cc                                     ;;; prints
  (lambda (escape)                          ;;; inner: here
    (with-exception-handler                 ;;; outer: #[condition ---]
      (lambda (obj)                         ;;; returns
        (print "outer" obj)                 ;;; 12
        (escape 12))
      (lambda ()
        (with-exception-handler
          (lambda (obj)
            (print "inner" obj)
            'there)
          (lambda ()
            (print "returned" (raise 'here))))))))))

```

Here, the call to `raise` calls the inner exception handler, which returns, causing `raise` to re-raise a non-continuable exception to the outer exception handler. The outer exception handler then calls the escape continuation.

The following procedure provides a useful example of using the exception handling mechanism. Consider a simple definition of the procedure `configuration-option` which returns the value associated with a key where the key/value pairs are stored in an association list in a configuration file.

```

(define (configuration-option filename key)
  (cdr (assq key (call-with-input-file filename read))))

```

Possible things may go wrong with calling `configuration-option` including errors opening the file, errors reading from the file (file may be corrupt), error in `assq` since what's read may not be an association list, and error in `cdr` since the key may not be in the association list. Handling all error possibilities is tedious and error prone. Exceptions provide a clean way of solving the problem. Instead of guarding against all possible errors, we install a handler that suppresses all errors and returns a default value if things go wrong. Error handling for `configuration-option` may be added as follows:

```
(define (configuration-option filename key default)
  (define (getopt)
    (cdr (assq key (call-with-input-file filename read))))
  (call/cc
    (lambda (k)
      (with-exception-handler
        (lambda (_) (k default))
        getopt))))))
```

Chapter 3

The (ikarus) library

In addition to the libraries listed in the R⁶RS standard, Ikarus contains the (ikarus) library which provides additional useful features. The (ikarus) library is a composite library—it exports a superset of all the supported bindings of R⁶RS. While not all of the exports of (ikarus) are documented at this time, this chapter attempts to describe a few of these useful extensions. Extensions to Scheme’s lexical syntax are also documented.

`#!ikarus`

reader syntax

Ikarus extends Scheme’s lexical syntax (R⁶RS Chapter 4) in a variety of ways including:

- end-of-file marker, `#!eof` (page 27)
- gensym syntax, `#{gensym}` (page 35)
- graph syntax, `#nn= #nn#` (page 40)

The syntax extensions are made available by default on all input ports, until the `#!r6rs` token is read. Thus, reading the `#!r6rs` token disables all extensions to the lexical syntax on the specific port, and the `#!ikarus` enables them again.

If you are writing code that is intended to be portable across different Scheme implementations, we recommend adding the `#!r6rs` token to the top of every script

and library that you write. This allows Ikarus to alert you when using non-portable features. If you're writing code that's intended to be Ikarus-specific, we recommend adding the `#!ikarus` token in order to get an immediate error when your code is run under other implementations.

port-mode **procedure**
(port-mode ip)

The `port-mode` procedure accepts an input port as an argument and returns one of `r6rs-mode` or `ikarus-mode` as a result. All input ports initially start in the `ikarus-mode` and thus accept Ikarus-specific reader extensions. When the `#!r6rs` token is read from a port, its mode changes to `ikarus-mode`.

```
> (port-mode (current-input-port))
ikarus-mode
> #!r6rs (port-mode (current-input-port))
r6rs-mode
> #!ikarus (port-mode (current-input-port))
ikarus-mode
```

set-port-mode! **procedure**
(set-port-mode! ip mode)

The `set-port-mode!` procedure modifies the lexical syntax accepted by subsequent calls to read on the input port. The mode is a symbol which should be one of `r6rs-mode` or `ikarus-mode`. The effect of setting the port mode is similar to that of reading the `#!r6rs` or `#!ikarus` from that port.

```
> (set-port-mode! (current-input-port) 'r6rs-mode)
> (port-mode (current-input-port))
r6rs-mode
```

#!eof**reader syntax**

The end-of-file marker, `#!eof`, is an extension to the R⁶RS syntax. The primary utility of the `#!eof` marker is to stop the reader (e.g. `read` and `get-datum`) from reading the rest of the file.

```
#!/usr/bin/env scheme-script
(import (ikarus))
<some code>
(display "goodbye\n")

#!eof
<some junk>
```

The `#!eof` marker also serves as a datum in Ikarus, much like `#t` and `#f`, when it is found inside other expressions.

```
> (eof-object)
#!eof
> (read (open-input-string ""))
#!eof
> (read (open-input-string "#!eof"))
#!eof
> (quote #!eof)
#!eof
> (eof-object? '#!eof)
#t
> #!r6rs #!eof
Unhandled exception
Condition components:
  1. &error
  2. &who: tokenize
  3. &message: "invalid syntax: #!e"
> #!ikarus #!eof
$
```

3.1 Parameters

Parameters in Ikarus¹ are intended for customizing the behavior of a procedure during the dynamic execution of some piece of code. Parameters are first class entities (represented as procedures) that hold the parameter value. A parameter procedure accepts either zero or one argument. If given no arguments, it returns the current value of the parameter. If given a single argument, it must set the state to the value of the argument. Parameters replace the older concept of using starred `*global*` customization variables. For example, instead of writing:

```
(define *screen-width* 72)
```

and then mutating the variable `*screen-width*` with `set!`, we could wrap the variable `*screen-width*` with a `screen-width` parameter as follows:

```
(define *screen-width* 72)
(define screen-width
  (case-lambda
    [()] *screen-width*]
    [(x) (set! *screen-width* x)]))
```

The value of `screen-width` can now be passed as argument, returned as a value, and exported from libraries.

make-parameter

procedure

```
(make-parameter x)
(make-parameter x f)
```

As parameters are common in Ikarus, the procedure `make-parameter` is defined to model the common usage pattern of parameter construction.

`(make-parameter x)` constructs a parameter with `x` as the initial value. For example, the code above could be written succinctly as:

¹Parameters are found in many Scheme implementations such as Chez Scheme and MzScheme.

```
(define screen-width (make-parameter 72))
```

`(make-parameter x f)` constructs a parameter which filters the assigned values through the procedure `f`. The initial value of the parameter is the result of calling `(f x)`. Typical uses of the filter procedure include checking some constraints on the passed argument or converting it to a different data type. The `screen-width` parameter may be constructed more robustly as:

```
(define screen-width
  (make-parameter 72
    (lambda (w)
      (assert (and (integer? w) (exact? w)))
      (max w 1))))
```

This definition ensures, through `assert`, that the argument passed is an exact integer. It also ensures, through `max` that the assigned value is always positive.

parameterize

syntax

```
(parameterize ([lhs* rhs*] ...) body body* ...)
```

Parameters can be assigned to by simply calling the parameter procedure with a single argument. The `parameterize` syntax is used to set the value of a parameter within the dynamic extent of the `body body* ...` expressions.

The `lhs* ...` are expressions, each of which must evaluate to a parameter. Such parameters are not necessarily constructed by `make-parameter`—any procedure that follows the parameters protocol works.

The advantage of using `parameterize` over explicitly assigning to parameters (same argument applies to global variables) is that you're guaranteed that whenever control exits the body of a `parameterize` expression, the value of the parameter is reset back to what it was before the body expressions were entered. This is true even in the presence of `call/cc`, errors, and exceptions.

The following example shows how to set the text property of a terminal window. The parameter `terminal-property` sends an ANSI escape sequence to the terminal

whenever the parameter value is changed. The use of `terminal-property` within `parameterize` changes the property before `(display "RED!")` is called and resets it back to normal when the body returns.

```
(define terminal-property
  (make-parameter "0"
    (lambda (x)
      (display "\x1b;[")
      (display x)
      (display "m")
      x)))

(display "Normal and ")
(parameterize ([terminal-property "41;37"])
  (display "RED!"))
(newline)
```

3.2 Local library imports

`import` **syntax**
(import import-spec* ...)

The `import` keyword which is exported from the (`ikarus`) library can be used anywhere definitions can occur: at a script body, library's top-level, or in internal definitions context. The syntax of the local `import` form is similar to the `import` that appears at the top of a library or a script form, and carries with it the same restrictions: no identifier name may be imported twice unless it denotes the same identifier; no identifier may be both imported and defined; and imported identifiers are immutable.

Local `import` forms are useful for two reasons: (1) they minimize the namespace clutter that usually occurs when many libraries are imported at the top level, and (2) they limit the scope of the import and thus help modularize a library's dependencies.

Suppose you are constructing a large library and at some point you realize that one of your procedures needs to make use of some other library for performing a specific task. Importing that library at top level makes it available for the entire library. Consequently, even if that library is no longer used anywhere in the code (say when the code that uses it is deleted), it becomes very hard to delete the import without first examining the entire library body for potential usage leaks. By locally importing a library into the appropriate scope, we gain the ability to delete the `import` form when the procedure that was using it is deleted.

3.3 Local modules

This section is not documented yet. Please refer to Section 10.5 of Chez Scheme User's Guide [2], Chapter 3 of Oscar Waddel's Ph.D Thesis [8], and its POPL99 paper [9] for details on using the `module` and `import` keywords. Ikarus's internal module system is similar in spirit to that of Chez Scheme.

<code>module</code>	<code>syntax</code>
<code>(module M definitions ... expressions ...)</code>	
<code>(module definitions ... expressions ...)</code>	

<code>import</code>	<code>syntax</code>
<code>(import M)</code>	

3.4 Gensyms

Gensym stands for a *generated symbol*—a fresh symbol that is generated at run time and is guaranteed to be *not* eq? to any other symbol present in the system. Gensyms are useful in many applications including expanders, compilers, and interpreters when generating an arbitrary number of unique names is needed.

Ikarus is similar to Chez Scheme in that the readers (including the read procedure) and writers (including write and pretty-print) maintain the read/write invariance on gensyms. When a gensym is written to an output port, the system automatically generates a random unique identifier for the gensym. When the gensym is read back through the `#{gensym}` read syntax, a new gensym is *not* regenerated, but instead, it is looked up in the global symbol table.

A gensym's name is composed of two parts: a *pretty* string and a *unique* string. The Scheme procedure `symbol->string` returns the pretty string of the gensym and not its unique string. Gensyms are printed by default as `#{pretty-string unique-string}`.

gensym	procedure
(gensym)	
(gensym string)	
(gensym symbol)	

The procedure `gensym` constructs a new gensym. If passed no arguments, it constructs a gensym with no pretty name. The pretty name is constructed when and if the pretty name of the resulting gensym is needed. If `gensym` is passed a string, that string is used as the pretty name. If `gensym` is passed a symbol, the pretty name of the symbol is used as the pretty name of the returned gensym. See `gensym-prefix` (page 42) and `gensym-count` (page 43) for details.

```
> (gensym)
#{g0 |y0zf>G1FvcTJE0xw|}
> (gensym)
#{g1 |U%X&sF6kX!YC8LW=|}
> (eq? (gensym) (gensym))
#f
```

(gensym string) constructs a new gensym with string as its pretty name. Similarly, (gensym symbol) constructs a new gensym with the pretty name of symbol, if it has one, as its pretty name.

```
> (gensym "foo")
#{foo |>Vg0l1CM&$dSvRN=|}
> (gensym 'foo)
#{foo |!TqQLmtw2hoEYfU>|}
> (gensym (gensym 'foo))
#{foo |N2C>500>C?OR0UBU|}
```

gensym?

procedure

(gensym? x)

The gensym? predicate returns #t if its argument is a gensym, and returns #f otherwise.

```
> (gensym? (gensym))
#t
> (gensym? 'foo)
#f
> (gensym? 12)
#f
```

gensym->unique-string

procedure

(gensym->unique-string gensym)

The gensym->unique-string procedure returns the unique name associated with the gensym argument.

```
> (gensym->unique-string (gensym))
"Yukro1LMgP?%ElcR"
```

```

#{gensym} reader syntax
#{unique-name}
#{pretty-name unique-name}
#:pretty-name

```

Ikarus's read and write procedures extend the lexical syntax of Scheme by the ability to read and write gensyms using one of the three forms listed above.

`#{unique-name}` constructs, at read time, a gensym whose unique name is the one specified. If a gensym with the same unique name already exists in the system's symbol table, that gensym is returned.

```

> ' #{some-long-name}
#{g0 |some-long-name|}
> (gensym? ' #{some-long-unique-name})
#t
> (eq? ' #{another-unique-name} ' #{another-unique-name})
#t

```

The two-part `#{pretty-name unique-name}` gensym syntax is similar to the syntax shown above with the exception that if a new gensym is constructed (that is, if the gensym did not already exist in the symbol table), the pretty name of the constructed gensym is set to `pretty-name`.

```

> ' #{foo unique-identifier}
#{foo |unique-identifier|}
> ' #{unique-identifier}
#{foo |unique-identifier|}
> ' #{bar unique-identifier}
#{foo |unique-identifier|}

```

The `#:pretty-name` form constructs, at read time, a gensym whose pretty name is `pretty-name` and whose unique name is fresh. This form guarantees that the resulting gensym is not `eq?` to any other symbol in the system.

```

> ' #:foo

```

```
#{foo |j=qTGlEwS/Zlp2Dj|}
> (eq? '#:foo '#:foo)
#f
```

generate-temporaries

example

The (rnrs syntax-case) library provides a generate-temporaries procedure, which takes a syntax object (representing a list of things) and returns a list of fresh identifiers. Using gensym, that procedure can be defined as follows:

```
(define (generate-temporaries* stx)
  (syntax-case stx ()
    [(x* ...)
     (map (lambda (x)
            (datum->syntax #'unimportant
              (gensym
               (if (identifier? x)
                   (syntax->datum x)
                   't))))
          #'(x* ...))]))
```

The above definition works by taking the input `stx` and destructuring it into the list of syntax objects `x* ...`. The inner procedure maps each `x` into a new syntax object (constructed with `datum->syntax`). The datum is a gensym, whose name is the same name as `x` if `x` is an identifier, or the symbol `t` if `x` is not an identifier. The output of `generate-temporaries*` generates names similar to their input counterpart:

```
> (print-gensym #f)
> (generate-temporaries* #'(x y z 1 2))
(#<syntax x> #<syntax y> #<syntax z> #<syntax t> #<syntax t>)
```

3.5 Printing

pretty-print **procedure**

```
(pretty-print datum)
(pretty-print datum output-port)
```

The procedure `pretty-print` is intended for printing Scheme data, typically Scheme programs, in a format close to how a Scheme programmer would write it. Unlike `write`, which writes its input all in one line, `pretty-print` inserts spaces and new lines in order to produce more pleasant output.

```
(define fact-code
  '(letrec ([fact (lambda (n) (if (zero? n) 1 (* n (fact (- n 1)))))]
    (fact 5)))

> (pretty-print fact-code)
(letrec ((fact
          (lambda (n) (if (zero? n) 1 (* n (fact (- n 1))))))
  (fact 5))
```

The second argument to `pretty-print`, if supplied, must be an output port. If not supplied, the `current-output-port` is used.

Limitations: As shown in the output above, the current implementation of `pretty-print` does not handle printing of square brackets properly.

pretty-width **parameter**

```
(pretty-width)
(pretty-width n)
```

The parameter `pretty-width` controls the number of characters after which the `pretty-print` starts breaking long lines into multiple lines. The initial value of

`pretty-width` is set to 60 characters, which is suitable for most terminals and printed material.

```
> (parameterize ([pretty-width 40])
  (pretty-print fact-code))
(letrec ((fact
  (lambda (n)
    (if (zero? n)
        1
        (* n (fact (- n 1)))))))
  (fact 5))
```

Note that `pretty-width` does not guarantee that the output will not extend beyond the specified number. Very long symbols, for examples, cannot be split into multiple lines and may force the printer to go beyond the value of `pretty-width`.

format **procedure**
 (format fmt-string args ...)

The procedure `format` produces a string formatted according to `fmt-string` and the supplied arguments. The format string contains markers in which the string representation of each argument is placed. The markers include:

"~s" instructs the formatter to place the next argument as if the procedure `write` has printed it. If the argument contains a string, the string will be quoted and all quotes and backslashes in the string will be escaped. Similarly, characters will be printed using the `#\x` notation.

"~a" instructs the formatter to place the next argument as if the procedure `display` has printed it. Strings and characters are placed as they are in the output.

"~b" instructs the formatter to convert the next argument to its binary (base 2) representation. The argument must be an exact number. Note that the `#b` numeric prefix is not produced in the output.

"~o" is similar to "~b" except that the number is printed in octal (base 8).

"~x" is similar to "~b" except that the number is printed in hexadecimal (base 16).


```
(print-graph #t)
(print-graph #f)
```

The graph notation is a way of marking and referencing parts of a data structure and, consequently, creating shared and cyclic data structures at read time instead of resorting to explicit mutation at run time. The `#n=` marks the following data structure with mark n , where n is a nonnegative integer. The `#n#` references the data structure marked n . Marks can be assigned and referenced in any order but each mark must be assigned to exactly once in an expression.

```
> (let ([x '#0=(1 2 3)])
    (eq? x '#0#))
#t
> (let ([x '#0#] [y '#0=(1 2 3)])
    (eq? x y))
#t
> (eq? (cdr '(12 . #1#)) '#1=(1 2 3))
#t
> (let ([x '#1=(#1# . #1#)])
    (and (eq? x (car x))
         (eq? x (cdr x))))
#t
```

The `print-graph` parameter controls how the writers (e.g. `pretty-print` and `write`) handle shared and cyclic data structures. In Ikarus, all writers detect cyclic data structures and they all terminate on all input, cyclic or otherwise.

If the value of `print-graph` is set to `#f` (the default), then the writers does not attempt to detect shared data structures. Any part of the input that is shared is printed as if no sharing is present. If the value of `print-graph` is set to `#t`, all sharing of data structures is marked using the `#n=` and `#n#` notation.

```
> (parameterize ([print-graph #f])
    (let ([x (list 1 2 3 4)])
      (pretty-print (list x x x))))
((1 2 3 4) (1 2 3 4) (1 2 3 4))
```

```

> (parameterize ([print-graph #t])
  (let ([x (list 1 2 3 4)])
    (pretty-print (list x x x))))
(#0=(1 2 3 4) #0# #0#)

> (parameterize ([print-graph #f])
  (let ([x (list 1 2)])
    (let ([y (list x x x x)])
      (set-car! (last-pair y) y)
      (pretty-print (list y y))))))
(#0=((1 2) (1 2) (1 2) #0#) #0#)

> (parameterize ([print-graph #t])
  (let ([x (list 1 2)])
    (let ([y (list x x x x)])
      (set-car! (last-pair y) y)
      (pretty-print (list y y))))))
(#0=(#1=(1 2) #1# #1# #0#) #0#)

```

print-gensym

parameter

```

(print-gensym)
(print-gensym #t)
(print-gensym #f)
(print-gensym 'pretty)

```

The parameter `print-gensym` controls how gensyms are printed by the various writers.

If the value of `print-gensym` is `#f`, then gensym syntax is suppressed by the writers and only the gensyms' pretty names are printed. If the value of `print-gensym` is `#t`, then the full `#{pretty unique}` syntax is printed. Finally, if the value of `print-gensym` is the symbol `pretty`, then gensyms are printed using the `#:pretty` notation.

```

> (parameterize ([print-gensym #f])
  (pretty-print (list (gensym) (gensym))))

```

```
(g0 g1)

> (parameterize ([print-gensym #t])
  (pretty-print (list (gensym) (gensym))))
({g2 |KR1M2&CTt1<B0n/m|} #{g3 |FBAb&7NC6&=c82!0|})

> (parameterize ([print-gensym 'pretty])
  (pretty-print (list (gensym) (gensym))))
(#{g4 #:g5})
```

The initial value of `print-gensym` is `#t`.

gensym-prefix

parameter

```
(gensym-prefix)
(gensym-prefix string)
```

The parameter `gensym-prefix` specifies the string to be used as the prefix to generated pretty names. The default value of `gensym-prefix` is the string "g", which causes generated strings to have pretty names in the sequence `g0`, `g1`, `g2`, etc.

```
> (parameterize ([gensym-prefix "var"]
  [print-gensym #f])
  (pretty-print (list (gensym) (gensym) (gensym))))
(var0 var1 var2)
```

Beware that the `gensym-prefix` controls how pretty names are generated, and has nothing to do with how `gensym` constructs a new `gensym`. In particular, notice the difference between the output in the first example with the output of the examples below:

```
> (pretty-print
  (parameterize ([gensym-prefix "var"] [print-gensym #f])
    (list (gensym) (gensym) (gensym))))
(g3 g4 g5)
```

```
> (let ([ls (list (gensym) (gensym) (gensym))])
      (parameterize ([gensym-prefix "var"] [print-gensym #f])
        (pretty-print ls)))
(var5 var6 var7)
```

gensym-count**parameter**

```
(gensym-count)
(gensym-count n)
```

The parameter `gensym-count` determines the number which is attached to the `gensym-prefix` when gensyms' pretty names are generated. The initial value of `gensym-count` is 0 and is incremented every time a pretty name is generated. It might be set to any non-negative integer value.

```
> (let ([x (gensym)])
      (parameterize ([gensym-count 100] [print-gensym #f])
        (pretty-print (list (gensym) x (gensym)))))
(g100 g101 g102)
```

Notice from all the examples so far that pretty names are generated in the order at which the gensyms are printed, not in the order in which gensyms were created.

3.6 Tracing

trace-define

syntax

```
(trace-define (name . args) body body* ...)
(trace-define name expression)
```

The trace-define syntax is similar to define except that the bound value, which must be a procedure, becomes a traced procedure. A traced procedure prints its arguments when it is called and prints its values when it returns.

```
> (trace-define (fact n)
    (if (zero? n) 1 (* n (fact (- n 1)))))
> (fact 5)
|(fact 5)
| (fact 4)
| |(fact 3)
| | (fact 2)
| | |(fact 1)
| | | (fact 0)
| | | 1
| | |1
| | 2
| |6
| 24
|120
120
```

The tracing facility in Ikarus preserves and shows tail recursion and distinguishes it from non-tail recursion by showing tail calls starting at the same line in which their parent was called.

```
> (trace-define (fact n)
    (trace-define (fact-aux n m)
      (if (zero? n) m (fact-aux (- n 1) (* n m))))
    (fact-aux n 1))
```

```

> (fact 5)
|(fact 5)
|(fact-aux 5 1)
|(fact-aux 4 5)
|(fact-aux 3 20)
|(fact-aux 2 60)
|(fact-aux 1 120)
|(fact-aux 0 120)
|120
120

```

Moreover, the tracing facility interacts well with continuations and exceptions.

```

> (call/cc
  (lambda (k)
    (trace-define (loop n)
      (if (zero? n)
          (k 'done)
          (+ (loop (- n 1)) 1)))
    (loop 5)))
|(loop 5)
| (loop 4)
| |(loop 3)
| | (loop 2)
| | |(loop 1)
| | | (loop 0)
done

```

trace-lambda

syntax

```
(trace-lambda name args body body* ...)
```

The `trace-lambda` macro is similar to `lambda` except that the resulting procedure is traced: it prints the arguments it receives and the results it returns.

make-traced-procedure**procedure**`(make-traced-procedure name proc)`

The procedure `make-traced-procedure` takes a name (typically a symbol) and a procedure. It returns a procedure similar to `proc` except that it traces its arguments and values.

```

> (define (fact n)
  (if (zero? n)
      (lambda (k) (k 1))
      (lambda (k)
        ((fact (- n 1))
         (make-traced-procedure `(k ,n)
                                (lambda (v)
                                  (k (* v n))))))))))
> (call/cc
  (lambda (k)
    ((fact 5) (make-traced-procedure 'K k))))
|((k 1) 1)
|((k 2) 1)
|((k 3) 2)
|((k 4) 6)
|((k 5) 24)
|(K 120)
120

```

3.7 Timing

This section describes some of Ikarus's timing facilities which may be useful for benchmarking and performance tuning.

time **syntax**
(time expression)

The `time` macro performs the following: it evaluates `expression`, then prints a summary of the run time statistics, then returns the values returned by `expression`. The run-time summary includes the number of bytes allocated, the number of garbage collection runs, and the time spent in both the mutator and the collector.

```
> (let (
      ;;; 10 million
      (define ls (time (vector->list (make-vector 10000000))))
      (time (append ls ls))
      (values))
  running stats for (vector->list (make-vector 10000000)):
  3 collections
  672 ms elapsed cpu time, including 547 ms collecting
  674 ms elapsed real time, including 549 ms collecting
  120012328 bytes allocated
  running stats for (append ls ls):
  4 collections
  1536 ms elapsed cpu time, including 1336 ms collecting
  1538 ms elapsed real time, including 1337 ms collecting
  160000040 bytes allocated
```

Note: The output listed above is *just a sample* that was taken at some point on some machine. The output on your machine at the time you read this may vary.

time-it**procedure**

(time-it who thunk)

The procedure `time-it` takes a datum denoting the name of the computation and a thunk (i.e. a procedure with no arguments), invokes the thunk, prints the stats, and returns the values obtained from invoking the thunk. If the value of `who` is non-`false`, `who` is used when displaying the run-time statistics. If the value of `who` is `#f`, then no name for the computation is displayed.

```
> (time-it "a very fast computation"
    (lambda () (values 1 2 3)))
running stats for a very fast computation:
  no collections
  0 ms elapsed cpu time, including 0 ms collecting
  0 ms elapsed real time, including 0 ms collecting
  56 bytes allocated
1
2
3

> (time-it #f (lambda () 12))
running stats:
  no collections
  0 ms elapsed cpu time, including 0 ms collecting
  0 ms elapsed real time, including 0 ms collecting
  32 bytes allocated
12
```

Chapter 4

Contributed Libraries

We try to keep Ikarus Scheme small and its complexity manageable. Libraries that are not an essential part of Ikarus are not included in the Ikarus proper, instead, they are distributed with Ikarus in source form. Such libraries may be written specifically for Ikarus, or they may be portable libraries that can be used in Ikarus. SRFIs or other libraries contributed by members of the Scheme community belong to this section.

Using contributed libraries is no different from using any of the built-in libraries—all one has to do is add the library name to the `import` clause and the rest is done by the system.

If you have written a useful R⁶RS library and wish for it to be available for a wider audience, contact us and we would be delighted to include information about it in the next release of Ikarus.

Note: Contributed libraries may have bugs on their own or may exhibit bugs in Ikarus itself. If you have a problem using any of these libraries, please try to resolve the issue by contacting the library author first. Do not hesitate to file a bug on Ikarus if you believe that Ikarus is at fault.

Library files

The contributed libraries are installed in your system when Ikarus was installed. By default, running the configure script installs the contributed libraries into the `/usr/local/lib/ikarus` directory. If a `--prefix DIR` argument was supplied to configure, then the libraries are installed in the `DIR/ikarus/lib` directory.

You may install additional libraries into the Ikarus library directory. Doing so makes them available for `import` into other libraries and script regardless of where the importing code is located or the current directory in which it is executed.

Defining `IKARUS_LIBRARY_PATH`

There may be situations in which you may wish to install your own libraries into a different location. For example, you may not have sufficient administrative privileges to write to the system directory, or you may wish to keep your own libraries separate from the standard libraries. Whatever the reason is, you can store your library files in any location you want and set up the `IKARUS_LIBRARY_PATH` environment variable to point to these locations. The value of `IKARUS_LIBRARY_PATH` is a colon-separated list of directories in which Ikarus will search.

For example, suppose your script imports the `(streams derived)` library. First, Ikarus will map the library name to the file path `streams/derived.ss`. Suppose that Ikarus was installed using the `--prefix /usr/local` configuration option, and suppose further that the value of `IKARUS_LIBRARY_PATH` is set by the user to be `/home/john/ikarus-libraries:/home/john/srfis`. Ikarus will search in the following locations in sequence until it finds the file it is looking for.

```
./streams/derived.ss  
/home/john/ikarus-libraries/streams/derived.ss  
/home/john/srfis/streams/derived.ss  
/usr/local/lib/ikarus/streams/derived.ss
```

Warning: The current behavior of Ikarus regarding the IKARUS_LIBRARY_PATH is preliminary and is likely to change in future releases.

The method in which the value of IKARUS_LIBRARY_PATH is defined is typically shell dependant. If you use GNU Bash, you typically set the values of environment variables in the ~/.bash_profile or ~/.bashrc file by adding the following lines:

```
IKARUS_LIBRARY_PATH=/path/to/some/directory:/and/another
export IKARUS_LIBRARY_PATH
```

4.1 (srfi *)

Ported by: Derick Eddington

Scheme Requests For Implementation (SRFIs) for R6RS/Ikarus can be found at:
<https://code.launchpad.net/~ikarus-srfi-team/ikarus-libraries/srfi>

Currently provided:

- SRFI-0: (srfi feature-expand)
- SRFI-1: (srfi lists)
- SRFI-2: (srfi and-let)
- SRFI-6: (srfi string-ports)
- SRFI-8: (srfi receive)
- SRFI-9: (srfi records)
- SRFI-11: (srfi let-values)
- SRFI-13: (srfi strings)
- SRFI-14: (srfi char-set)
- SRFI-16: (srfi case-lambda)
- SRFI-19: (srfi time)
- SRFI-23: (srfi error-reporting)
- SRFI-26: (srfi specialize-procedures)
- SRFI-27: (srfi random)
- SRFI-31: (srfi rec)
- SRFI-37: (srfi args-fold)
- SRFI-39: (srfi parameters)

- SRFI-41: (srfi streams) by Phil Bewig
- SRFI-42: (srfi eager-comprehensions)
- SRFI-43: (srfi vectors)
- SRFI-67: (srfi compare)
- SRFI-78: (srfi lightweight-testing)

To install (you'll need a new version of the Bazaar revision control system):

```
$ bzr checkout --lightweight http://bazaar.launchpad.net/  
~ikarus-srfi-team/ikarus-libraries/srfi
```

After you do the above, you'll get a new `srfi/` directory in the directory you ran the checkout command, and that parent directory of the new `srfi/` directory needs to be in your `IKARUS_LIBRARY_PATH`, so that attempts to import (`srfi ---`) will look in the directory containing the `srfi/` directory.

You can stay up-to-date by changing directory into your `srfi/` and doing:

```
$ bzr update
```

4.2 (math number-theory)

Provided by: Jens Axel Søgaard

URL: <https://code.launchpad.net/~soegaard/ikarus-libraries/soegaard>

This library contains number theory code that I have written over a long period. The code began as an experiment. I grabbed a book on number theory from the shelf (*Elementary Number Theory* by Gareth A. Jones and J. Mary Jones) and began illustrating each definition and each theorem with Scheme code. The first half of the source code is thus a well commented mix of definitions, theorems and code.

The second half contains more sophisticated algorithms mostly of from the excellent book *Modern Computer Algebra* by Joachim von zur Gathen and Jürgen Gerhard. The algorithms for factorizing large integers come from this book.

Finally there are some definitions of special functions, mostly inspired by the problems of the Euler Project.

4.3 (r6rs-clos)

Provided by: Christian Sloma

URL: <https://launchpad.net/r6rs-clos>

R⁶RS-clos is a port of tiny-clos to the latest (6th) revision of the language standard for scheme. It uses the library system that is new in R⁶RS to structure the code based on functionality (bootstrap of core classes and generic functions, actual implementation of the standard protocols, class layout and slot access ...).

The homepage for now is <https://launchpad.net/r6rs-clos>, where my current development branch can be found.

4.4 SRFI-41: (streams)

The (streams), (streams primitive), and (streams derived) libraries are written by Philip L. Bewig as the reference implementation for SRFI-41. See <http://srfi.schemers.org/srfi-41/srfi-41.html> for more details. The following abstract is excerpted from the SRFI document.

Abstract

Streams, sometimes called lazy lists, are a sequential data structure containing elements computed only on demand. A stream is either null or is a pair with a stream in its cdr. Since elements of a stream are computed only when accessed, streams can be infinite. Once computed, the value of a stream element is cached in case it is needed again.

Streams without memoization were first described by Peter Landin in 1965. Memoization became accepted as an essential feature of streams about a decade later. Today, streams are the signature data type of functional programming languages such as Haskell.

This Scheme Request for Implementation describes two libraries for operating on streams: a canonical set of stream primitives and a set of procedures and syntax derived from those primitives that permits convenient expression of stream operations. They rely on facilities provided by R⁶RS, including libraries, records, and error reporting. To load both stream libraries, say:

```
(import (streams))
```

Chapter 5

Missing Features

Ikarus does not fully conform to R⁶RS yet. Although it implements more than 94% of R⁶RS's macros and procedures, some are still missing. This section summarizes the set of missing features and procedures.

- Numeric tower is complete except for complex numbers.
Consequences:
 - Reader does not recognize complex number notation (e.g. $5-7i$).
 - Procedures that may construct complex numbers from non-complex arguments may signal an error or return an incorrect value (e.g., `(sqrt -1)` should return $+i$ instead of signaling an error).
- The procedure `equal?` may not terminate on `equal?` infinite (circular) input.
- `number->string` does not accept the third argument (precision). Similarly, `string->number` and the reader do not recognize the `|p` notation.

5.1 List of missing R⁶RS procedures

The following procedures are missing from (rnrs base):

angle make-polar make-rectangular

The following procedures are missing from (rnrs unicode):

string-downcase string-titlecase string-upcase
 string-normalize-nfc string-normalize-nfd
 string-normalize-nfkc string-normalize-nfkd

The following procedures are missing from (rnrs arithmetic bitwise):

bitwise-ior bitwise-xor bitwise-if bitwise-copy-bit-field
 bitwise-reverse-bit-field bitwise-rotate-bit-field

The following procedures are missing from (rnrs arithmetic fixnum):

fxreverse-bit-field fxrotate-bit-field

The following procedures are missing from (rnrs hashtables):

make-eqv-hashtable make-hashtable equal-hash
 hashtable-hash-function hashtable-equivalence-function

The following procedures are missing from (rnrs io ports):

string->bytevector bytevector->string
 port-has-port-position? port-position
 port-has-set-port-position!? set-port-position!
 make-custom-binary-input/output-port
 make-custom-textual-input/output-port
 open-file-input/output-port

Bibliography

- [1] Philip L. Bewig. Scheme request for implementation 41: Streams. 2007.
- [2] R. Kent Dybvig. *Chez Scheme Version 7 User's Guide*. Cadence Research Systems, 2005.
- [3] R. Kent Dybvig, David Eby, and Carl Bruggeman. Don't stop the BiBOP: Flexible and efficient storage management for dynamically-typed languages. Technical Report 400, Indiana University, March 1994.
- [4] Abdulaziz Ghuloum and R. Kent Dybvig. Generation-friendly Eq hash tables. In *Proceedings of the 2007 Workshop on Scheme and Functional Programming*, pages 27–35. Université Laval Technical Report DIUL-RT-0701, 2007.
- [5] Abdulaziz Ghuloum and R. Kent Dybvig. Implicit phasing for R6RS libraries. In *ICFP '07: Proceedings of the 2007 ACM SIGPLAN international conference on Functional programming*, pages 303–314, New York, NY, USA, 2007. ACM.
- [6] Michael Sperber, R. Kent Dybvig, Matthew Flatt, and Anton Van Straaten (Editors). Revised⁶ report on the algorithmic language Scheme. 2007.
- [7] Michael Sperber, R. Kent Dybvig, Matthew Flatt, and Anton Van Straaten (Editors). Revised⁶ report on the algorithmic language Scheme–standard libraries. 2007.
- [8] Oscar Waddell. *Extending the Scope of Syntactic Abstraction*. PhD thesis, Indiana University Computer Science Department, August 1999.
- [9] Oscar Waddell and R. Kent Dybvig. Extending the scope of syntactic abstraction. In *Conference Record of POPL'99: The 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 203–213, January 1999.

Index

- #!eof, [27](#)
- #!ikarus, [25](#)
- #:pretty reader syntax, [35](#)
- #{pretty unique} reader syntax, [35](#)
- #{unique} reader syntax, [35](#)
- Boot files, [7](#)
- command-line, [7](#)
- Command-line switches, [7](#)
- Examples
 - generate-temporaries, [36](#)
 - Hello World, [14](#)
- format, [38](#)
- fprintf, [39](#)
- generate-temporaries, [36](#)
- gensym, [33](#)
- #{gensym}, [35](#)
- gensym->unique-string, [34](#)
- gensym-count, [43](#)
- gensym-prefix, [42](#)
- gensym?, [34](#)
- IKARUS_LIBRARY_PATH, [50](#)
- import, [31](#), [32](#)
- Invoke, [16](#)
- make-parameter, [28](#)
- make-traced-procedure, [46](#)
- module, [32](#)
- parameterize, [29](#)
- port-mode, [26](#)
- pretty-print, [37](#)
- pretty-width, [37](#)
- print-gensym, [41](#)
- print-graph, [39](#)
- printf, [39](#)
- R⁶RS Script, [7](#)
 - Import, [13](#)
- set-port-mode!, [26](#)
- SRFI-41: Streams, [56](#)
 - (streams), [56](#)
 - (streams primitive), [56](#)
 - (streams derived), [56](#)
- time, [47](#)
- time-it, [48](#)
- trace-define, [44](#)
- trace-lambda, [45](#)