

Technical Report



Proceedings of the 2020 Scheme and Functional Programming Workshop

Edited by Baptiste Saleil and Michael D. Adams

January 2021

Computer Science and Engineering technical reports
published by University of Michigan are available at:

<https://cse.engin.umich.edu/research/technical-reports-publications/>

Preface

This report aggregates the papers presented at the twenty-first annual Scheme and Functional Programming Workshop, hosted on August 28th, 2020, online and co-located with the twenty-fifth International Conference on Functional Programming. The Scheme and Functional Programming Workshop is held every year to provide an opportunity for researchers and practitioners using Scheme and related functional programming languages like Racket, Clojure, and Lisp, to share research findings and discuss the future of the Scheme programming language. Seven papers and three lightning talks were submitted to the workshop, and each submission was reviewed by three members of the program committee. After deliberation, four papers and three lightning talks were accepted to the workshop. In addition to the four papers and three lightning talks presented,

- Martin Henz and Tobias Wrigstad gave an invited keynote speech entitled *SICP JS: Ketchup on Caviar?*
- Bohdan Khomtchouk and Jonah Fleishhacker gave an invited keynote speech entitled *21st Century Lisp in Academic Research and Pedagogy*.

Thank you to all the presenters, panelists, participants, and members of the program committee.

Program Committee

Michael D. Adams, University of Michigan (Program Co-Chair)
Baptiste Saleil, IBM Canada (Program Co-Chair)
Maxime Chevalier-Boisvert, Université de Montréal
Ryan Culpepper, Czech Technical University
Kimball Germane, University of Utah
Yukiyoshi Kameyama, University of Tsukuba
Andy Keep, Cisco Systems, Inc
Julien Pagès, Université de Montréal
Alexey Radul, Massachusetts Institute of Technology

Steering Committee

Will Byrd, University of Alabama at Birmingham
Will Clinger, Northeastern University
Marc Feeley, Université de Montréal
Dan Friedman, Indiana University
Olin Shivers, Northeastern University

Contents

Keynote: SICP JS: Ketchup on Caviar?	1
Keynote: 21st Century Lisp in Academic Research and Pedagogy	2
Paper: Clotho: A Racket Library for Parametric Randomness	3
Paper: Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set	14
Paper: Running Scheme On Bare Metal (Experience Report)	51
Paper: Scheme for Scientific Computing	66
Talk: On Teaching Type Systems as Macros	77
Talk: Designing a Programming Environment Based on the Program Design Recipe	78
Talk: Programming with Petri Nets to Reason about Concurrency	79

SICP JS: Ketchup on Caviar?

Martin Henz
National University of Singapore
henz@comp.nus.edu.sg

Tobias Wrigstad
Uppsala University
tobias.wrigstad@it.uu.se

With its minimalism, the language Scheme is well suited, if not designed, for teaching the structure and interpretation of computer programs (SICP) to freshmen computer science students, and Harold Abelson and Gerald Jay Sussman made use of the language in their eponymous book, whose second edition was published in 1996. The presenters applied the same minimalism to JavaScript, by identifying four sublanguages just expressive enough for the first four chapters of SICP, and named the languages Source §1, 2, 3 and 4. (There turned out to be no need for a sublanguage for chapter 5 of SICP.) Due to changes introduced to JavaScript with ECMAScript 2015, the Source languages are similar enough to Scheme for a relatively close adaptation of SICP to JavaScript. The resulting book by Abelson and Sussman as original authors, and by the presenters as adapters, is available online, including a side-by-side comparison.

We encountered the following issues during the adaptation due to the differences between the Source languages and Scheme, and briefly sketch here how they are resolved in SICP JS.

The distinction between statements and expressions, and the use of `return` is probably the most significant change from SICP to SICP JS. A notable consequence is the need to wrap return values in data structures in 4.1.1 and 4.1.3 in order to distinguish `x => { return x; }` from `x => { x; }` the latter of which returns undefined in JavaScript. We faithfully implement JavaScript's return statements in chapters 4 and 5, such that control can return to the caller from anywhere in the function body. This leads to several significant changes in these chapters, compared to the original. As a benefit, SICP JS helps readers understand statement-oriented languages such as Java and Python better.

Both Scheme and JavaScript (in strict mode, introduced in ECMAScript 5) employ lexical scoping. The Source languages only use JavaScript's `const` and `let` (introduced in ECMAScript 2015) and avoid JavaScript's `var`. The treatment of the scope of variables in chapter 4 and 5 becomes more uniform in SICP JS compared to SICP, as a result of consistently applying a treatment of `const` and `let` akin to Scheme's derived expression `letrec`.

The absence of Scheme's homoiconicity might at the surface be considered a major obstacle to adapting SICP to languages with a conventional syntax. However, SICP already hides the concrete syntax of programs behind an abstraction layer, which greatly simplifies the JavaScript adaptation. The introduction of an explicit parser suffices for adapting chapter 4 (including section 4.4 on logic programming), and the

controller instructions in chapter 5 of SICP JS enjoy a syntax similar to SICP, through the use of constructors, which fit naturally into section 5.2.3. On the negative side, the lack of macros and our restriction to a JavaScript-compatible parser required significant changes to and occasionally replacement of exercises in chapter 4.

The audience is welcome to inspect SICP JS by visiting <https://source-academy.github.io/sicp>. A comparison edition lets the reader inspect the changes and compare them line-by-line with the original. The presentation will leave ample time for discussion.

The presentation will also cover the Source Academy, an online learning environment for programming, developed by and for students at the National University of Singapore, which implements the four Source languages along with several variants and extensions introduced in SICP.

21st Century Lisp in Academic Research and Pedagogy

Bohdan Khomtchouk
University of Chicago
bohdan@uchicago.edu

Jonah Fleishhacker
University of Chicago
jfleishhacker@uchicago.edu

Lisp-family languages (LFLs) continue to influence and inform modern programming language design and applications. In bioinformatics and computational biology, LFLs have successfully been applied to high-performance computing, database curation, systems biology, drug discovery, computational chemistry and nanotechnology, among much more. Furthermore, dialects such as Racket, with its language-oriented programming and streamlined development environment, or Clojure, with its Java ecosystem and comprehensive open-source libraries, present exciting possibilities in

pedagogy, particularly in introductory computer science and functional programming courses. We present an educational perspective on beginning one's journey in computer science with LFLs, and the advantages unique to this pedagogical approach. Distinct advantages of using LFLs for students can be found in fundamental theoretical concepts of computer science and functional programming, along with practical software development in industry settings.

Clotho: A Racket Library for Parametric Randomness

Pierce Darragh
pdarragh@cs.utah.edu
University of Utah
Salt Lake City, UT, USA

William Gallard Hatch
william@hatch.uno
University of Utah
Salt Lake City, UT, USA

Eric Eide
eeide@cs.utah.edu
University of Utah
Salt Lake City, UT, USA

Abstract

Programs such as simulators and fuzz testers often use randomness to walk through a large state space in search of interesting paths or outcomes. These explorations can be made more efficient by employing heuristics that “zero-in” on paths through the state space that are more likely to lead to interesting solutions. Given one path that exhibits a desired property, it may be beneficial to generate and explore similar paths to determine if they produce similarly interesting results. When the random decisions made during this path exploration can be manipulated in such a way that they correspond to discrete structural changes in the result, we call it *parametric randomness*.

Many programming languages, including Racket, provide only simple randomness primitives, making the implementation of parametric randomness somewhat difficult. To address this deficiency, we present Clotho: a Racket library for parametric randomness, designed to be both easy to use and flexible. Clotho supports multiple strategies for using parametric randomness in Racket applications without hassle.

1 Introduction

There are many applications in which a developer may want to use pseudo-random number generators (PRNGs) to explore a given search space while using the results of previous explorations to inform choices in subsequent navigation of the space. Examples include:

- Generating many random programs that share a common attribute.
- Producing sentences from a grammar with a common prefix.
- Walking a large graph, such as that of a social network, without changing an initial portion of the walk.
- Implementing a genetic algorithm.
- Modeling multiple, similar paths in a simulation.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Scheme '20, August 28, 2020, Jersey City, NJ, USA
© 2020 Copyright held by the owner/author(s).

Many mainstream programming languages only provide very simple randomness primitives and leave the more complex uses of these functions up to developers to implement on a per-case basis. This can be tedious and is prone to error.

While working on a random program generator (§2.1), we found ourselves in need of a system for manipulating the outcomes of randomness functions in a predictable manner. We wanted to “record” a sequence of randomly generated values, modify that sequence in some way, and then feed the modified recording back to our system to get a new—but similar—sequence of randomly generated values. Crucially, unchanged portions of the recording must produce the same results in subsequent executions as they did in the original generation. We call this process *parametric randomness*.

We define parametric randomness as a kind of random value generation that is amenable to predictable external manipulation. For example, consider a random value sequence generated without any manipulation: [4, 8, 15, 16, 23, 42]. After generating this initial sequence, one could employ parametric randomness to produce new sequences that are similar to the original:

- [4, 8, 15, 16, 17]
- [4, 8, 15, 16, 23, 43]
- [4, 8, 15, 16, 23, 19, 68]
- [4, 8, 12, 16, 23, 42]
- [4, 3, 15, 16, 37, 42]

Imagine that in each of these sequences, the values correspond to the choices made by a program that randomly explores a decision tree. This collection of sequences therefore represents multiple explorations of the decision tree. The paths exhibit some similar properties, but may lead to fundamentally different outcomes by the search program.

We have developed a Racket library, Clotho, that enables developers to easily engage in this style of search-space exploration with parametric randomness.¹ Our library implements the following functionality:

1. All the existing randomness functionality of `racket/base` and `racket/random`, which define the Racket standard library’s randomness functions.
2. Convenience functions for generating common values using parametric random generation functions (e.g.,

¹The library can be installed from the official Racket package catalog via `raco pkg install clotho`. Alternatively, the source code can be downloaded at <https://gitlab.flux.utah.edu/xsmith/clotho>.

Booleans, integers, Unicode characters, and Unicode strings).

3. Support for parametric randomness by specifying a seed sequence.
4. A macro for wrapping functions that use Racket’s built-in randomness functions, to ensure the use of parametric randomness outside of the functions specified in our library.
5. A macro for supporting randomness abstraction (§5.2).

In this paper we explain the motivation, design, and implementation of Clotho with examples along the way. We present:

- Motivation for parametric randomness, and background on why Racket’s existing randomness functionality is insufficient for our purposes. (§2)
- A high-level overview of the functionality provided by Clotho. (§3)
- Explanation of the under-the-hood implementation that powers Clotho’s parametric generation functionality. (§4)
- Detailed examples that illustrate the capabilities of Clotho. (§5)
- Discussion on the limitations of, and potential future work for, Clotho. (§6)

2 Background

2.1 Motivation

We developed Clotho as part of an implementation of a random program generation tool [Hatch et al. 2020]. *Random program generation* is the process of automatically creating whole programs without human input. This technique has proven especially useful in the domain of testing programming language compilers and interpreters, where human-written tests can miss edge cases or are otherwise insufficient to trigger bugs in the implementation [Padhye et al. 2019; Yang et al. 2011].

When we use our tool, we sometimes want to generate programs that are similar to those that the tool has generated previously. One way to achieve this is with a *parametric generator* [Padhye et al. 2019]: a generator that inputs a sequence that encodes the “random” choices that the generator will make. The generator processes the input sequence and outputs a new test case. The crucial characteristic of a parametric generator is that simple (e.g., bit-level) modifications to the input sequence result in structural changes to the generated output in a relatively predictable manner. Essentially, each primitive element of the input sequence is a parameter that can be adjusted to modify the output.

If we consider that a random program generator is a function that randomly walks a path in the decision tree of possible output programs, then a parametric generator is one that exposes its decisions as parameters that can be tuned. Coupled with an external metric for recognizing “interesting”

```

1 (require racket/random)
2
3 (random-seed 0)
4 (random) ;; => 0.8571568490678037
5 (random) ;; => 0.6594215608573717
6 (random) ;; => 0.205654820840853
7 (random-seed 0)
8 (random) ;; => 0.8571568490678037
9 (random) ;; => 0.6594215608573717
10 (parameterize
11   ([current-pseudo-random-generator
12    (make-pseudo-random-generator)])
13   (random-seed 0)
14   (random) ;; => 0.8571568490678037
15   (random)) ;; => 0.6594215608573717
16 (random) ;; => 0.205654820840853

```

Figure 1. Randomness in Racket using `random-seed`, where comments show the result of each `random` call. The PRNG created in the body of the `parameterize` expression has no impact on the PRNG that exists externally.

generator outputs—e.g., test cases that extend code coverage of the compiler or interpreter under test—a parametric random program generator can be driven to generate more interesting programs over time.

Clotho arose from our implementation of random program generation in Racket, but we have made Clotho a standalone package because it can be used in other domains as well.

2.2 What Racket Offers

Racket provides a number of functions for generating random values in its standard library, as well as some methods for directly manipulating the current source of randomness (a pseudo-random number generator, or PRNG) to make the outputs of randomness functions manipulable. However, we found Racket’s built-in functionality to be lacking in expressive capability on its own.

The most fundamental of Racket’s randomness functions is `random`. The `random` function, when called without any arguments, produces an inexact number in the interval $[0, 1)$ with uniform probability.

Behind the scenes, Racket uses a system-wide PRNG to generate random values on demand. This PRNG is a *parameter*² in Racket, which means that its value can be dynamically re-bound in a local context with a `parameterize` expression. The parameter’s name is `current-pseudo-random-generator`, which we will call CPRG for short, and it conforms to the type predicate `pseudo-random-generator?`. The CPRG is instantiated automatically at run time without

²The term “parameter” is unfortunately overloaded in this paper by necessity. When referring to “Racket parameters” or the `parameterize` form, we mean the Racket-specific concept of a parameter as explained here: <https://docs.racket-lang.org/reference/parameters.html>.

any configuration, and it is used implicitly in all of the standard library randomness functions, so the average user never needs to interact with it directly.

For our purposes, though, direct interaction is necessary. Fortunately, Racket provides some mechanisms to manipulate the `CRNG`, which we could potentially use to induce parametric behavior.

The first mechanism lies in the `random-seed` function. This is a side-effecting function that takes as argument an integer in the interval $[0, 2^{31})$, and then uses that value to seed the `CRNG`. The values returned by subsequent calls to any of the randomness functions, such as `random`, are determined by whatever seed was passed to `random-seed`. This effectively means that the randomness can be manipulated indirectly by choosing seed values.

The second mechanism is to create an entirely new `PRNG` by calling either `make-pseudo-random-generator` (which takes no arguments and automatically creates a new `PRNG` seeded by the current system time) or `vector->pseudo-random-generator` (which takes a specially formatted vector as argument and produces a new `PRNG` object from an algorithm performed on those values). This new `PRNG`, which also conforms to `pseudo-random-generator?`, can then be used as the Racket-wide `CRNG` by using a `parameterize` expression. Within the body of this expression, all randomness is handled by the newly created `PRNG`.

One can see an example of some of these mechanisms in action in Figure 1. Unfortunately, there are some problems with these approaches:

1. They are unwieldy to use for manipulating randomness repeatedly.
2. Only the seeds are manipulable. It is quite difficult to implement a system using the provided mechanisms where we want to be able to manipulate *some* of the randomly generated values output by a `PRNG`.
3. Subsequently, there is no way to parameterize the random generation in the manner needed by a parametric generator.

Because we want to use parametric randomness to explore state spaces, it is points 2 and 3 that are the main concerns for us (though 1 is relevant in terms of library design). We want to be able to “replay” a sequence of random generations up to a point, and then deviate. Racket’s provided functionality does not make this easy.

One might consider the possibility of simply replacing the `CRNG` with a custom value that conforms to the `pseudo-random-generator?` predicate. However, `pseudo-random-generator?` is not open to external implementation: the only values that conform are those produced by either `make-pseudo-random-generator` or `vector->pseudo-random-generator`. This severely limits the ability of developers to implement custom random-generation mechanisms.

At this point, we have a choice: do we develop an entirely new `PRNG` system separate from Racket’s existing functionality, or do we attempt to wrap what Racket provides?

Although it may be tempting to implement everything fresh, we chose to wrap Racket’s existing randomness functions. Many existing libraries depend on these functions (such as the `math/distributions` module in the standard library), and a completely custom `PRNG` solution would jeopardize support for these libraries within our random program-generation tool. Supporting these libraries is important to us (because we want to use them without implementing their functionality ourselves), so our decision is made for us: we wrap Racket’s existing randomness functionality to support the parametricity we desire.

3 Design

Clotho’s client-facing interface has three parts:

1. A `current-pseudo-random-generator-like` parameter that controls random generation and provides an interface for users to manipulate random generation as needed.
2. Two macros for enabling advanced functionality.
3. Various convenience functions to make random generation simpler, e.g., `random-int` and `random-bool`.

This section focuses on items 1 and 2, as the convenience functions of 3 are not interesting on their own. Where convenience functions are used in examples in this paper, they will be summarized appropriately on a case-by-base basis.

3.1 The Parameter and Its Maker

The underlying functionality of Clotho is managed by the `current-random-source` parameter, which we abbreviate `CRS` hereafter. All Clotho randomness functions must be called within a context in which the `CRS` has been parameterized:

```
(parameterize
  ([current-random-source ...])
  ...)
```

A new `CRS` is created by using the `make-random-source` function. This function can accept arguments in a few forms.

When called with *zero arguments*, `make-random-source` functions very similarly to Racket’s built-in `make-pseudo-random-generator` function for creating `PRNGs`. Essentially, it will generate a new randomly seeded `random-source?` that is suitable for parameterization.

Alternatively, `make-random-source` can be called with *an integer argument*. In this case, that integer is treated as a random seed value. This initializes the generated `random-source?` deterministically (i.e., using the same seed value repeatedly will produce identical results each time).

Lastly, *a byte string argument* can be supplied. Clotho views this byte string as a sequence of four-byte integers. The first such integer is used internally and will be explained

Table 1. Clotho’s external API, leaving out most convenience functions.

Function/Value Name	Brief Explanation
<code>current-random-source</code>	Parameter responsible for generating random values.
<code>random-source?</code>	Type predicate that tests whether the argument is a valid <code>current-random-source</code> parameter.
<code>make-random-source</code>	Generates a new <code>random-source?</code> .
<code>current-random-source-initialized?</code>	Tests if <code>current-random-source</code> is initialized.
<code>assert-current-random-source-initialized!</code>	Raise an error if <code>current-random-source</code> is not initialized.
<code>get-current-random-source-byte-string</code>	Returns the byte string from <code>current-random-source</code> .
<code>wrap-external-randomness</code>	Wraps randomness functions defined outside Clotho.
<code>with-new-random-source</code>	Wraps a region with a new <code>current-random-source</code> .
<code>random</code>	Like Racket’s <code>random</code> , but using <code>current-random-source</code> .

later (§4). All subsequent integers are used to deterministically generate random values. (When a byte string is supplied that is not divisible into four-byte segments, it is padded with 0-value bytes at the end.)

3.2 The Macros

Clotho provides two macros to help with some of the more common advanced usage scenarios.

The `wrap-external-randomness` form enables the use of externally defined randomness functions within the parametric framework of Clotho. These “external” functions consult Racket’s `cprg` directly (§3.5). This macro allows Clotho to be compatible with any such existing functions or libraries with very little effort on the part of the Clotho user.

The `with-new-random-source` form is a shorter way of parameterizing the `current-random-source`. Specifically, it uses the `current-random-source` to generate a new seed value, then creates a *new* `random-source?` and installs it as the `current-random-source` parameter. This enables *randomness abstraction*, which is elaborated upon in §5.

3.3 The #langs

In addition to the forms described above, Clotho provides a few `#langs`. They are:

- `#lang clotho/racket/base`: Provides all of the bindings of `racket/base`, but without any randomness functionality.
- `#lang clotho`: Provides all of the bindings of Clotho, as well as all the bindings of `clotho/racket/base`. This is likely to be the most useful for most people.
- `#lang clotho/stateful`: In addition to providing the bindings of `#lang clotho`, this `#lang` also initializes a global `current-random-source` so that random generation can be performed imperatively, similar to how `#lang racket/base` works.

Of course, these can also be used as simple `require` forms. We find that interactions in the Racket REPL are greatly improved by starting the session with `(require clotho/`

`stateful)`, because this enables easy execution of randomness functions without parameterizing the CRS each time. There are additional bindings provided in the module to interact with the random source, which are described in Clotho’s documentation.

3.4 API Summary

Table 1 summarizes Clotho’s exposed interface. The table omits Clotho’s convenience functions for obtaining random data of various types, but they are straightforward: e.g., `random-bool` returns a random Boolean value, `random-int` produces a random signed integer value, and so on. Clotho also provides `random-ref`, which works identically to the Racket-provided function of the same name: i.e., when given a list, it returns an element of that list selected at random with a uniform distribution.

The `get-current-random-source-byte-string` function returns a byte string that encodes the history of values that have been returned by the CRS. This byte string can be used as-is to initialize a new `random-source?` that will replay the recorded values, assuming that the same sequence of calls is made to draw random values from the new source. Alternatively, one can use a *mutation* of the byte string to create a `random-source?` that will produce a *modified* sequence of values.

3.5 An Example

Figure 2 shows the core functionality in action. In the left-hand column, the example defines a `card` struct and some lists describing the suits and values that a card can have. The `random-card` function randomly selects a suit and a value by using the `random-ref` function. Note that the `current-random-source` parameter has not been seen yet; that will come later. The `make-deck` function builds an ordered deck of cards, and the `random-deck` function returns a randomly shuffled deck. We will discuss its use of the `wrap-external-randomness` macro at the end of this section.

Now look the right-hand column of Figure 2. The first segment of code shows the CRS being parameterized with a

```

1 #lang clotho
2 (require racket/list)
3
4 (struct card (suit value) #:transparent)
5 (define card-suits
6   '(♠ ♥ ♣ ♦))
7 (define card-values
8   '(A 2 3 4 5 6 7 8 9 10 J Q K))
9
10 (define (random-card)
11   (card (random-ref card-suits)
12         (random-ref card-values)))
13
14 (define (make-deck)
15   (for*/list ([suit card-suits]
16              [value card-values])
17             (card suit value)))
18
19 (define (random-deck)
20   (wrap-external-randomness
21    (shuffle (make-deck))))
22
23 (define gcrsbs ;; alias to simplify later code
24   get-current-random-source-byte-string)
25
26 ;; code continues in the next column ->
27 (define-values
28   (cv bl) ;; "Card Value" and "Byte List"
29   (parameterize ([current-random-source
30                  (make-random-source)])
31                 (values
32                  (random-card)
33                  (bytes->list (gcrsbs)))))
34 (print cv) ;; -> (card '♠ 10)
35 (print bl) ;; -> '(196 156 203 55
36                  ;; 232 115 4 248
37                  ;; 19 113 78 202)
38
39 (define inbs ;; "INput Byte String"
40   (list->bytes
41    (list-update bl 7 add1)))
42
43 (define-values
44   (ncv nbl) ;; "New CV" and "New BL"
45   (parameterize ([current-random-source
46                  (make-random-source inbs)])
47                 (values
48                  (random-card)
49                  (bytes->list (gcrsbs)))))
50 (print ncv) ;; -> (card '♠ 10)
51 (print nbl) ;; -> '(196 156 203 55
52                  ;; 232 115 4 249
53                  ;; 19 113 78 202)

```

Figure 2. A Clotho example. Modifying the byte string obtained from a random source leads to a different random-card result.

fresh random source. From within this parameterized region, a new random card and the CRS’s byte string are returned.

Let us imagine that we want to repeat our invocation of `random-card`, but we want to modify the suit that comes out. We can do this by modifying the byte string that we obtained after our initial invocation. As shown, we can use the modified byte string to create a new random source in the second `parameterize` call. The result of our second call is as we had hoped: a new random card is returned, and its suit is changed while its value is unchanged.

Selecting the byte to modify to effect this change is (relatively) straightforward in this example. Clotho uses the first four bytes of the byte string for internal initialization (§4). After that, each call to `random-ref` is associated with four bytes of the byte string. The call to `random-ref` that determines the suit is the first invocation of a randomness function within the parameterization of the CRS, so the bytes that affect its outcome are at indices 4–7 of the byte string. The code in Figure 2 creates a new input byte string by incrementing the seventh byte of the original byte string by 1.

Finally, we return to `random-deck`, which uses the `wrap-external-randomness` macro (§3.2). When this macro is invoked, it consumes four bytes from the CRS’s byte string to seed and parameterize a Racket-wide cprg. This is necessary for interfacing with Racket functions, e.g., `shuffle`,

that do not use Clotho but instead use Racket’s standard randomness functions. By having Clotho instantiate a new cprg before calling those “external” randomness functions, Clotho ensures that the values returned by those functions are determined by the seed that Clotho supplies—and thus, the outcomes of those functions can be reproduced. In summary, within the body of `wrap-external-randomness`, the outcomes of all external randomness functions are determined by a single seed supplied by Clotho. When it is possible to do so, we recommend wrapping calls to external randomness functions individually, so that each call will correspond to a unique portion of Clotho’s byte string.

4 Implementation

At its core, Clotho’s functionality is provided through a struct type that is never directly exposed to the user. This struct is the type of the value managed by the `current-random-source` parameter mentioned in previous sections. The definition of the struct is quite simple:

```

(struct random-source-struct
  ([bts #:mutable]
   [idx #:mutable]
   [prg]
   [add #:mutable]))

```

```

1 #lang clotho
2 ;; minesweeper.rkt
3
4 (require racket/list)
5 (provide (all-defined-out))
6
7 (define (play-game n mines)
8   (define (play-moves covered exposed moves)
9     (if (empty? covered)
10        `(win ,(reverse moves))
11        (let* ([move (random n)]
12              [moves (cons move moves)])
13          (cond
14            [(memq move mines)
15             `(lose ,(reverse moves))]
16            [(memq move exposed)
17             `(illegal ,(reverse moves))]
18            [else
19             (play-moves (remq move covered)
20                        (cons move exposed)
21                        moves)]))))
22   (play-moves
23     (remq* mines (range n)) empty empty))
24
25 (define gcrsbs
26   get-current-random-source-byte-string)
27
28 ;; code continues in the next column ->
29 (define (mutate-bytes input-bytes move-num)
30   (define target-index (* 4 move-num))
31   (define original-int
32     (integer-bytes->integer
33      input-bytes #f (system-big-endian?)
34      target-index (+ 4 target-index)))
35   (define mutated-int (+ 1 original-int))
36   (integer->integer-bytes
37    mutated-int 4 #f (system-big-endian?)
38    input-bytes target-index))
39
40 (define (solve-game n mines [src-bts (bytes)]
41       [outcomes (list)])
42   (define-values (outcome outcome-bytes)
43     (parameterize
44      ([current-random-source
45       (make-random-source src-bts)])
46      (values (play-game n mines)
47              (gcrsbs))))
48   (let ([outcomes (cons outcome outcomes)])
49     (if (eq? (first outcome) 'win)
50         (values (reverse outcomes)
51                 outcome-bytes)
52         (solve-game n mines
53                    (mutate-bytes
54                     outcome-bytes
55                     (length (second outcome)))
56                    outcomes))))

```

Figure 3. A simple, one-dimensional Minesweeper game and a naive mutational solver.

The types of these fields are as follows:

- `bts`: bytes?
- `idx`: integer?
- `prg`: pseudo-random-generator?
- `add`: list of integer?

A new `random-source-struct` is created using a byte string. If no byte string is supplied, a new byte string consisting of four 0-value bytes is created. This byte string is stored in `bts`. An index value, `idx`, is kept to point to the next four-byte segment of the byte string to use for random generation. The first four bytes of the byte string are used to seed a new, Racket-standard PRNG that is stored in the `prg` field—which is why the byte string must always have at least four bytes. A list of temporarily stored integers is kept in `add`, explained in more detail below.

The struct’s design enables two forms of value generation: one using the byte string to “replay” values, and the other producing new values using the PRNG stored in `prg`. They work together seamlessly. When a value needs to be generated, the next four bytes from the `bts` byte string are taken and used; if there are no bytes remaining to be used, a new value is generated from the PRNG.

When a segment of bytes is used for generating a value, they are not returned directly as the result of a randomness function. Consider the following code segment:

```

(if (random-bool)
    (random-int)
    (random-bool))

```

Assume that we run this code and observe the values #t and 42 being produced. We obtain the byte string that caused these results, manipulate it, create a new CRS and rerun the code. If the result of the first call to `random-bool` in the new run is #f, Clotho must *reinterpret* the bytes that previously determined the result from `random-int`: now those bytes must determine the result of the second invocation of `random-bool`. To avoid implementing its own conversions from bytes to return values, Clotho interprets four-byte segments of the `bts` byte string as (integer) seed values for PRNGs, which it creates on demand. When a random value is needed, it consumes the next four bytes from the `bts` byte string to produce an integer, seeds a new CPRG with that integer, and invokes the appropriate randomness function, which calls `random` from the Racket standard library. The CPRG determines the value returned by `random`.

If the `bts` byte string is exhausted when a value is requested, the PRNG stored in `prg` is used to generate a new

value. The value generated by the PRNG is then used to seed a CPRG and call the appropriate randomness function, as just described. However, the seed for the CPRG is stored in the add list until a client calls `get-current-random-source-byte-string` to obtain the byte string. When this happens, the add's values are reversed and appended to the `bts` byte string, and then that (now extended) byte string is returned. The add-list reduces the time complexity of intermediate random generations, because the `bytes-append` function can be very costly. By storing unserialized generated values in a list, we avoid incurring significant overhead.

The index, `idx`, is used to keep place while drawing values from the `bts` byte string. When the byte string is exhausted, the `idx` is set to `#f`.

The `random-source-struct`, along with its automatically provided functions, is kept private from even the rest of Clotho. An API is provided to the rest of the library that ensures certain conditions are maintained:

1. The `bts` byte string must contain at least four bytes, to be used for seeding the source's prg PRNG.
2. The `idx` index value corresponds to the head of the next four-byte sequence to be read from the `bts` byte string.
3. If all of the bytes have been read from the `bts` byte string, the `idx` index value is set to `#f`.
4. Whenever the call is made to extract the `bts` byte string from the struct, the add add-list must be reversed, converted to bytes, and appended to the byte string. (The add list must then be cleared.)
5. When a new value is requested and the `idx` index is `#f`, the prg PRNG is used to generate a new value.

These conditions ensure that random generation works as explained in §3.

5 Using Clotho

We have explained the design and implementation of Clotho and provided some small examples to show its use, but how can it be used to accomplish parametric generation? In this section, we give a brief example client in the form of a simplified Minesweeper game and show how to write a naive mutational fuzzer to force a win.

(Note: The non-figure code in this section is meant to be read additively. The output of `print` is shown in a comment to the right of the call. We convert byte strings to lists of byte-as-integers using `bytes->list` so the values are easier to read.)

5.1 Playing a Game of Minesweeper

Figure 3 contains the implementation of the game and its fuzzer. The `play-game` function contains the core game logic, which we will mostly gloss over here except to point out the use of the random function used to make a guess on the board.

An example call to `play-game` might look like this:

```
(require "minesweeper.rkt")
(define-values
  (r bs) ;; "Result" and "Byte String"
  (parameterize
    ([current-random-source
      (make-random-source)]
     (values
      (play-game 5 '(2 3))
      (bytes->list (gcrsbs)))))

  (print r) ;; -> '(illegal (4 0 4))
  (print bs) ;; -> '(125 35 151 62
                   ;; 0 0 0 0
                   ;; 1 236 216 117
                   ;; 33 15 40 66))
```

This plays a game of Minesweeper consisting of 5 cells, with mines hidden in cells 2 and 3. The result of the game and the resulting byte string from the generated `current-random-source` are returned. In this game, the player made an illegal move by attempting to expose the 4 cell twice. However, the first two moves (guessing 4 and then 0) were legal, so our player was on the right track! Let's use Clotho to modify this game so the player wins.

There are 16 bytes in the returned byte string `bs`. The first four bytes of `bs` are devoted to creating a PRNG (§4), and each call to `random` after that adds an additional four bytes to the byte string. This means that $(16 - 4) \div 4 = 3$ moves were made—which lines up correctly with the output list of moves we saw: `'(4 0 4)`. Since the last move in the sequence is the one that caused a failure, we want to try mutating bytes 12–15 in the byte string (remembering that the string is zero-indexed).

A simple mutation to try (which is used by the `mutate-bytes` function in Figure 3) is to increment those four bytes by 1. We can use the mutated byte string to build a new `current-random-source` and see what the outcome is:

```
(define inbs ;; "INput Byte String"
  (mutate-bytes (list->bytes bs) 3))
(define-values
  (nr nbs)
  (parameterize
    ([current-random-source
      (make-random-source inbs)]
     (values
      (play-game 5 '(2 3))
      (bytes->list (gcrsbs)))))

  (print nr) ;; -> '(win (4 0 1))
  (print nbs) ;; -> '(125 35 151 62
                    ;; 0 0 0 0
                    ;; 1 236 216 117
                    ;; 33 15 40 67))
```

Hooray! With that change, the player made the correct final guess and won the game.

```

1 (define (play-games game-count n mines)
2   (for/list ([_ (range game-count)])
3     (with-new-random-source
4       (play-game n mines))))
5
6 (define (solve-games game-count n mines
7                 [src-bts (bytes)]
8                 [results (list)])
9   (define-values (outcomes outcome-bytes)
10    (parameterize
11      ([current-random-source
12       (make-random-source src-bts)]
13       (values (play-games game-count n mines)
14              (gcrsbs))))
15
16   (define lost-game-index
17     (index-where
18      outcomes
19      (lambda (outcome)
20        (not (eq? 'win (first outcome))))))
21
22   (let ([results (cons outcomes results)])
23     (if lost-game-index
24         (solve-games
25          game-count n mines
26          (mutate-bytes outcome-bytes
27                       (add1 lost-game-index))
28          results)
29         (values (reverse results)
30                outcome-bytes))))

```

Figure 4. An abstracted solver for multiple Minesweeper games played in a series.

The `solve-game` function (Figure 3) packages up this iterative refinement process. Here, the `current-random-source` is parameterized (using an empty byte string by default) and a game is played. If the game is won, the outcome is returned along with the byte string. Otherwise, the byte string is mutated and `solve-game` calls itself recursively with the new byte string and the list of outcomes accumulated so far.

5.2 Playing Multiple Minesweeper Games

Nobody wants to play just one game of Minesweeper. Let's expand our example to play multiple consecutive games and solve them all!

A simple approach is to write a `play-games` function that uses `for/list` to call `play-game` multiple times and accumulate the results. However, this raises a question: at what level do we parameterize the `current-random-source`? We could parameterize it outside of the `for/list`, using a single random source for all the games. Alternatively, we could do the parameterization *inside* the `for/list`, using a new random source for each game.

Let's try them both out and see what works best!

We begin by implementing the first method where all of the games are parameterized together:

```

(define (play-games game-count n mines)
  (parameterize
    ([current-random-source
     (make-random-source)])
    (values
     (for/list ([_ (in-range game-count)])
       (play-game n mines))
     (gcrsbs))))

```

The return values are a list of the game results and the byte string obtained after the last game is played.

While straightforward to implement, this can cause unpredictable effects during later mutation (such as that implemented by our `solve-game` function). Each game in this example can consume 1–3 random values, depending on how many moves the player makes. If mutating a move in one of the earlier games of the sequence causes that game to complete in a different number of moves than it did previously, subsequent games will play differently than they did before. Often, this is undesirable.

Let us instead try moving the parameterization of the `current-random-source` to a per-game level:

```

(define (play-games game-count n mines)
  (for/list ([_ (in-range game-count)])
    (parameterize
      ([current-random-source
       (make-random-source)])
      (cons
       (play-game n mines)
       (gcrsbs)))))

```

In this version of `play-games`, the return value is a list of pairs of game results with each game's corresponding byte string. This means each game's randomness is independent from the rest, but the cost is a more complicated output—one that does not naturally play well with tools that operate on a single byte string!

There is actually a third option: use *both*, but make the inner parameterization depend on a value in the parent parameterization. Or, to put it in code:

```

(define (play-games game-count n mines)
  (parameterize
    ([current-random-source
     (make-random-source)])
    (values
     (for/list ([_ (in-range game-count)])
       (parameterize
         ([current-random-source
          (make-random-source
           (random-seed-value))]
          (play-game n mines))
         (gcrsbs)))))

```

(Note that this code uses the `random-seed-value` function, which is a convenience function that generates a value

suitable for use as a random seed in either the CRS or the Racket CPRG.)

This `play-games` returns two values: a list of the results of the games, and a single top-level byte string. However, this byte string is different from the ones we've seen so far. Previously, the 4-byte segments of the byte strings (after the initial 4-byte segment reserved for internal use) corresponded to individual moves made during the game. But in this latest byte string, each 4-byte segment corresponds to an entire game. Essentially, we have abstracted the randomness: instead of fine-grained control over each randomness function, we now have coarse-grained control over the meta-randomness function `play-game`.

To make it easier to use *randomness abstraction*, Clotho provides the `with-new-random-seed` macro (§3.2). Called without any arguments, it functions identically to the inner `parameterize` function in the previous code segment, creating a new `current-random-source` using a seed value generated from the parent random source.

Figure 4 shows an example of implementing the code in this way. The `play-games` function plays multiple games of Minesweeper with identical inputs, each executed within its own `with-new-random-source` body. This function is called from within `solve-games`, which mutates bytes according to which games have resulted in a loss. The `solve-games` function has its own parameterization of the CRS (which functions as the top-level parameterization in this code), and `play-games` uses the macro to abstract the randomness within each call to `play-game`.

Below, we show the output of one execution of this function as seen in the Racket REPL, with the `illegal` symbol abbreviated to `ill` and columns vertically aligned for readability. The output can be rather long because multiple games are played many times, so we show only an excerpt of the output that illustrates the point.

```
> (define-values
  (results outcome-bytes)
  (solve-games 3 5 '(2 3) #f))

... ;; <output lines removed>
((lose (4 1 3)) (ill (0 1 0)) (lose (0 3)))
((lose (2)) (ill (0 1 0)) (lose (0 3)))
((win (1 0 4)) (ill (0 1 0)) (lose (0 3)))
((win (1 0 4)) (ill (4 1 4)) (lose (0 3)))
((win (1 0 4)) (ill (4 4)) (lose (0 3)))
((win (1 0 4)) (lose (2)) (lose (0 3)))
((win (1 0 4)) (ill (4 0 0)) (lose (0 3)))
((win (1 0 4)) (lose (3)) (lose (0 3)))
((win (1 0 4)) (ill (4 4)) (lose (0 3)))
((win (1 0 4)) (lose (4 2)) (lose (0 3)))
((win (1 0 4)) (win (4 0 1)) (lose (0 3)))
((win (1 0 4)) (win (4 0 1)) (lose (2)))
((win (1 0 4)) (win (4 0 1)) (lose (3)))
((win (1 0 4)) (win (4 0 1)) (ill (4 4)))
... ;; <output lines removed>
```

From this excerpt, one can observe how modifying a game earlier in the series does not affect the outcomes of subsequent games. If either of the second or third games had initially resulted in a win, they would not be affected by the modifications to the earlier games.

5.3 Xsmith

A major motivation in the creation of Clotho has been to aid in feedback-directed fuzzing. We have performed some preliminary work using Clotho for fuzzing by using it in conjunction with Xsmith [Hatch et al. 2020] and AFL [Zalewski 2020] to fuzz an implementation of the Lua programming language [Jerusalimschy et al. 1996]. In the future we intend to do further work on feedback-directed fuzzing using Clotho and Xsmith.

While Clotho was designed with guiding Xsmith program generators in mind, it could be used for other structured data generators, such as QuickCheck [Claessen and Hughes 2000] or Racket's `data/enumerate` package [New 2020].

6 Discussion

In building Clotho, we found some limitations and uncovered potential future directions of investigation.

6.1 Limitations

Because Racket's `pseudo-random-generator?` type is closed (i.e., cannot be implemented by an external source), interfacing with libraries that use Racket's built-in randomness functions can be awkward. We introduced the `wrap-external-randomness` macro to address this, but it is not an ideal solution: it requires the user to wrap every call to any function defined in an external library that uses randomness. This can be somewhat tedious. Clotho provides its own `clotho/math/distributions` library that automatically finds all top-level bindings in Racket's `math/distributions` library and wraps them in our macro, allowing them to be used in code that uses Clotho. A significant (though perhaps not wholly detrimental) caveat of this is that a Clotho client is unable to exercise as fine-grained control over random generation as it would be able to if it implemented the functions for itself. When external functions are wrapped, the granularity of Clotho's control is at the level of calls to the external library, rather than at the level of calls to individual randomness functions (§3.5).

Another issue lies with the `with-new-random-source` function, and can be seen in §5.2. This macro provides an abstraction layer for randomness, but abstraction comes with a cost: one can no longer make fine-grained adjustments! In the example shown, the use of `with-new-random-source` precludes the solver from manipulating specific moves, instead requiring it to iterate on the seed for the PRNG that is used for entire games. This loss of granularity can cause a significant decrease in solver efficiency, because the mutations

may take the solver further away from the solution before they bring it closer. However, despite this limitation, we feel that there are situations that warrant the use of randomness abstraction, and so we leave `with-new-random-source` in the Clotho library.

6.2 Future Work

In working around the limitations of `with-new-random-source`, we have started to speculate about the use of a different datatype for capturing random generation. Currently, randomness is encoded into a single byte string. Instead, one can imagine perhaps using a list of bytes, which itself may contain sub-lists of bytes. Using this data structure, `with-new-random-source` could represent the abstracted `current-random-source`'s encoding as a sub-list in the parent context's `current-random-source`. This would allow for parameterizing abstracted regions with fine granularity. Clotho could also include a number of mutation functions specifically intended to aid in the modification of these sub-lists, which would further improve a developer's ability to manage randomness in an application.

Another improvement that could be made to Clotho's data representation would be to use data from the input byte string directly as return values from Clotho's randomness functions, only turning to a PRNG when really necessary. This would allow solvers to more directly manipulate the byte string, which may improve efficiency in finding interesting solutions. However, it is not straightforward, due to the need for input bytes to potentially be consumed by *any* randomness function (§4).

7 Related Work

We created Clotho because we want to use it in Xsmith [Hatch et al. 2020], our tool for creating random program generators. We want the ability to control all of the choices that Xsmith makes while generating a program—an idea that we borrowed from the Zest fuzzing system [Padhye et al. 2019]. Zest explores the state space of a system under test (SUT) by invoking a random test-case generator and using code-coverage feedback from the SUT to tune subsequent generation. Zest requires the test-case generator to be built such that its “random” choices are determined by a bit sequence that is input to the generator. Zest provides bit sequences to the generator in order to create test cases; when a test case triggers new code coverage, Zest mutates the corresponding bit sequence to create new inputs to the test-case generator. The key insight of Zest is that bit-level manipulations to the generator's input produce structural changes in the resulting test cases: “Zest converts a random-input generator into an equivalent deterministic *parametric generator*” [Padhye et al. 2019, p. 332]. Clotho enables similar functionality in

Xsmith by providing an easy-to-use library for making random generation parametric in the same way as Zest, while also providing additional benefits.

Other systems have also manipulated the input of a test-case generator toward increasing the code coverage of a SUT. Crowbar [Dolan and Preston 2017], for example, is a testing library for OCaml that leverages AFL [Zalewski 2020] to (1) generate bit-level inputs that Crowbar turns into structured test cases and (2) measure coverage within the SUT. The DeepState unit-testing library for C and C++ supports coverage-directed fuzzing and can initialize the state of the SUT using bits from a provided input sequence [Goodman and Groce 2018].

Some prior work has manipulated the input to a test-case generator not toward increasing the code coverage of a SUT, but instead toward finding small test cases that trigger a behavior of interest in the SUT, e.g., a program crash or other bug. For example, Seq-Reduce [Regehr et al. 2012] relied on the Csmith program generator and aimed to “automate most or all of the work required to reduce bug-triggering test cases for C compilers.” Seq-Reduce would first run Csmith to generate a program in Csmith's the normal way (using a PRNG), but recording the random decisions that were made during generation. Seq-Reduce would then repeatedly run Csmith, each time trying to discover a new decision sequence that would yield a new, smaller program that preserves the “interesting behavior” of the original. Regehr et al. [2012] concluded that Seq-Reduce was not very effective in general: changes early in the decision sequence would greatly impact the program generator, making it unlikely that the newly generated test case would preserve the behavior of interest. In Clotho, one can mitigate this issue by using the `with-new-random-source` macro to organize test-case generation into subparts, each of which draws from an independently seeded source of values.

Hypothesis [MacIver and Donaldson 2020] is a more recent test-case generator that implements “internal test-case reduction,” i.e., the idea of manipulating the random decisions made during test-case generation, toward coaxing the generator into producing small test cases that preserve a property of interest. Starting from an input “choice sequence” that yields an interesting but large test case, Hypothesis applies heuristics that simplify the choice sequence and yield smaller test cases. In contrast to the conclusion reached by Regehr et al. for Seq-Reduce, MacIver and Donaldson found that internal reduction with Hypothesis often produced good results. For this reason, we speculate that Clotho may be useful for adding internal test-case reduction to Hypothesis-like tools written in Racket.

Wingate et al. [2011] present a method for providing parametric randomness in the implementations of probabilistic programming languages. Their method creates a naming scheme for each program trace that accesses a source of random data, i.e., a random function or variable. While running,

each access to random data triggers a database lookup, keyed by the name of the current program trace. Values from the database are returned as the results of random-data accesses, and new random values are added to the database as necessary for new trace names. The database may be preserved and altered for future executions to provide parametric deterministic randomness. Unlike Clotho, the technique described by Wingate et al. requires a whole-program transformation to track and name each program trace.

8 Conclusion

We have introduced Clotho, a Racket library that provides parametric randomness where you need it. In addition to its own functions, Clotho wraps and re-provides all of the randomness functionality of the `racket/base`, `racket/random`, and `math/distributions` modules in the Racket standard library—making parametric randomness accessible without conflict. Clotho provides mechanisms for supporting other randomness functions on a case-by-case basis. Because Clotho provides support for mixing parametric randomness with traditional PRNG-based randomness, it is flexible for a wide range of use cases.

We built Clotho as part of a random program generator in Racket, but its potential use is far more general than that. We hope that Clotho will prove useful to other developers seeking to manipulate randomly generated values within their programs.

Acknowledgments

We thank the anonymous Scheme '20 reviewers for their valuable comments on this work. Their feedback helped us to greatly improve this paper. This material is based upon work supported by the National Science Foundation under Grant Number 1527638.

References

- Koen Claessen and John Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *Proceedings of the 5th ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. 268–279. <https://doi.org/10.1145/351240.351266>
- Stephen Dolan and Mindy Preston. 2017. Testing with Crowbar. In *Proceedings of the OCaml Users and Developers Workshop*. https://ocaml.org/meetings/ocaml/2017/extended-abstract__2017__stephen-dolan_mindy-preston__testing-with-crowbar.pdf
- Peter Goodman and Alex Groce. 2018. DeepState: Symbolic Unit Testing for C and C++. In *Proceedings of BAR 2018: Workshop on Binary Analysis Research*. https://www.ndss-symposium.org/wp-content/uploads/2018/07/bar2018_9_Goodman_paper.pdf
- William Gallard Hatch, Pierce Darragh, and Eric Eide. 2020. Xsmith software repository. <https://gitlab.flux.utah.edu/xsmith/xsmith>
- Roberto Jerusalimschy, Luiz Henrique de Figueiredo, and Waldemar Celes Filho. 1996. Lua—An Extensible Extension Language. *Software: Practice and Experience* 26, 6 (June 1996), 635–652. [https://doi.org/10.1002/\(SICI\)1097-024X\(199606\)26:6<635::AID-SPE26>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1097-024X(199606)26:6<635::AID-SPE26>3.0.CO;2-P)
- David R. MacIver and Alastair F. Donaldson. 2020. Test-Case Reduction via Test-Case Generation: Insights From the Hypothesis Reducer. In *34th European Conference on Object-Oriented Programming (ECOOP '20)*. <https://doi.org/10.4230/LIPIcs.ECOOP.2020.13>
- Max S. New. 2020. data/enumerate library documentation. <https://docs.racket-lang.org/data/Enumerations.html>
- Rohan Padhye, Caroline Lemieux, Koushik Sen, Mike Papadakis, and Yves Le Traon. 2019. Semantic Fuzzing with Zest. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA '19)*. 329–340. <https://doi.org/10.1145/3293882.3330576>
- John Regehr, Yang Chen, Pascal Cuoq, Eric Eide, Chucky Ellison, and Xuejun Yang. 2012. Test-Case Reduction for C Compiler Bugs. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '12)*. 335–346. <https://doi.org/10.1145/2254064.2254104>
- David Wingate, Andreas Stuhlmüller, and Noah Goodman. 2011. Lightweight Implementations of Probabilistic Programming Languages Via Transformational Compilation. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS) (Proceedings of Machine Learning Research)*, Geoffrey Gordon, David Dunson, and Miroslav Dudík (Eds.), Vol. 15. 770–778. <http://proceedings.mlr.press/v15/wingate11a.html>
- Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. 2011. Finding and Understanding Bugs in C Compilers. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '11)*. 283–294. <https://doi.org/10.1145/1993498.1993532>
- Michal Zalewski. 2020. AFL software repository. <https://github.com/google/AFL>

Solving SICP

An Experience Report on Solving the World's Most Famous Programming Problem Set

Vladimir Nikishkin
vladimir.nikishkin@gmail.com

Abstract

This report is written as a post-mortem of a project that has, perhaps, been the author's most extensive personal project: creating a complete and comprehensive solution to one of the most famous programming problem sets in the modern computer science curriculum *Structure and Interpretation of Computer Programs* by Abelson, G. J. Sussman, and J. Sussman.

It measures exactly:

- How much effort SICP requires (729 hours 19 minutes (over eight months), 292 sessions).
- How many computer languages it involves (6).
- How many pieces of software are required (9).
- How much communication with peers is needed.

It suggests:

- A practical software-supported task management procedure for solving coursework.
- Several improvements, on the technical side, to any hard skills teaching process.
- Several improvements, on the social side, to any kind of teaching process.

The solution is published online (the source code and pdf file):

- <http://gitlab.com/Lockywolf/chibi-sicp> (Nikishkin 2020)

This report (and the data in the appendix) can be applied immediately as:

- A single-point estimate of the SICP problem set difficulty.
- A class handout aimed at increasing students' motivation to study.
- A data source for a study of learning patterns among adult professionals aiming for continuing education.
- An “almost ready” protocol for a convenient problem-set solution procedure, which produces artefacts that can be later used as a student portfolio.
- An “almost ready”, and “almost convenient” protocol for measuring time consumption of almost any problem set expressible in a digital form.

Additionally, a time-tracking data analysis can be reproduced interactively in the org-mode version of this

report. (See: Appendix: Emacs Lisp code for data analysis)

CCS Concepts • **Social and professional topics** → **Computing education programs**; **CS1**; **Computer science education**; **Software engineering education**; **Computational science and engineering education**; • **Computing methodologies** → **Symbolic and algebraic manipulation**; *Philosophical/theoretical foundations of artificial intelligence*; • **Software and its engineering** → **Software organization and properties**; **Software system structures**; **Abstraction, modeling and modularity**; **Software architectures**.

Keywords scheme, r7rs, teaching, programming, literate programming, fortran, reproducible research, emacs lisp

1 Introduction

Programming language textbooks are not a frequent object of study, as they are expected to convey existing knowledge. However, teaching practitioners, when they face the task of designing a computer science curriculum for their teaching institution, have to base their decisions on something. An “ad-hoc” teaching method, primarily based on studying some particular programming language fashionable at the time of selection, is still a popular choice.

There have been attempts to approach course design with more rigour. The “Structure and Interpretation of Computer Programs” was created as a result of such an attempt. SICP was revolutionary for its time, and perhaps can be still considered revolutionary nowadays. Twenty years later, this endeavour was analysed by Felleisen in a paper “Structure and Interpretation of Computer Science Curriculum” (Felleisen, Findler, Flatt, and Krishnamurthi 2004). He then reflected upon the benefits and drawbacks of the **deliberately designed** syllabus from a pedagogical standpoint. He proposed what he believes to be a pedagogically superior successor to the first generation of **deliberate** curriculum. (See: “How to Design Programs” (*HTDP*) Felleisen, Findler, Flatt, and Krishnamurthi 2018)

Leaving aside the pedagogical quality of the textbook (as the author is not a practising teacher), this report touches a different (and seldom considered!) aspect of a computer science (and in general, any other subject's) curriculum. That is, precisely, how much work is required to pass a particular course.

This endeavour was spurred by the author's previous experience of learning about partial differential equations through a traditional paper-and-pen based approach, only mildly augmented with a time-tracking software. But even such a tiny augmentation already exposed an astonishing disparity between a declared laboriousness of a task and the empirically measured time required to complete it.

The author, therefore, decided to build upon the previous experience and to try and design as smooth, manageable, and measurable approach to performing university coursework, as possible. A computer science subject provided an obvious choice.

The solution was planned, broken down into parts, harnessed with a software support system, and executed in a timely and measured manner by the author, thus proving that the chosen goal is doable. The complete measured data are provided. Teaching professionals may benefit from it when planning coursework specialised to their requirements.

More generally, the author wants to propose a comprehensive reassessment of university teaching in general,

based on empirical approaches (understanding precisely how, when, and what each party involved in the teaching process does), in order to select the most efficient (potentially even using an optimisation algorithm) strategy when selecting a learning approach for every particular student.

2 Solution approach

The author wanted to provide a solution that would satisfy the following principles:

1. Be complete.
2. Be a reasonably realistic model of a solution process as if executed by the intended audience of the course – that is, freshman university students with little programming experience.
3. Be done in a “fully digital”, “digitally native” form.
4. Be measurable.

These principles need an explanation.

2.1 Completeness

2.1.1 Just solve all of the exercises

The author considers completeness to be an essential property of every execution of a teaching syllabus.

In simple words, what does it mean “to pass a course” or “to learn a subject” at all? How exactly can one formalise the statement “I know calculus”? Even simpler, what allows a student to say “I have learnt everything that was expected in a university course on calculus”?

It would be a good idea to survey teachers, students, employers, politicians and random members of the community to establish what it means for them that a person “knows a subject”.

Following are some potential answers to these questions:

- Passing an oral examination.
- Passing a written examination.
- Passing a project defence committee questioning.
- Completing a required number of continuous assessment (time-limited) tasks.
- Completing coursework.
- Attending a prescribed number of teaching sessions (lectures or tutorials).
- Reading a prescribed amount of prescribed reading material.

Any combination of these can also be chosen to signify the “mastering” of a subject, but the course designer is then met with a typical goal-attainment, multi-objective optimisation problem (Gembicki and Haines 1975); such problems are still usually solved by reducing the multiple goals to a single, engineered goal.

Looking at the list above from a “Martian point of view” (Berne 1973), we will see that all the goals listed

above are reducible to a single “completing coursework” goal. “Completing coursework” is not reducible to any of those specific sub-goals in general, so the “engineered goal” may take the shape of a tree-structured problem set (task/subtask). “Engineered” tasks may include attending tutorials, watching videos and writing feedback.

Moreover, thinking realistically, doing coursework often is the only way that a working professional can study without altogether abandoning her job.

Therefore, choosing a computer science textbook that is known primarily for the problem set that comes with it, even more than for the actual text of the material, was a natural choice.

However, that is not enough, because even though “just solving all of the exercises” may be the most measurable and the most necessary learning outcome, is it sufficient?

As the author intended to “grasp the skill” rather than just “pass the exercises”, he initially considered inventing additional exercises to cover parts of the course material not covered by the original problem set.

For practical reasons (in order for the measured data to reflect the original book’s exercises), in the “reference solution” referred to in this report’s bibliography, the reader will not find exercises that are not a part of the original problem set.

The author, however, re-drew several figures from the book, representing those types of figures that are not required to be drawn by any of the exercises.

This was done in order to “be able to reproduce the material contained in the book from scratch at a later date”. This was done only for the cases for which the author considered the already available exercises insufficient. The additional figures did not demand a large enough amount of working time to change the total difficulty estimate noticeably.

2.1.2 A faithful imitation of the university experience

One common objection to the undertaken endeavour may be the following. In most universities (if not all), it is not necessary to solve all exercises in order to complete a course. This is often true, and especially true for mathematics-related courses (whose problem books usually contain several times more exercises than reasonably cover the course content). The author, however, considers SICP exercises not to be an example of such a problem set. The exercises cover the course material with minimal overlap, and the author even considered adding several more for the material that the exercises did not fully cover.

Another objection would be that a self-study experience cannot faithfully imitate a university experience at all because a university course contains tutorials and demonstrations as crucial elements. Problem-solving

methods are “cooked” by teaching assistants and delivered to the students in a personalised manner in those tutorials.

This is indeed a valid argument. However, teaching assistants may not necessarily come from a relevant background; they are often recruited from an available pool and not explicitly trained. For such cases, the present report may serve as a crude estimate of the time needed for the teaching assistants to prepare for the tutorials.

Furthermore, many students choose not to attend classes at all either because they are over-confident, or due to high workload. For these groups, this report may serve similarly as a crude estimate.

Moreover, prior research suggests that the learning outcome effect of class attendance on the top quartile (by grade) of the students is low. (Clair 1999 and Kooker 1976)

For the student groups that benefit most from tutorials, this report (if given as a recommended reading for the first lesson) may serve as additional evidence in favour of attendance.

Additionally, nothing seems to preclude recording videos of tutorials and providing them as a supplementary material at the subsequent deliveries of the course. The lack of interactivity may be compensated for by a large amount of the material (such as the video recordings of questions and answers) accumulated through many years and a well-functioning query system.

2.1.3 Meta-cognitive exercises

It is often underestimated how much imbalance there is between a teacher and a pupil. The teacher not only better knows the subject of study – which is expected – but is also deciding how and when a student is going to study. This is often overlooked by practitioners, who consider themselves simply as either as sources of knowledge or, even worse, as only the examiners. However, it is worth considering the whole effect that a teacher has on the student’s life. In particular, a student has no other choice than to trust the teacher on the choice of exercises. A student will likely mimic the teacher’s choice of tools used for the execution of a solution.

The main point of the previous paragraph is that teaching is not only the process of data transmission. It is also the process of metadata transmission, the development of meta-cognitive skills. (See Ku and Ho 2010) Therefore, meta-cognitive challenges, although they may very well be valuable contributions to the student’s “thinking abilities”, deserve their own share of consideration when preparing a course.

Examples of meta-cognitive challenges include:

- Non-sequentiality of material and exercises, so that earlier exercises are impossible to solve without first solving later ones.
- The incompleteness of the treatise.
- The terseness of the narrative.
- Lack of modern software support.
- Missing difficulty/hardness estimation for tasks.
- The vastly non-uniform difficulty of the problems.

An additional challenge to the learning process is the lack of peer support. There have been attempts by learning institutions to encourage peer support among students, but the successfulness of those attempts is unclear. Do students really help each other in those artificially created support groups? Inevitably, communication in this those groups will not be limited only to the subject of study. To what extent does this side-communication affect the learners?

A support medium is even more critical for adult self-learners, who do not get even those artificial support groups created by the school functionaries and do not get access to teaching assistance.

It should be noted that the support medium (a group chat platform, or a mailing list) choice, no matter how irrelevant to the subject itself it may be, is a significant social factor. This is not to say that a teacher should create a support group in whatever particular social medium that happens to be fashionable at the start of the course. This is only to say that **deliberate effort** should be spent on finding the best support configuration.

In the author's **personal experience**:

- The #scheme Freenode channel was used as a place to ask questions in real-time. #emacs was also useful.
- <http://StackOverflow.com> was used to ask asynchronous questions.
- The Scheme Community Wiki <http://community.schemewiki.org> was used as reference material.
- The author emailed some prominent members of the Scheme community with unsolicited questions.
- The author was reporting errors in the documents generated by the Scheme community process.
- The author was asking for help on the Chibi-Scheme mailing list.
- There was also some help from the Open Data Science Slack chat.
- There was also some help from the Closed-Circles data science community.
- There was also some help from the rulinux@conference.jabber.ru community.
- There was also some help from the Shanghai Linux User Group.
- There was also some help from the <http://www.dxdy.ru> scientific forum.

- There was also some help from the Haskell self-study group in Telegram.

It should be noted that out of those communities, only the Open Data Science community, and a small Haskell community reside in “fashionable” communication systems.

The summary of the community interaction is under the “meta-cognitive” exercises section because the skill of finding people who can help you with your problems is one of the most useful soft skills and one of the hardest to teach. Moreover, the very people who can and may answer questions are, in most situations, not at all obliged to do so, so soliciting an answer from non-deliberately-cooperating people is another cognitive exercise that is worth covering explicitly in a lecture.

Repeating the main point of the previous paragraph in other words: human communities consist of rude people. Naturally, no-one can force anyone to bear rudeness, but no-one can force anyone to be polite, either. The meta-cognitive skill of extracting valuable knowledge from willing but rude people is critical but seldom taught.

The author considers it vital to convey to students, as well as to teachers, the following idea: it is not the fashion, population, easy availability, promotion, and social acceptability of the support media that matters. Unfortunately, it is not even the technological sophistication, technological modernity or convenience; it is the availability of information and the availability of people who can help.

Support communication was measured by the following:

- Scheme-system related email threads in the official mailing list: **28**.
- Editor/IDE related email threads + bug reports: **16**.
- Presentation/formatting related email threads: **20**.
- Syllabus related email threads: **3**.
- Documentation related email threads (mostly obsolete link reports): **16**.
- IRC chat messages: **2394** #scheme messages initiated by the author (the number obtained by simple filtering by the author's nickname).
- Software packages re-uploaded to Software Forges: **2** (recovered from original authors' personal archives).

The author did not collect measures of other communication means.

2.1.4 Figures to re-typeset

Several figures from SICP were re-drawn using a textual representation. The choice of figures was driven by the idea that someone who successfully completed the book should also be able to re-create the book material and therefore should know how to draw similar diagrams.

Therefore, those were chosen to be representative of the kinds of figures not required to be drawn by any exercise.

The list of re-drawn figures:

- 1.1 Tree representation, showing the value of each sub-combination.
- 1.2 Procedural decomposition of the `sqrt` program.
- 1.3 A linear recursive process.
- 2.2 Box-and-pointer representation of `(cons 1 2)`.
- 2.8 A solution to the eight-queens puzzle.
- 3.32 The integral procedure viewed as a signal-processing system.
- 3.36 An RLC circuit.
- 5.1 Data paths for a Register Machine.
- 5.2 Controller for a GCD Machine.

2.2 Behaviour modelling, reenactment and the choice of tools

2.2.1 The author's background

On starting the project, the author already possessed a PhD in Informatics, although not in software engineering. This gave an advantage over a first-year undergraduate student. However, to a large extent, the author still resembled a newbie, as he never before used a proudly functional programming language, and had never used any programmers' editor other than Notepad++. Another noticeable difference was that the author could type fast without looking at a keyboard (so-called touch-typing). This skill is taught at some U.S.A. high schools but is still not considered mandatory all over the world.

NOTE: This whole report depends heavily on the fact that the author had learnt how to touch-type, and can do it relatively quickly. Without the skill of fast touch-typing, almost all of the measurements are meaningless, and the choice of tools may seem counter-intuitive or even arbitrary.

The goal the author had was slightly ambiguous, in the sense that the intention was to model (reenact) an "idealised" student, that is the one that does not exist, in the sense that the author decided to:

- Perform all exercises honestly, no matter how hard they be or how much time they take.
- Solve all exercises without cheating; this did not prohibit consulting other people's solutions without direct copying.
- Try to use the tools that may have been available at the disposal of the students in 1987, although possibly the most recent versions.
- Try to follow the "Free Software/Open Source/Unix-way" approach as loosely formulated by well-known organisations, as closely as possible.
- Try to prepare a "problem set solution" in a format that may be potentially presentable to a university teacher in charge of accepting or rejecting it.

While the first three principles turned out to be almost self-fulfilling, the last one turned out to be more involved.

The author's personal experience with university-level programming suggested that, on average, the largest amount of time is spent on debugging input and output procedures. The second-largest amount is usually dedicated to inventing test cases for the code. The actual writing of the substantive part of the code comes only third.

It is known that SICP had been intended as a deliberately created introductory course. The author assumed that a large part of the syllabus would be dedicated to solving the two most common difficulties described above. This assumption turned out to not be the case. Rather than solving them, SICP just goes around them, enforcing a very rigid standard on the input data instead.

While not originally designed for such a treatment, SICP's approach greatly simplified formatting the ready-to-submit coursework solution as a single file with prose, code blocks, input blocks, and figures interleaved (a so-called "notebook" format.)

The ambiguity characteristic comes from the need to find a balance between the two "more realistic" mental models of student behaviour. One would be representing a "lazy" student, who would be only willing to work enough to get a passing score. This model would be responsible for saving time and choosing the tools that would possess the least possible incompatibility with the assessment mechanism. The other would be the model of an "eager" student, who would be willing to study the material as deeply as possible, possibly never finishing the course, and would be responsible for the quality of learning and for choosing the best tools available. The idea of two different types of motivation is to some extent similar to the "Theory X and theory Y" approach proposed by McGregor (McGregor 1960).

Let us try to imagine being an "ideal student", a mixture of the two models described above, and make the decisions as if the imaginary student would be doing them. Informally this can be summarised as "I will learn every tool that is required to get the job done to the extent needed to get the job done, but not the slightest bit more". (There exist far more sophisticated models of student behaviour, most of them mathematical, see e.g. Hlosta, Herrmannova, Vachova, Kuzilek, Zdrahal, and Wolff 2018, however, a simple mental model was deemed sufficient in this particular case.)

2.2.2 The tools

The final choice of tools turned out to be the following:

- **Chibi-Scheme** as the scheme implementation
- **srfi-159** as a petty-printing tool
- **srfi-27** as a random bits library

srfi-18 as a threading library
(chibi time) as a timing library
(chibi ast) (not strictly necessary) macro expansion tool
(chibi process) for calling ImageMagick
GNU Emacs as the only IDE
org-mode as the main editing mode and the main planning tool
f90-mode as a low-level coding adaptor
geiser turned out to be not ready for production use, but still useful for simple expressions evaluation
magit as the most fashionable GUI for git
gfortran as the low-level language
PlantUML as the principal diagramming language
TikZ + luaLaTeX as the secondary diagramming language
Graphviz as a tertiary diagramming language
ImageMagick as the engine behind the “picture language” chapter
git as the main version control tool
GNU diff, bash, grep as the tools for simple text manipulation

Figure 1. List of tools required to solve SICP

Chibi-Scheme was virtually the only scheme system claiming to fully support the latest Scheme standard, r7rs-large (Red Edition), so there was no other choice. This is especially true when imagining a student unwilling to go deeper into the particular curiosities of various schools of thought, responsible for creating various partly-compliant Scheme systems. Several libraries (three of which were standardised, and three of which were not) were used to ensure the completeness of the solution. Effectively, it is not possible to solve all the exercises using only the standardised part of the Scheme language. Even Scheme combined with standardised extensions is not enough. However, only one non-standard library was strictly required: **(chibi process)**, which served as a bridge between Scheme and the graphics toolkit.

git is not often taught in schools. The reasons may include the teachers’ unwillingness to busy themselves with something deemed trivial or impossible to get by without, or due to them being overloaded with work. However, practice often demonstrates that students still too often graduate without yet having a concept of file version control, which significantly hinders work efficiency. Git was chosen because it is, arguably, the most widely used version-control system.

ImageMagick turned out to be the easiest way to draw images consisting of simple straight lines. There is still no standard way to connect Scheme applications to applications written in other languages. Therefore,

by the principle of minimal extension, ImageMagick was chosen, as it required **just a single** non-standard Scheme procedure. Moreover, this procedure (a simple synchronous application call) is likely to be the most standard interoperability primitive invented. Almost all operating systems support applications executing other applications.

PlantUML is a code-driven implementation of the international standard of software visualisation diagrams. The syntax is straightforward and well documented. The PlantUML-Emacs interface exists and is relatively reliable. The textual representation conveys the hacker spirit and supports easy version control. UML almost totally dominates the software visualisation market, and almost every university programming degree includes it to some extent. It seemed, therefore very natural (where the problem permitted) to solve the “diagramming” problems of the SICP with the industry-standard compliant diagrams.

Graphviz was used in an attempt to use another industry standard for solving diagramming problems not supported by the UML. The **dot** package benefits from being fully machine-parsable and context-independent even more than UML. However, it turned out to be not as convenient as expected.

TikZ is practically the only general-purpose, code-driven drawing package. So, when neither UML nor Graphviz managed to embed the complexity of the models diagrammed properly, TikZ ended up being the only choice. Just as natural an approach could be to draw everything using a graphical tool, such as Inkscape or Adobe Illustrator. The first problem with the images generated by such tools, though, is that they are hard to manage under version control. The second problem is that it was desirable to keep all the product of the course in one digital artefact (i.e., one file). Single-file packaging would reduce confusion caused by the different versions of the same code, make searching more straightforward, and simplify the presentation to a potential examiner.

gfortran, or GNU Fortran, was the low-level language of choice for the last two problems in the problem set. The reasons for choosing this not very popular language were:

- The author already knew the C language, so compared to an imaginary first-year student, would have an undue advantage if using C.
- Fortran is low-level enough for the purposes of the book.
- There is a free/GPL implementation of Fortran.
- Fortran 90 already existed by the time SICP 2nd ed. was published.

GNU Unix Utilities the author did not originally intend to use these, but **diff** turned out to be extremely

effective for illustrating the differences between generated code pieces in Chapter 5. Additionally, in some cases, they were used as a universal glue between different programs.

GNU Emacs is, de facto, the most popular IDE among Scheme users, the IDE used by the Free Software Foundation founders, likely the editor used when writing SICP, also likely to be chosen by an aspiring freshman to be the most “hacker-like” editor. It is, perhaps, the most controversial choice, as the most likely IDE to be used by freshmen university students, in general, would be Microsoft Visual Studio. Another popular option would be Dr.Racket, which packages a component dedicated to supporting solving SICP problems. However, Emacs turned out to be having the best support for a “generic Lisp” development, even though its support for Scheme is not as good as may be desired. The decisive victory point ended up being the org-mode (discussed later). Informally speaking, entirely buying into the Emacs platform ended up being a substantial mind-expanding experience. The learning curve is steep, however.

As mentioned above, the main point of this report is to supply the problem execution measures for public use. Later sections will elaborate on how data collection about the exercise completion was performed, using org-mode’s time-tracking facility. The time-tracking data in the section 8 do not include learning Emacs or org-mode. However, some data about these activities were collected nevertheless:

Reading the Emacs Lisp manual required **10** study sessions of total length 32 hours 40 minutes. Additional learning of Emacs without reading the manual required 59 hours 14 minutes.

2.3 Org-mode as a universal medium for reproducible research

Org-mode helps to resolve dependencies between exercises. SICP provides an additional challenge (metacognitive exercise) in that its problems are highly dependent on one another. As an example, problems from Chapter 5 require solutions to the successfully solved problems of Chapter 1. A standard practice of modern schools is to copy the code (or other forms of solution) and paste it into the solution of a dependent exercise. However, in the later parts of SICP, the solutions end up requiring tens of pieces of code written in the chapters before. Sheer copying would not just blow up the solution files immensely and make searching painful; it would also make it extremely hard to propagate the fixes to the bugs discovered by later usages back into the earlier solutions.

People familiar with the work of Donald Knuth will recognise the similarity of org-mode with his WEB system and its web2c implementation. Another commonly

used WEB-like system is Jupyter Notebook (See Project Jupyter Developers 2019).

Org-mode helps package a complete student’s work into a single file. Imagine a case in which student needs to send his work to the teacher for examination. Every additional file that a student sends along with the code is a source of potential confusion. Even proper file naming, though it increases readability, requires significant concentration to enforce and demands that the teacher dig into peculiarities that will become irrelevant the very moment after he signs the work off. Things get worse when the teacher has not just to examine the student’s work, but also to test it (which is a typical situation with computer science exercises.)

Org-mode can be exported into a format convenient for later revisits. Another reason to carefully consider the solution format is the students’ future employability. This problem is not unfamiliar to the Arts majors, who have been collecting and arranging “portfolios” of their work for a long time. However, STEM students generally do not understand the importance of a portfolio. A prominent discussion topic in job interviews is, “What have you already done?”. Having a portfolio, in a form easily presentable during an interview, may be immensely helpful to the interviewee.

A potential employer is almost guaranteed not to have any software or equipment to run the former student’s code. Even the student himself would probably lack a carefully prepared working setup at the interview. Therefore, the graduation work should be “stored”, or “canned” in a format as portable and time-resistant as possible.

Unsurprisingly, the most portable and time-resistant format for practical use is plain white paper. Ideally, the solution (after being examined by a teacher) should be printable as a report. Additionally, the comparatively (in relation to the full size of SICP) small amount of work required to turn a solution that is “just enough to pass” into a readable report would be an emotional incentive for the students to carefully post-process their work. Naturally, “plain paper” is not a very manageable medium nowadays. The closest manageable approximation is PDF. So, the actual “source code” of a solution should be logically and consistently exportable into a PDF file. Org-mode can serve this purpose through the PDF export backend.

Org-mode has an almost unimaginable number of use cases. (For example, this report has been written in org-mode.) While the main benefits of using org-mode for the coursework formatting was the interactivity of code execution, and the possibility of export, another benefit that appeared almost for free was minimal-overhead time-tracking (human performance profiling.) Although this initially appeared as a by-product of choosing a specific

tool, the measures collected with the aid of org-mode is the main contribution of this report.

The way org-mode particulars were used is described in the next section, along with the statistical summary.

2.4 Different problem types

SICP's problems can be roughly classified into the following classes:

- Programming problems in Scheme without input.
- Programming problems in Scheme with input (possibly running other programs).
- Programming problems in Scheme with graphical output.
- Programming problems in a “low-level language of your choice”.
- Mathematical problems.
- Standard-fitting drawing exercises.
- Non-standard drawing exercises.
- Essays.

Wonderfully absent are the problems of the data analysis kind.

This section will explain how these classes of problem can be solved in a “single document mode”.

Essays is the most straightforward case. The student can just write the answer to the question below the heading corresponding to a problem. Org-mode provides several minimal formatting capabilities that are enough to cover all the use cases required.

Mathematical problems require that a \TeX -system be present on the student machine, and employ org-mode's ability to embed \TeX ' mathematics, along with previews, directly into the text. The author ended up conducting almost zero pen-and-paper calculations while doing SICP's mathematical exercises.

Programming exercises in Scheme are mostly easily formatted as org-mode “babel-blocks”, with the output being pasted directly into the document body, and updated as needed.

Programming exercises in Scheme with input require a little bit of effort to make them work correctly. It is sometimes not entirely obvious whether the input should be interpreted as verbatim text, or as executable code. Ultimately, it turned out to be possible to format all the input data as either “example” or “code” blocks, feed them into the recipient blocks via an “:stdin” block directive and present all the test cases (different inputs) and test results (corresponding outputs) in the same document.

Programming exercises in a low-level language required wrapping the low-level language code into “babel” blocks, and the result of combining those into a “shell” block. This introduces an operating system dependency.

However, GNU Unix Utilities are widespread enough to consider this not a limitation.

Programming exercises with graphical output turned out to be the trickiest part from the software suite perspective. Eventually, a Scheme-system (chibi) dependent wrapper around the ImageMagick graphics manipulation tool was written. Org-mode has a special syntax for the inclusion of graphic files, so the exercise solutions were generating the image files and pasting the image inclusion code into the org buffer.

Standard drawing exercises illustrate a problem that is extremely widespread, but seldom well understood, perhaps because people aiming to solve it usually do not come from the programming community. Indeed, there are several standard visual conventions for industrial illustrations and diagramming, including UML, ArchiMate, SDL, and various others. Wherever a SICP figure admitted a standard-based representation, the author tried to use that standard to express the answer to the problem. The PlantUML code-driven diagramming tool was used most often, as its support for UML proved to be superior to the alternatives. The org-plantuml bridge made it possible to solve these problems in the manner similar to the coding problems – as “org-babel” blocks.

Non-standard drawing exercises, the most prominent of those requiring drawing environment diagrams (debugging interfaces), were significantly more challenging. When a prepared mental model (i.e. an established diagramming standard) was absent, that diagram had to be implemented from scratch in an improvised way. The TikZ language proved to have enough features to cover the requirements of the book where PlantUML was not enough. It required much reading of the manual and an appropriate level of familiarity with \TeX .

3 Time analysis, performance profiling and graphs

This section deals with explaining exactly how the working process was organised and later shows some aggregated work measures that have been collected.

3.1 Workflow details and profiling

The execution was performed in the following way:

At the start of the work, the outline-tree corresponding to the book subsection tree was created. Most leaves are two-state **TODO**-headings. (Some outline leaves correspond to sections without problems, and thus are not **TODO**-styled.)

TODO-heading is a special type of an org-mode heading, that exports its state (TODO/DONE) to a simple database, which allows monitoring of the overall TODO/DONE ratio of the document.

Intermediate levels are not **TODO**-headings, but they contain the field representing the total ratio of **DONE** problems in a subtree.

The top-level ratio is the total number of finished problems divided by the total number of problems.

An example of the outline looks the following:

```
* SICP [385/404]
** Chapter 1: Building abstractions ... [57/61]
*** DONE Exercise 1.1 Interpreter result
    CLOSED: [2019-08-20 Tue 14:23]...
*** DONE Exercise 1.2 Prefix form
    CLOSED: [2019-08-20 Tue 14:25]
#+begin_src scheme :exports both :results value
  (/ (+ 5 4 (- 2 (- 3 (+ 6 (/ 4 5))))))
    (* 3 (- 6 2) (- 2 7)))
#+end_src

#+RESULTS:
: -37/150
...
```

Figure 2. Execution file example

When work is clearly divided into parts and, for each unit, its completion status is self-evident, the visibility of completeness creates a sense of control in the student. The “degree of completeness of the whole project”, available at any moment, provides an important emotional experience of “getting close to the result with each completed exercise”.

Additional research is needed on how persistent this emotion is in students and how much it depends on the uneven distribution of difficulty or the total time consumption. There is, however, empirical evidence that even very imprecise, self-measured KPIs do positively affect the chance of reaching the goal. (See: VanWormer, French, Pereira, and Welsh 2008)

From the author’s personal experience, uneven distribution of difficulties at the leaf-level tasks is a major demotivating factor. However, the real problems we find in daily life are not of consistent difficulty, and therefore managing an uneven distribution of difficulty is a critical meta-cognitive skill. Partitioning a large task into smaller ones (*not necessarily* in the way suggested by the book) may be a way to tackle this problem. Traces of this approach are visible through the “reference” solution PDF.

The problems were executed almost sequentially. Work on the subsequent problem was started immediately after the previous problem had been finished.

Out of more than 350 exercises, only 13 were executed out of order (See section 3.2). Sequentiality of problems is essential for proper time accounting because the total

time attributed to a problem is the sum of durations of all study sessions between the end of the problem considered and the end of the previous problem. It is not strictly required for the problem sequence to be identical to the sequence proposed by the book, but it is important that, if a problem is postponed, the study sessions corresponding to the initial attempt to solve this problem be somehow removed from the session log dataset.

In this report, study sessions corresponding to the initial attempts of solving out of order problems were simply ignored. This has not affected the overall duration measures much because those sessions were usually short.

Sequentiality is one of the weakest points of this report. It is generally hard to find motivation to work through a problem set sequentially. SICP does enforce sequentiality for a large share of problems by making the later problems depend on solutions of the previous ones, but this “dependence coverage” is not complete.

As the most straightforward workaround, the author may once again suggest dropping the initial attempts of solving the out-of-order problems from the data set entirely. This should be relatively easy to do because the student (arguably) is likely to decide whether to continue solving the problem or to postpone it within one study session. This study session may then be appropriately trimmed.

The author read the whole book before starting the project. The time to read the prose could also be included in project’s total time consumption, but the author decided against it. In fact, when approached from the viewpoint of completing the exercises, material given in the book appeared to have nothing in common with the perception created by only reading the text.

A deliberate effort was spent on avoiding closing a problem at the same time as closing the study session.

The reason for this is to exploit the well-known tricks (See: Adler and Kounin 1939):

- “When you have something left undone, it is easier to make yourself start the next session.”
- Even just reading out the description of a problem makes the reader start thinking about how to solve it.

The data come in two datasets, closely related.

Dataset 1: Exercise completion time was recorded using a standard org-mode closure time tracking mechanism. (See Appendix: Full data on the exercise completion times.) For every exercise, completion time was recorded as an org-mode time-stamp, with minute-scale precision.

Dataset 2: Study sessions were recorded in a separate org-mode file in the standard org-mode time interval standard (two time-stamps):

```
"BEGIN_TIME -- END_TIME".
```

(See Appendix: Full data on the study sessions.)

During each study session, the author tried to concentrate as much as possible, and to do only the activities related to the problem set. These are not limited to just writing the code and tuning the software setup. They include the whole “package” of activities leading to the declaration of the problem solved. These include, but are not limited to, reading or watching additional material, asking questions, fixing bugs in related software, and similar activities.

Several software problems were discovered in the process of making this solution. These problems were reported to the software authors. Several of those problems were fixed after a short time, thus allowing the author to continue with the solution. For a few of the problems, workarounds were found. None of the problems prevented full completion of the problem set.

The author found it very helpful to have a simple dependency resolution tool at his disposal. As has been mentioned above, SICP’s problems make heavy use of one another. It was therefore critical to find a way to re-use code within a single org-mode document. Indeed org’s WEB-like capabilities («noweb»-links) proved to be sufficient. Noweb-links is a method for verbatim inclusion of a code block into other code blocks. In particular, Exercise 5.48 required inclusion of **58** other code blocks into the final solution block. Pure copying would not suffice because SICP exercises often involve the evaluation of the code written before (in the previous exercises) by the code written during the execution of the current exercise. Therefore, later exercises are likely to expose errors in the earlier exercises’ solutions.

3.2 Out-of-order problems and other measures

The following figure presents some of the aggregated measurements on solving of the problem set.

- **729** hours total work duration.
- **2.184** hours mean time spent on solving one problem.
- **0.96** hours was required for the dataset median problem.
- **94.73** hours for the hardest problem: writing a Scheme interpreter in a low-level language.
- **652** study sessions.
- **1.79** study sessions per problem on average.
- **>78000**-lines long .org file (**>2.6** megabytes) (5300 pages in a PDF).
- **1** median number of study sessions required to solve a single problem. The difference of almost 2 with the average hints that the few hardest problems required significantly more time than typical ones.
- **13** problems were solved out of order:
 - “Figure 1.1 Tree representation...”

- “Exercise 1.3 Sum of squares.”
- “Exercise 1.9 Iterative or recursive?”
- “Exercise 2.45 Split.”
- “Exercise 3.69 Triples.”
- “Exercise 2.61 Sets as ordered lists.”
- “Exercise 4.49 Alyssa’s generator.”
- “Exercise 4.69 Great-grandchildren.”
- “Exercise 4.71 Louis’ simple queries.”
- “Exercise 4.79 Prolog environments.”
- “Figure 5.1 Data paths for a Register Machine.”
- “Exercise 5.17 Printing labels.”
- “Exercise 5.40 Maintaining a compile-time environment.”

Figure 3. Aggregated measures of problem set execution

Thirteen problems were solved out-of-order. This means that those problems may have been the trickiest (although not necessarily the hardest.)

3.3 Ten hardest problems by raw time

Exercise	Days Spent	Spans Sessions	Minutes Spent
Exercise 2.46 <code>make-vect</code> .	2.578	5	535
Exercise 4.78 Non-deterministic queries.	0.867	6	602
Exercise 3.28 Primitive <code>orgate</code> .	1.316	2	783
Exercise 4.79 Prolog environments.	4.285	5	940
Exercise 3.9 Environment structures.	21.030	10	1100
Exercise 4.77 Lazy queries.	4.129	9	1214
Exercise 4.5 <code>cond</code> with arrow.	12.765	7	1252
Exercise 5.52 Making a compiler for Scheme.	22.975	13	2359
Exercise 2.92 Add, mul for different variables.	4.556	11	2404
Exercise 5.51 EC-evaluator in low-level language.	28.962	33	5684

Figure 4. The ten hardest problems

It is hardly unexpected that writing a Scheme interpreter in a low-level language (**Exercise 5.51**) turned out to be the most time-consuming problem of all the problem set. After all, it required learning an entirely new language from scratch. In the author’s case, the low-level language happened to be Fortran 2018. Learning Fortran

up to the level required is a relatively straightforward, albeit time-consuming.

Exercise 5.52, a compiler for Scheme, implicitly required that the previous exercise be solved already, as the runtime support code is shared between these two problems. All of the compiled EC-evaluator turned out to be just a single (very long) Fortran function.

Exercise 2.29 proves that it is possible to create significantly difficult exercises even without introducing the concept of mutation into the curriculum. This problem bears the comment from the SICP authors, “This is not easy!”. Indeed, the final solution contained more than eight hundred lines of code, involved designing an expression normalisation algorithm from scratch, and required twenty-five unit tests to ensure consistency. It is just a huge task.

Exercise 4.5 is probably one of those exercises that would benefit most from a Teaching Assistant’s help. In fact, the exercise itself is not that hard. The considerable workload comes from the fact that, in order to test that the solution is correct, a fully working interpreter is required. Therefore, this exercise, in fact, includes reading the whole of Chapter 4 and assembling the interpreter. Furthermore, the solution involves a lot of list manipulation, which is itself inherently error-prone if using only the functions already provided by SICP.

Exercise 4.77 required heavy modification of the codebase that had already been accumulated. It is likely to be the most architecture-intensive exercise of the book, apart from the exercise requiring a full rewrite of the backtracking engine of Prolog in a non-deterministic evaluator (**Exercise 4.78**). The code is very hard to implement incrementally, and the system is hardly testable until the last bit is finished. Furthermore, this exercise required the modification of the lowest-level data structures of the problem domain and modifying all the higher-level functions accordingly.

Exercise 4.79, is, in fact, an open-ended problem. The author considers it done, but the task is formulated so vaguely that it opens up an almost infinite range of possible solutions. This problem can hence consume any amount of time.

Exercise 3.9 required implementing a library for drawing environment diagrams. It may seem a trivial demand, as environment diagramming is an expected element of a decent debugger. However, the Scheme standard does not include many debugging capabilities. Debugging facilities differ among different Scheme implementation, but even those are usually not visual enough to generate the images required by the book. There exists an EnvDraw library (and its relatives), but the author failed to embed any of them into easily publishable Scheme

code. It turned out to be more straightforward to implement drawing diagrams as TikZ pictures in embedded L^AT_EX-blocks.

The time spent on **Exercise 3.28** includes the assembly of the whole circuit simulation code into a working system. The time required actually to solve the problem was comparatively short.

The same can be said about **Exercise 2.46**, which required writing a bridge between a Scheme interpreter and a drawing system. The exercise itself is relatively easy.

To sum up this section, the most laborious exercises in the book are the ones that require a student to:

- implement language features that are “assumed to be given”;
- assemble scattered code fragments into a working program;
- solve problems that have little to no theoretical coverage in the book.

In total, the ten most challenging problems account for 280 hours of work which is more than a third of the full problem set workload.

3.4 Minutes spent per problem




Figure 5. Minutes spent per problem

This graph is probably the most representative of the whole problem set. As expected, the last few problems turned out to be among the hardest. The second part of the course turned out to be more time-consuming than the first one.

3.5 Days spent per problem

The figure depicts the number of days (Y-axis) a problem (enumerated by the X-axis coordinate) was loaded in the author’s brain. In simple words, it is the number of days

that the state of “trying to solve a problem number X” spanned.

This measure is less justified than the “high concentration” time presented on the figure in the previous section. However, it may nevertheless be useful for encouraging students who get demotivated when spending a long “high concentration” session on a problem with no apparent success. Naturally, most (but not all) problems are solvable within one session (one day).




Figure 6. Days spent per problem

The second spike in the distribution can be attributed to general tiredness while solving such as huge problem set and a need for a break. The corresponding spike on the graph of the study sessions is less prominent.

3.6 Study sessions per problem




Figure 7. Study sessions per problem

A “session” may be defined as a period of high concentration when the student is actively trying to solve a

problem and get the code (or essay) written. This graph presents the number of sessions (Y-axis) spent on each problem (enumerated by the X-axis), regardless of the session length.

When a student goes on a vacation, the problem, presumably, remains loaded in the student’s brain. However, periodic “assaults” in the form of study sessions may be necessary to feed the subconscious processing with the new data.

During vacation time, there should be a spike on the “days per problem” graph, but not the “sessions per problem graph”. This can be seen on the second spike in the “days per problem” graph, which has its counterpart on the “sessions per problem” graph. The counterpart is much shorter.

3.7 Difficulty histogram (linear)

The linearly-scaled difficulty histogram depicts how many problems (Y-axis) require up to “bin size” hours for solution. Naturally, most of the exercises are solvable within one to three hours.




Figure 8. Difficulty distribution (linear)

3.8 Difficulty histogram (logarithmic)

The logarithmically-scaled difficulty histogram depicts how many problems (Y-axis) require up to 2^X hours for solution. It is very interesting to observe that the histogram shape resembles a uni-modal distribution. It is hard to think of a theoretical foundation on which to base assumptions for the distribution law. Prior research, however, may imply that the distribution is log-normal. (See Crow and Shimizu 2018)




Figure 9. Difficulty distribution (logarithmic)

4 Conclusion and Further Work

4.1 Conclusion

As follows immediately from the introduction, this report is essentially a single-point estimate of the difficulty distribution of a university-level problem set.

As far as the author knows, this is the first such a complete difficulty breakdown of a university-level problem set in existence.

As has been mentioned in section 3.2, the complete execution of the problem set required 729 hours. In simple words, this is a very long time. If a standard working day is assumed to have the length of 8 hours, the complete solution would require 91 days, or 14 weeks, or 3.5 months.

In the preface to the second edition, the authors claim that a redacted version (e.g. dropping the logical programming part, the part dedicated to the implementation of the register machine simulator, and most of the compiler-related sections) of the course can be covered in one semester. This statement is in agreement with the numbers presented in this report. Nevertheless, as the teachers would probably not want to assign every problem in the book to the student, they would need to make a selection based on both the coverage of the course topics and the time required. The author hopes that this report can provide an insight into the difficulty aspect.

On the other hand, the author would instead recommend opting for a two-semester course. If several of the hardest problems (i.e. problems discussed in section 3.3) are left out, the course can be fitted into two 300-hour modules. Three hundred hours per semester-long course

matches the author’s experience of studying partial differential equations at the Moscow Institute of Physics and Technology.

Another important consideration is the amount of time that instructors require to verify solutions and to write feedback for the students. It is reasonable to assume that marking the solutions and writing feedback would require the same amount of time (within an order of magnitude) as the amount needed to solve the problem set, since every problem solution would have to be visited by a marker at least once. For simplicity, the author assumes that writing feedback would require 72 hours per student.

This parameter would then be multiplied by the expected number of students per group, which may vary between institutions, but can be lower-bounded by 5. Therefore the rough estimate would be $\text{const} \cdot 72 \cdot 5 \approx 360$ hours, or 45 full working days (2 months). This duration is hardly practicable for a lone teacher, even if broken down over two semesters. (Each requiring 180 hours.) On the other hand, if the primary teacher is allowed to hire additional staff for marking, the problem becomes manageable again. One of the applications of this report may be as supporting evidence for lead instructors (professors) asking their school administration for teaching assistants.

4.2 Further work

The field of difficulty assessment (especially with the computer-based tools) of university courses still offers a lot to investigate. As far as the author of this report knows, this is the first exhaustive difficulty assessment of a university course. (This is not to say that SICP has not been successfully solved in full before. Various solutions can be found on many well-known software forges.)

The first natural direction of research would then be expanding the same effort towards other problem sets and other subjects.

On the other hand, this report is just a single point estimate, and therefore extremely biased. It may be a significant contribution if the same problem set (or indeed parts or even single problems of it) be solved by different people following the same protocol.

The provision of the solution protocol, the software setup and the time-tracking procedure, is deemed by the author to be a contribution of this report.

Professors teaching such a course are encouraged to show this report to their students and to suggest executing the problem set required along the lines of the protocol given here.

Another research direction could be towards finding an optimal curriculum design beyond the areas covered by SICP. It should not be unexpected if the students decide not to advance further in the course as long as

their personal difficulty assessment exceeds a certain unknown threshold. In other words, the author suspects that, at some point, the students may feel an emotion that may be expressed as, “I have been solving this for too long, and see little progress; I should stop.”

It would be interesting to measure such a threshold and to suggest curriculum design strategies that aim to minimise course drop-out. Such strategies may include attempts at hooking into students’ intrinsic motivation (and proper measurements of the execution process may provide an insight on where it is hidden), as well as better designing an extrinsic motivation toolset (e.g. finding better KPIs for rewards and penalties, and proper measures should be helpful in this approach as well).

It would be interesting to observe whether the students who follow the protocol (and see their progress after each session) are more or less likely to drop the course than those who do not. This could constitute a test of intrinsic motivation in line with the self-determination theory of Deci and Ryan (see Ryan and Deci 2017).

Another important direction may be the development and formalisation of coursework submission formats, in order to facilitate further collection of similar data on this or other problem sets.

4.3 Informal review

This section contains the author’s personal view on the problem set and the questions it raises.

The author (Vladimir Nikishkin), enjoyed doing it. On the other hand, it is hard to believe that teaching this course to first-year undergraduate students can easily be made successful. It is unlikely that a real-world student can dedicate seven hundred hours to a single subject, even if the subject is broken down into two semesters without significant support (the more so, recalling that 25 years has passed since the second edition was released, during which time the world of programming has expanded enormously.) Even if such a student is found, he would probably have other subjects in the semester, as well as the need to attend classes and demonstrations.

Admittedly, out of almost four hundred exercises, the author cannot find a single superfluous one. Even more, the author had to add some extra activities in order to cover several topics better. Every exercise teaches some valuable concept and nudges the student into thinking more deeply.

The course could have been improved in the area of garbage collection and other memory management topics. Indeed, the main cons-memory garbage collector is explained with sufficient detail to implement it, but several other parts of the interpreter memory model are left without explanation. Very little is said about efficiently storing numbers, strings and other objects.

There is not very much information about a rational process of software development. While this is not fundamental knowledge, but it would be helpful to undergraduates.

The last two exercises amount to one-fifth of the whole work. It was entirely unexpected to see a task to be completed in a language other than Scheme after having already finished most of the exercises.

Probably the biggest drawback of the book is the absence of any conclusion. Indeed, the book points the reader’s attention into various directions by means of an extensive bibliography. However, the author, as a willing student, would like to see a narrativised overview of the possible future directions.

4.4 Informal recommendations

If the author may, by virtue of personally experiencing this transformative experience, give a few suggestions to university curriculum designers, they would be the following:

- Deliberately teach students to use \TeX , and especially well technically harnessed \TeX (using a professional text editor, additional supportive software, such as syntax checkers, linters, and documentation lookup systems).

This is often considered to be a meta-cognitive exercise to be solved by the students, but the author’s personal experience is not reassuring in this aspect. Very few students, and even professionals, use \TeX efficiently. It took more than **50** hours just to refresh the skill of using \TeX that the author had already learnt, in order to write a thesis.

- Deliberately teach students to touch-type. This may not be necessary in the regions where touch-typing is included in the standard high school curriculum, but poor touch-typing skills are still a major problem in most parts of the world.
- Deliberately teach students to read software manuals. Indeed, much modern software has manuals built-in piece-wise right into the software itself. Often reading the whole manual is not required to perform the task. However, doing the reading at least once (i.e. reading **some** manual from the first page to the last), is a very enlightening experience, and additionally useful in teaching how to assess the time needed to grasp the skill of using a piece of software. As a by-product, this experience may help the students to write better manuals for their own software.
- Teach students to use a timer when doing homework, even if it is not an org-mode timer. A realistic assessment of how much effort things actually take is a paradigm-shifting experience.

- When writing a book on any subject, start from designing exercises, and afterwards write the text that helps to develop the skills required to solve those. Reading SICP without doing the exercises proved to be almost useless for this project, which was done two years after the first reading.
- Consider introducing elements of industrial illustration standards (UML, ArchiMate) into the teaching flow of an introductory programming course. Courses created to deliberately cover these standards typically suffer from being disconnected from the problem domain. (Few people would like to draw a yet another model of an ATM machine.) Introductory programming provides a surrogate domain that can be mapped onto the diagrams relatively easily and is unlikely to cause rejection.

5 Materials

This section attempts to provide a complete list of materials used in the process of the problem set solution. It is not to be confused with the list of materials used in the preparation of this Experience Report.

5.1 Books

- Structure and Interpretation of Computer Programs 2nd Ed. (Abelson, G. J. Sussman, and J. Sussman 1996)
- Structure and Interpretation of Computer Programs 1st Ed. (Abelson and G. J. Sussman 1985)
- Modern Fortran Explained 2018. (Metcalf, Reid, and Cohen 2018)
- Revised⁷ Report on Algorithmic Language Scheme. (Shinn, Cowan, Gleckler, et al. 2013)
- Logic Programming: A Classified Bibliography. (Balbin and Lecot 1985)
- Chibi-Scheme Manual. (Shinn 2018)
- TikZ Manual. (Tantau 2019)
- PlantUML Manual. (PlantUML Developers 2019a)
- UML Weekend Crash Course. (Pender 2002)
- GNU Emacs Manual. (Stallman et al. 2020c)
- GNU Emacs Lisp Reference Manual. (Stallman et al. 2020b)
- GNU Emacs Org-Mode Manual. (Dominik 2010)
- Debugging With GDB. (Stallman et al. 2020a)
- Implementations of Prolog. (J. A. Campbell 1984)

5.2 Software

- GNU Emacs. (Free Software Foundation 2019)
- Org-mode for Emacs. (Dominik et al. 2019)
- Chibi-Scheme. (Shinn et al. 2019)
- MIT/GNU Scheme. [For for portability checks]. (T. Campbell et al. 2019)
- Geiser. (Ruiz et al. 2020)

- GNU Debugger (GDB). (Free Software Foundation 2020)
- luaLaTeX/T_EX Live. (TeX User Groups 2019)
- TikZ/PGF. (Tantau et al. 2019)
- PlantUML. (PlantUML Developers 2019b)
- Graphviz. (Ellson et al. 2016)
- Slackware Linux 14.2-current. (Volkerding et al. 2019)

5.3 Papers

- Revised Report on the Propagator Model. (Radul and G. J. Sussman 2011)
- On Implementing Prolog In Functional Programming. (Carlsson 1984)
- eu-Prolog, Reference Manual. (Kohlbecker 1984)

References

- [1] Harold Abelson and Gerald J. Sussman. 1985. *Structure and Interpretation of Computer Programs*. (1st edition). MIT Press. ISBN: 0-262-51036-7 (cited on page 15).
- [2] Harold Abelson, Gerald J. Sussman, and Julia Sussman. 1996. *Structure and Interpretation of Computer Programs*. (2nd edition). MIT Press. ISBN: 0-262-01153-0 (cited on pages 1, 15).
- [3] Dan L. Adler and Jacob S. Kounin. 1939. Some factors operating at the moment of resumption of interrupted tasks. *The Journal of Psychology*, 7, 2, (April 1, 1939), 255–267. ISSN: 0022-3980. DOI: 10.1080/00223980.1939.9917632. Retrieved 05/17/2020 from <https://doi.org/10.1080/00223980.1939.9917632> (cited on page 9).
- [4] Isaac Balbin and Koenraad Lecot. 1985. *Logic Programming*. Springer Netherlands. ISBN: 978-94-009-5044-3. DOI: 10.1007/978-94-009-5044-3. <http://dx.doi.org/10.1007/978-94-009-5044-3> (cited on page 15).
- [5] Eric Berne. 1973. *What Do You Say After You Say Hello?* Bantam Books, New York (cited on page 2).
- [6] John A. Campbell, editor. 1984. *Implementations of Prolog*. Ellis Horwood/Halsted Press/Wiley. ISBN: 0-85312-675-5 (cited on page 15).
- [7] Taylor Campbell et al. 2019. MIT/GNU Scheme. Version 9.2. (2019). <https://www.gnu.org/software/mit-scheme/> (cited on page 15).
- [8] Mats Carlsson. 1984. On implementing Prolog in functional programming. *New Gener. Comput.*, 2, 4, 347–359. DOI: 10.1007/BF03037326. <https://doi.org/10.1007/BF03037326> (cited on page 15).
- [9] Karen L. St. Clair. 1999. A case against compulsory class attendance policies in higher education. *Innovative Higher Education*, 23, 3, 171–180. ISSN:

- 0742-5627. DOI: 10.1023/a:1022942400812. <http://dx.doi.org/10.1023/A:1022942400812> (cited on page 3).
- [10] Edwin L. Crow and Kunio Shimizu. 2018. *Lognormal Distributions: Theory and Applications*. Routledge, (May 2018). ISBN: 978-0-203-74866-4. DOI: 10.1201/9780203748664. <http://dx.doi.org/10.1201/9780203748664> (cited on page 12).
- [11] Carsten Dominik. 2010. *The Org-Mode 7 Reference Manual: Organize Your Life with GNU Emacs*. with contributions by David O'Toole, Bastien Guerry, Philip Rooke, Dan Davison, Eric Schulte, and Thomas Dye. Network Theory, UK (cited on page 15).
- [12] Carsten Dominik et al. 2019. Org-mode. Version 9.0. (2019). <https://www.orgmode.org> (cited on page 15).
- [13] John Ellson et al. 2016. Graphviz. Version 2.40.1. (December 25, 2016). <https://graphviz.org> (cited on page 15).
- [14] Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. 2018. *How to Design Programs: an Introduction to Programming and Computing*. The MIT Press, Cambridge, Massachusetts. ISBN: 978-0-262-53480-2 (cited on page 2).
- [15] Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. 2004. The structure and interpretation of the computer science curriculum. *Journal of Functional Programming*, 14, (July 2004), 365–378. DOI: 10.1017/S0956796804005076 (cited on page 2).
- [16] Free Software Foundation. 2020. GNU debugger. Version 9.1. (2020). <https://www.gnu.org/software/gdb/> (cited on page 15).
- [17] Free Software Foundation. 2019. GNU Emacs. Version 26.1. (2019). <https://www.gnu.org/software/emacs/> (cited on page 15).
- [18] Floyd W. Gembicki and Yacov Y. Haimes. 1975. Approach to performance and sensitivity multiobjective optimization: the goal attainment method. *IEEE Transactions on Automatic Control*, 20, 6, (December 1975), 769–771. ISSN: 0018-9286. DOI: 10.1109/tac.1975.1101105. <http://dx.doi.org/10.1109/TAC.1975.1101105> (cited on page 2).
- [19] Martin Hlosta, Drahomira Herrmannova, Lucie Vachova, Jakub Kuzilek, Zdenek Zdrahal, and Anika Wolff. 2018. Modelling student online behaviour in a virtual learning environment. *arXiv preprint arXiv:1811.06369* (cited on page 5).
- [20] Eugene Kohlbecker. 1984. eu-Prolog, Reference Manual and Report. Technical report. University of Indiana (Bloomington), Computer Science Department, (April 1984) (cited on page 15).
- [21] E. W. Kooker. 1976. Changes in grade distributions associated with changes in class attendance policies. *Psychology: A Journal of Human Behavior*, 13, 56–57 (cited on page 3).
- [22] Kelly Y. L. Ku and Irene T. Ho. 2010. Metacognitive strategies that enhance critical thinking. *Metacognition and Learning*, 5, 3, (July 2010), 251–267. DOI: 10.1007/s11409-010-9060-6. <https://doi.org/10.1007/s11409-010-9060-6> (cited on page 3).
- [23] Douglas McGregor. 1960. Theory X and theory Y. *Organization theory*, 358, 374 (cited on page 5).
- [24] Michael Metcalf, John Reid, and Malcolm Cohen. 2018. *Modern Fortran Explained*. Oxford University Press, (October 2018). ISBN: 978-0-19-881189-3. DOI: 10.1093/oso/9780198811893.001.0001. <http://dx.doi.org/10.1093/oso/9780198811893.001.0001> (cited on page 15).
- [25] Vladimir Nikishkin. 2020. A full solution to the structure and interpretation of computer programs. (May 30, 2020). <https://gitlab.com/Lockywolf/chibisicp/> (cited on page 1).
- [26] Thomas Pender. 2002. *UML Weekend Crash Course*. Hungry Minds, Indianapolis, IN. ISBN: 978-0-7645-4910-6 (cited on page 15).
- [27] PlantUML Developers. 2019. Drawing UML with plantuml. PlantUML language reference guide. (September 1, 2019). <http://plantuml.com/guide> (cited on page 15).
- [28] PlantUML Developers. 2019. PlantUML. (September 1, 2019). <http://plantuml.net> (cited on page 15).
- [29] Project Jupyter Developers. 2019. Jupyter Notebook: a server-client application that allows editing and running notebook documents via a web browser. P.O. Box 90596, Austin, TX 78709: NumFOCUS Foundation, (2019). <https://jupyter.org/> (cited on page 7).
- [30] Alexey Radul and Gerald J. Sussman. 2011. Revised report on the propagator model. (March 16, 2011). <http://groups.csail.mit.edu/mac/users/gjs/propagators/> (cited on page 15).
- [31] Jose A. O. Ruiz et al. 2020. Geiser. Version 0.11.2. (2020). <http://www.nongnu.org/geiser/index.html> (cited on page 15).
- [32] Richard M Ryan and Edward L Deci. 2017. *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Publications (cited on page 14).
- [33] Eric Schulte, Dan Davison, Thomas Dye, and Carsten Dominik. 2012. A multi-language computing environment for literate programming and reproducible research. *Journal of Statistical Software*, 46, 3, (January 2012), 1–24. ISSN: 1548-7660. <http://www.jstatsoft.org/v46/i03>.

[34] Alex Shinn. 2018. Chibi-Scheme. Minimal scheme implementation for use as an extension language. (January 25, 2018). <http://synthcode.com/scheme/chibi/> (cited on page 15).

[35] Alex Shinn et al. 2019. Chibi-Scheme. Version 0.8.0. (2019). <http://synthcode.com/scheme/chibi/> (cited on page 15).

[36] Alex Shinn, John Cowan, Arthur A. Gleckler, et al., editors. 2013. *Revised⁷ Report on the Algorithmic Language Scheme*. <https://small.r7rs.org/> (cited on page 15).

[37] Richard Stallman et al. 2020. Debugging with GDB. The GNU source-level debugger. <https://sourceware.org/gdb/download/onlinedocs/gdb/index.html> (cited on page 15).

[38] Richard Stallman et al. 2020. GNU Emacs Lisp Reference Manual. <https://www.gnu.org/software/emacs/manual/pdf/elisp.pdf> (cited on page 15).

[39] Richard Stallman et al. 2020. GNU Emacs Manual. <https://www.gnu.org/software/emacs/manual/pdf/emacs.pdf> (cited on page 15).

[40] Till Tantau. 2019. The TikZ and PGF Packages. Manual for version 3.1.5b-42-ga7cccca0. (May 19, 2019). <https://pgf-tikz.github.io/pgf/pgfmanual.pdf> (cited on page 15).

[41] Till Tantau et al. 2019. Portable graphics format. Version 3.1.3. (May 9, 2019). <https://github.com/pgf-tikz/pgf> (cited on page 15).

[42] TeX User Groups. 2019. TeX Live. Version 2019. (2019). <https://tug.org/texlive/> (cited on page 15).

[43] Jeffrey J. VanWormer, Simone A. French, Mark A. Pereira, and Ericka M. Welsh. 2008. The impact of regular self-weighing on weight management: a systematic literature review. *International Journal of Behavioral Nutrition and Physical Activity*, 5, 1, 54. DOI: 10.1186/1479-5868-5-54. <https://doi.org/10.1186/1479-5868-5-54> (cited on page 9).

[44] Patric Volkerding et al. 2019. Slackware® Linux. Version 14.2-current. Slackware Linux, Inc., (2019). <http://www.slackware.com/> (cited on page 15).

6 Appendix: Analysed data on problem difficulty

For the code used to generate the tables in the following sections, see: Appendix: Emacs Lisp code for data analysis.

6.1 Analysed time consumption

No	Exercise Name	Days Spent	Spans Sessions	Minutes Spent
----	---------------	------------	----------------	---------------

1	Exercise 1.1 Interpreter result	1.211	2	459
2	Exercise 1.2 Prefix form	0.001	1	2
3	Figure 1.1 Tree representation, showing the value of each su	0.007	1	10
4	Exercise 1.4 Compound expressions	0.003	1	4
5	Exercise 1.5 Ben's test	0.008	1	11
6	Exercise 1.6 If is a special form	0.969	2	118
7	Exercise 1.7 Good enough?	0.949	3	436
8	Exercise 1.8 Newton's method	0.197	2	193
9	Exercise 1.10 Ackermann's function	3.038	2	379
10	Exercise 1.11 Recursive vs iterative	0.037	1	54
11	Exercise 1.12 Recursive Pascal's triangle	0.012	1	17
12	Exercise 1.13 Fibonacci	0.092	1	132
13	Exercise 1.9 Iterative or recursive?	3.722	2	65
14	Exercise 1.14 count-change	1.038	2	50
15	Exercise 1.15 sine	0.267	2	195
16	Exercise 1.16 Iterative exponentiation	0.032	1	46
17	Exercise 1.17 Fast multiplication	0.019	1	28
18	Exercise 1.18 Iterative multiplication	0.497	2	23
19	Exercise 1.19 Logarithmic Fibonacci	1.374	2	93
20	Exercise 1.20 GCD applicative vs normal	0.099	1	142
21	Exercise 1.21 smallest-divisor	0.027	1	39
22	Exercise 1.22 timed-prime-test	0.042	1	61
23	Exercise 1.23 (next test-divisor)	0.383	2	5
24	Exercise 1.24 Fermat method	0.067	1	96
25	Exercise 1.25 expmod	0.051	1	74
26	Exercise 1.26 square vs mul	0.003	1	4

27	Exercise 1.27 Carmichael numbers	0.333	2	102	56	Exercise 2.10 div-interval-better	0.010	1	15
28	Exercise 1.28 Miller-Rabin	0.110	1	158	57	Exercise 2.11 mul-interval-nine-cases	0.052	1	75
29	Exercise 1.29 Simpson's integral	0.464	2	68	58	Exercise 2.12 make-center-percent	0.393	2	43
30	Exercise 1.30 Iterative sum	0.030	2	10	59	Exercise 2.13 formula for tolerance	0.003	1	5
31	Exercise 1.31 Product	0.028	1	40	60	Exercise 2.14 parallel-resistors	0.047	1	68
32	Exercise 1.32 Accumulator	0.017	1	24	61	Exercise 2.15 better-intervals	0.007	1	10
33	Exercise 1.33 filtered-accumulate	0.092	1	133	62	Exercise 2.16 interval-arithmetic	0.002	1	3
34	Exercise 1.34 lambda	0.006	1	8	63	Exercise 2.17 last-pair	0.966	2	89
35	Exercise 1.35 fixed-point	0.265	2	87	64	Exercise 2.18 reverse	0.006	1	9
36	Exercise 1.36 fixed-point-with-dampening	0.035	1	50	65	Exercise 2.19 coin-values	0.021	1	30
37	Exercise 1.37 cont-frac	0.569	2	348	66	Exercise 2.20 dotted-tail notation	0.311	2	156
38	Exercise 1.38 euler constant	0.000	1	0	67	Exercise 2.21 map-square-list	0.013	1	19
39	Exercise 1.39 tan-cf	0.025	1	36	68	Exercise 2.22 wrong list order	0.007	1	10
40	Exercise 1.40 newtons-method	0.205	2	6	69	Exercise 2.23 for-each	0.006	1	9
41	Exercise 1.41 double-double	0.010	1	15	70	Exercise 2.24 list-plot-result	0.111	2	75
42	Exercise 1.42 compose	0.004	1	6	71	Exercise 2.25 caddr	0.037	1	54
43	Exercise 1.43 repeated	0.019	1	27	72	Exercise 2.26 append cons list	0.011	1	16
44	Exercise 1.44 smoothing	0.099	2	142	73	Exercise 2.27 deep-reverse	0.433	2	40
45	Exercise 1.45 nth-root	0.056	1	80	74	Exercise 2.28 fringe	0.026	1	37
46	Exercise 1.46 iterative-improve	0.033	1	48	75	Exercise 2.29 mobile	0.058	1	83
47	Exercise 2.1 make-rat	1.608	2	109	76	Exercise 2.30 square-tree	0.100	2	122
48	Exercise 2.2 make-segment	0.024	1	34	77	Exercise 2.31 tree-map square tree	0.019	1	27
49	Exercise 2.3 make-rectangle	2.183	2	174	78	Exercise 2.32 subsets	0.010	1	15
50	Exercise 2.4 cons-lambda	0.007	1	10	79	Exercise 2.33 map-append-length	0.375	2	96
51	Exercise 2.5 cons-pow	0.041	1	59	80	Exercise 2.34 horners-rule	0.006	1	8
52	Exercise 2.6 Church Numerals	0.024	1	34	81	Exercise 2.35 count-leaves-accumulate	0.011	1	16
53	Exercise 2.7 make-interval	0.019	1	28	82	Exercise 2.36 accumulate-n	0.006	1	9
54	Exercise 2.8 sub-interval	0.124	1	58	83	Exercise 2.37 matrix-*-vector	0.017	1	24
55	Exercise 2.9 interval-width	0.006	1	8					

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

84	Exercise 2.38 fold-left	0.372	2	65	109	Exercise 2.61 sets as ordered lists	0.004	1	6
85	Exercise 2.39 reverse fold-right fold-left	0.005	1	7	110	Exercise 2.63 tree->list (binary search tree)	0.078	1	113
86	Exercise 2.40 unique-pairs	0.029	1	42	111	Exercise 2.64 balanced-tree	2.740	3	106
87	Exercise 2.41 triple-sum	2.195	2	57	112	Exercise 2.65 tree-union-set	9.785	2	47
88	Figure 2.8 A solution to the eight-queens puzzle.	0.001	1	2	113	Exercise 2.66 tree-lookup	0.035	1	50
89	Exercise 2.42 k-queens	3.299	2	122	114	Exercise 2.67 Huffman decode a simple message	0.303	3	108
90	Exercise 2.43 slow k-queens	0.019	1	28	115	Exercise 2.68 Huffman encode a simple message	0.023	1	33
91	Exercise 2.46 make-vector	2.578	5	535	116	Exercise 2.69 Generate Huffman tree	0.608	2	160
92	Exercise 2.47 make-frame	0.083	1	10	117	Exercise 2.70 Generate a tree and encode a song	0.072	2	57
93	Exercise 2.48 make-segment	0.054	1	78	118	Exercise 2.71 Huffman tree for frequencies 5 and 10	0.258	2	202
94	Exercise 2.49 segments->painter applications	0.294	2	139	119	Exercise 2.72 Huffman order of growth	0.050	2	26
95	Exercise 2.50 flip-horiz and rotate270 and rotate180	0.019	1	27	120	Exercise 2.73 data-driven-deriv	0.605	2	189
96	Exercise 2.51 below	1.801	4	524	121	Exercise 2.74 Insatiable Enterprises	0.410	4	171
97	Exercise 2.44 up-split	1.169	2	89	122	Exercise 2.75 make-from-mag-ang message passing	0.019	1	28
98	Exercise 2.45 split	0.113	2	23	123	Exercise 2.76 types or functions?	0.003	1	5
99	Exercise 2.52 modify square-limit	0.450	2	58	124	Exercise 2.77 generic-algebra-magnitude	0.772	3	190
100	Exercise 2.53 quote introduction	0.008	1	11	125	Exercise 2.78 Ordinary numbers for Scheme	0.212	2	67
101	Exercise 2.54 equal? implementation	0.050	1	72	126	Exercise 2.79 generic-equality	1.786	2	28
102	Exercise 2.55 quote quote	0.000	1	0	127	Exercise 2.80 Generic arithmetic zero?	0.056	1	80
103	Exercise 2.56 differentiation-exponentiation	0.393	2	65	128	Exercise 2.81 coercion to-itself	0.749	3	330
104	Exercise 2.57 differentiate-three-sum	0.560	3	147	129	Exercise 2.82 three-argument-coercion	0.433	2	230
105	Exercise 2.58 infix-notation	0.112	1	161	130	Exercise 2.83 Numeric Tower and (raise)	0.717	3	116
106	Exercise 2.59 union-set	0.277	2	6					
107	Exercise 2.60 duplicate-set	0.012	1	17					
108	Exercise 2.62 ordered-union-set (ordered list)	0.973	2	14					

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

131	Exercise 2.84 Using <code>raise</code> (<code>raise-type</code>) in <code>apply-generic</code>	0.865	2	135	152	Exercise 3.8 Right-to-left vs Left-to-right	0.026	1	38
132	Exercise 2.85 Dropping a type	3.089	5	507	153	Exercise 3.9 Environment structures	21.030	10	1100
133	Exercise 2.86 Compound complex numbers	0.274	2	108	154	Exercise 3.10 Using <code>let</code> to create state variables	4.933	2	138
134	Exercise 2.87 Generalized zero?	0.919	4	389	155	Exercise 3.11 Internal definitions	0.994	2	219
135	Exercise 2.88 Subtraction of polynomials	0.646	3	50	156	Exercise 3.12 Drawing <code>append!</code>	2.966	3	347
136	Exercise 2.89 Dense term-lists	0.083	1	120	157	Exercise 3.13 <code>make-cycle</code>	0.010	1	14
137	Exercise 2.90 Implementing dense polynomials as a separate p	0.400	2	148	158	Exercise 3.14 <code>mystery</code>	0.385	2	77
138	Exercise 2.91 Division of polynomials	0.111	2	103	159	Exercise 3.15 <code>set-to-wow!</code>	1.942	3	117
139	Exercise 2.92 Ordering of variables so that addition and mul	4.556	11	964	160	Exercise 3.16 <code>count-pairs</code>	0.171	1	118
140	Exercise 2.93 Rational polynomials	0.378	3	198	161	Exercise 3.17 Real <code>count-pairs</code>	0.029	1	42
141	Exercise 2.94 Greatest-common-divisor for polynomials	0.091	1	131	162	Exercise 3.18 Finding cycles	0.012	1	17
142	Exercise 2.95 Illustrate the non-integer problem	0.450	2	149	163	Exercise 3.19 Efficient finding cycles	0.934	2	205
143	Exercise 2.96 Integerizing factor	0.325	2	275	164	Exercise 3.20 Procedural <code>set-car!</code>	0.633	2	121
144	Exercise 2.97 Reduction of polynomials	0.201	1	140	165	Exercise 3.21 queues	0.021	1	30
145	Exercise 3.1 accumulators	0.425	2	53	166	Exercise 3.22 procedural queue	0.294	2	67
146	Exercise 3.2 make-monitored	0.027	1	39	167	Exercise 3.23 dequeue	0.049	2	71
147	Exercise 3.3 password protection	0.010	1	14	168	Exercise 3.24 tolerant tables	0.780	3	33
148	Exercise 3.4 call-the-cops	0.010	1	15	169	Exercise 3.25 multilevel tables	2.103	2	486
149	Exercise 3.5 Monte-Carlo	0.528	2	98	170	Exercise 3.26 binary tree table	0.013	1	18
150	Exercise 3.6 reset a <code>prng</code>	0.479	2	68	171	Exercise 3.27 memoization	0.802	2	2
151	Exercise 3.7 Joint accounts	0.059	1	85	172	Exercise 3.28 primitive or-gate	1.316	2	783
					173	Exercise 3.29 Compound or-gate	0.001	1	2
					174	Exercise 3.30 ripple-carry adder	0.009	1	13
					175	Exercise 3.31 Initial propagation	0.013	1	18
					176	Exercise 3.32 Order matters	0.007	1	10
					177	Exercise 3.33 averager constraint	9.460	3	198

178	Exercise 3.34 Wrong squarer	0.042	1	61	201	Exercise 3.57 exponential additions fibs	0.007	1	10
179	Exercise 3.35 Correct squarer	0.012	1	17	202	Exercise 3.58 Cryptic stream	0.010	1	14
180	Exercise 3.36 Connector environment diagram	3.319	3	263	203	Exercise 3.59 power series	0.422	2	30
181	Exercise 3.37 Expression-based constraints	0.037	1	53	204	Exercise 3.60 mul-series	0.048	1	69
182	Exercise 3.38 Timing	0.061	1	88	205	Exercise 3.61 power-series-inversion	0.087	1	126
183	Exercise 3.39 Serializer	1.266	4	269	206	Exercise 3.62 div-series	0.006	1	8
184	Exercise 3.40 Three parallel multiplications	5.973	3	332	207	Exercise 3.63 sqrt-stream	0.299	2	8
185	Exercise 3.41 Better protected account	4.229	2	97	208	Exercise 3.64 stream-limit	1.546	2	55
186	Exercise 3.42 Saving on serializers	0.023	1	33	209	Exercise 3.65 approximating logarithm	0.039	1	56
187	Exercise 3.43 Multiple serializations	0.040	1	58	210	Exercise 3.66 lazy pairs	0.515	2	107
188	Exercise 3.44 Transfer money	0.005	1	7	211	Exercise 3.67 all possible pairs	0.010	1	14
189	Exercise 3.45 new plus old serializers	0.004	1	6	212	Exercise 3.68 pairs-louis	0.012	1	17
190	Exercise 3.46 broken test-and-set!	0.007	1	10	213	Exercise 3.70 merge-weighted	0.522	2	188
191	Exercise 3.47 semaphores	1.044	2	53	214	Exercise 3.71 Ramanujan numbers	0.035	1	51
192	Exercise 3.48 serialized-exchange deadlock	0.022	1	31	215	Exercise 3.72 Ramanujan 3-numbers	0.901	2	187
193	Exercise 3.49 When numbering accounts doesn't work	0.008	1	11	216	Figure 3.32	0.022	1	32
194	Exercise 3.50 stream-map multiple arguments	0.317	3	96	217	Exercise 3.73 RC-circuit	0.090	1	130
195	Exercise 3.51 stream-show	0.007	1	10	218	Exercise 3.74 zero-crossings	0.153	1	221
196	Exercise 3.52 streams with mind-boggling	0.034	1	49	219	Exercise 3.75 filtering signals	0.056	1	81
197	Exercise 3.53 stream power of two	0.016	1	23	220	Exercise 3.76 stream-smooth	0.073	2	36
198	Exercise 3.54 mul-streams	0.005	1	7	221	Exercise 3.77	0.038	1	55
199	Exercise 3.55 streams partial-sums	0.013	1	18	222	Exercise 3.78 second order differential equation	0.039	1	56
200	Exercise 3.56 Hamming's streams-merge	0.015	1	21	223	Exercise 3.79 general second-order ode	0.007	1	10
					224	Figure 3.36	0.058	1	84
					225	Exercise 3.80 RLC circuit	0.013	1	19
					226	Exercise 3.81 renerator-in-streams	0.040	1	57
					227	Exercise 3.82 streams Monte-Carlo	0.378	2	57

228	Exercise 4.1 list-of-values ordered	0.437	2	14	250	Exercise 4.23 Analysing sequences	0.005	1	7
229	Exercise 4.2 application before assignments	0.021	1	30	251	Exercise 4.24 Analysis time test	0.022	1	32
230	Exercise 4.3 data-directed eval	0.030	1	43	252	Exercise 4.25 lazy factorial	0.034	1	49
231	Exercise 4.4 eval-and and eval-or	0.035	1	50	253	Exercise 4.26 unless as a special form	0.313	1	451
232	Exercise 4.5 cond with arrow	12.765	7	1252	254	Exercise 4.27 Working with mutation in lazy interpreters	0.515	2	112
233	Exercise 4.6 Implementing let	0.019	1	27	255	Exercise 4.28 Eval before applying	0.005	1	7
234	Exercise 4.7 Implementing let*	0.046	1	66	256	Exercise 4.29 Lazy evaluation is slow without memoization	0.035	1	50
235	Exercise 4.8 Implementing named let	0.070	1	101	257	Exercise 4.30 Lazy sequences	0.153	2	74
236	Exercise 4.9 Implementing until	0.928	3	102	258	Exercise 4.31 Lazy arguments with syntax extension	0.092	2	112
237	Exercise 4.10 Modifying syntax	14.168	3	462	259	Exercise 4.32 streams versus lazy lists	0.503	2	87
238	Exercise 4.11 Environment as a list of bindings	4.368	2	194	260	Exercise 4.33 quoted lazy lists	0.097	2	103
239	Exercise 4.12 Better abstractions for setting a value	0.529	2	120	261	Exercise 4.34 printing lazy lists	0.219	3	205
240	Exercise 4.13 Implementing <code>make-unbound!</code>	0.550	2	149	262	Exercise 4.50 The <code>ramb</code> operator	0.813	4	266
241	Exercise 4.14 meta map versus built-in map	0.004	1	6	263	Exercise 4.35 <code>an-integer-between</code> and Pythagorean triples	0.103	2	138
242	Exercise 4.15 The <code>halts?</code> predicate	0.018	1	26	264	Exercise 3.69 triples	0.115	2	85
243	Exercise 4.16 Simultaneous internal definitions	0.162	2	177	265	Exercise 4.36 infinite search for Pythagorean triples	0.011	1	16
244	Exercise 4.17 Environment with simultaneous definitions	0.036	1	52	266	Exercise 4.37 another method for triples	0.035	1	51
245	Exercise 4.18 Alternative scanning	0.018	1	26	267	Exercise 4.38 Logical puzzle - Not same floor	0.027	1	39
246	Exercise 4.19 Mutual simultaneous definitions	0.220	2	96	268	Exercise 4.39 Order of restrictions	0.003	1	5
247	Exercise 4.20 letrec	0.206	2	195	269	Exercise 4.40 People to floor assignment	0.019	1	28
248	Exercise 4.21 Y-combinator	0.013	1	18	270	Exercise 4.41 Ordinary Scheme to solve the problem	0.072	1	103
249	Exercise 4.22 Extending evaluator to support <code>let</code>	1.768	3	144					

271	Exercise 4.42 The liars puzzle	0.503	1	81	297	Exercise 4.72 interleave-stream	0.002	1	3
272	Exercise 4.43 Problematical Recreations	0.052	1	75	298	Exercise 4.73 flatten-stream delays	0.006	1	8
273	Exercise 4.44 Non-deterministic eight queens	0.074	1	106	299	Exercise 4.67 loop detector	0.251	1	361
274	Exercise 4.45 Five parses	0.186	3	145	300	Exercise 4.68 reverse rule	0.686	2	321
275	Exercise 4.46 Order of parsing	0.007	1	10	301	Exercise 4.69 great grandchildren	0.080	2	65
276	Exercise 4.47 Parse verb phrase by Louis	0.013	1	18	302	Exercise 4.71 Louis' simple queries	0.134	2	69
277	Exercise 4.48 Extending the grammar	0.037	1	1	303	Exercise 4.74 Alyssa's streams	0.044	1	64
278	Exercise 4.49 Alyssa's generator	0.031	1	45	304	Exercise 4.75 unique special form	0.055	1	79
279	Exercise 4.51 Implementing permanent-set!	0.030	1	43	305	Exercise 4.76 improving and	0.797	2	438
280	Exercise 4.52 if-fail	0.063	1	91	306	Figure 5.2 Controller for a GCD Machine	0.167	3	124
281	Exercise 4.53 test evaluation	0.005	1	7	307	Exercise 5.1 Register machine plot	0.020	1	29
282	Exercise 4.54 analyze-require	0.468	2	31	308	Figure 5.1 Data paths for a Register Machine	0.599	2	115
283	Exercise 4.55 Simple queries	0.258	2	372	309	Exercise 5.2 Register machine language description of Exerci	0.006	1	8
284	Exercise 4.56 Compound queries	0.018	1	26	310	Exercise 5.3 Machine for sqrt using Newton Method	0.306	2	286
285	Exercise 4.57 custom rules	0.147	3	112	311	Exercise 5.4 Recursive register machines	1.001	4	274
286	Exercise 4.58 big shot	0.025	1	36	312	Exercise 5.5 Hand simulation for factorial and Fibonacci	0.110	1	158
287	Exercise 4.59 meetings	0.031	1	45	313	Exercise 5.6 Fibonacci machine extra instructions	0.011	1	16
288	Exercise 4.60 pairs live near	0.016	1	23	314	Exercise 5.7 Test the 5.4 machine on a simulator	0.458	2	133
289	Exercise 4.61 next-to relation	0.008	1	11	315	Exercise 5.8 Ambiguous labels	0.469	1	160
290	Exercise 4.62 last-pair	0.033	1	48	316	Exercise 5.9 Prohibit (op)s on labels	0.017	1	25
291	Exercise 4.63 Genesis	0.423	2	40	317	Exercise 5.10 Changing syntax	0.011	1	16
292	Figure 4.6 How the system works	0.022	1	31	318	Exercise 5.11 Save and restore	0.619	3	186
293	Exercise 4.64 broken outranked-by	0.065	1	94					
294	Exercise 4.65 second-degree subordinates	0.012	1	17					
295	Exercise 4.66 Ben's accumulation	0.013	1	18					
296	Exercise 4.70 Cons-stream delays its second argument	0.167	3	79					

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

319	Exercise 5.12 Data paths from controller	0.424	2	183	340	Exercise 5.32 symbol-lookup optimization	0.052	1	75
320	Exercise 5.13 Registers from controller	0.470	2	101	341	Exercise 5.33 compiling factorial-alt	0.753	2	267
321	Exercise 1.3 Sum of squares	1.044	1	6	342	Exercise 5.34 compiling iterative factorial	0.169	1	243
322	Exercise 5.14 Profiling	0.347	2	57	343	Exercise 5.35 Decompilation	0.022	1	32
323	Exercise 5.15 Instruction counting	0.052	1	75	344	Exercise 5.36 Order of evaluation	0.845	4	256
324	Exercise 5.16 Tracing execution	0.058	1	83	345	Exercise 5.37 preserving	0.135	1	194
325	Exercise 5.18 Register tracing	0.631	2	90	346	Exercise 5.38 open code primitives	0.914	3	378
326	Exercise 5.19 Breakpoints	0.149	1	215	347	Exercise 5.41 find-variable	0.028	1	40
327	Exercise 5.17 Printing labels	0.001	1	1	348	Exercise 5.39 lexical-address-lookup	0.044	1	64
328	Exercise 5.20 Drawing a list "(#1=(1 . 2) #1)"	0.189	2	139	349	Exercise 5.42 Rewrite compile-variable and ~compile-assign	0.679	2	118
329	Exercise 5.21 Register machines for list operations	0.617	2	115	350	Exercise 5.40 maintaining a compile-time environment	0.085	2	101
330	Exercise 5.22 append and append! as register machines	0.047	1	68	351	Exercise 5.43 Scanning out defines	0.249	3	261
331	Exercise 5.23 Extending EC-evaluator with let and cond	0.862	4	363	352	Exercise 5.44 open code with compile-time environment	0.020	1	29
332	Exercise 5.24 Making cond a primitive	0.160	2	199	353	Exercise 5.45 stack usage analysis for a factorial	0.528	1	61
333	Exercise 5.25 Normal-order (lazy) evaluation	1.010	4	342	354	Exercise 5.46 stack usage analysis for fibonacci	0.017	1	25
334	Exercise 5.26 Explore tail recursion with factorial	0.195	2	26	355	Exercise 5.47 calling interpreted procedures	0.049	1	71
335	Exercise 5.27 Stack depth for a recursive factorial	0.008	1	11	356	Exercise 5.48 compile-and-run	1.020	3	264
336	Exercise 5.28 Interpreters without tail recursion	0.028	1	40	357	Exercise 5.49 read-compile-execute-print loop	0.015	1	22
337	Exercise 5.29 Stack in tree-recursive Fibonacci	0.015	1	21	358	Exercise 4.77 lazy queries	4.129	9	1214
338	Exercise 5.30 Errors	0.615	3	147	359	Exercise 5.50 Compiling the metacircular evaluator	0.007	1	10
339	Exercise 5.31 a preserving mechanism	0.417	2	161					

360	Exercise 4.78 non-deterministic queries	0.867	6	602
361	Exercise 5.51 Translating the EC-evaluator into a low-level	28.962	33	5684
362	Exercise 5.52 Making a compiler for Scheme	22.975	13	2359
363	Exercise 4.79 prolog environments	4.285	5	940

6.2 Time consumption histogram linear

Bin Lower Bound (Minutes)	N. tasks
0.	301
177.625	38
355.25	14
532.875	2
710.5	1
888.125	2
1065.75	2
1243.375	1
1421.	0
1598.625	0
1776.25	0
1953.875	0
2131.5	0
2309.125	1
2486.75	0
2664.375	0
2842.	0
3019.625	0
3197.25	0
3374.875	0
3552.5	0
3730.125	0
3907.75	0
4085.375	0
4263.	0
4440.625	0
4618.25	0
4795.875	0
4973.5	0
5151.125	1

6.3 Time consumption histogram logarithmic

Bin Lower Bound (Minutes)	N. tasks
1	2
2	6
4	15
8	41
16	55
32	67
64	85
128	52
256	29
512	6
1024	3
2048	1
4096	1

7 Appendix: Full data on the study sessions.

This section lists the data on each study session in the "BEGIN_TIMESTAMP-END_TIMESTAMP:duration" format.

The earliest time stamp also marks the beginning of the whole project.

[2020-05-10 Sun 14:39]-[2020-05-10 Sun 18:00]	3:21
[2020-05-09 Sat 19:13]-[2020-05-09 Sat 22:13]	3:00
[2020-05-09 Sat 09:34]-[2020-05-09 Sat 14:34]	5:00
[2020-05-08 Fri 21:45]-[2020-05-08 Fri 23:17]	1:32
[2020-05-08 Fri 18:30]-[2020-05-08 Fri 21:18]	2:48
[2020-05-06 Wed 10:12]-[2020-05-06 Wed 11:09]	0:57
[2020-05-05 Tue 12:11]-[2020-05-06 Wed 00:00]	11:49
[2020-05-04 Mon 18:20]-[2020-05-05 Tue 00:30]	6:10
[2020-05-04 Mon 14:02]-[2020-05-04 Mon 17:43]	3:41
[2020-05-03 Sun 21:03]-[2020-05-03 Sun 22:02]	0:59
[2020-04-30 Thu 09:28]-[2020-04-30 Thu 11:23]	1:55
[2020-04-29 Wed 20:00]-[2020-04-29 Wed 23:25]	3:25
[2020-04-28 Tue 22:55]-[2020-04-29 Wed 00:11]	1:16
[2020-04-28 Tue 21:00]-[2020-04-28 Tue 22:50]	1:50
[2020-04-27 Mon 20:09]-[2020-04-27 Mon 22:09]	2:00
[2020-04-26 Sun 20:10]-[2020-04-26 Sun 23:52]	3:42
[2020-04-21 Tue 11:01]-[2020-04-21 Tue 12:26]	1:25
[2020-04-13 Mon 11:40]-[2020-04-13 Mon 11:55]	0:15
[2020-04-11 Sat 11:50]-[2020-04-11 Sat 15:50]	4:00
[2020-04-10 Fri 09:50]-[2020-04-10 Fri 14:26]	4:36
[2020-04-09 Thu 19:50]-[2020-04-09 Thu 23:10]	3:20
[2020-04-09 Thu 09:55]-[2020-04-09 Thu 13:00]	3:05
[2020-04-08 Wed 22:50]-[2020-04-08 Wed 23:55]	1:05
[2020-04-08 Wed 18:30]-[2020-04-08 Wed 21:11]	2:41
[2020-04-08 Wed 09:15]-[2020-04-08 Wed 12:15]	3:00
[2020-04-07 Tue 20:46]-[2020-04-07 Tue 23:37]	2:51
[2020-04-07 Tue 09:41]-[2020-04-07 Tue 11:57]	2:16
[2020-04-06 Mon 18:58]-[2020-04-06 Mon 21:20]	2:22
[2020-04-06 Mon 12:09]-[2020-04-06 Mon 14:15]	2:06
[2020-04-05 Sun 11:30]-[2020-04-05 Sun 15:11]	3:41
[2020-04-04 Sat 22:08]-[2020-04-04 Sat 22:45]	0:37
[2020-04-04 Sat 17:54]-[2020-04-04 Sat 20:50]	2:56
[2020-04-04 Sat 17:24]-[2020-04-04 Sat 17:41]	0:17
[2020-04-04 Sat 15:15]-[2020-04-04 Sat 16:10]	0:55
[2020-04-03 Fri 20:22]-[2020-04-03 Fri 22:21]	1:59
[2020-04-01 Wed 13:05]-[2020-04-01 Wed 15:05]	2:00
[2020-03-29 Sun 13:05]-[2020-03-29 Sun 22:05]	9:00

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2020-03-28 Sat 13:04]-[2020-03-28 Sat 22:04] 9:00	[2020-02-24 Mon 10:08]-[2020-02-24 Mon 12:39] 2:31
[2020-03-26 Thu 20:20]-[2020-03-26 Thu 23:33] 3:13	[2020-02-23 Sun 19:20]-[2020-02-23 Sun 20:48] 1:28
[2020-03-26 Thu 10:43]-[2020-03-26 Thu 14:39] 3:56	[2020-02-23 Sun 12:52]-[2020-02-23 Sun 16:45] 3:53
[2020-03-24 Tue 20:00]-[2020-03-24 Tue 23:50] 3:50	[2020-02-22 Sat 21:35]-[2020-02-23 Sun 00:25] 2:50
[2020-03-24 Tue 09:10]-[2020-03-24 Tue 12:34] 3:24	[2020-02-22 Sat 19:59]-[2020-02-22 Sat 21:03] 1:04
[2020-03-23 Mon 19:56]-[2020-03-23 Mon 23:06] 3:10	[2020-02-22 Sat 12:20]-[2020-02-22 Sat 18:35] 6:15
[2020-03-23 Mon 10:23]-[2020-03-23 Mon 13:23] 3:00	[2020-02-21 Fri 20:55]-[2020-02-22 Sat 00:30] 3:35
[2020-03-23 Mon 09:06]-[2020-03-23 Mon 10:56] 1:50	[2020-02-21 Fri 17:30]-[2020-02-21 Fri 18:51] 1:21
[2020-03-22 Sun 18:46]-[2020-03-22 Sun 22:45] 3:59	[2020-02-21 Fri 10:40]-[2020-02-21 Fri 16:40] 6:00
[2020-03-22 Sun 12:45]-[2020-03-22 Sun 13:46] 1:01	[2020-02-20 Thu 17:00]-[2020-02-20 Thu 23:33] 6:33
[2020-03-21 Sat 19:07]-[2020-03-21 Sat 21:35] 2:28	[2020-02-20 Thu 14:43]-[2020-02-20 Thu 15:08] 0:25
[2020-03-17 Tue 19:11]-[2020-03-17 Tue 22:11] 3:00	[2020-02-20 Thu 10:05]-[2020-02-20 Thu 13:54] 3:49
[2020-03-15 Sun 09:10]-[2020-03-15 Sun 12:41] 3:31	[2020-02-19 Wed 21:35]-[2020-02-20 Thu 00:36] 3:01
[2020-03-14 Sat 23:01]-[2020-03-14 Sat 23:54] 0:53	[2020-02-19 Wed 19:50]-[2020-02-19 Wed 21:30] 1:40
[2020-03-14 Sat 20:46]-[2020-03-14 Sat 23:01] 2:15	[2020-02-19 Wed 13:34]-[2020-02-19 Wed 18:15] 4:41
[2020-03-14 Sat 20:39]-[2020-03-14 Sat 20:46] 0:07	[2020-02-19 Wed 11:10]-[2020-02-19 Wed 13:34] 2:24
[2020-03-14 Sat 17:23]-[2020-03-14 Sat 20:39] 3:16	[2020-02-18 Tue 21:05]-[2020-02-19 Wed 00:27] 3:22
[2020-03-14 Sat 12:00]-[2020-03-14 Sat 15:53] 3:53	[2020-02-18 Tue 19:02]-[2020-02-18 Tue 20:13] 1:11
[2020-03-13 Fri 20:01]-[2020-03-13 Fri 23:01] 3:00	[2020-02-18 Tue 16:58]-[2020-02-18 Tue 18:36] 1:38
[2020-03-13 Fri 09:20]-[2020-03-13 Fri 11:58] 2:38	[2020-02-18 Tue 10:55]-[2020-02-18 Tue 15:21] 4:26
[2020-03-12 Thu 20:30]-[2020-03-12 Thu 23:29] 2:59	[2020-02-17 Mon 19:20]-[2020-02-18 Tue 00:12] 4:52
[2020-03-11 Wed 12:12]-[2020-03-11 Wed 13:18] 1:06	[2020-02-17 Mon 15:20]-[2020-02-17 Mon 18:00] 2:40
[2020-03-11 Wed 10:45]-[2020-03-11 Wed 11:09] 0:24	[2020-02-17 Mon 14:17]-[2020-02-17 Mon 15:09] 0:52
[2020-03-11 Wed 09:15]-[2020-03-11 Wed 10:45] 1:30	[2020-02-16 Sun 21:21]-[2020-02-17 Mon 00:52] 3:31
[2020-03-10 Tue 20:22]-[2020-03-11 Wed 00:09] 3:47	[2020-02-16 Sun 20:03]-[2020-02-16 Sun 20:14] 0:11
[2020-03-10 Tue 09:08]-[2020-03-10 Tue 13:44] 4:36	[2020-02-16 Sun 19:00]-[2020-02-16 Sun 19:30] 0:30
[2020-03-09 Mon 22:28]-[2020-03-09 Mon 23:32] 1:04	[2020-02-16 Sun 16:06]-[2020-02-16 Sun 18:38] 2:32
[2020-03-09 Mon 09:08]-[2020-03-09 Mon 11:59] 2:51	[2020-02-16 Sun 12:59]-[2020-02-16 Sun 14:37] 1:38
[2020-03-08 Sun 18:30]-[2020-03-08 Sun 21:29] 2:59	[2020-02-16 Sun 10:30]-[2020-02-16 Sun 12:22] 1:52
[2020-03-08 Sun 16:51]-[2020-03-08 Sun 18:08] 1:17	[2020-02-15 Sat 22:10]-[2020-02-15 Sat 23:52] 1:42
[2020-03-08 Sun 13:50]-[2020-03-08 Sun 15:36] 1:46	[2020-02-15 Sat 21:01]-[2020-02-15 Sat 21:50] 0:49
[2020-03-08 Sun 11:56]-[2020-03-08 Sun 13:28] 1:32	[2020-02-15 Sat 15:03]-[2020-02-15 Sat 18:34] 3:31
[2020-03-07 Sat 18:00]-[2020-03-07 Sat 21:36] 3:36	[2020-02-14 Fri 18:53]-[2020-02-15 Sat 04:33] 9:40
[2020-03-07 Sat 11:35]-[2020-03-07 Sat 16:09] 4:34	[2020-02-13 Thu 16:15]-[2020-02-13 Thu 17:21] 1:06
[2020-03-06 Fri 17:37]-[2020-03-06 Fri 21:48] 4:11	[2020-02-13 Thu 00:12]-[2020-02-13 Thu 01:45] 1:33
[2020-03-06 Fri 13:11]-[2020-03-06 Fri 14:16] 1:05	[2020-02-12 Wed 18:36]-[2020-02-12 Wed 22:30] 3:54
[2020-03-06 Fri 09:42]-[2020-03-06 Fri 12:39] 2:57	[2020-02-12 Wed 13:16]-[2020-02-12 Wed 14:55] 1:39
[2020-03-05 Thu 16:54]-[2020-03-05 Thu 21:34] 4:40	[2020-02-12 Wed 08:37]-[2020-02-12 Wed 12:20] 3:43
[2020-03-05 Thu 08:58]-[2020-03-05 Thu 13:24] 4:26	[2020-02-11 Tue 18:51]-[2020-02-11 Tue 21:54] 3:03
[2020-03-04 Wed 19:51]-[2020-03-04 Wed 22:51] 3:00	[2020-02-11 Tue 04:30]-[2020-02-11 Tue 08:09] 3:39
[2020-03-04 Wed 11:33]-[2020-03-04 Wed 12:31] 0:58	[2020-02-10 Mon 06:42]-[2020-02-10 Mon 07:28] 0:46
[2020-03-04 Wed 09:32]-[2020-03-04 Wed 11:01] 1:29	[2020-02-06 Thu 15:42]-[2020-02-06 Thu 22:08] 6:26
[2020-03-03 Tue 19:13]-[2020-03-03 Tue 21:46] 2:33	[2020-02-01 Sat 15:05]-[2020-02-01 Sat 15:36] 0:31
[2020-03-03 Tue 12:20]-[2020-03-03 Tue 14:58] 2:38	[2020-01-23 Thu 17:06]-[2020-01-23 Thu 18:51] 1:45
[2020-03-03 Tue 09:13]-[2020-03-03 Tue 11:57] 2:44	[2020-01-22 Wed 20:53]-[2020-01-22 Wed 21:05] 0:12
[2020-03-02 Mon 18:30]-[2020-03-02 Mon 18:50] 0:20	[2020-01-22 Wed 13:40]-[2020-01-22 Wed 20:20] 6:40
[2020-03-02 Mon 12:01]-[2020-03-02 Mon 14:43] 2:42	[2020-01-21 Tue 15:33]-[2020-01-21 Tue 16:57] 1:24
[2020-03-02 Mon 09:02]-[2020-03-02 Mon 11:30] 2:28	[2020-01-17 Fri 19:13]-[2020-01-17 Fri 23:00] 3:47
[2020-03-01 Sun 19:07]-[2020-03-01 Sun 21:25] 2:18	[2020-01-11 Sat 10:56]-[2020-01-11 Sat 18:24] 7:28
[2020-03-01 Sun 17:50]-[2020-03-01 Sun 18:41] 0:51	[2020-01-10 Fri 22:20]-[2020-01-10 Fri 23:56] 1:36
[2020-03-01 Sun 11:09]-[2020-03-01 Sun 15:15] 4:06	[2020-01-10 Fri 09:40]-[2020-01-10 Fri 13:20] 3:40
[2020-02-29 Sat 21:30]-[2020-02-29 Sat 22:16] 0:46	[2020-01-09 Thu 20:10]-[2020-01-09 Thu 22:15] 2:05
[2020-02-29 Sat 12:48]-[2020-02-29 Sat 19:17] 6:29	[2020-01-09 Thu 08:50]-[2020-01-09 Thu 09:55] 1:05
[2020-02-28 Fri 20:21]-[2020-02-28 Fri 23:10] 2:49	[2020-01-08 Wed 19:21]-[2020-01-09 Thu 00:42] 5:21
[2020-02-28 Fri 18:26]-[2020-02-28 Fri 19:22] 0:56	[2020-01-08 Wed 09:20]-[2020-01-08 Wed 18:12] 8:52
[2020-02-28 Fri 11:55]-[2020-02-28 Fri 12:02] 0:07	[2020-01-07 Tue 16:31]-[2020-01-07 Tue 18:31] 2:00
[2020-02-27 Thu 09:20]-[2020-02-27 Thu 10:57] 1:37	[2020-01-07 Tue 08:55]-[2020-01-07 Tue 12:49] 3:54
[2020-02-26 Wed 20:47]-[2020-02-26 Wed 23:44] 2:57	[2020-01-06 Mon 22:30]-[2020-01-06 Mon 23:31] 1:01
[2020-02-26 Wed 12:07]-[2020-02-26 Wed 13:40] 1:33	[2020-01-06 Mon 09:20]-[2020-01-06 Mon 11:56] 2:36
[2020-02-26 Wed 09:29]-[2020-02-26 Wed 11:00] 1:31	[2020-01-04 Sat 20:25]-[2020-01-04 Sat 21:09] 0:44
[2020-02-25 Tue 19:18]-[2020-02-25 Tue 22:51] 3:33	[2020-01-04 Sat 09:37]-[2020-01-04 Sat 13:22] 3:45
[2020-02-25 Tue 09:01]-[2020-02-25 Tue 10:42] 1:41	[2020-01-03 Fri 21:13]-[2020-01-03 Fri 23:59] 2:46
[2020-02-24 Mon 19:23]-[2020-02-25 Tue 00:15] 4:52	[2020-01-03 Fri 18:13]-[2020-01-03 Fri 19:13] 1:00
[2020-02-24 Mon 13:00]-[2020-02-24 Mon 13:36] 0:36	[2020-01-03 Fri 12:08]-[2020-01-03 Fri 14:12] 2:04

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2020-01-02 Thu 09:35]-[2020-01-02 Thu 11:58] 2:23	[2019-10-19 Sat 19:25]-[2019-10-19 Sat 20:45] 1:20
[2019-12-29 Sun 02:12]-[2019-12-29 Sun 05:42] 3:30	[2019-10-19 Sat 16:12]-[2019-10-19 Sat 18:47] 2:35
[2019-12-26 Thu 16:59]-[2019-12-26 Thu 19:51] 2:52	[2019-10-17 Thu 19:18]-[2019-10-17 Thu 22:55] 3:37
[2019-12-23 Mon 05:03]-[2019-12-23 Mon 05:31] 0:28	[2019-10-17 Thu 09:30]-[2019-10-17 Thu 11:42] 2:12
[2019-12-23 Mon 03:02]-[2019-12-23 Mon 04:03] 1:01	[2019-10-16 Wed 14:52]-[2019-10-16 Wed 14:59] 0:07
[2019-12-22 Sun 16:51]-[2019-12-22 Sun 18:40] 1:49	[2019-10-16 Wed 09:08]-[2019-10-16 Wed 10:08] 1:00
[2019-12-21 Sat 19:23]-[2019-12-22 Sun 00:19] 4:56	[2019-10-15 Tue 22:35]-[2019-10-15 Tue 23:30] 0:55
[2019-12-20 Fri 14:10]-[2019-12-20 Fri 17:11] 3:01	[2019-10-15 Tue 19:30]-[2019-10-15 Tue 21:40] 2:10
[2019-12-19 Thu 23:20]-[2019-12-19 Thu 23:38] 0:18	[2019-10-15 Tue 09:10]-[2019-10-15 Tue 12:56] 3:46
[2019-12-18 Wed 10:47]-[2019-12-18 Wed 12:47] 2:00	[2019-10-14 Mon 19:51]-[2019-10-14 Mon 23:10] 3:19
[2019-12-09 Mon 10:47]-[2019-12-09 Mon 13:21] 2:34	[2019-10-14 Mon 15:57]-[2019-10-14 Mon 17:23] 1:26
[2019-12-08 Sun 17:47]-[2019-12-09 Sun 00:28] 6:41	[2019-10-12 Sat 20:05]-[2019-10-12 Sat 21:33] 1:28
[2019-12-07 Sat 16:07]-[2019-12-07 Sat 23:15] 7:08	[2019-10-12 Sat 15:56]-[2019-10-12 Sat 16:07] 0:11
[2019-12-06 Fri 19:04]-[2019-12-06 Fri 20:54] 1:50	[2019-10-12 Sat 10:31]-[2019-10-12 Sat 12:31] 2:00
[2019-12-04 Wed 18:06]-[2019-12-05 Thu 00:42] 6:36	[2019-10-11 Fri 19:55]-[2019-10-11 Fri 22:34] 2:39
[2019-12-04 Wed 12:36]-[2019-12-04 Wed 13:05] 0:29	[2019-10-11 Fri 17:55]-[2019-10-11 Fri 19:28] 1:33
[2019-12-03 Tue 22:18]-[2019-12-03 Tue 23:27] 1:09	[2019-10-11 Fri 14:35]-[2019-10-11 Fri 14:47] 0:12
[2019-12-03 Tue 21:21]-[2019-12-03 Tue 22:18] 0:57	[2019-10-11 Fri 09:10]-[2019-10-11 Fri 11:10] 2:00
[2019-12-03 Tue 12:40]-[2019-12-03 Tue 15:25] 2:45	[2019-10-10 Thu 20:26]-[2019-10-10 Thu 21:48] 1:22
[2019-12-02 Mon 20:06]-[2019-12-02 Mon 23:30] 3:24	[2019-10-10 Thu 17:26]-[2019-10-10 Thu 19:40] 2:14
[2019-12-01 Sun 22:07]-[2019-12-02 Mon 01:06] 2:59	[2019-10-10 Thu 12:15]-[2019-10-10 Thu 14:37] 2:22
[2019-12-01 Sun 18:59]-[2019-12-01 Sun 19:59] 1:00	[2019-10-10 Thu 08:50]-[2019-10-10 Thu 11:29] 2:39
[2019-11-30 Sat 14:19]-[2019-11-30 Sat 15:15] 0:56	[2019-10-09 Wed 20:16]-[2019-10-09 Wed 20:55] 0:39
[2019-11-29 Fri 20:07]-[2019-11-29 Fri 21:24] 1:17	[2019-10-09 Wed 16:46]-[2019-10-09 Wed 17:55] 1:09
[2019-11-29 Fri 11:51]-[2019-11-29 Fri 12:10] 0:19	[2019-10-09 Wed 11:27]-[2019-10-09 Wed 13:38] 2:11
[2019-11-28 Thu 09:30]-[2019-11-28 Thu 15:00] 5:30	[2019-09-29 Sun 17:01]-[2019-09-29 Sun 17:23] 0:22
[2019-11-26 Tue 09:15]-[2019-11-26 Tue 12:57] 3:42	[2019-09-27 Fri 08:56]-[2019-09-27 Fri 10:20] 1:24
[2019-11-25 Mon 10:35]-[2019-11-25 Mon 13:02] 2:27	[2019-09-26 Thu 21:25]-[2019-09-26 Thu 23:38] 2:13
[2019-11-20 Wed 12:08]-[2019-11-20 Wed 14:29] 2:21	[2019-09-25 Wed 21:55]-[2019-09-25 Wed 22:18] 0:23
[2019-11-20 Wed 09:25]-[2019-11-20 Wed 11:32] 2:07	[2019-09-25 Wed 12:20]-[2019-09-25 Wed 15:12] 3:02
[2019-11-19 Tue 11:45]-[2019-11-19 Tue 14:42] 2:57	[2019-09-25 Wed 09:20]-[2019-09-25 Wed 11:25] 2:05
[2019-11-13 Wed 20:52]-[2019-11-13 Wed 22:25] 1:33	[2019-09-24 Tue 22:10]-[2019-09-24 Tue 23:16] 1:06
[2019-11-12 Tue 19:47]-[2019-11-12 Tue 21:14] 1:27	[2019-09-24 Tue 12:05]-[2019-09-24 Tue 13:49] 1:44
[2019-11-12 Tue 09:30]-[2019-11-12 Tue 11:49] 2:19	[2019-09-24 Tue 01:17]-[2019-09-24 Tue 02:15] 0:58
[2019-11-11 Mon 21:03]-[2019-11-11 Mon 23:03] 2:00	[2019-09-23 Mon 21:26]-[2019-09-23 Mon 22:57] 1:31
[2019-11-10 Sun 21:45]-[2019-11-10 Sun 23:25] 1:40	[2019-09-22 Sun 14:52]-[2019-09-22 Sun 18:51] 3:59
[2019-10-31 Thu 09:20]-[2019-10-31 Thu 11:07] 1:47	[2019-09-21 Sat 16:50]-[2019-09-21 Sat 17:55] 1:05
[2019-10-30 Wed 10:35]-[2019-10-30 Wed 13:55] 3:20	[2019-09-21 Sat 12:31]-[2019-09-21 Sat 15:44] 3:13
[2019-10-29 Tue 22:35]-[2019-10-30 Wed 00:13] 1:38	[2019-09-20 Fri 22:05]-[2019-09-21 Sat 00:05] 2:00
[2019-10-29 Tue 09:33]-[2019-10-29 Tue 11:33] 2:00	[2019-09-20 Fri 14:38]-[2019-09-20 Fri 17:20] 2:42
[2019-10-28 Mon 21:52]-[2019-10-29 Tue 00:14] 2:22	[2019-09-20 Fri 11:42]-[2019-09-20 Fri 12:48] 1:06
[2019-10-28 Mon 18:23]-[2019-10-28 Mon 19:23] 1:00	[2019-09-19 Thu 21:14]-[2019-09-20 Fri 00:33] 3:19
[2019-10-28 Mon 09:07]-[2019-10-28 Mon 15:10] 6:03	[2019-09-19 Thu 09:15]-[2019-09-19 Thu 11:14] 1:59
[2019-10-27 Sun 20:44]-[2019-10-28 Mon 00:48] 4:04	[2019-09-18 Wed 20:55]-[2019-09-18 Wed 23:25] 2:30
[2019-10-27 Sun 14:17]-[2019-10-27 Sun 15:42] 1:25	[2019-09-17 Tue 22:05]-[2019-09-17 Tue 22:56] 0:51
[2019-10-27 Sun 12:15]-[2019-10-27 Sun 13:33] 1:18	[2019-09-14 Sat 14:20]-[2019-09-14 Sat 16:57] 2:37
[2019-10-26 Sat 13:53]-[2019-10-26 Sat 14:10] 0:17	[2019-09-12 Thu 09:31]-[2019-09-12 Thu 10:36] 1:05
[2019-10-26 Sat 10:15]-[2019-10-26 Sat 10:58] 0:43	[2019-09-11 Wed 22:40]-[2019-09-12 Thu 01:41] 3:01
[2019-10-25 Fri 15:12]-[2019-10-25 Fri 17:55] 2:43	[2019-09-11 Wed 12:11]-[2019-09-11 Wed 15:16] 3:05
[2019-10-25 Fri 09:10]-[2019-10-25 Fri 09:59] 0:49	[2019-09-11 Wed 09:19]-[2019-09-11 Wed 11:49] 2:30
[2019-10-24 Thu 22:23]-[2019-10-25 Fri 00:05] 1:42	[2019-09-10 Tue 20:60]-[2019-09-10 Tue 23:35] 2:35
[2019-10-24 Thu 18:45]-[2019-10-24 Thu 21:21] 2:36	[2019-09-10 Tue 16:30]-[2019-09-10 Tue 19:35] 3:05
[2019-10-24 Thu 09:03]-[2019-10-24 Thu 10:47] 1:44	[2019-09-10 Tue 14:30]-[2019-09-10 Tue 14:41] 0:11
[2019-10-23 Wed 21:24]-[2019-10-24 Wed 23:49] 2:25	[2019-09-10 Tue 10:27]-[2019-09-10 Tue 11:27] 1:00
[2019-10-23 Wed 09:09]-[2019-10-23 Wed 10:55] 1:46	[2019-09-09 Mon 09:29]-[2019-09-09 Mon 12:45] 3:16
[2019-10-22 Tue 22:35]-[2019-10-23 Wed 00:13] 1:33	[2019-09-08 Sun 23:07]-[2019-09-09 Mon 00:46] 1:39
[2019-10-22 Tue 19:10]-[2019-10-22 Tue 21:38] 2:28	[2019-09-08 Sun 15:10]-[2019-09-08 Sun 15:57] 5:57
[2019-10-22 Tue 09:18]-[2019-10-22 Tue 12:02] 2:44	[2019-09-06 Fri 12:05]-[2019-09-06 Fri 13:40] 1:35
[2019-10-21 Mon 23:39]-[2019-10-21 Mon 23:49] 0:10	[2019-09-04 Wed 20:01]-[2019-09-04 Wed 23:19] 3:18
[2019-10-21 Mon 17:23]-[2019-10-21 Mon 18:28] 1:05	[2019-09-04 Wed 17:01]-[2019-09-04 Wed 20:00] 2:59
[2019-10-21 Mon 09:05]-[2019-10-21 Mon 13:58] 4:53	[2019-09-04 Wed 09:12]-[2019-09-04 Wed 12:12] 3:00
[2019-10-20 Sun 23:27]-[2019-10-21 Mon 00:00] 0:33	[2019-09-03 Tue 19:40]-[2019-09-04 Wed 01:20] 5:40
[2019-10-20 Sun 19:32]-[2019-10-20 Sun 20:23] 0:51	[2019-09-03 Tue 11:12]-[2019-09-03 Tue 14:46] 3:34
[2019-10-20 Sun 12:55]-[2019-10-20 Sun 14:45] 1:50	[2019-09-03 Tue 10:00]-[2019-09-03 Tue 10:39] 0:39

```
[2019-09-02 Mon 19:55]-[2019-09-03 Tue 00:00]|4:05
[2019-09-02 Mon 09:53]-[2019-09-02 Mon 13:37]|3:44
[2019-09-01 Sun 19:10]-[2019-09-02 Mon 00:46]|5:36
[2019-08-31 Sat 11:21]-[2019-08-31 Sat 11:44]|0:23
[2019-08-30 Fri 19:21]-[2019-08-30 Fri 23:49]|4:28
[2019-08-30 Fri 15:21]-[2019-08-30 Fri 16:11]|0:50
[2019-08-29 Thu 14:10]-[2019-08-29 Thu 15:16]|1:06
[2019-08-25 Sun 14:15]-[2019-08-25 Sun 21:55]|7:40
[2019-08-22 Thu 15:01]-[2019-08-22 Thu 19:39]|4:38
[2019-08-22 Thu 09:12]-[2019-08-22 Thu 13:30]|4:18
[2019-08-21 Wed 21:15]-[2019-08-22 Thu 00:17]|3:02
[2019-08-21 Wed 12:21]-[2019-08-21 Wed 14:39]|2:18
[2019-08-20 Tue 10:57]-[2019-08-20 Tue 15:04]|4:07
[2019-08-19 Mon 09:19]-[2019-08-19 Mon 13:32]|4:13
```

8 Appendix: Full data on the exercise completion times.

This section lists the data on the minute each exercise was considered complete. (Local time.) For statistical purposes the beginning of each exercise is considered to be the completion time of the previous one. For the first exercise, the beginning time is *[2019-08-19 Mon 09:19]*.

```
Figure 1.1 Tree with the values of subcombinations
[2019-08-20 Tue 14:35]
Exercise 1.1 Interpreter result
[2019-08-20 Tue 14:23]
Exercise 1.2 Prefix form
[2019-08-20 Tue 14:25]
Exercise 1.3 Sum of squares
[2020-02-28 Fri 12:01]
Exercise 1.4 Compound expressions
[2019-08-20 Tue 14:39]
Exercise 1.5 Ben's test
[2019-08-20 Tue 14:50]
Exercise 1.6 If is a special form
[2019-08-21 Wed 14:05]
Exercise 1.7 Good enough?
[2019-08-22 Thu 12:52]
Exercise 1.8 Newton's method
[2019-08-22 Thu 17:36]
Exercise 1.9 Iterative or recursive?
[2019-08-29 Thu 15:14]
Exercise 1.10 Ackermann's function
[2019-08-25 Sun 18:31]
Exercise 1.11 Recursive vs iterative
[2019-08-25 Sun 19:25]
Exercise 1.12 Recursive Pascal's triangle
[2019-08-25 Sun 19:42]
Exercise 1.13 Fibonacci
[2019-08-25 Sun 23:04]
Exercise 1.14 ~count-change~
[2019-08-30 Fri 16:09]
Exercise 1.15 ~sine~
[2019-08-30 Fri 22:34]
Exercise 1.16 Iterative exponentiation
[2019-08-30 Fri 23:20]
Exercise 1.17 Fast multiplication
[2019-08-30 Fri 23:48]
Exercise 1.18 Iterative multiplication
[2019-08-31 Sat 11:43]
Exercise 1.19 Logarithmic Fibonacci
```

```
[2019-09-01 Sun 20:42]
Exercise 1.20 GCD applicative vs normal
[2019-09-01 Sun 23:04]
Exercise 1.21 ~smallest-divisor~
[2019-09-01 Sun 23:43]
Exercise 1.22 ~timed-prime-test~
[2019-09-02 Mon 00:44]
Exercise 1.23 ~test-divisor~
[2019-09-02 Mon 09:56]
Exercise 1.24 Fermat method
[2019-09-02 Mon 11:32]
Exercise 1.25 ~expmod~
[2019-09-02 Mon 12:46]
Exercise 1.26 ~square~ vs ~mul~
[2019-09-02 Mon 12:50]
Exercise 1.27 Carmichael numbers
[2019-09-02 Mon 20:50]
Exercise 1.28 Miller-Rabin
[2019-09-02 Mon 23:28]
Exercise 1.29 Simpson's integral
[2019-09-03 Tue 10:36]
Exercise 1.30 Iterative sum
[2019-09-03 Tue 11:19]
Exercise 1.31 Product
[2019-09-03 Tue 11:59]
Exercise 1.32 Accumulator
[2019-09-03 Tue 12:23]
Exercise 1.33 ~filtered-accumulate~
[2019-09-03 Tue 14:36]
Exercise 1.34 lambda
[2019-09-03 Tue 14:44]
Exercise 1.35 Fixed-point
[2019-09-03 Tue 21:05]
Exercise 1.36 Fixed-point-with-dampening
[2019-09-03 Tue 21:55]
Exercise 1.37 Cont-frac
[2019-09-04 Wed 11:35]
Exercise 1.38 Euler constant
[2019-09-04 Wed 11:35]
Exercise 1.39 Tan-cf
[2019-09-04 Wed 12:11]
Exercise 1.40 Newtons-method
[2019-09-04 Wed 17:06]
Exercise 1.41 Double-double
[2019-09-04 Wed 17:21]
Exercise 1.42 Compose
[2019-09-04 Wed 17:27]
Exercise 1.43 Repeated
[2019-09-04 Wed 17:54]
Exercise 1.44 Smoothing
[2019-09-04 Wed 20:17]
Exercise 1.45 Nth root
[2019-09-04 Wed 21:37]
Exercise 1.46 ~iterative-improve~
[2019-09-04 Wed 22:25]
Exercise 2.1 ~make-rat~
[2019-09-06 Fri 13:00]
Exercise 2.2 ~make-segment~
[2019-09-06 Fri 13:34]
Exercise 2.3 ~make-rectangle~
[2019-09-08 Sun 17:58]
Exercise 2.4 ~cons~ lambda
[2019-09-08 Sun 18:08]
Exercise 2.5 ~cons~ pow
```

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2019-09-08 Sun 19:07]	[2019-09-12 Thu 00:50]
Exercise 2.6 Church Numerals	Exercise 2.38 ~fold-left~
[2019-09-08 Sun 19:41]	[2019-09-12 Thu 09:45]
Exercise 2.7 ~make-interval~	Exercise 2.39 Reverse ~fold-right~ ~fold-left~
[2019-09-08 Sun 20:09]	[2019-09-12 Thu 09:52]
Exercise 2.8 ~sub-interval~	Exercise 2.40 ~unique-pairs~
[2019-09-08 Sun 23:07]	[2019-09-12 Thu 10:34]
Exercise 2.9 ~interval-width~	Exercise 2.41 ~triple-sum~
[2019-09-08 Sun 23:15]	[2019-09-14 Sat 15:15]
Exercise 2.10 Div interval better	Figure 2.8 A solution to the eight-queens puzzle
[2019-09-08 Sun 23:30]	[2019-09-14 Sat 15:17]
Exercise 2.11 Mul interval nine cases	Exercise 2.42 k-queens
[2019-09-09 Mon 00:45]	[2019-09-17 Tue 22:27]
Exercise 2.12 ~make-center-percent~	Exercise 2.43 Slow k-queens
[2019-09-09 Mon 10:11]	[2019-09-17 Tue 22:55]
Exercise 2.13 Formula for tolerance	Exercise 2.44 ~up-split~
[2019-09-09 Mon 10:16]	[2019-09-23 Mon 22:54]
Exercise 2.14 Parallel resistors	Exercise 2.45 ~split~
[2019-09-09 Mon 11:24]	[2019-09-24 Tue 01:37]
Exercise 2.15 Better intervals	Exercise 2.46 ~make-vect~
[2019-09-09 Mon 11:34]	[2019-09-20 Fri 12:48]
Exercise 2.16 Interval arithmetic	Exercise 2.47 ~make-frame~
[2019-09-09 Mon 11:37]	[2019-09-20 Fri 14:48]
Exercise 2.17 ~last-pair~	Exercise 2.48 ~make-segment~
[2019-09-10 Tue 10:48]	[2019-09-20 Fri 16:06]
Exercise 2.18 ~reverse~	Exercise 2.49 ~segments->painter~ applications
[2019-09-10 Tue 10:57]	[2019-09-20 Fri 23:10]
Exercise 2.19 Coin values	Exercise 2.50 ~flip-horiz~ ~rotate270~ ~rotate180~
[2019-09-10 Tue 11:27]	[2019-09-20 Fri 23:37]
Exercise 2.20 Dotted-tail notation	Exercise 2.51 ~below~
[2019-09-10 Tue 18:55]	[2019-09-22 Sun 18:50]
Exercise 2.21 Map square list	Exercise 2.52 Modify square-limit
[2019-09-10 Tue 19:14]	[2019-09-24 Tue 12:25]
Exercise 2.22 Wrong list order	Exercise 2.53 Quote introduction
[2019-09-10 Tue 19:24]	[2019-09-24 Tue 12:36]
Exercise 2.23 ~for-each~	Exercise 2.54 ~equal?~ implementation
[2019-09-10 Tue 19:33]	[2019-09-24 Tue 13:48]
Exercise 2.24 List plot result	Exercise 2.55 Quote quote
[2019-09-10 Tue 22:13]	[2019-09-24 Tue 13:48]
Exercise 2.25 ~caddr~	Exercise 2.56 Differentiation exponentiation
[2019-09-10 Tue 23:07]	[2019-09-24 Tue 23:14]
Exercise 2.26 ~append~ ~cons~ ~list~	Exercise 2.57 Differentiate three sum
[2019-09-10 Tue 23:23]	[2019-09-25 Wed 12:40]
Exercise 2.27 Deep reverse	Exercise 2.58 ~infix-notation~
[2019-09-11 Wed 09:47]	[2019-09-25 Wed 15:21]
Exercise 2.28 Fringe	Exercise 2.59 ~union-set~
[2019-09-11 Wed 10:24]	[2019-09-25 Wed 22:00]
Exercise 2.29 Mobile	Exercise 2.60 ~duplicate-set~
[2019-09-11 Wed 11:47]	[2019-09-25 Wed 22:17]
Exercise 2.30 ~square-tree~	Exercise 2.61 Sets as ordered lists
[2019-09-11 Wed 14:11]	[2019-09-26 Thu 21:44]
Exercise 2.31 Tree-map square tree	Exercise 2.62 ~ordered-union-set~ (ordered list)
[2019-09-11 Wed 14:38]	[2019-09-26 Thu 21:38]
Exercise 2.32 Subsets	Exercise 2.63 ~tree->list~ (binary search tree)
[2019-09-11 Wed 14:53]	[2019-09-26 Thu 23:37]
Exercise 2.33 Map append length	Exercise 2.64 Balanced tree
[2019-09-11 Wed 23:53]	[2019-09-29 Sun 17:22]
Exercise 2.34 Horner's rule	Exercise 2.65 ~tree-union-set~
[2019-09-12 Thu 00:01]	[2019-10-09 Wed 12:13]
Exercise 2.35 ~count-leaves-accumulate~	Exercise 2.66 Tree-lookup
[2019-09-12 Thu 00:17]	[2019-10-09 Wed 13:03]
Exercise 2.36 ~accumulate-n~	Exercise 2.67 Huffman decode a simple message
[2019-09-12 Thu 00:26]	[2019-10-09 Wed 20:20]
Exercise 2.37 ~matrix-*-vector~	Exercise 2.68 Huffman encode a simple message

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2019-10-09 Wed 20:53]	[2019-10-29 Tue 11:17]
Exercise 2.69 Generate Huffman tree	Exercise 3.4 Call-the-cops
[2019-10-10 Thu 11:28]	[2019-10-29 Tue 11:32]
Exercise 2.70 Generate a tree and encode a song	Exercise 3.5 Monte-Carlo
[2019-10-10 Thu 13:11]	[2019-10-30 Wed 00:12]
Exercise 2.71 Huffman tree for 5 and 10	Exercise 3.6 reset a prng
[2019-10-10 Thu 19:22]	[2019-10-30 Wed 11:42]
Exercise 2.72 Huffman order of growth	Exercise 3.7 Joint accounts
[2019-10-10 Thu 20:34]	[2019-10-30 Wed 13:07]
Exercise 2.73 Data-driven ~deriv~	Exercise 3.8 Right-to-left vs Left-to-right
[2019-10-11 Fri 11:05]	[2019-10-30 Wed 13:45]
Exercise 2.74 Insatiable Enterprises	Exercise 3.9 Environment structures
[2019-10-11 Fri 20:56]	[2019-11-20 Wed 14:28]
Exercise 2.75 ~make-from-mag-ang~ message passing	Exercise 3.10 ~let~ to create state variables
[2019-10-11 Fri 21:24]	[2019-11-25 Mon 12:52]
Exercise 2.76 Types or functions?	Exercise 3.11 Internal definitions
[2019-10-11 Fri 21:29]	[2019-11-26 Tue 12:44]
Exercise 2.77 Generic algebra magnitude	Exercise 3.12 Drawing ~append!~
[2019-10-12 Sat 16:01]	[2019-11-29 Fri 11:55]
Exercise 2.78 Ordinary numbers for Scheme	Exercise 3.13 ~make-cycle~
[2019-10-12 Sat 21:06]	[2019-11-29 Fri 12:09]
Exercise 2.79 Generic equality	Exercise 3.14 ~mystery~
[2019-10-14 Mon 15:58]	[2019-11-29 Fri 21:23]
Exercise 2.80 Generic arithmetic zero?	Exercise 3.15 ~set-to-wow!~
[2019-10-14 Mon 17:18]	[2019-12-01 Sun 19:59]
Exercise 2.81 Coercion to itself	Exercise 3.16 ~count-pairs~
[2019-10-15 Tue 11:16]	[2019-12-02 Mon 00:05]
Exercise 2.82 Three argument coercion	Exercise 3.17 Real ~count-pairs~
[2019-10-15 Tue 21:40]	[2019-12-02 Mon 00:47]
Exercise 2.83 Numeric Tower and (raise)	Exercise 3.18 Finding cycles
[2019-10-16 Wed 14:53]	[2019-12-02 Mon 01:04]
Exercise 2.84 ~raise-type~ in ~apply-generic~	Exercise 3.19 Efficient finding cycles
[2019-10-17 Thu 11:39]	[2019-12-02 Mon 23:29]
Exercise 2.85 Dropping a type	Exercise 3.20 Procedural ~set-car!~
[2019-10-20 Sun 13:47]	[2019-12-03 Tue 14:40]
Exercise 2.86 Compound complex numbers	Exercise 3.21 Queues
[2019-10-20 Sun 20:22]	[2019-12-03 Tue 15:10]
Exercise 2.87 Generalized zero?	Exercise 3.22 Procedural queue
[2019-10-21 Mon 18:25]	[2019-12-03 Tue 22:13]
Exercise 2.88 Subtraction of polynomials	Exercise 3.23 Dequeue
[2019-10-22 Tue 09:55]	[2019-12-03 Tue 23:24]
Exercise 2.89 Dense term-lists	Exercise 3.24 Tolerant tables
[2019-10-22 Tue 11:55]	[2019-12-04 Wed 18:07]
Exercise 2.90 Dense polynomials as a package	Exercise 3.25 Multilevel tables
[2019-10-22 Tue 21:31]	[2019-12-06 Fri 20:35]
Exercise 2.91 Division of polynomials	Exercise 3.26 Binary tree table
[2019-10-23 Wed 00:11]	[2019-12-06 Fri 20:53]
Exercise 2.92 Add, mul for different variables	Exercise 3.27 Memoization
[2019-10-27 Sun 13:32]	[2019-12-07 Sat 16:08]
Exercise 2.93 Rational polynomials	Exercise 3.28 Primitive or-gate
[2019-10-27 Sun 22:36]	[2019-12-08 Sun 23:43]
Exercise 2.94 GCD for polynomials	Exercise 3.29 Compound or-gate
[2019-10-28 Mon 00:47]	[2019-12-08 Sun 23:45]
Exercise 2.95 Non-integer problem	Exercise 3.30 Ripple-carry adder
[2019-10-28 Mon 11:35]	[2019-12-08 Sun 23:58]
Exercise 2.96 Integerizing factor	Exercise 3.31 Initial propagation
[2019-10-28 Mon 19:23]	[2019-12-09 Mon 00:16]
Exercise 2.97 Reduction of polynomials	Exercise 3.32 Order matters
[2019-10-29 Tue 00:12]	[2019-12-09 Mon 00:26]
Exercise 3.1 Accumulators	Exercise 3.33 Averager constraint
[2019-10-29 Tue 10:24]	[2019-12-18 Wed 11:29]
Exercise 3.2 Make-monitored	Exercise 3.34 Wrong squarer
[2019-10-29 Tue 11:03]	[2019-12-18 Wed 12:30]
Exercise 3.3 Password protection	Exercise 3.35 Correct squarer

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2019-12-18 Wed 12:47]
Exercise 3.36 Connector environment diagram
[2019-12-21 Sat 20:27]
Exercise 3.37 Expression-based constraints
[2019-12-21 Sat 21:20]
Exercise 3.38 Timing
[2019-12-21 Sat 22:48]
Exercise 3.39 Serializer
[2019-12-23 Mon 05:11]
Exercise 3.40 Three parallel multiplications
[2019-12-29 Sun 04:32]
Exercise 3.41 Better protected account
[2020-01-02 Thu 10:02]
Exercise 3.42 Saving on serializers
[2020-01-02 Thu 10:35]
Exercise 3.43 Multiple serializations
[2020-01-02 Thu 11:33]
Exercise 3.44 Transfer money
[2020-01-02 Thu 11:40]
Exercise 3.45 New plus old serializers
[2020-01-02 Thu 11:46]
Exercise 3.46 Broken test-and-set!
[2020-01-02 Thu 11:56]
Exercise 3.47 Semaphores
[2020-01-03 Fri 12:59]
Exercise 3.48 Serialized-exchange deadlock
[2020-01-03 Fri 13:30]
Exercise 3.49 When numbering does not work
[2020-01-03 Fri 13:41]
Exercise 3.50 ~stream-map~ multiple arguments
[2020-01-03 Fri 21:18]
Exercise 3.51 ~stream-show~
[2020-01-03 Fri 21:28]
Exercise 3.52 Streams with mind-boggling
[2020-01-03 Fri 22:17]
Exercise 3.53 Stream power of two
[2020-01-03 Fri 22:40]
Exercise 3.54 ~mul-streams~
[2020-01-03 Fri 22:47]
Exercise 3.55 Streams partial-sums
[2020-01-03 Fri 23:05]
Exercise 3.56 Hamming's streams-merge
[2020-01-03 Fri 23:26]
Exercise 3.57 Exponential additions fibs
[2020-01-03 Fri 23:36]
Exercise 3.58 Cryptic stream
[2020-01-03 Fri 23:50]
Exercise 3.59 Power series
[2020-01-04 Sat 09:58]
Exercise 3.60 ~mul-series~
[2020-01-04 Sat 11:07]
Exercise 3.61 ~power-series-inversion~
[2020-01-04 Sat 13:13]
Exercise 3.62 ~div-series~
[2020-01-04 Sat 13:21]
Exercise 3.63 ~sqrt-stream~
[2020-01-04 Sat 20:32]
Exercise 3.64 ~stream-limit~
[2020-01-06 Mon 09:38]
Exercise 3.65 Approximating logarithm
[2020-01-06 Mon 10:34]
Exercise 3.66 Lazy pairs
[2020-01-06 Mon 22:55]
Exercise 3.67 All possible pairs

[2020-01-06 Mon 23:09]
Exercise 3.68 ~pairs-louis~
[2020-01-06 Mon 23:26]
Exercise 3.69 ~triples~
[2020-02-17 Mon 20:10]
Exercise 3.70 ~merge-weighted~
[2020-01-07 Tue 11:58]
Exercise 3.71 Ramanujan numbers
[2020-01-07 Tue 12:49]
Exercise 3.72 Ramanujan 3-numbers
[2020-01-08 Wed 10:27]
Figure 3.32 Integral-signals
[2020-01-08 Wed 10:59]
Exercise 3.73 RC-circuit
[2020-01-08 Wed 13:09]
Exercise 3.74 Zero-crossings
[2020-01-08 Wed 16:50]
Exercise 3.75 Filtering signals
[2020-01-08 Wed 18:11]
Exercise 3.76 ~stream-smooth~
[2020-01-08 Wed 19:56]
Exercise 3.77 Streams integral
[2020-01-08 Wed 20:51]
Exercise 3.78 Second order differential equation
[2020-01-08 Wed 21:47]
Exercise 3.79 General second-order ode
[2020-01-08 Wed 21:57]
Figure 3.36
[2020-01-08 Wed 23:21]
Exercise 3.80 RLC circuit
[2020-01-08 Wed 23:40]
Exercise 3.81 Generator-in-streams
[2020-01-09 Thu 00:37]
Exercise 3.82 Streams Monte-Carlo
[2020-01-09 Thu 09:42]
Exercise 4.1 ~list-of-values~ ordered
[2020-01-09 Thu 20:11]
Exercise 4.2 Application before assignments
[2020-01-09 Thu 20:41]
Exercise 4.3 Data-directed eval
[2020-01-09 Thu 21:24]
Exercise 4.4 ~eval-and~ and ~eval-or~
[2020-01-09 Thu 22:14]
Exercise 4.5 ~cond~ with arrow
[2020-01-22 Wed 16:36]
Exercise 4.6 Implementing let
[2020-01-22 Wed 17:03]
Exercise 4.7 Implementing let*
[2020-01-22 Wed 18:09]
Exercise 4.8 Implementing named let
[2020-01-22 Wed 19:50]
Exercise 4.9 Implementing until
[2020-01-23 Thu 18:06]
Exercise 4.10 Modifying syntax
[2020-02-06 Thu 22:08]
Exercise 4.11 Environment as a list of bindings
[2020-02-11 Tue 06:58]
Exercise 4.12 Better abstractions setting value
[2020-02-11 Tue 19:40]
Exercise 4.13 Implementing ~make-unbound!~
[2020-02-12 Wed 08:52]
Exercise 4.14 Meta map versus built-in map
[2020-02-12 Wed 08:58]
Exercise 4.15 The ~halts?~ predicate

Paper - Solving SICP: An Experience Report on Solving the World's Most Famous Programming Problem Set

[2020-02-12 Wed 09:24]
Exercise 4.16 Simultaneous internal definitions
[2020-02-12 Wed 13:17]
Exercise 4.17 Environment for internal definitions
[2020-02-12 Wed 14:09]
Exercise 4.18 Alternative scanning
[2020-02-12 Wed 14:35]
Exercise 4.19 Mutual simultaneous definitions
[2020-02-12 Wed 19:52]
Exercise 4.20 ~letrec~
[2020-02-13 Thu 00:49]
Exercise 4.21 Y-combinator
[2020-02-13 Thu 01:07]
Exercise 4.22 Extending evaluator to support ~let~
[2020-02-14 Fri 19:33]
Exercise 4.23 Analysing sequences
[2020-02-14 Fri 19:40]
Exercise 4.24 Analysis time test
[2020-02-14 Fri 20:12]
Exercise 4.25 Lazy factorial
[2020-02-14 Fri 21:01]
Exercise 4.26 ~unless~ as a special form
[2020-02-15 Sat 04:32]
Exercise 4.27 Mutation in lazy interpreters
[2020-02-15 Sat 16:54]
Exercise 4.28 Eval before applying
[2020-02-15 Sat 17:01]
Exercise 4.29 Lazy eval slow without memoization
[2020-02-15 Sat 17:51]
Exercise 4.30 Lazy sequences
[2020-02-15 Sat 21:32]
Exercise 4.31 Lazy arguments with syntax extension
[2020-02-15 Sat 23:44]
Exercise 4.32 Streams versus lazy lists
[2020-02-16 Sun 11:49]
Exercise 4.33 Quoted lazy lists
[2020-02-16 Sun 14:09]
Exercise 4.34 Printing lazy lists
[2020-02-16 Sun 19:25]
Exercise 4.35 Pythagorean triples
[2020-02-17 Mon 17:25]
Exercise 4.36 Infinite Pythagorean triples
[2020-02-17 Mon 20:26]
Exercise 4.37 Another method for triples
[2020-02-17 Mon 21:17]
Exercise 4.38 Logical puzzle - Not same floor
[2020-02-17 Mon 21:56]
Exercise 4.39 Order of restrictions
[2020-02-17 Mon 22:01]
Exercise 4.40 People to floor assignment
[2020-02-17 Mon 22:29]
Exercise 4.41 Ordinary Scheme floor problem
[2020-02-18 Tue 00:12]
Exercise 4.42 The liars puzzle
[2020-02-18 Tue 12:16]
Exercise 4.43 Problematical Recreations
[2020-02-18 Tue 13:31]
Exercise 4.44 Nondeterministic eight queens
[2020-02-18 Tue 15:17]
Exercise 4.45 Five parses
[2020-02-18 Tue 19:45]
Exercise 4.46 Order of parsing
[2020-02-18 Tue 19:55]
Exercise 4.47 Parse verb phrase by Louis

[2020-02-18 Tue 20:13]
Exercise 4.48 Extending the grammar
[2020-02-18 Tue 21:06]
Exercise 4.49 Alyssa's generator
[2020-02-18 Tue 21:51]
Exercise 4.50 The ~ramb~ operator
[2020-02-17 Mon 14:56]
Exercise 4.51 Implementing ~permanent-set!~
[2020-02-18 Tue 22:34]
Exercise 4.52 ~if-fail~
[2020-02-19 Wed 00:05]
Exercise 4.53 Test evaluation
[2020-02-19 Wed 00:12]
Exercise 4.54 ~analyze-require~
[2020-02-19 Wed 11:26]
Exercise 4.55 Simple queries
[2020-02-19 Wed 17:38]
Exercise 4.56 Compound queries
[2020-02-19 Wed 18:04]
Exercise 4.57 Custom rules
[2020-02-19 Wed 21:36]
Exercise 4.58 Big shot
[2020-02-19 Wed 22:12]
Exercise 4.59 Meetings
[2020-02-19 Wed 22:57]
Exercise 4.60 Pairs live near
[2020-02-19 Wed 23:20]
Exercise 4.61 Next-to relation
[2020-02-19 Wed 23:31]
Exercise 4.62 Last-pair
[2020-02-20 Thu 00:19]
Exercise 4.63 Genesis
[2020-02-20 Thu 10:28]
Figure 4.6 How the system works
[2020-02-20 Thu 10:59]
Exercise 4.64 Broken outranked-by
[2020-02-20 Thu 12:33]
Exercise 4.65 Second-degree subordinates
[2020-02-20 Thu 12:50]
Exercise 4.66 Ben's accumulation
[2020-02-20 Thu 13:08]
Exercise 4.67 Loop detector
[2020-02-20 Thu 23:20]
Exercise 4.68 Reverse rule
[2020-02-21 Fri 15:48]
Exercise 4.69 Great grandchildren
[2020-02-21 Fri 17:43]
Exercise 4.70 Cons-stream delays second argument
[2020-02-20 Thu 17:08]
Exercise 4.71 Louis' simple queries
[2020-02-21 Fri 20:56]
Exercise 4.72 ~interleave-stream~
[2020-02-20 Thu 17:11]
Exercise 4.73 ~flatten-stream~ delays
[2020-02-20 Thu 17:19]
Exercise 4.74 Alyssa's streams
[2020-02-21 Fri 22:00]
Exercise 4.75 ~unique~ special form
[2020-02-21 Fri 23:19]
Exercise 4.76 Improving ~and~
[2020-02-22 Sat 18:27]
Exercise 4.77 Lazy queries
[2020-03-14 Sat 15:42]
Exercise 4.78 Non-deterministic queries

[2020-03-15 Sun 12:40]
Exercise 4.79 Prolog environments
[2020-05-10 Sun 17:59]
Figure 5.1 Data paths for a Register Machine
[2020-02-23 Sun 13:18]
Figure 5.2 Controller for a GCD Machine
[2020-02-22 Sat 22:27]
Exercise 5.1 Register machine plot
[2020-02-22 Sat 22:56]
Exercise 5.2 Register machine Exercise 5.1
[2020-02-23 Sun 13:26]
Exercise 5.3 Machine for `sqrt`, Newton Method
[2020-02-23 Sun 20:47]
Exercise 5.4 Recursive register machines
[2020-02-24 Mon 20:49]
Exercise 5.5 Manual factorial and Fibonacci
[2020-02-24 Mon 23:27]
Exercise 5.6 Fibonacci machine extra instructions
[2020-02-24 Mon 23:43]
Exercise 5.7 Test the 5.4 machine on a simulator
[2020-02-25 Tue 10:42]
Exercise 5.8 Ambiguous labels
[2020-02-25 Tue 21:58]
Exercise 5.9 Prohibit (op)s on labels
[2020-02-25 Tue 22:23]
Exercise 5.10 Changing syntax
[2020-02-25 Tue 22:39]
Exercise 5.11 Save and restore
[2020-02-26 Wed 13:30]
Exercise 5.12 Data paths from controller
[2020-02-26 Wed 23:40]
Exercise 5.13 Registers from controller
[2020-02-27 Thu 10:57]
Exercise 5.14 Profiling
[2020-02-28 Fri 20:21]
Exercise 5.15 Instruction counting
[2020-02-28 Fri 21:36]
Exercise 5.16 Tracing execution
[2020-02-28 Fri 22:59]
Exercise 5.17 Printing labels
[2020-02-29 Sat 17:43]
Exercise 5.18 Register tracing
[2020-02-29 Sat 14:07]
Exercise 5.19 Breakpoints
[2020-02-29 Sat 17:42]
Exercise 5.20 Drawing a list `(#1=(1 . 2) #1)`
[2020-02-29 Sat 22:15]
Exercise 5.21 Register machines list operations
[2020-03-01 Sun 13:03]
Exercise 5.22 `append` and `append!` as machines
[2020-03-01 Sun 14:11]
Exercise 5.23 EC-evaluator with `let` and `cond`
[2020-03-02 Mon 10:52]
Exercise 5.24 Making `cond` a primitive
[2020-03-02 Mon 14:42]
Exercise 5.25 Normal-order (lazy) evaluation
[2020-03-03 Tue 14:57]
Exercise 5.26 Tail recursion with `factorial`
[2020-03-03 Tue 19:38]
Exercise 5.27 Stack depth for recursive factorial
[2020-03-03 Tue 19:49]
Exercise 5.28 Interpreters without tail recursion
[2020-03-03 Tue 20:29]
Exercise 5.29 Stack in tree-recursive Fibonacci

[2020-03-03 Tue 20:50]
Exercise 5.30 Errors
[2020-03-04 Wed 11:35]
Exercise 5.31 a `preserving` mechanism
[2020-03-04 Wed 21:36]
Exercise 5.32 Symbol-lookup optimization
[2020-03-04 Wed 22:51]
Exercise 5.33 Compiling `factorial-alt`
[2020-03-05 Thu 16:55]
Exercise 5.34 Compiling iterative factorial
[2020-03-05 Thu 20:58]
Exercise 5.35 Decompilation
[2020-03-05 Thu 21:30]
Exercise 5.36 Order of evaluation
[2020-03-06 Fri 17:47]
Exercise 5.37 `preserving`
[2020-03-06 Fri 21:01]
Exercise 5.38 Open code primitives
[2020-03-07 Sat 18:57]
Exercise 5.39 `lexical-address-lookup`
[2020-03-07 Sat 20:41]
Exercise 5.40 Compile-time environment
[2020-03-08 Sun 15:02]
Exercise 5.41 `find-variable`
[2020-03-07 Sat 19:37]
Exercise 5.42 Compile variable and assignment
[2020-03-08 Sun 12:59]
Exercise 5.43 Scanning out defines
[2020-03-08 Sun 21:00]
Exercise 5.44 Open code compile-time environment
[2020-03-08 Sun 21:29]
Exercise 5.45 Stack usage for `factorial`
[2020-03-09 Mon 10:09]
Exercise 5.46 Stack usage for `fibonacci`
[2020-03-09 Mon 10:34]
Exercise 5.47 Calling interpreted procedures
[2020-03-09 Mon 11:45]
Exercise 5.48 `compile-and-run`
[2020-03-10 Tue 12:14]
Exercise 5.49 `read-compile-execute-print` loop
[2020-03-10 Tue 12:36]
Exercise 5.50 Compiling the metacircular evaluator
[2020-03-14 Sat 15:52]
Exercise 5.51 EC-evaluator in low-level language
[2020-04-13 Mon 11:45]
Exercise 5.52 Making a compiler for Scheme
[2020-05-06 Wed 11:09]

9 Appendix: Emacs Lisp code for data analysis

This section included the Emacs Lisp code used to analyse the data above. The code is directly executable in the org-mode version of the report. Interested readers reading the PDF version are advised to consult the org-mode version.

```
(require 'org-element)
(cl-labels (

; lexical-defun
(decorate-orgtable (tbl)
  (seq-concatenate
```

```

    'string
    "("
    "| Exercise | Days | Sessions | Minutes |"
    (char-to-string ?\n)
    "|- + - + - + - |"
    (format-orgtable tbl)
    ")")
  )

  ; lexical-defun
  (format-orgtable (list-of-lists)
  (apply
  #'seq-concatenate
  (cons
  'string
  (seq-map
  (lambda (x) (format-table-line x))
  list-of-lists)))
  )

  ; lexical-defun
  (format-table-line (line)
  (seq-concatenate 'string
  (char-to-string ?\n)
  "|")
  (substring
  (car line)
  0
  (min 60 (seq-length (car line))))
  "|")
  (format "%3.3f"(caddr line))
  "|")
  (format "%3d" (nth 4 line))
  "|")
  (format "%3.3f" (nth 6 line))
  "|")
  )

  ;; lexical-defun
  (get-study-sessions-data ()
  (save-excursion
  (org-babel-goto-named-src-block
  "study-sessions-data")
  (seq-map (lambda (x)
  (list
  (org-time-string-to-seconds
  (substring-no-properties
  x
  3
  23))
  (org-time-string-to-seconds
  (substring-no-properties
  x
  26
  46))
  ))
  (seq-subseq
  (split-string
  (org-element-property
  :value
  (org-element-at-point))
  "\n")
  0
  -1)))

```

```

  )

  ;; lexical-defun
  (get-task-sequence-data ()
  (save-excursion
  (org-babel-goto-named-src-block
  "completion-times-data")
  (let ((exercise-index 0))
  (seq-mapn
  (lambda (nam dat)
  (setq exercise-index
  (+ 1 exercise-index))
  (list nam dat exercise-index))
  (apply #'seq-concatenate
  (cons 'list
  (seq-map-indexed
  (lambda (x idx)
  (if (= 0 (mod idx 2))
  (list x)
  nil))
  (seq-subseq
  (split-string
  (org-element-property
  :value (org-element-at-point))
  "\n")
  0
  -1))))
  (apply #'seq-concatenate
  (cons 'list
  (seq-map-indexed
  (lambda (x idx)
  (if (= 1 (mod idx 2))
  ;(print x)
  (list x)
  nil))
  (seq-subseq
  (split-string
  (org-element-property
  :value (org-element-at-point))
  "\n")
  0
  -1)))))))
  )

  ;; lexical-defun
  (sort-task-seq (task-seq)
  (seq-sort
  (lambda (x y)
  (if (org-time< (cadr x)
  (cadr y))
  t
  nil))
  task-seq)
  )

  ;; lexical-defun
  (find-out-of-order-tasks (task-seq)
  (seq-reduce
  (lambda (acc next-elem)
  (if (org-time<
  (cadr next-elem) (cadr acc))
  (list (+ 1 (car acc))
  (cadr next-elem)
  (cons (caddr acc) (caddr acc))

```

```

        next-elem)
      (list (car acc)
            (cadr next-elem)
            (caddr acc) next-elem)))
  task-seq
  (list 0 "2019-08-19 Mon 09:19" (list) (list)))
)

;; lexical-defun
(find-spanning-sessions-and-duration
 (prev-time-stamp
  next-time-stamp
  study-sessions)
 (seq-reduce
  (lambda (acc next-session)
    (let ((session-start (car next-session))
          (session-end (cadr next-session)))
      (cond ((<= session-end prev-time-stamp)
             acc)
            ((<= next-time-stamp session-start)
             acc)
            (t (list (+ (car acc) 1)
                     (+ (cadr acc)
                        (cond ((and (<= prev-time-stamp session-start)
                                   (<= session-end next-time-stamp))
                              (- session-end prev-time-stamp))
                            ((and (<= session-start prev-time-stamp)
                                   (<= prev-time-stamp session-end)
                                   (<= session-end next-time-stamp))
                              (- session-end prev-time-stamp))
                            ((and (<= prev-time-stamp session-start)
                                   (<= session-start next-time-stamp)
                                   (<= next-time-stamp session-end))
                              (- next-time-stamp session-start))
                            ((and (<= session-start prev-time-stamp)
                                   (<= next-time-stamp session-end))
                              (- next-time-stamp prev-time-stamp))
                              (t 0))))))))))
  study-sessions
  (list 0 0)))

;; lexical-defun
(summarize-list (sorted-task-seq study-sessions)
 (cadr (seq-reduce
  (lambda (acc next-elem)
    (let ((prev-time-stamp (car acc))
          (retval (cadr acc))
          (next-time-stamp
           (org-time-string-to-seconds
            (cadr next-elem)))
          (exercise-name (car next-elem))
          (exercise-index (caddr next-elem)))
      (let ((spans-sessions
             (find-spanning-sessions-and-duration
              prev-time-stamp
              next-time-stamp
              study-sessions)))
          (list next-time-stamp
                (cons
                 (list exercise-name
                       :spent-time-calendar-days
                       (/ (-
                          next-time-stamp
                          prev-time-stamp)

```

```

        (* 60 60 24))
      :spans-sessions
      (if (not (eq 0 (car spans-sessions)))
          (car spans-sessions)
          (error
           "Fix time: %s, spans-sessions=%s"
           next-elem
           spans-sessions))
      :spent-time-net-minutes
      (/ (cadr spans-sessions) 60)
      :original-index
      exercise-index
      retval))))))
  sorted-task-seq
  (list
   (org-time-string-to-seconds
    "2019-08-19 Mon 09:19")
   ()))
)

(r-h 1)
(seq-reverse (seq-subseq 1 0)))

;; lexical-defun
(make-logarithmic-histogram (astrotime-list)
 (let* ((numbins
         (ceiling
          (log (+ 1.0
                 (seq-reduce
                  #'max
                  (seq-map
                   (lambda (x) (nth 6 x))
                   (r-h astrotime-list))
                  0))
          2))))
        (aset acc
              nbin
              (+ 1 (aref acc nbin)))
        acc)
  (r-h astrotime-list)
  (make-vector numbins 0)))

;; lexical-defun
(make-linear-histogram (astrotime-list)
 (let* ((numbins 32)
        (binsize
         (ceiling
          (/ (seq-reduce
              #'max
              (seq-map
               (lambda (x) (nth 6 x))
               (r-h astrotime-list))
              0)
            numbins ))))
        (seq-reduce
         (lambda (acc elem)

```

```

    (let* ((hardness (nth 6 elem))
          (nbin (floor (/ hardness binsize))))
      (aset acc
            nbin
            (+ 1 (aref acc nbin)))
      acc))
  (r-h astrotime-list)
  (make-vector numbins 0))
)

;; lexical-defun
(sort-by-hardness (astrotime-list)
;; 6 is the hardness index
(seq-sort (lambda (x y)
           (let* ((hardness-x (nth 6 x))
                  (hardness-y (nth 6 y)))
             (if (< hardness-x hardness-y)
                 t
                 nil)))
          astrotime-list)
)

;; lexical-defun
(sort-by-nsessions (astrotime-list)
;; 4 is the nsessions index
(seq-sort (lambda (x y)
           (let* ((nses-x (nth 4 x))
                  (nses-y (nth 4 y)))
             (if (< nses-x nses-y)
                 t
                 nil)))
          astrotime-list)
)

;; lexical-defun
(sort-by-original-index (astrotime-list)
;; 8 is the original index
(seq-sort (lambda (x y)
           (let* ((oidx-x (nth 8 x))
                  (oidx-y (nth 8 y)))
             (if (< oidx-x oidx-y)
                 t
                 nil)))
          astrotime-list)
)

) ;; end cl-labels defuns

(let* (
  ;; lexical-define
  (study-sessions (get-study-sessions-data))

  ;; lexical-define
  (task-seq (get-task-sequence-data))

  ;; lexical-define
  (sorted-task-seq (sort-task-seq task-seq))

  ;; lexical-define
  (out-of-order-tasks
   (find-out-of-order-tasks task-seq))

  ;; lexical-define

```

```

(astrotime-list
 (summarize-list
  sorted-task-seq
  study-sessions))

;; lexical-define
(problems-sorted-by-completion-time
 (seq-reverse astrotime-list))

;; lexical-define
(logarithmic-histogram
 (make-logarithmic-histogram astrotime-list))

;; lexical-define
(linear-histogram
 (make-linear-histogram astrotime-list))

;; lexical-define
(problems-sorted-by-hardness
 (sort-by-hardness astrotime-list))

;; lexical-define
(problems-sorted-by-nsessions
 (sort-by-nsessions astrotime-list))

;; lexical-define
(problems-sorted-by-original-index
 (sort-by-original-index astrotime-list))
)

(princ (char-to-string ?\()))
(pp "Amount of the out-of-order-problems: ")
(princ (char-to-string ?\()))
(pp (number-to-string
     (car out-of-order-tasks)))
(princ (char-to-string ?\n))

(pp "Out-of-order problems :")
(princ (char-to-string ?\n))
(pp (caddr out-of-order-tasks))
(princ (char-to-string ?\n))

(pp "Task summary (completion time):")
(princ (char-to-string ?\n))
(princ
 (decorate-orgtable
  (seq-subseq
   problems-sorted-by-completion-time
   0 3)))
(princ (char-to-string ?\n))

(princ (char-to-string ?\n))
(pp "Task summary (original-index):")
(princ (char-to-string ?\n))
;; (pp (seq-subseq
;;      problems-sorted-by-original-index 0 2))
(princ
 (decorate-orgtable
  (seq-subseq
   problems-sorted-by-original-index
   0 3)))

```

```
(princ (char-to-string ?\n))

;; Hardest 10 problems
(princ (char-to-string ?\n))
(pp "Hardest 10 problems (raw):")
(princ (char-to-string ?\n))
;; (pp (seq-subseq
;; problems-sorted-by-original-index 0 2))
(princ
 (decorate-orgtable
  (seq-subseq
   problems-sorted-by-hardness
   -10)))
(princ (char-to-string ?\n))

;; Hardest 10 problems
(princ (char-to-string ?\n))
(pp "Hardest 10 problems (sessions):")
(princ (char-to-string ?\n))
;; (pp (seq-subseq
;; problems-sorted-by-original-index 0 2))
(princ
 (decorate-orgtable
  (seq-subseq
   problems-sorted-by-nsessions
   -10)))
(princ (char-to-string ?\n))

(princ (char-to-string ?\n))
(pp "Logarithmic histogram:")
;; Make a logarithmic histogram
(princ (char-to-string ?\n))

(pp logarithmic-histogram)
(princ (char-to-string ?\n))

(pp "Linear histogram:")
(princ (char-to-string ?\n))
;; Make a linear histogram
(pp linear-histogram)
(princ (char-to-string ?\n))

(pp "Median difficulty:")
(princ (char-to-string ?\n))

(pp
 (nth
  (floor (/ (seq-length
            problems-sorted-by-hardness)
           2))
  problems-sorted-by-hardness))

(pp "Median n-sessions:")
(princ (char-to-string ?\n))

(pp
 (nth
  (floor (/ (seq-length
            problems-sorted-by-nsessions)
           2))
  problems-sorted-by-nsessions))
(princ (char-to-string ?\)))
))
```

10 Table of Contents

CONTENTS	
Abstract	1
1 Introduction	2
2 Solution approach	2
2.1 Completeness	2
2.1.1 Just solve all of the exercises	2
2.1.2 A faithful imitation of the university experience	3
2.1.3 Meta-cognitive exercises	3
2.1.4 Figures to re-typeset	4
2.2 Behaviour modelling, reenactment and the choice of tools	5
2.2.1 The author's background	5
2.2.2 The tools	5
2.3 Org-mode as a universal medium for reproducible research	7
2.4 Different problem types	8
3 Time analysis, performance profiling and graphs	8
3.1 Workflow details and profiling	8
3.2 Out-of-order problems and other measures	10
3.3 Ten hardest problems by raw time	10
3.4 Minutes spent per problem	11
3.5 Days spent per problem	11
3.6 Study sessions per problem	12
3.7 Difficulty histogram (linear)	12
3.8 Difficulty histogram (logarithmic)	12
4 Conclusion and Further Work	13
4.1 Conclusion	13
4.2 Further work	13
4.3 Informal review	14
4.4 Informal recommendations	14
5 Materials	15
5.1 Books	15
5.2 Software	15
5.3 Papers	15
6 Appendix: Analysed data on problem difficulty	17
6.1 Analysed time consumption	17
6.2 Time consumption histogram linear	25
6.3 Time consumption histogram logarithmic	25
7 Appendix: Full data on the study sessions.	25
8 Appendix: Full data on the exercise completion times.	28
9 Appendix: Emacs Lisp code for data analysis	33
10 Table of Contents	37
Contents	37

Running Scheme On Bare Metal (Experience Report)

Samuel Yvon
DIRO
Université de Montréal
Canada
samuel.yvon@umontreal.ca

Marc Feeley
DIRO
Université de Montréal
Canada
feeley@iro.umontreal.ca

Abstract

Programming language implementations have features such as threads, memory management, type safety, and REPLs that duplicate some of the work done by the underlying operating system. The objective of hosting a language on bare metal is to unify the language implementation and operating system to have a lean system with no redundancy for running applications.

In this paper we report on our experience implementing this idea for the Scheme language, specifically an extension of the Gambit Scheme system on the commodity x86-32 platform. Our system, Mimoso, is written mostly in Scheme and supports the execution of Scheme applications with access to some essential devices (keyboard, terminal, bitmap screen, serial links, hard disks).

We report on the main problems we have encountered and how the use of some of the advanced features of the Gambit Scheme system eased the development. One of the main difficulties was the handling of hardware interrupts in a garbage collected language. We also discuss the benefits of using a REPL for experimentation and continuation passing style in the device drivers.

CCS Concepts • Software and its engineering → Operating systems;

Keywords Operating Systems, Scheme, Dynamic Languages

1 Introduction

An operating system (OS) offers an abstraction over computer hardware that makes software portable over an assortment of hardware platforms and configurations. This abstraction provides applications with a common interface to safely use the hardware resources, such as managing the use of memory (RAM and disk) and multiplexing the execution of concurrent processes on the processor(s). These services are implemented by the OS through a tight coupling with the hardware using drivers specific to each available device. It is essential to interact at a low level of abstraction with the components of the computer, often requiring processor specific instructions that are centered around communicating with peripherals.

A programming language offers abstractions similar to the ones provided by an OS: threads, memory management, filesystem operations, type safety, etc. Typically the programming language's abstractions are implemented over the ones provided by the OS. This layering is suboptimal because some work is duplicated and the hardware-enforced protections and privilege levels implemented by the OS, which have a run time cost, are in principle not needed when using a safe language. A language implemented *on bare metal* removes the need for a host OS by integrating its functionality in the runtime system. In a sense, the language implementation becomes the OS and *applications* are simply modules run in separate threads.

Our goal is to explore this design for Scheme on commodity hardware. We are interested in determining the feasibility of a bare metal implementation mostly written in Scheme and whether Scheme's high-level nature, dynamic type checking and garbage collection are hindrances. We think that Scheme also offers benefits for a better and faster development process. In particular, we conjecture that the read-eval-print loop (REPL) facilitates the kind of tinkering encountered when developing device drivers that improves upon the traditional edit-compile-run cycle.

The idea of a bare metal language implementation is not new in the realm of Lisp-like languages. Early Lisp machines such as MIT's CONS and CADR machines relied on custom hardware [21] for performance. Our work is closer to more recent efforts that use Common Lisp on commodity hardware, such as Mezzano [8] and ChrysaLisp [19].

In order to investigate our research questions, we have implemented Mimoso, an extension to the standard Gambit Scheme system that runs on commodity x86-32 PCs with no separate OS. Our use of Gambit is motivated by the fact that it has good performance and it provides a rich set of features, including preemptive multithreading, yet its source code has very few OS dependencies. Moreover, Gambit's fully featured REPL allows investigating its usefulness in our context.

While some parts of the implementation are in C and assembly for reasons detailed in Section 2, most of the device-interfacing code is written in Scheme. This allows us to use the memory management, threading support and libraries of the Gambit runtime system, which are for the most part also written in Scheme.

This paper is organized as follows. Section 2 provides an overview of Mimosa’s general architecture. Section 3 provides some basic insight into the Gambit runtime system that helps one understand some of the design choices. Section 4 reports on the development process itself. Section 5 compares implementation in C versus Scheme, explaining the pros and cons of the experience. Section 6 presents related work. Section 7 adds some information about the codebase along with links to the source code. Our plans to improve Mimosa are presented in Section 8.

2 Architecture

2.1 General Structure

Mimosa’s code base has evolved from an x86-32 OS, MINOS, originally developed in C++ and x86 assembly language for an undergraduate OS course. Mimosa is divided in two layers of abstraction over the machine. The bottom layer provides basic C functions roughly corresponding to a subset of the libc C library and the top layer provides the Scheme procedures of the Gambit runtime library, some of which are implemented using the bottom layer. For brevity we call these layers the *C kernel* and the *Scheme kernel*.

The main device drivers are implemented in Scheme and reside in the Scheme kernel. Figure 1 depicts this construction. Interrupts from the hardware devices are received by the C kernel, which propagates them with associated data to the drivers in the Scheme kernel. An abstraction over machine I/O instructions allows the drivers to communicate directly with the devices, frequently bypassing the lower C kernel entirely.

Note that for historical and practical reasons, the C kernel uses a few features of the C++ language but is mostly compatible with the C language, aside from the use of function overloading. Also note that currently the Scheme kernel is mostly Scheme code that has been compiled to C code by the Gambit compiler.

While the current version of Mimosa only works on the 86-32 bits architecture, support for 64 bits would not be impossible to add. We expect a transition to 64 bit to require some additional configuration of the chipset upon boot. Of course, the Gambit executable would need to be compiled for the new architecture, something already supported.

2.2 Boot Process

When the PC boots up, the bootloader finds and executes a second stage loader, which loads from disk the C kernel using the BIOS and executes it. Once it is initialized and the x86 processor has been switched to 32 bit protected mode, it loads the Scheme kernel in several parts. First, the Gambit runtime system (RTS) is loaded into memory from disk and is executed. The RTS then loads various Scheme libraries and proceeds to load the Scheme kernel’s initialization routine, which loads the Scheme drivers. Each driver performs the

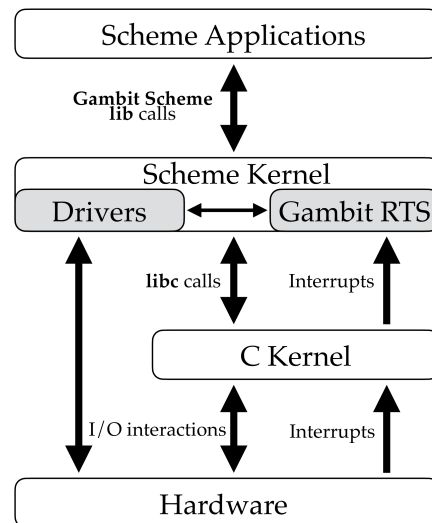


Figure 1. Mimosa’s kernel structure

required setup of its hardware device. In the case of the hard disk (ATA) driver, a call is made to the C kernel to disable its hard disk driver so that all future disk accesses will be managed by the Scheme driver.

2.3 Special Considerations

By design, Mimosa does not implement virtual memory since applications are running within the Gambit environment. The C kernel thus manages the free RAM as one large heap. To enable Scheme code to have the same access to the hardware as a normal kernel, all code runs with the same privilege level and so the entire system can be considered to operate in *kernel level*. In the future, when Mimosa is fully bootstrapped, the Gambit RTS will be statically linked with the C kernel, but currently, to improve modularity, it is loaded from disk by the C kernel. The implementation of a dynamic program loader is avoided by compiling the RTS to a fixed address, and using a table of function pointers at a fixed address to call the C kernel functions from the RTS.

2.4 Usage

Once Mimosa is booted, the operator at the keyboard is presented with a REPL giving interactive access to the primordial thread to evaluate commands, load files and modules from disk, edit files and show the status of the running threads. Figure 3 is a capture of the PC’s screen with a sample interaction. Another thread runs a secondary REPL on the PC’s serial link, allowing interactive access from an external development environment. This is particularly useful for experimentation when Mimosa is run in a virtual machine such as QEMU [5] that is configured to feed the serial link from a network connection, allowing the serial link REPL to

be reached easily using telnet, nc, shell scripts and similar utilities.

2.5 Modules and Drivers

Mimosa includes modules and drivers for essential features:

1. **Video:** accesses the screen bitmap to draw rectangles and text, display the mouse cursor, and perform *BIT-BLIT* using direct memory-mapped I/O in 16 color 640x480 mode.
2. **Term:** emulates a terminal that manages a rectangular window on the screen and responds to some VT100 escape sequences, providing a familiar interface to programs, in particular the primary REPL's line editor.
3. **Keyboard:** receives data packets from the PS/2 interface, converts them to a stream of key codes that feeds the terminal's input.
4. **UART:** manages the standard 4 serial links and connects one to the secondary REPL's input and output.
5. **Threading:** both kernels implement preemptive multithreading and synchronisation mechanisms; the Scheme kernel uses first-class continuations and the C kernel has a traditional implementation with separate stacks and assembly code context switching (kept until Mimosa's full bootstrap).
6. **IDE:** provides an implementation for the ATA disk driver in programmed-input-output mode (PIO).
7. **Disk Caching:** caches disk sectors accessed by the IDE driver to speed up accesses and synchronize writes to the disks.
8. **FAT32:** both kernels implement the FAT32 file system operations with disk caching (the C kernel driver is needed to load the Scheme kernel from disk until Mimosa's full bootstrap).
9. **RTC:** provides access to the system's real-time clock, providing access to the current date and time as well as regular interrupts to track time and implement preemptive multithreading.

3 Gambit Runtime System

Mimosa's Scheme kernel relies heavily on features of Gambit's standard RTS. In this section, we detail some key features to provide context for subsequent sections.

3.1 Machine Code Generation

The source code of the Gambit RTS has some parts, such as the garbage collector, hand written in C, but the largest part is Scheme code compiled to C by the Gambit compiler. Mimosa is built "ahead-of-time" by using the Gambit compiler to compile its Scheme modules to C and using a C compiler to finish the compilation process and link with the other C files. The bulk of Mimosa and Gambit's RTS are therefore translated to machine code via a C compiler. However, for maximum flexibility, the Scheme kernel contains the Gambit

compiler and related modules, including the experimental backend generating x86 machine code. The code generation module can be used independently by a Scheme program running in Mimosa to generate x86 machine code programmatically, more or less at the same level of abstraction as using an assembler on a desktop OS.

3.2 Garbage Collection

The RTS uses garbage collection (GC) to manage the memory allocated to Scheme objects. The Scheme heap is itself allocated incrementally in 512 KB sections obtained by calling the `malloc` function of the C kernel's `libc`. This allows the RTS to grow and shrink the heap at the end of a GC cycle proportionately to the amount of live objects. The RTS is configured so that the free space in the Scheme heap accounts for 50% of the heap. A blocking GC cycle is triggered when free space has been exhausted. Each GC cycle causes a pause that is proportional to the space occupied by live objects, which is about 9 MB when Mimosa is first booted (a pause of 70 milliseconds on modern hardware).

Both stop-and-copy (S&C) and mark-and-sweep (M&S) algorithms are used. The S&C algorithm, which moves objects, manages small objects that typically account for most of the Scheme heap. The M&S algorithm, which does not move objects, manages large objects and objects that are accessed by both Scheme and C code. These non-movable objects contain a reference count that is used by C code to mark the objects as reachable from the "C world" (so that they are not reclaimed erroneously when the "Scheme world" no longer references them).

3.3 Foreign Function Interface (FFI)

The `c-lambda` special form defines a Scheme procedure that when called will execute a C function (or literal C code) after converting the Scheme parameters to their C representation, and on return converts the function result from its C representation to Scheme. Similarly the `c-define` special form does the reverse, that is it defines a C procedure that when called from C will execute Scheme code.

The `c-lambda` form is useful to drop down to a lower-level of abstraction to perform operations that are easier to express in C. Moreover, the C code can contain an "asm" statement to execute specific machine instructions that aren't available in C, for example executing x86 `inb` and `outb` to read and write bytes to the processor's I/O ports, which is necessary to control some peripherals, and the x86 `cli` instruction that causes the hardware interrupts to be postponed until an `sti` instruction is executed.

3.4 Interrupts and Polling

Mimosa's C kernel communicates with the Scheme kernel mainly by notifying it when hardware events occur.

Gambit's standard RTS has a mechanism to set Scheme handlers for various asynchronous events, also known as

interrupts. Two standard interrupts are *user interrupts* which correspond to the user typing CTRL-C to stop the currently running thread and start a REPL, and *heartbeat interrupts* generated by a periodic timer at a high frequency to implement preemptive context switching of the Scheme threads (200 Hz when using Mimosa).

New types of interrupts can be added easily to the RTS. There is a predefined *high level interrupt* that is unused by the standard RTS and is used by Mimosa as explained in Section 4.2. The occurrence of an interrupt is signaled from C code through the `__raise_interrupt` function of the RTS which takes the interrupt number as its sole argument.

Given that interrupt handlers are Scheme code, it is necessary that Scheme code be interrupted at safe points, for example not inside a sequence of instructions that performs an allocation which could leave the RTS in an intermediate, inconsistent state. However, the call to `__raise_interrupt` may happen asynchronously in a low-level interrupt handler, so the handling of the interrupt in Scheme must be delayed until a safe point is reached. Of course, it is desirable for interrupts to be handled with minimal delay. To ensure that interrupts occur in a timely manner at safe points, `__raise_interrupt` only registers that the interrupt occurred and the Gambit compiler inserts throughout the generated C code some run time checks that detect the presence of a registered interrupt (*polling points* [14]). The appropriate Scheme interrupt handler(s) are called at these polling points in the currently running thread through a Scheme non-tail call. Stack area overflows are also detected at polling points and handled by extending the stack area, possibly triggering a GC, and execution is continued.

There are various ways to postpone interrupt processing. Hardware level interrupt processing can be postponed via the `x86 cli` and `sti` instructions, by setting a bit in the Programmable Interrupt Controller's interrupt mask register, and for some peripherals by changing a bit in any of its control registers that disable the signaling of interrupts. A global mask that controls whether polling points in any thread detect the presence of registered interrupts can also be modified from Scheme code by calling `##disable-interrupts!` and `##enable-interrupts!`. There is also a per-thread interrupt mask parameter object that is useful to prevent the polling points of a specific thread from processing a specific set of interrupts (possibly all). Finally, the generation of polling points by the Gambit compiler in a specific block of Scheme code can be disabled with the `(declare (not interrupts-enabled))` declaration which is useful to implement short critical sections at no cost.

3.5 Threading and I/O

The Gambit RTS implements SRFI 21 [16], a real-time extension of SRFI 18 [15] (Multithreading support) that adds the concepts of thread priority and priority inheritance that prevents priority inversion problems. The thread system is

completely implemented in Scheme using first-class continuations to capture the state of suspended threads. Data types for mutexes and condition variables are available for inter-thread synchronisation. All the synchronisation operations have an optional timeout parameter that limits the waiting time when that is appropriate.

The implementation of Scheme ports and the Scheme thread scheduler are fully integrated so an I/O request that cannot be fulfilled immediately will cause the current Scheme thread to be suspended until the operation can succeed, or a timeout is reached if specified.

Of particular relevance to Mimosa are vector-ports. A vector-port is a bidirectional port created with the procedure `open-vector` or `open-vector-pipe` that accumulates objects written to it using the `write` procedure, and delivers these objects in the same order when the `read` procedure is called. When the vector-port is empty the thread calling `read` will be suspended until another thread calls `write` to add an object. A limit on the amount of buffering can also be set so that when the limit is reached the writing thread will block until another thread frees space by reading data. This FIFO data structure is useful for loosely coupling producer and consumer threads that produce and consume data at variable rates.

4 Development Process

The implementation of Scheme on bare metal with the Gambit compiler as a run time environment faced a number of issues. The purpose of this section is to explain the main issues and choices that affected the final architecture.

Our plan was to first get the C kernel to the point where it was able to execute C programs compiled as *flat binaries*. Only then did the work on the Scheme kernel begin, with the intent to gradually replace all the functionality of the C kernel by a Scheme equivalent in the final system (an ongoing effort). We encountered and fixed the following issues:

1. The Gambit RTS requires a C library.
2. Interrupts were sometimes dropped.
3. Exclusive access to the hardware is sometimes required.

4.1 Custom C Library

Gambit's RTS is designed to be hosted on top of a minimal subset of the standard C library; in all about 70 functions are needed like `malloc`, `fopen`, `setitimer`, etc., half of which are `math.h` functions. Normally the build process automatically detects the features available on the OS that go beyond this subset to support advanced features like networking, subprocesses, dynamic library loading, etc. For Mimosa, a custom `libc` was written on top of the existing MINOS code base. To avoid implementing a program linker/loader, the

`libc` functions are called through a table of function pointers that is placed at a fixed address and the loaded program is also placed at a fixed address.

4.2 Dropped Interrupts

An important milestone was compiling the Gambit interpreter, which includes the Gambit RTS, and loading it from disk using the C kernel. This required very few changes to its source code and most of the work went into fixing issues in `libc` that were discovered in the process. With the interpreter, Scheme source code could now be loaded from disk using Gambit's standard module system mechanisms.

This allowed starting the implementation of the Scheme kernel drivers using the interpreter, with the source code stored on the file system. We observed that the system would regularly go unresponsive to repeated keyboard input during noticeable GC pauses. The abnormally long pauses were eventually traced to a poor internal encoding of the interpreted code that caused a high heap occupancy, and consequently long GC pauses. This issue was fixed (see Section 5), but it confirmed that interrupt handling in Scheme is problematic when using blocking GC.

Switching to an incremental GC algorithm, while an interesting approach to explore, is outside the scope of this project and the shorter pauses may still be too long for interrupts that need a quick processing, such as high baud rate serial links. Instead, a double two level queuing system is used, as explained in the next sections.

4.3 Analysis of Interrupt Management

To motivate this design, let us explore the types of interrupts one might encounter. On the x86 platform, until the acknowledgment of a particular interrupt is performed, the processor will not raise it again even when it occurs again. This mechanism allows us to control the rhythm of interruptions and thus delay their processing until the data has reached the Scheme kernel. We partition interrupts into two classes, the *flow controlled* and *flow uncontrolled* interrupts.

The flow controlled interrupts, such as the ATA ones, are interrupts in which the delay between reception and acknowledgment is of no consequence for correctness, as the data needs to be processed completely before more can come in. These interrupts are received in the C layer, but their processing is delayed until the Scheme layer receives it.

In contrast, flow uncontrolled interrupts are interrupts where the data is available fully upon interrupt reception and there is no way to limit the rate of reception. Consequently, the delay between reception and acknowledgment is critical because responding too slowly will cause data loss. In our system, we consider PS2 and UART interrupts to be of this kind. Therefore, a small layer of C partially processes them and stores the available data, waiting for the Scheme layer to complete their processing.

Flow controlled interrupts are placed in a set, since we will not encounter more than one at a given time, and flow uncontrolled are placed in a static queue, with their associated data.

Since we choose the size of the set to be big enough for all flow controlled interrupts to fit, a garbage collection pass will never result in lost data, since their treatment can be delayed as required. However, flow uncontrolled interrupts will start to fill the static queue if garbage collection passes are too long. This is because the Scheme kernel does not empty the queues during GC pauses. This can result in dropped interrupts. Quick processing is critical, as otherwise their source will also start dropping interrupts. The UART controller, for instance, might be unable to stop the sending device from further transmissions that would overflow the UART hardware receiving buffers when flow control wiring is missing (a common situation). Therefore, it appears that having a garbage collected language to handle interrupts will always yield this dilemma between data loss on the sending device or in the system, the balance of which is ruled by the size of the static queue. Choosing the optimal size pertains more to engineering, as it hinges on the needs of the user. Notice that even with a faster garbage collection solution, the problem will still exist as fast bursts of interrupts can still fill the queue.

An alternate solution would be the use of a subset of Scheme, with a manual memory management model or where memory allocation is forbidden. The use of this restricted Scheme code would allow the handling of flow uncontrolled interrupts during a garbage collection cycle. We discuss this more in depth in section 5.2.

4.4 Exclusive Access

Some drivers require exclusive access to resources with no simultaneous access by another module, thread or interrupt handler. There is a need to inhibit interrupts, and thread context switches at these times. In a uni-processor OS, this is typically achieved by clearing the processor's interrupt mask flag. In Mimosa, hardware interrupts are only one of the possible sources of interrupts and context switches are possible for other reasons than heartbeat interrupts, for example when a thread blocks reading an empty vectorport. Consequently, using the processor's interrupt mask flag only works in some contexts, and in general the higher level mechanisms explained in Section 3.4 need to be used.

Another context where this issue occurs is the need to protect the execution of Scheme threads that are part of Mimosa's implementation, such as the *idle* and *notification pump* threads. The processing of *user interrupts* (CTRL-C) must be inhibited while these threads are running otherwise deadlocks will occur (stopping the *notification pump* thread and starting a REPL will freeze the system because that thread is essential in the delivery of keyboard input). This

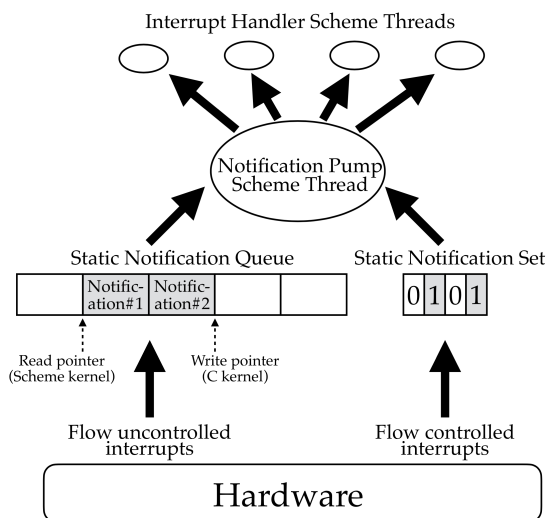


Figure 2. Interrupts queues in Mimosa

can be done by using the per-thread interrupt mask available in the Gambit RTS.

5 Using Scheme

This section highlights critical parts of the development process related to the use of Scheme. Scheme’s use throughout the project is compared with the use of C and assembly. We discuss expressing low-level primitives, OS services and runtime system services, and also discuss the impact of using a REPL.

5.1 Adapting Scheme

The development of low level software often requires using specific machine instructions that would not be otherwise directly accessible to an application. For example the `inb/inw/outb/outw` instructions are used on x86 to communicate with I/O ports. Similarly, direct memory access is a feature often used to read/write structures located in memory and communicate with memory mapped devices. Such constructs are understandably not provided in standard Scheme since userspace privileges normally forbid such instructions to applications on a typical OS.

Since the RTS runs in *kernel level* without constraints, we can use the Gambit compiler’s x86 code generation module within the RTS to create a Scheme wrapper over any instruction sequence needed. This approach is shown in Listing 1, which provides access to the `inb` instruction using a Scheme procedure with a machine code definition.

The Scheme kernel currently only uses direct memory access to communicate with the C kernel through shared memory structures, such as the notification queue. This can easily be implemented using Gambit’s FFI as shown in Listing 2.

```
(define inb
  (asm
    (lambda (cgc)
      (x86-mov  cgc (x86-edx) (x86-mem 4 (x86-esp)))
      (x86-sar  cgc (x86-edx) (x86-imm-int 2))
      (x86-mov  cgc (x86-eax) (x86-imm-int 0))
      (x86-in-dx cgc (x86-al))
      (x86-shl  cgc (x86-eax) (x86-imm-int 2))
      (x86-ret  cgc))))
```

Listing 1. The `inb` procedure. The `cgc` variable denotes the *code generation context*, an example of the usefulness of Gambit’s compiler features to extend Scheme.

```
(define fetch-u8
  (c-lambda ((pointer unsigned-int8) int32)
    unsigned-int8
    "___return(*(___arg1 + ___arg2));"))
```

Listing 2. Procedure to read bytes from memory defined using the FFI.

The `fetch-u8` procedure receives two parameters, a foreign object byte pointer and an integer offset, and fetches the byte at that offset from the address indicated by the pointer. When the special `#f` value is passed as the first parameter (representing the null pointer, i.e. address 0), the second parameter acts as the absolute address of the byte to fetch.

A common C/C++ idiom to fill memory structures with data is to cast the memory address into a pointer to the structure, providing an easy access to values at different offsets. However, support for that idiom is limited in Scheme. Memory blocks, such as sectors read by a disk, are represented by byte vectors in our implementation. To provide a similar behavior while respecting Scheme conventions, we use macros to unpack a vector into a memory structure and vice-versa. This results in a fairly dense macro that creates *C-like structures*, providing methods to pack and unpack these structures in a vector according to the size of each field. This has the obvious disadvantage that memory must be duplicated in two locations, as opposed to using pointers to the memory area. Our usage is limited to fairly small structures, but a possible solution to the data duplication is to abstract memory access as procedures over the vector itself. We mainly use this pattern in the FAT32 driver to access the file system data structures contained in memory. On a smaller scale, since data is represented as byte vectors, we created utilities to recreate wider integers from an endianness dependent sequence of bytes. Using macros allowed us to simplify the boilerplate code and quickly reach feature parity with C in that regard.

All of these tools allow us to access raw memory, I/O ports and other low level mechanisms without having to break the language’s semantics, providing the missing features

required without much difficulty. Of course, these adaptations are required because Scheme lacks these constructs in the first place. However, the tools provided by the language makes the adaptation fairly easy. Moreover, the extensible nature of the language makes the inclusion of those foreign concepts fully transparent to the developer, as these procedures do not suffer any limitations. The unusual context of having the full language available inside of its compiler certainly helps in providing tooling to extend on demand the capacity of the language.

5.2 Minimizing the impact of Garbage Collection

As previously discussed, GC created noticeable pauses and made the system less responsive. The use of a two layer notification queue helped avoid dropping interrupts due to the pauses but it did not resolve the root of the issue. By compiling device drivers and Gambit libraries, particularly the import system, we were able to reduce the duration of each pause from an average of 1.1 seconds to an average of 70 milliseconds¹, going from 379MB of heap usage to a mere 9MB.

While these pauses are hard to notice, a system with more allocated resources, perhaps running user applications and additional device drivers would eventually run into the same problems. Therefore, this does not solve the root of the issue but it brings down our memory footprint to a reasonable level for experimentation.

The impact of memory management is also a concern in other operating systems where applications use automatic GC such as Android [2]. However, a distinction to be made is that in Mimosa, the Scheme kernel itself is written in a language needing GC and the drivers are slowed down by GC cycles. The current execution model, where applications are loaded in the same runtime as the rest of the kernel, makes the system much more vulnerable to slow down since a badly behaving program can adversely affect the entire system.

We have yet to explore the idea of having multiple Gambit instances running. This is the approach taken by Android [1] to manage application memory. Such an architecture would require the inclusion of communication mechanisms since system services might be provided by a different virtual machine. This approach has the disadvantage of introducing overhead for system calls and a bigger memory usage. An alternative is the implementation of a higher level system for managing resources within a single run time. This is somewhat related to the approach taken by the J-Kernel[28].

The problems encountered with GC in this project have also been encountered in the implementation of LISP machines. However, the choice of the x86 platform prohibits us from hardware specific solutions described in [18]. Other GC

algorithms, such as Baker's real time garbage collection algorithm [4] provide constant bounds on pause time. Such an implementation could help in providing strict bounds on the GC pauses. Collectors based on the lifetime of objects such as introduced in [22] could also improve collection pauses since it optimizes for short lived objects.

5.2.1 Coexistence of Compiled and Interpreted Scheme Drivers

To include compiled drivers in Mimosa, we compiled modules with our custom Gambit implementation. The resulting C files were used to recreate another Gambit version that includes those compiled files as builtin libraries. This produces the final Scheme kernel used in the system. While this increases the build time, it enables us to run the system with a mix of compiled and interpreted libraries and drivers. As expected, the development cycle of compiled sources is much slower but both performance and memory footprint are improved.

5.3 Using Scheme for Writing Device Drivers

Many OS related algorithms make use of the imperative nature of C, C++ and assembly. In particular, interactions with peripheral devices are often performed in sequences of output instructions followed by sequences of input instructions. This sometimes conflicts with the functional interfaces provided by Scheme. In this section we discuss how we integrated and expressed routines in the Scheme codebase and the benefits of using Scheme as opposed to C.

5.3.1 Code Style

Some drivers, such as the keyboard driver, are easy to convert to Scheme. In essence, the driver acts as a transformation routine from a scancode to a printable character, making it fairly natural to convert from C to Scheme. However, this is not the case of all the device drivers we implemented. The hard disk access (ATA), real time clock (RTC) and FAT32 routines cannot be directly translated in most cases. Many interactions require looping over data, sending information to the controllers, waiting and performing more interactions. It would be possible to model all of these interactions with function composition, having each function perform recursive looping when needed. However, this would reduce readability and introduce complexity for tasks that would be easier to do in an imperative fashion. For instance, the ATA controller reset routines makes for a fairly dense function that requires many trivial checks that need to be performed sequentially. In these cases, we found using the `begin` and `for-each` special forms gave us enough power to express these operations easily at the cost of potentially less idiomatic Scheme code. It is however important to point out that rethinking many of the algorithms, especially in the case of the filesystem driver, mostly solved this issue.

¹On QEMU 4.2.0, without CPU passthrough

5.3.2 Error Handling

Operations involving hardware can fail for multiple reasons. The success of such operations must be communicated through the code to allow the system to react accordingly. The traditional C approach is to use integer codes to communicate error status. In the Scheme kernel, symbols are used as a replacement for numerical identifiers. Symbols provide a convenient way to provide a unique identifier that is both human-readable while debugging and easy for the machine to compare. This contrasts with integer error codes that are usually hard for a human to understand but easy for a machine to compare or strings that provide the opposite trade-off. This approach makes the REPL usage much easier as interaction with devices provides the user with a clear result of the operation, avoiding the need for a lookup while debugging, and allowing error states in structures to be left as-is.

5.3.3 Flow Control

A common C pattern used frequently in the original implementation of the FAT32 driver is the usage of a return value to indicate an error code and a pointer argument to transmit the actual return value. This is mainly used for operations that might fail for hardware reasons or failed memory allocations. This lends well to the imperative nature of C-like languages because we can perform a correctness check in an `if` statement and then use the result through the variable. Since raw pointers do not exist in Scheme, the problem boils down to having multiple return values. Scheme's dynamic typing allows us to return different types of arguments, allowing us to have procedures that sometimes evaluate to an error code, sometimes to a result. This is however a questionable practice since the semantics of the value changes but the type might not always change, making the detection of errors hard. We felt like the best solution was to use continuation passing style (CPS), a common functional programming pattern. Using CPS, we can specify success and failure continuations, allowing the code to naturally branch out. We typically use a success continuation that receives the actual value and an error continuation that receives an error code in the form of a symbol for the reasons mentioned above. This allows for elegant function composition and reuse throughout the driver code. Continuation passing style also makes use of Scheme's tail calling, reducing the impact on the stack, something that cannot be replicated by using function pointers in C, making this pattern more practical in Scheme. Listing 3 and 4 provide examples of these patterns.

Another variation on this pattern is used in the acquisition of hard disk sectors. The system provides a cache to speed up lookup and use of disk block. In the C kernel, the usage of a block often follows this pattern (Listing 5).

```
int data = 0;
error_code err;
if(HAS_ERROR(err =
    operation(var1, var2, &data))) {
    ... manipulate data ...
} else {
    ... handle err ...
}
```

Listing 3. Error checking pattern in C

```
(operation
  var1
  var2
  (lambda (data)
    ...manipulate data...)
  (lambda (err)
    ...handle err...))
```

Listing 4. Error checking and control flow in Scheme

```
disk* d = ...;
cache_block* cb = ...;
if (!ERROR(err =
    disk_cache_block_acquire(d, 0, &cb))) {
    rwmutex_readlock(cb->mut);
    uint8* data = cb->buf;
    ... manipulate data ...
    rwmutex_readunlock(cb->mut);
    release_error = disk_cache_block_release(cb);
}
```

Listing 5. Common C pattern for accessing cached blocks

Such a pattern introduces multiple points of failure, requiring the mutexes to be acquired and released. Block operations need to be limited to the lock's scope. Such a sequential list of tasks is also hard to integrate into Scheme code and is poor use of the language. Using CPS, all the lock operations are integrated in the block routine and the continuation receives the data vector as an argument. When the data manipulation is done, the lock and the block are released. An equivalent usage is given in Listing 6.

```
(with-sector
  d ; the disk
  0 ; the sector address
  MRW ; read-write
  (lambda (sector-vector)
    ... Data manipulation ...
  ))
```

Listing 6. Cached block (sectors) access in Scheme

The memory access control operations are implied in the `with-sector` procedure that encapsulates the critical section accesses. This greatly reduces the amount of duplicated logic throughout the code base by providing a single interface to access data.

This greatly simplifies safe access over shared structures compared to what the C language offers. This comfort is mainly given by the language's ability to manipulate procedures. Another example is found in the Scheme RTC, where system or run time interrupts can significantly affect the results of interactions with the clock. To safely perform an operation, the `atomic` macro is used (Listing 7), simplifying access and ensuring safety of the code.

```
(define-macro (atomic . exprs)
  (let ((r (gensym)))
    ` (begin
      ; Run Time interrupts
      (##disable-interrupts!)
      ; System interrupts
      (disable-interrupts)
      (let ((,r (begin ,@exprs)))
        (enable-interrupts)
        (##enable-interrupts!)
        ,r)
      )))
```

Listing 7. The atomic macro that allows safe access to the RTC

5.3.4 Continuation Passing Style for Asynchronous Operations

As we mentioned, the ability to manipulate procedures as first class objects is of great use in the development of systems. A common use for something akin to CPS in imperative languages is the use of callback functions in the context of asynchronous programming. For instance, JavaScript libraries extensively uses this pattern to specify the code that has to be executed when the operation completes but on a different timeline of execution. In general, event based programming can make use of continuations to simplify dispatching of events and their handling.

Clearly, asynchronous programming takes place in OS, even for simple drivers, due to the desire of avoiding polling (or busy waiting) and using interrupts as events. A more concrete example is the implementation of the ATA driver. We dive in this specific use case to illustrate the point.

There are many modes of execution for ATA drivers, but the specification requires all drives to support the programmed input-output mode (PIO) which is the supported implementation in Mimosa. Sending a command to a hard drive potentially implies physical movement of its parts and so it can take time before the data requested is available. The system must therefore wait until the information is available. Waiting for tasks to complete can generally be implemented

in one of two ways. Polling the hard drive (or any device in general) may be performed to wait until the device has completed the requested operation. However, polling wastes processor cycles that could be used for other tasks. The alternative approach of waiting for an interrupt to be received avoids this. However, the trade-off is the complexity required to manage reception of the events. Proper handling of the interrupts requires having state variables to keep track of what operation was requested first and where to store the received data. In the C kernel, in the case of the hard disk driver, all of this is managed through a handler function and an external queue that keeps track of the information and provides allocation and deallocation of queue data. In the Scheme layer, we make use of continuations to store all the desired information, and push the continuation to a command queue. This makes the code much simpler, since all of the information required for handling the interrupt is present while creating the command. Upon interrupt reception, the only required action is to call the first continuation in the queue since the continuation stores the information for us. Section A in the appendix lists the code for both the C and Scheme handling routines, focusing on the sector read operation. A more complete implementation of an OS could make use of this in many aspects of the system, such as process scheduling, synchronization of I/O operations and message passing between drivers amongst many use cases.

5.4 The REPL

The REPL allows interactive code experimentation by providing a more dynamic development cycle with immediate feedback. While REPLs are available in many language implementations, they are typically nonexistent or awkward in languages traditionally used for OS development. In this section, we discuss our usage of the REPL in the context of Mimosa's development.

Mimosa offers two independent REPLs once the Scheme kernel has completed loading: a keyboard driven on-screen REPL and a serial link driven REPL (which is accessible through a network connection when Mimosa is run on a virtual machine such as QEMU). Since the Scheme drivers are loaded through the RTS, the REPLs have access to all driver functionality, enabling the programmer to interactively perform tests on the implementation, live-repairing the implementation, and experimenting new features.

5.4.1 Debugging

With the Scheme REPL, tests can be performed by loading a test file and executing it in the REPL. Results can then be manipulated directly through the REPL. While a C programmer can add code to print the result of simple types to confirm expectations, the Scheme developer can go further and inspect more types of results and, most importantly, repair the problem without a time consuming edit-compile-reboot cycle. Furthermore, with the serial link REPL, the debugging

can benefit from the convenience of a separate development machine, such as driving automated testing with scripts running on the development machine.

The REPL also helps with debugging infrequently occurring issues that depend on a specific state of a complex transaction that is hard to reproduce in a new run. This is because the REPL can access all of the system's state information directly after a bug happens. This was particularly helpful in the development of the serial link driver to detect a transition to an invalid state that would freeze the data flow. After noticing with the REPL that manual calls to the *read* operation would still read characters, it was a simple matter to access the UART registers with the REPL and notice that some interrupt acknowledgment condition was not fulfilled, stalling the data flow.

5.4.2 Exploration

Device programming requires a good understanding of the available registers and data interfaces provided by the hardware. While documentation usually provides a detailed specification of how a device behaves, the REPL can provide empirical experience to help understand the specification.

This proved beneficial for testing the behavior of code about to be added to the system before committing to it. In effect, it allows the interactive development of code and the evaluation of the behavior of devices. This was useful when writing the UART and RTC drivers as it enabled experimenting with various interface procedures and verification of the expected output.

Of course, since the REPL runs on the system that is being developed, it is not immune to crashes of the system. Incorrect device manipulation might render the REPL less effective. For instance, bugs regarding the serial driver obviously affects the REPLs opened through the network. Since the system only supports one core of the processor, deadlocks or frozen devices might disable the main interface, making interaction impossible.

5.5 Evaluation of the Architecture

The way the system is organized provides restrictions and benefits. We evaluate our approach by pointing out limitations on our use of the Scheme language. While we sometimes focus on the specifics of the Scheme language, many of these remarks might apply to other hosted languages, arranged in a similar configuration.

Like mentioned in earlier parts, the processor interrupts must be passed between the two layers to enable proper reaction. The current system for transferring interrupts is illustrated in Figure 2 and relies on a shared memory section. To accommodate for the variable number of essential parameters per flow uncontrolled interrupt, variable length notifications frames are inserted in the buffer, adding to the complexity of the mechanism.

It is well known that using C can lead to erroneous manipulations, particularly regarding memory, since the language provides few guardrails. This is mainly due to fact that C code has an almost direct translation to machine code, with few run-time checks for operation safety. This is an attractive proposition if the goal is run-time performance, since the lack of checks implies less processor cycles wasted performing checks.

This contrasts with Scheme, where data manipulation involves type checking and input validation where necessary. In a Scheme program running on a mainstream OS, incorrect manipulation will result in a clean termination. In a situation where the operator has REPL access, he will be informed of the error and the debugging can start. Since we are running the system on top of a thin C layer to provide common routines, bugs in those routines will cause erroneous execution at the Scheme level. Since these issues can occur in an unrelated part of the system, coded in a different language, these bugs are typically hard to locate and can cause unpredictable effects, akin to illegal memory manipulations in C. These issues often freeze the Scheme kernel, rendering the REPL and other tools unusable. An option for avoiding these issues would be to program more of the system in Scheme. Having some parts of the kernel in a limited subset of Scheme would enable these parts to be independent of the C kernel, but would however mean limited use of the language, relying on some compilation mechanisms to bridge the gap between what is used in Scheme and what is available from the machine.

6 Related work

There have been a number of projects that have brought Scheme directly to the hardware. As mentioned in the introduction, this is not a new concept. As we noted, Lisp machines had a similar goal, providing a Lisp system and environment to the user, while allowing the Lisp programs to access the entire machine without restrictions. These computers often relied on hardware support to provide efficient operation; something we avoided here.

Many Scheme implementations target micro-controllers to provide an environment for embedded system development. While the goal is often to use Scheme as a development language, the hardware constraints provide an interesting challenge. PICBIT [17] and PICOBIT [26] are Scheme systems for Microchip's PIC micro-controllers, providing support for the limited nature of these 8 bit devices.

Using higher level languages to program hardware devices is of course not exclusive to Scheme. A much more complete Go system is discussed in [6] where the usage of high level languages for kernels is also evaluated. Their approach is similar to ours, as they load the Go run time and provide kernel features with a small compatibility layer.

Language	Excluding Scheme Libraries		Including Scheme Libraries	
	Files	LOC	Files	LOC
C/C++ Header	50	13806	63	14224
C	33	8488	45	13729
Scheme	17	4852	264	32935

Table 1. Code distribution of the project. The breakdown is detailed in Section 7.

The Rust language promises C-like performance without the dangers of manual memory management by enforcing a set of memory ownership rules. The Redox OS [12] is a micro-kernel developed in Rust, using the unusual memory management of Rust. The study of using Rust for a Unix-Like operating system is done in [23]. The increasing popularity of Rust in recent years has also spawned multiple projects at various level of completion [7, 9–11, 13].

Other projects leverage the REPL to provide a more convenient development environment. uLisp [20] is a Lisp implementation that runs on various micro-controllers, including Arduino boards. It provides a REPL over the micro-controller to interact with the hardware. uLisp mentions the educational benefits of using the REPL to learn programming and device programming since it allows for quick experimentation. Armpit Scheme [25] provides a Scheme REPL running on RISC micro-controllers, enabling manipulation of hardware and low-level constructs in Scheme. Loko Scheme [29] is an implementation of Scheme that compiles directly to bare metal, providing an interface to the hardware. As opposed to the aforementioned projects, Loko Scheme tries to be a more complete Scheme system, close in spirit to Mimosa.

The objectives and characteristics of an ideal Lisp operating system are discussed in [27] and provides many potential future goals for the Mimosa project, in particular in how GC is approached. In the same vein, Mezzano [8] is a fairly complete operating system written entirely in Common Lisp. ChrysaLisp [19] takes a different approach and is a virtualized OS that works in a hosted environment, with the goal of eventually moving to bare metal.

The idea of using high level languages to create applications that have full control over the machine is gaining traction, as the benefits of having a very small and specialized image can be enticing. Uni-kernels, such as presented in [24], are facing many of the issues we encountered here, especially regarding garbage collection.

7 Mimosa Source Code

The main repository for the project is available on Github at the following URL: <https://github.com/udem-dlteam/mimosa>.

In order to compare the size of the Scheme code base with the C code base, we used the `cloc` [3] utility to compare the amount of C code and Scheme code. We aggregated C++ sources and C sources, since our usage of C++ was

limited to C features including C++ function overloading for convenience. The breakdown was obtained on commit 73358d71 using version 1.82 of `cloc`. We excluded the `attic` directory since it does not contain any in-use code. We excluded the `fonts` directory, since it is more akin to binary data. We excluded the `utils` directory, since it does not contain code used by the operating system. Table 1 contains the breakdown obtained by running `cloc . --exclude-dir=attic,fonts,.vscode,util`. The table contains two breakdowns. One with the `archive-items` directory and one without. This directory contains Gambit’s Scheme libraries required for correct execution of the drivers, but is not code written strictly for this project. Irrelevant file types are excluded from the tables.

8 Conclusion

A Scheme runtime provides multiple features, most of which are not available in a language like C without the use of external libraries and tools. OS developers must forego such standard libraries since often times they rely on the OS itself. While we did have to reimplement some basic functionalities, the Gambit runtime provided many tools regarding threading, synchronization, data manipulation and high level operations over data structures. Gambit also provides compilation mechanisms, allowing us to transparently create lower level routines and interact with them without regard to their implementation. The Scheme REPL provides a convenient tool to interact with the code in a way that goes beyond traditional OS tools. The very nature of Scheme makes it very malleable, allowing the language to be adapted to various situations by extending it when necessary with the usage of macros and closures.

Of course, the run time impact is sometimes significant. The GC adversely impacts performance by freezing the system periodically. Mimosa requires a small layer of C code to support its functionalities. Because of the way interrupts are handled in Scheme, some devices that require immediate attention might not be properly managed.

We originally wondered whether using Scheme to write hardware facing code would provide a tangible benefit that would outweigh the penalty induced by the runtime environment. Of course, the constraints of some systems, such as a real-time operating system, are unsatisfied by our implementation that does not achieve the responsiveness required.

However, the features of Scheme make implementation of many constructs much easier, providing the flexibility required to create data manipulation routines of variable abstraction level over the hardware. Automatic memory management certainly relieves the programmer of a lot of effort at a certain price.

Mimosa is still a work in progress and we foresee additional experimentation. An incremental GC algorithm would provide interesting insight on the practicality of garbage collected languages for OS development. The segmentation of user programs into different runtimes would prevent badly behaving applications from slowing down the entire kernel. It would be interesting to see if continuations can be used to make fast inter-process calls to a kernel-provided library, a requirement if the applications are to be executed in their own environment. More short-term goals for Mimosa are:

1. Write the interrupt handling code in a restricted Scheme to allow some interrupt handling during GC cycles.
2. Eliminate the C kernel and use the Gambit x86 backend to remove the need for a C compiler.
3. Use the Actor model (communicating threads) for interrupt and event processing.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] 2020. [Android] Platform Architecture. <https://developer.android.com/guide/platform> Accessed: 2020-08-05.
- [2] 2020. Overview of memory management. <https://developer.android.com/topic/performance/memory-overview> Accessed: 2020-08-05.
- [3] AlDanial. 2020. Cloc. (2020). <https://github.com/AlDanial/cloc>
- [4] Henry G Baker Jr. 1978. List processing in real time on a serial computer. *Commun. ACM* 21, 4 (1978), 280–294.
- [5] Fabrice Bellard. 2005. QEMU, a fast and portable dynamic translator. In *USENIX Annual Technical Conference, FREENIX Track*, Vol. 41. 46.
- [6] Cody Cutler, M Frans Kaashoek, and Robert T Morris. 2018. The benefits and costs of writing a {POSIX} kernel in a high-level language. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 89–105.
- [7] Intermezz Developers. 2019. IntermezzOS. (2019). <https://github.com/intermezzos>
- [8] Mezzano Developers. 2020. Mezzano Common Lisp OS. (2020). <https://github.com/froggey/Mezzano>
- [9] Nebulet Developers. 2018. Nebulet. (2018). <https://github.com/nebulet/nebulet>
- [10] Quilt OS Developers. 2016. RustOS. (2016). <https://github.com/QuiltOS/QuiltOS>
- [11] Rux Developers. 2018. Rux. (2018). <https://source.that.world/source/rux/>
- [12] Redox Developers. 2020. Redox OS. (2020). <https://www.redox-os.org/docs/>
- [13] Tifflin Developers. 2020. Tifflin. (2020). https://github.com/thepowersgang/rust_os
- [14] Marc Feeley. 1993. Polling efficiently on stock hardware. In *Proceedings of the conference on Functional programming languages and computer architecture*. 179–187.
- [15] Marc Feeley. 2001. SRFI 18: Multithreading support. (2001). <https://srfi.schemers.org/srfi-18/srfi-18.html>
- [16] Marc Feeley. 2001. SRFI 21: Real-time multithreading support. (2001). <https://srfi.schemers.org/srfi-21/srfi-21.html>
- [17] Marc Feeley and Danny Dubé. 2003. PICBIT: A Scheme system for the PIC microcontroller. In *Proceedings of the Fourth Workshop on Scheme and Functional Programming*. Citeseer, 7–15.
- [18] Richard D Greenblatt, Thomas F Knight, John T Holloway, and David A Moon. 1980. A LISP machine. In *Proceedings of the fifth workshop on Computer architecture for non-numeric processing*. 137–138.
- [19] Chris Hinsley. 2020. (2020). <https://github.com/vygr/ChrysaLisp>
- [20] David Johnson-Davies. Unknown. uLisp. Project Website. (Unknown). <http://www.ulisp.com>
- [21] D Kevin Layer and Chris Richardson. 1991. Lisp Systems in the 1990s. *Commun. ACM* 34, 9 (1991), 48–57.
- [22] Henry Lieberman and Carl Hewitt. 1983. A real-time garbage collector based on the lifetimes of objects. *Commun. ACM* 26, 6 (1983), 419–429.
- [23] Alex Light. 2015. *Reenix: Implementing a unix-like operating system in rust*. Ph.D. Dissertation. Ph. D. dissertation, Master’s thesis, Brown University, Department of Computer Science.
- [24] Anil Madhavapeddy, Richard Mortier, Charalampos Rotsos, David Scott, Balraj Singh, Thomas Gazagnaire, Steven Smith, Steven Hand, and Jon Crowcroft. 2013. Unikernels: Library operating systems for the cloud. *ACM SIGARCH Computer Architecture News* 41, 1 (2013), 461–472.
- [25] Hubert Montas. 2006. Armpit Scheme. SourceForge Repository. (2006). <http://armpit.sourceforge.net/>
- [26] Vincent St-Amour and Marc Feeley. 2009. PICOBIT: a compact scheme system for microcontrollers. In *International Symposium on Implementation and Application of Functional Languages*. Springer, 1–17.
- [27] Robert Strandh. 2013. (2013). <http://metamodular.com/lispos.pdf>
- [28] Thorsten Von Eicken, Chi-Chao Chang, Grzegorz Czajkowski, Chris Hawblitzel, Deyu Hu, and Dan Spoonhower. 1999. J-Kernel: A capability-based operating system for Java. In *Secure Internet Programming*. Springer, 369–393.
- [29] Göran Weinholt. 2019-2020. Loko Scheme. GitLab Repository. (2019-2020). <https://scheme.fail/>

A Handling of IDE interrupts

```

void ide_irq(ide_controller* ctrl) {
    uint8 s;
    uint32 i;
    ide_cmd_queue_entry* entry;
    uint16 base;
    uint16* p = NULL;

    entry = &ctrl->cmd_queue[0]; // We only handle one operation at a time
    base = ide_controller_map[ctrl->id].base;

    cmd_type type = entry->cmd;

    if (type == cmd_read_sectors) {
        p = CAST(uint16*, entry->_.read_sectors.buf);
    } else if (type == cmd_write_sectors) {
        p = CAST(uint16*, entry->_.write_sectors.buf);
    } else if (type == cmd_flush_cache) {
        p = NULL;
    } else {
        panic(L"[IDE.CPP]_Unknown_command_type...");
    }

    s = inb(base + IDE_STATUS_REG);

    if (s & IDE_STATUS_ERR) {
        // #ifdef SHOW_DISK_INFO
        uint8 err = inb(base + IDE_ERROR_REG);
        term_write(cout, "***IDE_ERROR***\n");

        if (err & IDE_ERROR_BBK)
            term_write(cout, "Bad_block_mark_detected_in_sector's_ID_field\n");
        if (err & IDE_ERROR_UNC)
            term_write(cout, "Uncorrectable_data_error_encountered\n");
        if (err & IDE_ERROR_IDNF)
            term_write(cout, "Requested_sector's_ID_field_not_found\n");
        if (err & IDE_ERROR_ABRT)
            term_write(cout, "Command_aborted_(status_error_or_invalid_cmd)\n");
        if (err & IDE_ERROR_TK0NF)
            term_write(cout, "Track_0_not_found_during_recalibrate_command\n");
        if (err & IDE_ERROR_AMNF)
            term_write(cout, "Data_address_mark_not_found_after_ID_field\n");
        // #endif

        if (type == cmd_read_sectors) {
            entry->_.read_sectors.err = UNKNOWN_ERROR;
        } else if (type == cmd_write_sectors) {
            entry->_.write_sectors.err = UNKNOWN_ERROR;
        }
        condvar_mutexless_signal(entry->done);
        ide_cmd_queue_free(entry);
    } else if (type == cmd_read_sectors) {
        for (i = entry->_.read_sectors.count << (IDE_LOG2_SECTOR_SIZE - 1); i > 0; i--)
            *p++ = inw(base + IDE_DATA_REG);

        if (inb(base + IDE_ALT_STATUS_REG) & IDE_STATUS_DRQ) {
            entry->_.read_sectors.err = UNKNOWN_ERROR;
        } else {
            entry->_.read_sectors.err = NO_ERROR;
        }
        condvar_mutexless_signal(entry->done);
        ide_cmd_queue_free(entry);
    } else if (type == cmd_write_sectors) {
        if (entry->_.write_sectors.written < entry->_.write_sectors.count) {
            // Write the next sector to write
            for (uint16 i = 1 << (IDE_LOG2_SECTOR_SIZE - 1); i > 0; i--) {
                outw(*p++, base + IDE_DATA_REG);
            }
            entry->_.write_sectors.buf = p;
        }
    }
}

```

```

    entry->_.write_sectors.written++;
  } else {
    // This is the status interrupt
    if (inb(base + IDE_ALT_STATUS_REG) & IDE_STATUS_DRQ) {
      entry->_.write_sectors.err = UNKNOWN_ERROR;
    } else {
      entry->_.write_sectors.err = NO_ERROR;
    }
    condvar_mutexless_signal(entry->done);
    ide_cmd_queue_free(entry);
  }
} else if (type == cmd_flush_cache) {
  condvar_mutexless_signal(entry->done);
  ide_cmd_queue_free(entry);
}
}
}

```

Listing 8. Interrupt handling routine for ATA interrupts

```

(write ; continuation passed into the continuation queue
 (lambda ()
  (mutex-lock! mut)
  (let* ((status (inb stt-reg)))
    (if (mask status IDE-STATUS-ERR)
      (set! err (ide-handle-read-err cpu-port))
      (begin
        (for-each
         (lambda (i) (vector-set! word-vector i (inw data-reg)))
         (iota sz))
        (if (mask IDE-STATUS-DRQ (inb alt-reg))
          (set! err ERR-HWD))))
    ; Signal we are ready
    (condition-variable-signal! cv)
    (mutex-unlock! mut)))
  q)

; Interruption dispatching routine
(define (handle-ide-int controller-no)
 (let*
  ((ctrl (vector-ref IDE-CTRL-VECT controller-no))
   (q (ide-controller-continuations-queue ctrl))
   (cont (read q)))
  (and cont (cont))))

```

Listing 9. Handling of IRQ interrupts in Scheme (with the write continuation)

B The REPL environment

The image shows a terminal window titled 'console' on the left and a terminal window titled 'log' on the right. The 'console' window displays the Mimosa REPL environment with the following text:

```

Mimosa
> (define (set-pixels i) (store-u8 #f (+ #x0000 (* 80 (+ i 6)) 32) 255))
> (for-each set-pixels (iota 10))
> (import (srfi 132))
> (apropos 'list-sort)
"##" namespace:
  list-sort, list-sort!
"srfi/132#" namespace:
  list-sort, list-sort!, list-sorted?
> (list-sort <= '(3 1 4 1 5 9 2 6))
(1 1 2 3 4 5 6 9)
> (top)
*** THREAD LIST:
#<thread #1 primordial>  RUNNING PO
#<thread #2 notification-pump> WAITING #<condition-variable #3> 20s
#<thread #4 idle-thread>  RUNNING PO
#<thread #5 keyboard#interrupt-handler> WAITING #<condition-variable #6 #f>
#<thread #7 uart-pump-1>  WAITING #<input-output-port #8 (string)>
#<thread #9 uart-repl-1>  WAITING #<input-output-port #10 (string)>
#<thread #11 uart#interrupt-handler> WAITING #<condition-variable #12 #f>
>
  
```

The 'log' window displays the following system initialization logs:

```

log
Initializing Mimosa
CPU is GenuineIntel
family=6
model=5
stepping=3
has:
  Floating Point Unit
  Debug Extensions
  Page Size Extensions
  Time Stamp Counter
  Model Specific Registers
  Physical Address Extensions
  Machine Check Exception
  CMPXCHG8B instruction
  Local APIC
  Fast system call
  Page Global Enable
  Conditional MOVE
  Page Attribute Table
  MultiMedia Extensions
  FPUSAVE and FPURESTOR
  SSE instructions
  SSE2 instructions
Loading up disks...
Setting up IDE controllers...
Resetting the IDE...Reset of device 0
Reset of device 1
word 4f hi (should be 12B) = 12B
Maximum number of sectors that sh
all be transferred per interrupt on
READ/WRITE MULTIPLE commands = 16
Setting up IDE controller no: 1
Resetting the IDE...Reset of device 0
Reset of device 1
word 4f hi (should be 12B) = 0
Maximum number of sectors that sh
all be transferred per interrupt on
READ/WRITE MULTIPLE commands = 0
Setting up IDE controller no: 2
Setting up IDE controller no: 3
Loading up the virtual file system.
Unknown partition type: 0
Unknown partition type: 0
Enabling FS2
Enabling the terminal STDOUT bridge
Cleaning up communication memory space
Loading up LIBC
Enabling UART IRQ(s)MIMOSA V.2.0
Starting /dsk1/gambit/bin/asc.exe
The len of the gambit file is: 1069
1440
gambit file loaded...
Running program thread
  
```

Figure 3. The Mimosa "desktop", with the REPL on the left

Scheme for scientific computing

Francesco Montanari

Instituto de Física Teórica IFT-UAM/CSIC, Universidad Autónoma de Madrid

Madrid, Spain

francesco.montanari@uam.es

Abstract

Drawing from specific needs in physics and in machine learning, we review software engineering systems associated with a selection of Scheme implementations and dialects relevant for scientific computing: Chez Scheme, CHICKEN, Gambit and Racket. We address the needs of an *impatient schemer* who aims at profiting from currently available Scheme systems to solve research and data analysis problems. We examine aspects related to runtime performance, development tools (including availability of external libraries) and parallel computing. Based on two case studies, we first discuss the gap between the few Scheme numerical libraries and the extensive resources available for mainstream languages, which is a serious obstacle for prompt adoption of Scheme systems. Then we suggest that research projects building upon simple components can take advantage from Scheme-based languages to overcome expressiveness and efficiency limitations leading to cumbersome engineering practices in modern scientific computing. Further development of basic Scheme numerical libraries serving as common ground for the advancement of more specialized applications is desirable.

CCS Concepts • **Software and its engineering**; • **Applied computing** → **Physical sciences and engineering**;

1 Introduction

Modern scientific computing builds upon a plethora of software libraries addressing the most disparate needs. They often rely on low-level optimized programs written in runtime and memory efficient languages, wrapped by higher-level languages that allow more expressive algorithm design. Rapidly evolving software engineering practices are employed to distribute and run computationally intensive jobs. Still, fundamental research often needs to compose simple, well-understood parts into more or less complex systems, rather than profiting solely from large libraries with vast functionalities. Popular scientific computing frameworks can be cumbersome and even a limiting factor in software development, yet their preeminence is not necessarily justified on software engineering grounds. A simple, expressive language with mature and efficient implementations like Scheme [21, 51, 53] seems an underappreciated resource.

In this report we review aspects of Scheme-based languages relevant for scientific computing, drawing from specific needs in physics and in machine learning. While the work of a *patient seasoned schemer* meticulously building fully-featured state-of-the-art native libraries addressed to a wide diverse audience would be valuable, it is beyond the experience of this report. Instead, here we address the needs of an *impatient little schemer* who aims at immediately solving current research and data analysis problems. We focus on software engineering tools¹ available to tackle current physics and machine learning problems. We measure and discuss runtime performance of existing Scheme benchmark programs (including numeric tests focused on floating point arithmetic), examine support for developing and distributing libraries (highlighting currently available scientific computing software) and for parallel computing. Limits and advantages of such tools are further discussed in the context of two concrete use cases based on computer vision and cosmology, respectively.

A seminal application of Scheme to scientific computing has been the study of the chaotic planetary dynamics in the Solar system [55]. The project built a machine based on the Supercomputer Toolkit framework [9]. The Toolkit featured a low-level assembly lisp dialect addressing, among other features, multiprocessing. The Toolkit included a compiler converting high-level routines written in Scheme into efficient Toolkit programs. Two of the authors of the Toolkit and of the related astrophysics research work also argued in favor of the use of Scheme in physics education [56–58]. Describing a physics problem in terms of computer programs has the advantage of enforcing the use of an unambiguous notation, removing the educational obstacle of implicit context-dependent interpretation of mathematical expressions that are typical of physics notation. The authors show that Scheme is an extremely expressive language to implement programmatically functional physics notation and provide Scmutils [59], a collection of powerful numerical and algebraic packages written in MIT/GNU Scheme.

On the other hand, Scheme is not a common language [37, 60]. This implies fewer available libraries than mainstream choices like Python or R, difficulties in collaborating with colleagues and in integrating Scheme software into

¹We refer to, e.g., Abelson and Sussman [10], Dybvig [21], Felleisen et al. [22], Friedman et al. [25, 26, 27] for introductions to the Scheme programming language itself.

large collaborations' code bases. Furthermore, while the diversity in Scheme implementations and dialects is a valuable resource, it is challenging to profit from it in a holistic way: different Scheme implementations have different extensions to the standard specifications and rely on different software engineering tools (e.g., package managers).

Section 2 reviews popular practices employed in modern physics and in machine learning software development. Section 3, mainly addressed to the novice schemer (or to an experienced schemer interested in the perception of a computationally scientist about modern Scheme systems), reviews a few Scheme implementations and dialects together with the respective software development tools currently available. Based on two case studies, Section 4 discusses both limits and advantages about adopting Scheme for scientific computing. We conclude and examine future prospects in Section 5. Appendix A contains a reference to common SRFI libraries. Appendix B discusses two other Lisp dialects, Common Lisp and Clojure, in view of scientific computing.

2 Scientific Computing in physics and in machine learning

Modern software libraries employed in physics and in machine learning widely resort to languages such as C, C++ and Fortran whenever runtime is critical (also leveraging well-established high performance computing frameworks such as Message Passing Interface [24] and OpenMP [44]), or as legacy code. Most popular front ends, however, are written in dynamic languages with a simpler syntax, such as GNU Octave (or MATLAB), Mathematica, Python and R. For instance, the Numpy and Scipy Python libraries [62] rely on C, C++ and Fortran back ends. In the field of machine learning, PyTorch [45] and TensorFlow [7] exploit a C++/Python synergy to provide state-of-the-art libraries to process large amounts of data, easily accessible via a simple functional Python front ends (Keras [19] in the case of TensorFlow). A remarkable feature of Keras is the user interface consistency with other independent Python projects, including Numpy and Scikit-learn [15, 46], which ease both learning and composition of algorithms (this feature is not limited to Keras, for instance XGBoost [17] provides a wrapper around its main Gradient Boosting algorithm such that it can be called with the very same parameters as the analogous algorithm defined in Scikit-learn). Keras is also used at CERN for particle physics research, where another popular C++ data analysis framework, ROOT [14], is developed to process petabytes of data and provides a Python extension, PyROOT [11]. A further notable example for the adoption of this strategy are the software libraries publicly released by the LIGO collaboration that were used to analyze the first ever observation of gravitational waves from a binary black hole merger [8].

On the top of the rich and uniform library system, several software engineering solutions are used to distribute heavy

workload jobs on remote servers. Python distributions like Anaconda make it easy to install the required dependencies via package managers, while Docker and Singularity virtualization frameworks further focus on avoiding the dependency hell and on reproducibility. The Jupyter notebook is a popular interactive development interface that supports several languages and that is well suited to work also on remote machines. Code execution is combined with rich text and media in a literate programming environment. Intensive computations are comfortably launched from the interactive framework thanks to advanced libraries like Keras that can run both on CPU and GPU, requiring no changes to the user code.²

The low-level and high-level programming languages synergy is a resource when it is used to resort to well-tested and efficient legacy code. However, too often it is a necessity dictated by the lack of mathematical expressiveness of the low-level language, and the lack of efficiency of the high-level language. For instance, physics problems are in general independent of whether a given data structure is allocated on the heap or on the stack of a given machine. However, this a necessary design decision to avoid memory leaks in C code. This leads to a poor abstraction level when facing physics problems. More in general, the C imperative syntax makes it difficult to manipulate representations of mathematical functions. Conversely, Python is well-suited to design algorithms performing both numerical (see, e.g., the Numpy and Scipy front ends) and symbolic computations (see, e.g., SymPy [43]). However, the runtime inefficiency of the CPython interpreter frequently leads to cumbersome engineering solutions to embed C, C++, Fortran or Cython code. Furthermore, the CPython interpreter is not thread-safe and it only releases the Global Interpreter Lock (GIL) in I/O applications, so that again there is the need of leverage on external C code to bypass the GIL.³ There are libraries (e.g., the concurrent.futures standard library, or external projects like Mpi4py and Numba) that bypass the GIL, but they require objects to be serialized according to the Python pickle protocol, which is in general not guaranteed especially when dealing with wrappers of external C/C++ libraries.⁴

An ideal set of tools for scientific computing would include:

²This is important from the perspective of say a statistician or a physicist, since the problems they aim to solve are in general independent of the details of the machine used to carry out numerical computations, hence abstracting away such information allows to focus algorithms around the actual problem under consideration.

³There are Python implementations (e.g., Jython) that do not suffer from this limitation. However, Pypy (possibly the fastest Python implementation available) also has a GIL.

⁴Notably, in terms of functions, only top-level definitions can be pickle-serialized, which limits the expressiveness of Python by preventing the use of convenient internal functions, lambda expressions and lexical closures.

- A simple, expressive and well established language with efficient implementations.
- Tools to easily document, distribute and install libraries.
- A Foreign Function Interface (FFI) to interact with software compatible with the pervasive C language calling conventions. Relying on the FFI shouldn't be a necessity dictated by the limits of the language, but a choice (e.g., to build upon robust common libraries like BLAS and LAPACK for numerical linear algebra).
- Tools for parallel computing.
- Interactive development environment with insightful debugging and profiling tools.

Scheme systems offer a possibility to satisfy all of the above requirements.⁵

3 Overview of Scheme systems

In this section, mainly addressed to the novice schemer, we review four Scheme implementations and dialects in view of modern scientific computing needs including purely numerical applications, but also symbolic computations:⁶ Chez Scheme (<https://www.scheme.com>), CHICKEN (<https://www.call-cc.org>), Gambit (<http://gambitscheme.org>), Racket (<https://www.racket-lang.org/>).⁷ The Scheme standard language is defined in [53], but the *de facto* standards are the Revisedⁿ Report on the Algorithmic Language Scheme (RⁿRS) series of reports. Gambit and CHICKEN implement R⁵RS [40]. Chez Scheme implements R⁶RS [54] that, compared to R⁵RS, standardized important features like libraries. Racket, while being an almost superset of R⁶RS, does not aim at conforming to Scheme standards (but it does provide a R⁶RS module language). Support for the R⁷RS small report [51] is available via external packages and tools for Chez Scheme, CHICKEN and Racket, and there is ongoing work to implement R⁷RS in Gambit. All the languages are available under free software licenses and are easily installable from sources supporting several architectures, or from packages targeted at specific operating systems.

The languages have been selected based on community resources and active support, and bearing in mind runtime performance.⁸ We discuss aspects highlighted above in view of modern scientific computing needs. Table 1 summarizes

⁵The recent Julia programming language also aims at solving all of these issues. Notably, it also implements lisp-like hygienic macros, although the simpler syntax limits its scope compared to Scheme. On the other hand, in our experience runtime performance can be more subtle to achieve than in Scheme implementations here considered, but the language is still rapidly evolving.

⁶See [12] for a more generic discussion, and [38] for detailed differences between Scheme implementations.

⁷Here we consider the Racket CS variant built on Chez Scheme [23]. The regular Racket BC (before Chez Scheme) variant is built on C.

⁸The selection is somewhat arbitrary and unfair. It is beyond the scope of this report to review the large number of Scheme implementations currently available. A non-exhaustive list, which neglects among others an ever-growing number of Scheme-based domain specific languages, includes:

the comparison for the specific language versions here considered.

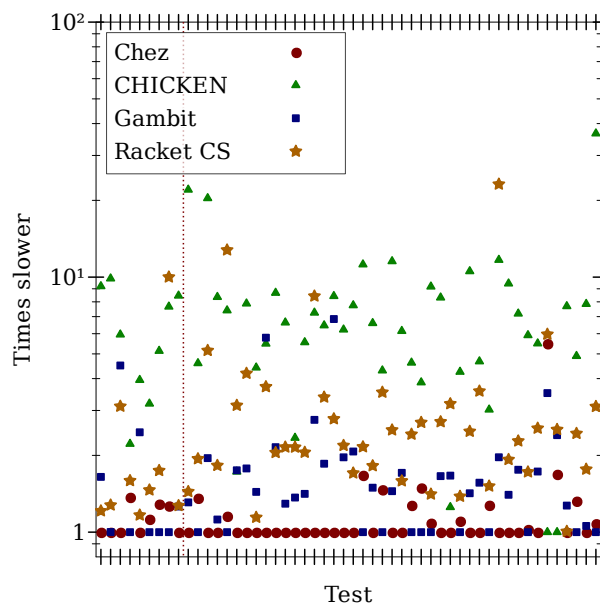


Figure 1. Speed measurements based on Larceny benchmark programs [2, 3]. Every tick on the horizontal axis corresponds to a benchmark test. For each test we show how much slower a given language is compared to the fastest one. Tests at the left of the vertical dotted line are numeric benchmarks.

Speed and reliability. Several speed benchmarks are available [1, 3, 5, 6]. The general approximate picture is that Chez Scheme and Gambit have very efficient compilers, comparable to other high-performance Lisp implementations like SBCL Common Lisp (i.e., runtime range within the same order of magnitude as C and C++, typically only up to $\lesssim 10$ times slower).

We measured runtime speed performance for Larceny benchmark programs [2, 3] on an Intel Core i7-8550U CPU @ 1.80GHz processor. We use the source code provided by the Scheme Benchmarks [3] project (in particular, using the same optimization flags).⁹ Figure 1 shows the results normalized to the fastest one for each test and for the Scheme implementations here considered.¹⁰ We are only interested

Bigloo, Bones, Chez Scheme, Chibi, CHICKEN, Cyclone, Foment, Gambit, Gerbil, Gauche, Guile, IronScheme, Kawa, Larceny, MIT/GNU Scheme, Mosh, Petite Chez Scheme, Picrin, Racket, Rhizome/Pi, RScheme, Sagittarius, Scheme48, Scheme 9 from Empty Space, Vicare, Ypsilon.

⁹Benchmarks are written in R⁷RS style. We used R⁷RS packages available via CHICKEN and Racket repositories and converted R⁷RS library forms to R⁶RS via Akku.scm for Chez Scheme. Source files are slightly adapted for Gambit to remove `(import (scheme ...))` statements.

¹⁰CHICKEN and Gambit provide a separate interpreter and compiler. Here we consider compiled programs.

Table 1. Summary of a few Scheme implementations and dialects in view of scientific computing. We also show tools supporting multiple platforms. The *Times slower* row roughly indicates how much slower a given language typically is compared to the fastest one. More precisely, it shows 25% and 75% quantile bounds for all speed benchmark tests considered in figure 1—see the text for important caveats that apply to interpret these numbers.

Feature	Chez Scheme 9.5	CHICKEN 5.0	Gambit 4.9	Racket 7.7 CS	Cross-platform
Times slower	1.0–1.1	4.4–8.4	1.0–1.8	1.7–3.1	
Package manager	—	chicken-install	(See footnote 14)	raco	Akku.scm
Repository	—	Eggs (~ 400 packages)	—	Racket Packages (~ 1400 packages)	Snow, Spheres, SRFI
Parallel computing	Native threads	Green threads (+ MPI extension)	Green threads	Futures and places (+ high-level extensions)	
IDE	REPL	REPL	REPL	REPL, DrRacket, IRacket	GNU Emacs
Debugging, profiling	Yes	Yes	Yes	Yes	
Documentation	ChezWEB, stex	chicken-doc	—	Scribble	Schmooz

in the broad overview involving diverse benchmark programs, for details about each separate test we refer to the Larceny [2] and Scheme Benchmarks [3] projects. The quantile ranges for each language reported in table 1 show that in most tests CHICKEN and Racket¹¹ have performance a few factors slower than Chez Scheme or Gambit. This picture is also roughly representative of the subset of numeric benchmarks targeting floating point arithmetic and reported at the left of the dotted line in figure 1. However, CHICKEN and Racket outperform the other implementations in a few cases. All of the languages considered are reliable in successfully completing all, or almost all of the benchmark tests analyzed (if an implementation failed completing a test, the corresponding mark is not shown in figure 1).

Here we only consider runtime benchmarks to identify a rough performance range, as many issues arise when interpreting them in terms of effective performance in real case specific scenarios. For instance, in the benchmark sets mentioned above, compiling options are kept constant. However, fine-tuning compiler and runtime options can speed up significantly a given program. Also, Racket provides a Typed Racket language that can boost performance in some cases. More importantly, besides runtime efficiency also development time and programming style are fundamental aspects. As seen, the expressiveness of Python makes it one of the most popular languages for scientific computing, yet the main CPython implementation can easily perform ~ 100 times slower than C, C++ or Fortran.¹² Having a Scheme

implementation that is only up to $\lesssim 10$ times slower than best-performing low-level languages is a significant improvement. Furthermore, the variance between programmers can be larger than the variance between languages [28, 48, 49]: an appropriate programming style is more important for overall efficiency than the particular language choice (provided they are in comparable runtime ranges). This suggests that adopting implementations like Chez Scheme or Gambit for scientific computing is a promising prospect to relegate the necessity of embedding low-level languages (e.g., C, C++, Fortran) to legacy code.

External libraries and package manager. CHICKEN and Racket have package managers to access the Eggs (currently containing ~ 400 packages) and Racket Packages (~ 1400 packages) repositories, respectively. The packaging procedure is fairly straightforward, dependencies are automatically solved, and correct building is continuously checked on several architectures. Concerning scientific computing tools, CHICKEN repository contains 11 libraries in its mathematical section, including an interface to the popular Fortran BLAS (Basic Linear Algebra Subprograms) library for low-level linear algebra operations, and a collection of basic statistics utilities.¹³ Racket built-in math library also includes, among other tools, linear algebra and statistics utilities. The Racket package repository lists 21 items classified as mathematical tools. For instance, `flomat` is a high-level interface to BLAS and LAPACK (Linear Algebra PACKAge—building itself upon BLAS, it provides a collection of higher-level linear algebra operations such as matrix factorization and eigenvalues computation) focused on double precision floating point numbers. Racket also has interfaces to GNU

¹¹We recall that here we consider the Racket variant built on Chez Scheme. We verified that the regular Racket variant built on C still leads to a similar quantile range as shown in table 1, but its performance is ~ 100 slower for the few benchmark programs that require capturing and invoking continuations.

¹²For some numerical applications such a performance is good enough. For instance, *cosmo.el* (<https://gitlab.com/montanari/cosmo-el>) is an Emacs package written entirely in Emacs Lisp (certainly not the best choice for numerical applications) that computes cosmographic quantities like numerically integrated cosmological distances. In general, however, algorithms

like numerical integration, solutions of ordinary differential equations, optimization or algebraic operations on large vectors, matrices or tensors can soon become prohibitive.

¹³CHICKEN 4 contains more than 20 mathematical libraries that may be easy to port to CHICKEN 5, including an interface to LAPACK and an automatic differentiation library.

multi-precision arithmetic GMP, some GNU Scientific Library (GSL) programs (numerical integration, random numbers), an excellent plotting library and a machine learning libraries collection (although the project is at an early stage and only includes Decision Trees and k-Nearest Neighbors algorithms). Among external Racket packages not listed on the official repositories, racket-cas is a simple Computer Algebra System. Still, for both CHICKEN and Racket the number of libraries directly useful for numerical computing is relatively scarce compared to, e.g., more of 10000 packages currently tagged as scientific tools on the Python Package Index (PyPI) repository.

Chez Scheme lacks of an established 3rd-party libraries package manager. Gambit refers to repositories of portable Scheme libraries.¹⁴

Besides implementation-specific solutions, the Scheme Requests for Implementation (SRFI) offer several libraries written in portable Scheme (appendix A gives a fast overview to common SRFIs). Also, many other portable Scheme packages retrieved from different sources¹⁵ are available via the Akku.scm package manager [30], which aims at supporting at once several R⁶RS and R⁷RS Scheme implementations, including Chez Scheme. For instance, R6RS-AD is a R⁶RS-compliant library to perform automatic differentiation for both forward and reverse mode [32] that can be easily installed with Akku.scm. Another project worth mentioning is GNU Guix, a cross-platform generic package manager with a GNU Guile¹⁶ front end and a focus on providing reproducible software deployment for high-performance computing.

Furthermore, all of the languages here considered have a FFI to call procedures written in C (or in languages that obey the same calling conventions as C).¹⁷ CHICKEN even allows one to easily embed C code into Scheme source files and has high-level extensions to the core FFI (including a Python FFI), and the Gambit compiler can easily create shared C libraries.

Parallel computing.¹⁸ Chez Scheme supports native Posix threads on several architectures, although it lacks of high-level parallel forms.¹⁹ Racket has support for parallelization

¹⁴ Gerbil Scheme (<https://cons.io/>) aims at providing a battery-included Gambit featuring a package manager, gxpkg (still, only a dozen of libraries are currently available).

¹⁵ E.g., Snow (<http://snow-fort.org/>) and Scheme Spheres (<http://www.schemespheres.org/>).

¹⁶ The recently released GNU Guile version 3.0 features a just-in-time compiler that speeds it up considerably in some benchmark compared to previous versions, possibly making it another interesting Scheme implementation for scientific computing.

¹⁷ The Darkart project (<https://guenchi.github.io/Darkart/>) is also developing binary interfaces to let Chez Scheme calling libraries of several languages at the forefront of scientific computing (Python, Julia, OCaml, etc.)

¹⁸ For a recent Scheme dialect created for modern GPU computing, see Harlan [35].

¹⁹ For instance, as a notable example GNU Guile provides `parallel`, `letpar`, `par-map` and `par-for-each` as parallel versions of `values`, `let`, `map` and `for-each`, respectively.

with built-in futures and places, and with MPI and user-friendly parallel mapping (`pmap`) libraries.

Gambit only supports green threads (managed entirely by Gambit's runtime, not taking advantage of multiple processors) via the Termite language. Also CHICKEN only supports green threads (and has a package for Termite-like concurrency), but parallel computing is available via a package wrapping the Message Passing Interface (MPI).

Editing, debugging and profiling. All the implementations considered here provide read-eval-print loop (REPL) shells for interactive development.²⁰ Chez Scheme and Gambit REPL prompts support multi-line editing, name completion, history of previously entered expressions (and automatic indentation for Chez Scheme). CHICKEN requires additional programs like `r1wrap` or Eggs extensions for comfortable editing via the REPL.

Code is more easily edited via external programs, Emacs being a popular editor with built-in modes and external packages supporting Scheme editing (e.g., Geiser and Paredit).²¹ Racket features an exceptional Integrated Development Environment (IDE), DrRacket, with graphical helpers and integrated debugger and macro stepper. Racket also features IRacket, a kernel for Jupyter that, together with DrRacket, is a user-friendly option for novel schemers not accustomed to Emacs.

Besides Racket, also Chez Scheme, CHICKEN and Gambit have built-in debuggers and introspection tools. Chez Scheme and Racket have a thorough profilers, while CHICKEN profiler is rather limited (e.g., it is at least not obvious how to trace procedures defined internally to other procedures).

Source documentation tools. Although Chez Scheme does not have an established documentation tool, `stex` (<https://github.com/dybvig/stex>) is a Latex-like language to write documentation (separate from source code) with built-in commands for Scheme code display and it has been used to produce, among others, the HTML version of the Chez Scheme User's Guide. `ChezWEB` [36] provides a literate programming environment for Chez Scheme.

CHICKEN and its packages documentation are available via a manual and a wiki. The `chicken-doc` extension makes all the documentation available via command line and from the interpreter, providing good accessibility to both the core language and to those Eggs extensions uploaded on the CHICKEN infrastructure.²²

Gambit does not provide documentation tools. However, Gambit and Chez Scheme are supported by SLIB (<https://people.csail.mit.edu/jaffer/SLIB>) that features a nice inline

²⁰ Superficially, the REPL is roughly analogous to the Python/IPython shell.

²¹ Emacs Scheme modes are in many aspects more interactive and powerful than the Jupyter notebooks familiar to Python developers, although less user-friendly.

²² There is also a CHICKEN extension (the Hahn egg) for in-source documentation similar to Doxygen, but it is not widely used.

documentation language, Schmooz. Racket relies on Scribble, a powerful language that is the standard to document Racket projects.

4 Case study

In this section we present two case studies based on experience developed in machine learning and in cosmology, respectively. The first example shows the difficulties that an *impatient schemer* may encounter in adopting Scheme, while the second one shows how Scheme is an excellent candidate to help replacing a cumbersome software development process.

Let us first stress that Common Lisp has been a more popular Lisp dialect in machine learning and physics applications (see Appendix B). We are aware of only a few applications of Scheme to research in natural sciences (see Section 1, or [50] where CHICKEN Scheme was used to prototype a domain specific language for describing computational models of neuronal ionic currents). In the particular context of physics, Lisp dialects in general have never been a popular choice [47, 64, 65].

4.1 Classification of American Sign Language hand gestures

The Sign Language MNIST dataset²³ can be used to identify the American Sign Language letter represented by a hand picture. This is an image classification problem and a typical workflow involves the following steps [18, 29]:

1. Preprocess the data to make them suitable for use in numerical algorithms. This includes appropriate vectorization and normalization of the input data.
2. Build a machine learning model and refine it fine-tuning its parameters.

While very high-level libraries (like Pandas for Python) are popular for visualization and manipulation of data tables, the tasks described in the first step are typically fairly straightforward and can be carried out with basic operations on arrays. For instance, Keras [19] represents data in terms of tensors (i.e., multidimensional Numpy arrays) that are then passed to back ends like TensorFlow. The first step could be easily solved in Scheme, although there is no evident advantage over, e.g., Numpy (built upon robust and efficient legacy C and Fortran code).

The second step involves complicated algorithms such as Convolutional Neural Networks (CNN).²⁴ Runtime speed

²³<https://www.kaggle.com/datamunge/sign-language-mnist>. The database webpage also lists several Jupyter notebooks illustrating different solutions.

²⁴In other machine learning problems dealing with non-perceptive data simpler shallow learning algorithms may be a better solution. However, state-of-the-art software rely on highly optimized libraries, like XGBoost [17] for Gradient Boosting, to process large data sets that make algebraic operations very expensive. Moreover, ensemble models often outperform single ones, which may require combining, e.g., shallow and deep learning methods from different libraries.

and memory efficiency are paramount. To get an idea about the order of magnitude involved in training most advanced models, the VGG16 network [52] (a popular CNN architecture) depends on $O(10^8)$ parameters. Algebraic operations on arrays involving these parameters are required to be carried out several times during training of the model, which may need hundreds of thousands of iterations. Apart from runtime and memory efficiency (and support for appropriate GPU hardware), modern libraries provide flexible and interactive front ends. For instance, given the costly training process, it is important to be able to tune its hyperparameters (e.g., the learning rate of gradient descent) and to stop training interactively. Besides allowing this, Keras provides a functional interface to build new complex deep learning architectures or, e.g., to define generators such that data is automatically processed at each training epoch (*augmenting* an image data set by applying random image transformations such as rotations, stretches, etc. is a common technique to contrast overfitting).

The best way we identify for an *impatient little schemer* to face the task here described is to interface Scheme with existing deep learning libraries. Given that this is likely to require non-trivial work to build a compatibility layer (e.g., via the FFI), there is no clear advantage over using directly existing popular and well supported libraries.

Let us stress that, while this case is representative of a wide range of applications, it is of course not a fair description of many other use cases. For instance, according to the 2019 Kaggle's report about the state of machine learning, the algorithms most used by data scientists overall are actually simple linear or logistic regression [39]. Scheme adoption may still be advantageous for these situations. Furthermore, given the high level of interactivity required in training processes, deep learning could still profit from the contribution of a *patient seasoned schemer*. In the next session we also discuss how Scheme expressiveness is particularly well suited to implement and manipulate expressions using automatic differentiation, at the base of fundamental backpropagation algorithms used to train efficiently neural networks [7, 45].

4.2 Modeling the Large Scale Structure of the universe

The large-scale distribution of galaxies in the Universe can be modeled based on cosmological perturbation theory. From the numerical point of view, this may involve the computation of hundreds of time-consuming (multidimensional) numerical integrals. A numerical project facing this task is composed by [20]:

- A core C++ library that deals with the numerically requiring integration algorithms. It also uses SWIG to automatize the building of a Python wrapper around the C++ library.

- A high-level Python library that profits from the C++ core wrapper and provides procedures corresponding directly to relevant physical statistical quantities. Users are supposed to operate through (and possibly extend) this high-level interface.

The code is suitable for massive parallelization. However, due to the GIL of the CPython interpreter (see section 2), neither built-in nor external Python libraries support the serialization of objects created via the automatic Python wrapper of the C++ back end.²⁵

The C++ library relies on design patterns that, although standard (e.g., factory methods), would be trivially replaced with much simpler solutions in Python. However, the latter option would have prohibitively slowed down the core module. Furthermore, part of the integrals solved by the C++ modules had to be manipulated symbolically to reach a tractable form. Given that the analytical expression involved hundreds of such integrals, this was carried out procedurally with the SymPy [43] Python library that performed mathematical simplification and returned C++ code.

A unified Scheme library, relying on the same external dependencies via the C FFI, would have considerably removed several layers of cumbersome software engineering solutions reducing precious development time. As discussed in the previous sections, Scheme is a flexible language to manipulate mathematical operations and benefits from highly efficient compilers. Contrary to the previous example, the project itself does not require particularly complicated algorithms and would have been within the reach of an *impatient little schemer*.

As an illustration of Scheme expressiveness, let us discuss more in details a simplified version of the symbolic computation mentioned above. A differential operator such as

$$P = D^2 + D, \quad (1)$$

where $D^n f$ is the n -th order derivative of a function f , can lead to cumbersome expressions when applied to functions like

$$f_{mn}(x) = j_m(x)j_n(x)x^2, \quad (2)$$

where $j_n(x)$ is the spherical Bessel function. Even in this simple case it is somewhat cumbersome, although straightforward, to compute by hand expressions like $Pf(x)$ or $P^2f(x)$. The actual case [20] involved hundreds of polynomials and trigonometric functions after symbolic simplification. Numerical (finite differences), symbolic or automatic differentiation can be used to carry out computations programmatically. However, the first option typically leads to large errors (much larger than working precision), and the second one becomes inefficient when leading to large symbolic expressions that take correspondingly long to evaluate. Automatic

²⁵A simple alternative solution applicable to this case is to manually launch several processes in parallel via shell scripts, a more involved one would require more tweaking of the automatically generated Python wrapper.

differentiation (AD) [32] algorithms compute derivatives at working precision of any function using at most a small constant factor more arithmetic operations than the original function, solving the issues related to the other strategies. AD libraries are available for C and C++, but they typically rely on involved implementations and require particular type definitions and redefinition (or compatibility layers) of functions not already implemented in terms of basic AD operators [4, 7, 16, 34, 41, 45, 63]. Due to these issues, the project under consideration [20] opted for more inefficient symbolic computation.²⁶ On the other hand, expressive AD libraries are available for Scheme. For instance, the R6RS-AD²⁷ is a relatively simple implementation written in portable R6RS code (hence easy to import, e.g., in Chez Scheme) of both *forward* and *reverse* (useful for machine learning) mode AD, also including convenience functions to compute gradients. The differential operator in equation 1 can be defined as follows (derivative-F provides the forward mode AD):

```
(import (AD))

(define (P proc)
  (define (diff-operator x)
    (+ ((derivative-F (derivative-F proc)) x)
      ((derivative-F proc) x)))
  diff-operator)
```

For simplicity, let us consider the $f_{01}(x)$ function involving spherical Bessel functions of order zero and one.²⁸

```
(define (spherical-j0 x)
  (if (= x 0) 1 (/ (sin x) x)))

(define (spherical-j1 x)
  (if (= x 0)
      0
      (- (/ (sin x) x x) (/ (cos x) x))))

(define (func x)
  (* (spherical-j0 x) (spherical-j1 x) x x))
```

Then, expressions like $Pf_{01}(x)$ and $P^2f_{01}(x) = P(Pf_{01}(x))$ are computed for, e.g., $x = 1$, respectively, as

²⁶Simple Computed Algebra Systems (CAS) are available for R6RS standards (<https://github.com/dharmatech/mpl>) and for Racket (racket-cas and rascas packages). Alternatively, Maxima [42] is an affirmed CAS written in Common Lisp that can output Lisp expression easily ported to Scheme.

²⁷<https://github.com/qobi/R6RS-AD>.
²⁸High-order spherical Bessel functions can be defined using recursion relations and asymptotic expansions leading to more involved implementations. We are not aware of a Scheme library providing spherical Bessel functions (or ordinary Bessel functions of half-integer order). However, the project here discussed [20] used an internal implementation due to issues that arise in popular libraries [4, 31] at large Bessel orders, so a Scheme implementation wouldn't have led to overhead in this direction. Furthermore, spherical Bessel functions are typically not included in AD libraries of other languages either and they need to be redefined in terms of basic AD operators and functions.

```
((P func) 1)
((P (P func)) 1)
```

The abstraction level reflects the mathematical expressions under consideration and it is comparable, e.g., with high-level Python AD packages like JAX [13], but it does not exclude the possibility of using directly high performance languages like Chez Scheme. The notation could be even further refined, e.g., to interpret multiplication of operators as composition, so that $(P (P \text{ func}))$ may alternatively be written as $((\text{expt } D \ 2) \text{ func})$ [56, 58, 59]. On more general grounds, Scheme shares only with other Lisp dialects the possibility to define sophisticated syntactic extensions useful, e.g., to reorder evaluation via new binding constructs not defined as core syntactic forms of a given language and, more broadly, to define convenient domain specific languages.

5 Conclusions

In this report we gave a coarse-grained review of modern scientific computing, drawing from needs in physics and in machine learning. We argued that Scheme is a simple and extremely expressive language well suited for numeric and symbolic computations. The adoption of efficient Scheme implementations in place of other popular high-level languages (like Python and R) may restrict the cumbersome need of embedding high-performance, low-level languages (e.g., C, C++ and Fortran) to legacy code.

Specifically, we compared Chez Scheme, CHICKEN, Gambit and Racket focusing on currently available software engineering tools. Chez Scheme and Gambit have speed competitive with low-level languages like C, but are supported by a relatively small community and lack of external libraries. CHICKEN and Racket also have good performance, typically only some factor (≤ 10) slower than Chez Scheme and Gambit, good community support and a large set of libraries (yet, numerical libraries are by far more scarce than for popular languages like Python). Racket provides the most user-friendly editing interface and documentation resources for novel schemers, together with the most comprehensive tools for scientific computing (including libraries for linear algebra, numerical integration, plotting, symbolic computations and for high-level parallel forms). Nothing forbids to take advantage from synergies among complementary systems. For instance, Chez Scheme could be employed in a computationally intensive program whose output is then analyzed via the Racket plot library.

The lack of numerical libraries is the main practical issue to promptly replace popular high-level languages like Python and R. Systematic efforts are required to wrap code written in other languages than Scheme for applications that rely on enormous libraries with several functionalities, providing consistent interfaces among different projects (e.g., Numpy, Keras, Scikit-learn for machine learning). As we discussed in a computer vision use case, this is not an optimal

perspective. Let us stress that while the patient work to build a Scheme library competitive with modern machine learning tools is besides the experience of this report, it would be a valuable effort. For instance, we argued that Scheme is better suited than languages like C and C++ to implement and use algorithms like automatic differentiation, at the base of modern neural network libraries [7, 45], and to manipulate mathematical expressions. Furthermore, machine learning benefits from interactive development with running training processes (e.g., to monitor and interrupt the training process or to change training hyperparameters) and Scheme systems provide competitive tools also in this direction. Scheme and its unique syntax transformers (still unchallenged by other programming languages than Lisp dialects) are also an interesting option, e.g., to boost current efforts [19] to implement reusable neural network architectures as composable programs in metalearning systems or to base machine learning on program synthesis (generating programs, for instance a neural network, by exploring a large space of possible programs).

In a use case based on a cosmology application we discussed how scientific research often requires to build simple components that may eventually be combined into more complex systems. In such situations Scheme systems already offer an excellent opportunity to rely at once on highly educational, expressive and efficient systems. Still, there is ample room for improvements. The most urgent advancement would seem to provide comprehensive base Scheme libraries (such as Numpy, Scipy and Sympy for Python) with support for modern parallel computing needs. This would provide common ground to ease and encourage development of further numerical applications, reducing the available libraries gap with popular high-level languages like Python and R. The scientific community could then profit more promptly from sophisticated but simple Scheme languages, overcoming issues related to popular but less expressive or efficient languages.

Acknowledgments

I thank Jason Hemann for encouraging this report and for insightful discussions, and the anonymous reviewers for constructive comments.

References

- [1] [n. d.]. Gambit Benchmarks. <https://www.iro.umontreal.ca/~gambit/bench.html>. Accessed: 2020-05-26.
- [2] [n. d.]. Larceny Benchmarks. <http://www.larcenists.org/benchmarksAboutR6.html>. Accessed: 2020-05-26.
- [3] [n. d.]. Scheme Benchmarks. <https://ecraven.github.io/r7rs-benchmarks/>. Accessed: 2020-05-26.
- [4] 1997–2019. Boost C++ Libraries. <http://www.boost.org/>. Accessed: 2020-05-26.
- [5] 2010. Racket Benchmarks. <https://blog.racket-lang.org/2010/01/benchmarks.html>. Accessed: 2020-05-26.

- [6] 2020. The Computer Language Benchmarks Game. <https://benchmarksgame-team.pages.debian.net/benchmarksgame/fastest/racket.html>. Accessed: 2020-05-26.
- [7] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [8] B.P. Abbott et al. 2016. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.* 116, 6 (2016), 061102. <https://doi.org/10.1103/PhysRevLett.116.061102> arXiv:gr-qc/1602.03837
- [9] H. Abelson, A. A. Berlin, J. Katzenelson, W. H. McAllister, G. J. Rozas, G. J. Sussman, and J. Wisdom. 1991. *The Supercomputer Toolkit: A General Framework for Special-purpose Computing*. Technical Report. MIT AI Lab. <ftp://publications.ai.mit.edu/ai-publications/pdf/AIM-1329.pdf>
- [10] H. Abelson and G.J. Sussman. 1996. *Structure and Interpretation of Computer Programs* (2nd ed.). MIT Press, Cambridge, MA, USA.
- [11] Kim Albertsson et al. 2018. Machine Learning in High Energy Physics Community White Paper. *J. Phys. Conf. Ser.* 1085, 2 (2018), 022008. <https://doi.org/10.1088/1742-6596/1085/2/022008> arXiv:physics.comp-ph/1807.02876
- [12] Andy Wingo. 2013. An opinionated guide to scheme implementations. <https://wingolog.org/archives/2013/01/07/an-opinionated-guide-to-scheme-implementations>. Accessed: 2020-05-26.
- [13] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, and Skye Wanderman-Milne. 2018. JAX: composable transformations of Python+NumPy programs. <http://github.com/google/jax>
- [14] Rene Brun and Fons Rademakers. 1996. ROOT - An Object Oriented Data Analysis Framework. <http://root.cern.ch/>, In AIHENP'96 Workshop, Lausanne. *Nucl. Inst. & Meth. in Phys. Res. A* 389, 81–86.
- [15] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 108–122.
- [16] Bob Carpenter, Matthew D. Hoffman, Marcus Brubaker, Daniel Lee, Peter Li, and Michael Betancourt. 2015. The Stan Math Library: Reverse-Mode Automatic Differentiation in C++. arXiv:cs.MS/1509.07164
- [17] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, New York, NY, USA, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [18] Francois Chollet. 2017. *Deep Learning with Python* (1st ed.). Manning Publications Co., USA.
- [19] François Chollet et al. 2015. Keras. <https://keras.io>.
- [20] Enea Di Dio, Ruth Durrer, Roy Maartens, Francesco Montanari, and Obinna Umeh. 2019. The Full-Sky Angular Bispectrum in Redshift Space. *JCAP* 04 (2019), 053. <https://doi.org/10.1088/1475-7516/2019/04/053> arXiv:astro-ph.CO/1812.09297
- [21] R. Kent Dybvig. 2009. *The Scheme Programming Language, 4th Edition* (4th ed.). The MIT Press.
- [22] Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. 2018. *How to Design Programs: An Introduction to Programming and Computing*. The MIT Press.
- [23] Matthew Flatt, Caner Derici, R. Kent Dybvig, Andrew W. Keep, Gustavo E. Massaccesi, Sarah Spall, Sam Tobin-Hochstadt, and Jon Zeppieri. 2019. Rebuilding Racket on Chez Scheme (Experience Report). *Proc. ACM Program. Lang.* 3, ICFP, Article 78 (July 2019), 15 pages. <https://doi.org/10.1145/3341642>
- [24] Message P Forum. 1994. *MPI: A Message-Passing Interface Standard*. Technical Report. USA.
- [25] D.P. Friedman, W.E. Byrd, O. Kiselyov, J. Hemann, D. Bibby, G.L. Steele, G.J. Sussman, and R.A. Kowalski. 2018. *The Reasoned Schemer*. MIT Press.
- [26] D.P. Friedman, M. Felleisen, and D. Bibby. 1996. *The Little Schemer*. MIT Press.
- [27] D.P. Friedman, M. Felleisen, D. Bibby, and G.L. Steele. 1996. *The Seasoned Schemer*. MIT Press.
- [28] Erann Gat. 2000. Point of view: Lisp as an alternative to Java. *intelligence* 11 (2000), 21–24. Issue 4. <https://doi.org/10.1145/355137.355142>
- [29] Aurélien Géron. 2017. *Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* (1st ed.). O'Reilly Media, Inc.
- [30] Göran Weinholt. 2019. Akku.scm. <https://akkuscm.org/>. Accessed: 2020-05-26.
- [31] Brian Gough. 2009. *GNU Scientific Library Reference Manual - Third Edition* (3rd ed.). Network Theory Ltd.
- [32] Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. 2015. Automatic differentiation in machine learning: a survey. *arXiv e-prints*, Article arXiv:1502.05767 (Feb. 2015), arXiv:1502.05767 pages. arXiv:1502.05767
- [33] Klaus Havelund, Mike Lowry, Seungjoon Park, Charles Pecheur, John Penix, Willem Visser, Jon L. White, and Recom Technologies. 2000. Formal Analysis of the Remote Agent Before and After Flight.
- [34] Robin J. Hogan. 2014. Fast Reverse-Mode Automatic Differentiation Using Expression Templates in C++. *ACM Trans. Math. Software* 40, 4 (jun 2014), 26:1–26:24. <http://doi.acm.org/10.1145/2560359>
- [35] Eric Holk, Ryan Newton, Jeremy Siek, and Andrew Lumsdaine. 2014. Region-based memory management for GPU programming languages. *ACM SIGPLAN Notices* 49 (12 2014), 141–155. <https://doi.org/10.1145/2714064.2660244>
- [36] Aaron W. Hsu. 2011. Hygienic Literate Programming: Lessons from ChezWEB. In *Workshop on Scheme and Functional Programming*.
- [37] IEEE Spectrum. 2019. The Top Programming Languages 2019. <https://spectrum.ieee.org/computing/software/the-top-programming-languages-2019>. Accessed: 2020-05-26.
- [38] John W. Cowan. 2016. Implementation Contrasts. <https://smallr7rs.org/wiki/ImplementationContrasts/>. Accessed: 2020-05-26.
- [39] Kaggle. 2019. Kaggle's State of Data Science and Machine Learning 2019. <https://www.kaggle.com/c/kaggle-survey-2019>
- [40] Richard Kelsey, William Clinger, and Jonathan Rees (eds.). 20007. *Revised⁶ Report on the Algorithmic Language Scheme*. Technical Report. <http://www.r6rs.org/>
- [41] Allan M. M. Leal et al. 2018. autodiff, a modern, fast and expressive C++ library for automatic differentiation. <https://autodiff.github.io> <https://autodiff.github.io>
- [42] Maxima. 2014. Maxima, a Computer Algebra System. Version 5.34.1. <http://maxima.sourceforge.net/>
- [43] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, Amit Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. 2017. SymPy: symbolic

computing in Python. *PeerJ Computer Science* 3 (Jan. 2017), e103. <https://doi.org/10.7717/peerj-cs.103>

[44] OpenMP Architecture Review Board. 2018. OpenMP Application Program Interface Version 5.0. <http://www.openmp.org/>

[45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>

[46] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.

[47] P. Prabhu, H. Kim, T. Oh, T. B. Jablin, N. P. Johnson, M. Zoufaly, A. Raman, F. Liu, D. Walker, Y. Zhang, S. Ghosh, D. I. August, J. Huang, and S. Beard. 2011. A survey of the practice of computational science. In *SC '11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–12.

[48] Lutz Prechelt. 1999. Technical opinion: Comparing Java vs. C/C++ efficiency differences to interpersonal differences. *Commun. ACM* 42 (1999), 109–112. Issue 10. <https://doi.org/10.1145/317665.317683>

[49] Lutz Prechelt. 2000. An empirical comparison of seven programming languages. *IEEE Computer* 33, 10 (October 2000), 23–29. <https://doi.org/10.1109/2.876288>

[50] Ivan Raikov and Erik De Schutter. 2012. The Layer-Oriented Approach to Declarative Languages for Biological Modeling. *PLoS computational biology* 8 (05 2012), e1002521. <https://doi.org/10.1371/journal.pcbi.1002521>

[51] Alex Shinn, John Cowan, and Arthur A. Gleckler (eds.). 2013. *Revised⁷ Report on the Algorithmic Language Scheme*. Technical Report. <http://www.r7rs.org/>

[52] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv e-prints*, Article arXiv:1409.1556 (Sept. 2014), arXiv:1409.1556 pages. [arXiv:cs.CV/1409.1556](http://arxiv.org/abs/1409.1556)

[53] IEEE Computer Society. 1991. IEEE Standard for the Scheme Programming Language. *IEEE Std. 1178-1990* (1991).

[54] Michael Sperber, R. Kent Dybvig, Matthew Flatt, and Anton Van Straaten (eds.). 2007. *Revised⁶ Report on the Algorithmic Language Scheme*. Technical Report. <http://www.r6rs.org/>

[55] G.J. Sussman and J. Wisdom. 1992. Chaotic Evolution of the Solar System. *Science* 257 (1992), 56–62. Issue 5066. <https://doi.org/10.1126/science.257.5066.56>

[56] G.J. Sussman and J. Wisdom. 2001. *Structure and Interpretation of Classical Mechanics*. MIT Press, Cambridge, MA, USA. https://mitpress.mit.edu/sites/default/files/titles/content/sicm_edition_2/book.html

[57] G.J. Sussman and J. Wisdom. 2002. *The Role of Programming in the Formulation of Ideas*. Technical Report. MIT AI Lab. <ftp://publications.ai.mit.edu/ai-publications/2002/AIM-2002-018.pdf>

[58] G.J. Sussman, J. Wisdom, and W. Farr. 2013. *Functional Differential Geometry*. MIT Press. <https://mitpress.mit.edu/books/functional-differential-geometry>

[59] G. J. Sussman. 1997–2019. Scmutils. <https://groups.csail.mit.edu/mac/users/gjs/6946/linux-install.htm>. Accessed: 2020-05-26.

[60] TIOBE. 2020. TIOBE Index for May 2020. <https://www.tiobe.com/tiobe-index/>. Accessed: 2020-05-26.

Table 2. Reference to common SRFIs. We group algorithms into a broad category and provide a short description for each SRFI, identified by a number.

Category	Description	SRFI
Basic libraries	Lists	1
	Vectors	133
	Strings	13 (and 130)
	Sorting	132
Sets and maps	Sets and bags	113 (or 1)
	Character sets	14
	Hash tables	125
Immutability	Immutable pairs/lists	116
	Random access pairs/lists	101
	Immutable dequeues	134
	Immutable text	135
Laziness	Generators	121
	Lazy sequences	127
	Streams	41
Miscellaneous	Boxes	111
	Comparators	128
	Comprehensions	42
	Ephemerals	124
	List queues	117
	Regular expressions	115
	Titlecase	129

[61] Didier Verna. 2006. Beating C in Scientific Computing Applications. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.150.640&rep=rep1&type=pdf>. In *European Conference on Object-Oriented Programming (ECOOP)*, Nantes, France.

[62] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

[63] A. Walther and A. Griewank. 2012. Getting started with ADOL-C. In *Combinatorial Scientific Computing*, U. Naumann and O. Schenk (Eds.). Chapman-Hall CRC Computational Science, Chapter 7, 181–202.

[64] Bebo White. 1991. *The Comparison and Selection of Programming High Energy Physics Applications*. Technical Report TN-91-6.

[65] I E Zacharov. 1990. Programming languages at LHC. CERN-ECP-90-019. CERN-ECP-90-19 (Dec 1990), 11 p. <https://doi.org/10.5170/CERN-1990-010-V-3.41>

A Reference to common SRFI libraries

Scheme Requests for Implementation (SRFI) provide Scheme users portable, useful code not necessarily present in typically minimalist language implementations. SRFIs are well documented on the dedicated website (<https://srfi.schemers.org/>), but to the novice schemer it may not be clear what is the best or most idiomatic solution for a given application. Table 2 reviews common SRFI libraries.²⁹ Most of them are basic algorithm that are also useful for numeric application.³⁰

B Other Lisp dialects

An obvious alternative to Scheme as a Lisp dialect is Common Lisp. Common Lisp has excellent, efficient (runtime usually of the same order as C or C++, only a few factors slower [6], see also Verna [61] for scientific computing benchmarks) implementations like the Steel Bank Common Lisp (SBCL) compiler (similar performance range as Chez Scheme or Gambit), a large community, an established package manager to easily access an extensive library repository (quicklisp) that receives frequent new contributions. A seminal application of Common Lisp to scientific computing is the real-time debug and fix of a race condition when the Deep Space 1 NASA mission was already operating [33]. A more recent application is the top score winning solution of the Higgs Boson challenge organized by CERN and Kaggle in 2014, entirely written in Common Lisp.³¹ However, Scheme is conceptually simpler than Common Lisp thanks to only one namespace for functions and variables (among other differences). E.g., Scheme only needs `define` for both variables and functions, while Common Lisp distinguishes between `defun`, `setf`, `defvar`, etc., and Scheme does not need a call to `funcall` for function application.

Another Lisp dialect worth considering is Clojure. It is actively maintained, the community is significantly larger than those around Scheme implementations (although relatively small compared to mainstream languages like Python) and provides excellent documentation and support. While being a battery-included modern language, it feels much simpler than Common Lisp. In some aspect its design choices may appear more immediate than Scheme, e.g., to Python developers (see for instance the implementation of hash tables and sets). Projects building and automatic dependency resolution are managed via Leiningen (build automation tool ideal for

medium size and large projects). Although there are interesting numerical projects like the Neanderthal library for linear algebra operations on CPU and GPU (based on BLAS and LAPACK standardization and optimized for Clojure), the number of native Clojure packages is not competitive with Python. However, running on the Java Virtual Machine, it is easy to access the large number of Java numerical libraries. E.g., the state-of-the-art TensorFlow library for deep learning provides a Java front end. (However, it should be noted that Java libraries are not as common in physics as C libraries.)

²⁹The list is strongly inspired by ongoing discussion for the large R⁷RS report.

³⁰For instance, list comprehensions are not implemented in the standard library of Scheme implementations discussed in the main text (except for Racket that does provide similar forms, like `for/list`). However, they are widely used in Python, hence SRFI 42 may serve as common ground for novel schemers. It is interesting to note i) the broader scope and flexibility of SRFI 42 compared to Python list comprehensions, and ii) the fact that in Python list comprehensions are preferable to other iterative forms for performance reason (not the case in Scheme).

³¹<https://github.com/melisgl/higgsml>.

On Teaching Type Systems as Macros (Lightning Talk)

Youyou Cong
Tokyo Institute of Technology
Tokyo, Japan
cong@c.titech.ac.jp

Naoya Furudono
Tokyo Institute of Technology
Tokyo, Japan
furudono.n.aa@m.titech.ac.jp

Hidehiko Masuhara
Tokyo Institute of Technology
Tokyo, Japan
masuhara@is.titech.ac.jp

ACM Reference Format:

Youyou Cong, Naoya Furudono, and Hidehiko Masuhara. 2020. On Teaching Type Systems as Macros (Lightning Talk). In *Scheme '20*. ACM, New York, NY, USA, 1 page.

Type systems are one of the fundamental things that everyone in the programming languages community must study. In our research group, we run a semester-long type systems seminar every two years, covering essential chapters of the traditional TaPL book [2]. This year, however, we are attempting a different approach, namely teaching type systems through programming in the TURNSTILE language of [1]. TURNSTILE is a Racket DSL for creating typed programming languages, built under the slogan "type systems as macros". More specifically, it allows the user to define typing rules

in the familiar, derivation-like syntax, and to reuse Racket's infrastructure for type checking and evaluation.

We report our experience in teaching type systems as macros to an undergraduate student who joined our research group this spring (the second author). In particular, we show how the student used TURNSTILE to build simple type-andeffect systems, and what difficulties the student had in the course of implementation.

References

- [1] CHANG, S., KNAUTH, A., AND GREENMAN, B. Type systems as macros. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages* (2017), pp. 694–705.
- [2] PIERCE, B. C. *Types and programming languages*. MIT press, 2002.

Scheme '20, August 28, 2020, Online
2020.

Designing a Programming Environment Based on the Program Design Recipe (Lightening Talk)

Junya Nose
Tokyo Institute of Technology
Tokyo, Japan
junya.nose@prg.is.titech.ac.jp

Youyou Cong
Tokyo Institute of Technology
Tokyo, Japan
cong@c.titech.ac.jp

Hidehiko Masuhara
Tokyo Institute of Technology
Tokyo, Japan
masuhara@is.titech.ac.jp

ACM Reference Format:

Junya Nose, Youyou Cong, and Hidehiko Masuhara. 2020. Designing a Programming Environment Based on the Program Design Recipe (Lightening Talk). In *Scheme '20*. ACM, New York, NY, USA, 1 page.

The *program design recipe*, introduced by Felleisen et al. [2018] in their textbook *How to Design Programs*, is a step-by-step procedure that solves a problem by programming. The design recipe encourages the programmer, after examining the problem statement, to analyze the data, create input-output examples, and develop a template, instead of immediately starting coding. This helps reduce various kinds of errors, such as non-exhaustive conditional expressions and infinite loops.

One problem in programming with the design recipe is the lack of error-checking support for steps other than coding. For instance, in DrRacket¹, the programmer cannot check the correctness of data definitions or templates, as they are not written as runnable Racket programs.

We propose a programming environment based on the design recipe. The environment covers the whole process of the design recipe, including data analysis and template construction, with an IDE-like user interface. We also create a domain-specific language that allows systematic design of the user interface and error checking algorithms. Although it is only partly implemented at this point, we hope to receive suggestions on the design of our environment, especially from the workshop participants who are teaching with the design recipe.

References

Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. 2018. *How to design programs: an introduction to programming and computing*. MIT Press.

¹<https://racket-lang.org/>

Scheme '20, August 28, 2020, Online
2020.

Programming with Petri Nets to Reason about Concurrency

Julien Lepiller
Yale University

ACM Reference Format:

Julien Lepiller. 2020. Programming with Petri Nets to Reason about Concurrency. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/nnnnnnn>. nnnnnnn

Functional programming languages make it easy to reason about data processing because the lack of side effects makes local reasoning possible. Introducing parallelism might break such reasoning as data is passed from a thread to another in a way that is not structurally obvious. Petri nets on the other hand are well suited to represent the global state of a parallel program [2, 5] and how different parts of it can interact. We propose a Guile library, `guile-petri`, that combines them to get the best of both worlds.

Previous work [1, 3, 4] propose to use petri nets either for a domain-specific application or with object-oriented languages. We advocate for a similar approach where a token represents a concrete value, transitions are Guile procedures and places are communication channels where tokens wait until they can be processed.

Using a library instead of a creating a new languages gives us the possibility of using both a graphical and a textual representation and edition tool. Contrary to visual programming approaches, a Petri net can be created or modified programmatically, giving more power to the programmer.

Acknowledgement

Part of this work was completed while working on the grant supported by the National Science Foundation under Grant No. CCF-1715387.

References

- [1] Lawrence Cabac, Daniel Moldt, and Heiko Rölke. 2003. A Proposal for Structuring Petri Net-Based Agent Interaction Protocols. In *Applications*

and Theory of Petri Nets 2003, Wil M. P. van der Aalst and Eike Best (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 102–120.

- [2] M. Minas and G. Frey. 2002. Visual PLC-programming using signal interpreted Petri nets. In *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, Vol. 6. 5019–5024 vol.6.
- [3] M. Usher and D. Jackson. 1998. A Petri net based visual programming language. In *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.98CH36218)*, Vol. 1. 107–112 vol.1.
- [4] Rüdiger Valk. 2001. *Concurrency in Communicating Object Petri Nets*. Springer Berlin Heidelberg, Berlin, Heidelberg, 164–195. https://doi.org/10.1007/3-540-45397-0_5
- [5] Rob van Glabbeek and Frits Vaandrager. 1987. Petri net models for algebraic theories of concurrency. In *PARLE Parallel Architectures and Languages Europe*, J. W. de Bakker, A. J. Nijman, and P. C. Treleaven (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 224–242.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Author Index

Cong, Youyou	77, 78
Darragh, Pierce	3
Eide, Eric	3
Feeley, Marc	51
Fleishhacker, Jonah	2
Furudono, Naoya	77
Hatch, William G.	3
Henz, Martin	1
Khomtchouk, Bohdan	2
Lepiller, Julien	79
Masuhara, Hidehiko	77, 78
Montanari, Francesco	66
Nikishkin, Vladimir	14
Nose, Junya	78
Wrigstad, Tobias	1
Yvon, Samuel	51