

# Compiling Concurrency Correctly Verifying Software Transactional Memory

Liyang HU

A Thesis submitted for the degree of Doctor of Philosophy

School of Computer Science

University of Nottingham

June 2012



## Abstract

Concurrent programming is notoriously difficult, but with multi-core processors becoming the norm, is now a reality that every programmer must face. Concurrency has traditionally been managed using low-level mutual exclusion *locks*, which are error-prone and do not naturally support the compositional style of programming that is becoming indispensable for today's large-scale software projects.

A novel, high-level approach that has emerged in recent years is that of *software transactional memory* (STM), which avoids the need for explicit locking, instead presenting the programmer with a declarative approach to concurrency. However, its implementation is much more complex and subtle, and ensuring its correctness places significant demands on the compiler writer.

This thesis considers the problem of formally verifying an implementation of STM. Utilising a minimal language incorporating only the features that we are interested in studying, we first explore various STM design choices, along with the issue of compiler correctness via the use of automated testing tools. Then we outline a new approach to concurrent compiler correctness using the notion of bisimulation, implemented using the Agda theorem prover. Finally, we show how bisimulation can be used to establish the correctness of a low-level implementation of software transactional memory.



# Acknowledgements

Many have had a part to play in this production, and I cannot hope to enumerate them exhaustively. Nevertheless, I would like to begin by thanking everyone at the Functional Programming Laboratory in Nottingham who have made it such an interesting place, academically and socially. Conor McBride deserves a special mention for his multitudes of infectious ideas that started me on this dependently-typed journey, as do Ulf Norell and Nils Anders Danielsson for the years they have put into Agda and its standard library that underpins a large part of this work.

There were plenty of ups and downs in the process. I am eternally grateful to my flatmate Rebecca who had a large part in maintaining my sanity, my muse Ana whose company kept my spirits up through those seemingly endless hours of writing, and my friend Tom for sharing his inexhaustible enthusiasm with me. The numerous thoughtful gifts from Star and Cosmo are also warmly received. I am very fortunate to have my parents, whose support made my aspirations possible. Thank you all.

I appreciate the effort of Andy Gordon and Thorsten Altenkirch in undertaking the rewardless task of my examination. Tsuru Capital—as well as being an excellent place to work—afforded me much time and flexibility with which to complete my corrections. Last but certainly not least, I would like to express my gratitude towards my supervisor Graham Hutton for his guidance and enduring patience through all these years, without whose insights and encouragement this thesis would certainly have found itself in perpetual limbo.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Background . . . . .	1
1.1.1	A Brief Note on Moore’s Law . . . . .	1
1.1.2	The End of the Free Lunch . . . . .	2
1.1.3	A Brief Look at Parallel and Concurrent Computing . . . . .	3
1.1.4	A Supercomputer on Every Desk . . . . .	4
1.2	Approaches to Concurrent Software . . . . .	5
1.2.1	Concurrency versus Parallelism . . . . .	5
1.2.2	Counting: easy as 0, 1, 2 . . . . .	6
1.2.3	Shared Memory and Mutual Exclusion . . . . .	7
1.2.4	Example: Deadlock . . . . .	8
1.2.5	Message Passing and Implicit Synchronisation . . . . .	10
1.2.6	Software Transactional Memory . . . . .	12
1.3	Thesis Overview . . . . .	14
1.4	Contributions . . . . .	15
<b>2</b>	<b>Software Transactional Memory</b>	<b>19</b>
2.1	Database Transactions . . . . .	19
2.2	Transactional Memory . . . . .	21
2.2.1	Hardware Transactional Memory . . . . .	22

# CONTENTS

2.2.2	Software Transactional Memory . . . . .	23
2.3	Implementing Transactions . . . . .	24
2.3.1	Log-Based Transactions . . . . .	25
2.3.2	Alternative Approaches to Atomicity . . . . .	26
2.4	Haskell and Sequential Computation . . . . .	27
2.4.1	Monads for Sequential Computation . . . . .	27
2.4.2	Modelling Mutable State . . . . .	28
2.4.3	Monadic Properties . . . . .	31
2.4.4	Input, Output and Control Structures . . . . .	33
2.5	Haskell and Concurrent Computation . . . . .	34
2.6	Haskell and Software Transactional Memory . . . . .	37
2.7	Conclusion . . . . .	43
<b>3</b>	<b>Semantics for Compiler Correctness</b>	<b>45</b>
3.1	Semantics . . . . .	45
3.1.1	Natural Numbers and Addition . . . . .	45
3.1.2	Denotational Semantics . . . . .	46
3.1.3	Big-Step Operational Semantics . . . . .	47
3.1.4	Small-Step Operational Semantics . . . . .	48
3.1.5	Modelling Sequential Computation with Monoids . . . . .	50
3.2	Equivalence Proofs and Techniques . . . . .	51
3.2.1	Rule Induction . . . . .	51
3.2.2	Proofs of Semantic Equivalence . . . . .	52
3.3	Compiler Correctness . . . . .	57
3.3.1	A Stack Machine and Its Semantics . . . . .	57
3.3.2	Compiler . . . . .	59
3.3.3	Compiler Correctness . . . . .	60
3.4	Conclusion . . . . .	64



<b>4</b>	<b>Randomised Testing in Haskell</b>	<b>65</b>
4.1	Executable Semantics . . . . .	66
4.1.1	Denotational Semantics . . . . .	66
4.1.2	Big-Step Operational Semantics . . . . .	66
4.1.3	Small-Step Operational Semantics . . . . .	68
4.1.4	Virtual Machine . . . . .	70
4.2	Randomised Testing with QuickCheck and HPC . . . . .	71
4.2.1	Generating Arbitrary Expressions . . . . .	72
4.2.2	Semantic Equivalence and Compiler Correctness . . . . .	74
4.2.3	Coverage Checking with HPC . . . . .	77
4.3	Conclusion . . . . .	79
<b>5</b>	<b>A Model of STM</b>	<b>81</b>
5.1	A Simple Transactional Language . . . . .	81
5.1.1	Syntax . . . . .	82
5.1.2	Transaction Semantics . . . . .	83
5.1.3	Process Soup Semantics . . . . .	87
5.2	A Simple Transactional Machine . . . . .	91
5.2.1	Instruction Set . . . . .	91
5.2.2	Compiler . . . . .	91
5.2.3	Implementing Transactions . . . . .	92
5.2.4	Virtual Machine . . . . .	96
5.3	Correctness of the Implementation . . . . .	101
5.3.1	Statement of Correctness . . . . .	101
5.3.2	Validation of Correctness . . . . .	103
5.4	Conclusion . . . . .	105

# CONTENTS

<b>6</b>	<b>Machine-Assisted Proofs in Agda</b>	<b>107</b>
6.1	Introduction to Agda . . . . .	107
6.1.1	Data and Functions . . . . .	108
6.1.2	Programs as Proofs and Types as Predicates . . . . .	109
6.1.3	Dependent Types . . . . .	110
6.1.4	Equality and its Properties . . . . .	113
6.1.5	Existentials and Dependent Pairs . . . . .	115
6.1.6	Reflexive Transitive Closure . . . . .	117
6.2	Agda for Compiler Correctness . . . . .	120
6.2.1	Syntax and Semantics . . . . .	120
6.2.2	Semantic Equivalence . . . . .	122
6.2.3	Stack Machine, Compiler, and Correctness . . . . .	124
6.3	Conclusion . . . . .	129
<b>7</b>	<b>Compiling Non-Determinism Correctly</b>	<b>131</b>
7.1	Existing Approach . . . . .	131
7.2	Related Work . . . . .	132
7.3	A Non-Deterministic Language . . . . .	134
7.3.1	Choice of Action Set . . . . .	136
7.4	Compiler, Virtual Machine and its Semantics . . . . .	136
7.5	Non-Deterministic Compiler Correctness . . . . .	137
7.6	Combined Machine and its Semantics . . . . .	139
7.7	Weak Bisimilarity . . . . .	141
7.8	The <b>elide-<math>\tau</math></b> Lemma . . . . .	145
7.9	Compiler Correctness . . . . .	147
7.9.1	Proof of <b>correctness</b> . . . . .	147
7.9.2	The <b>eval-left</b> Lemma . . . . .	149
7.9.3	The <b>eval-right</b> Lemma . . . . .	152

7.10	Conclusion . . . . .	156
<b>8</b>	<b>Compiling Concurrency Correctly</b>	<b>157</b>
8.1	The Fork Language . . . . .	157
8.1.1	Syntax and Virtual Machine . . . . .	157
8.1.2	Actions . . . . .	158
8.1.3	Semantics . . . . .	160
8.2	Combined Machines and Thread Soups . . . . .	161
8.3	Silent and Visible Transitions . . . . .	162
8.4	Bisimilarity . . . . .	165
8.5	Properties of Thread Soups . . . . .	166
8.5.1	Soups Concatenation Preserves Silent Transitions . . . . .	166
8.5.2	Partitioning Silent Transitions . . . . .	167
8.5.3	Partitioning a Non-Silent Transition . . . . .	167
8.5.4	Dissecting a Visible Transition . . . . .	168
8.5.5	Extracting the Active Thread . . . . .	169
8.6	The <b>elide-<math>\tau</math></b> Lemma . . . . .	170
8.7	Soup Concatenation Preserves Bisimilarity . . . . .	172
8.8	Compiler Correctness . . . . .	175
8.9	Conclusion . . . . .	177
<b>9</b>	<b>Transaction Correctness</b>	<b>179</b>
9.1	The Atomic Language . . . . .	179
9.1.1	Syntax . . . . .	179
9.1.2	Heaps and Variables . . . . .	180
9.1.3	Stop-the-World Semantics . . . . .	181
9.1.4	Transaction Logs and Consistency . . . . .	184
9.1.5	Log-Based Semantics . . . . .	187

## CONTENTS

9.2	Combined Semantics and Bisimilarity . . . . .	190
9.2.1	Combined Semantics . . . . .	190
9.2.2	Bisimilarity of Semantics . . . . .	191
9.2.3	Definition of Correctness . . . . .	193
9.3	Reasoning Transactionally . . . . .	193
9.3.1	Consistency-Preserving Transitions . . . . .	193
9.3.2	Heaps and Logs Equivalence . . . . .	196
9.3.3	Post-Commit Heap Equality . . . . .	199
9.4	Transaction Correctness . . . . .	200
9.4.1	Completeness of Log-Based Transactions . . . . .	201
9.4.2	Soundness of Log-Based Transactions . . . . .	203
9.5	Bisimilarity of Semantics . . . . .	209
9.5.1	Addition . . . . .	209
9.5.2	Right Evaluation . . . . .	210
9.5.3	Left Evaluation . . . . .	212
9.5.4	Transactions . . . . .	214
9.5.5	Putting It Together . . . . .	217
9.6	Conclusion . . . . .	217
<b>10</b>	<b>Conclusion</b>	<b>219</b>
10.1	Retrospection . . . . .	219
10.2	Summary of Contributions . . . . .	221
10.3	Directions for Further Research . . . . .	222

# Chapter 1

## Introduction

In this chapter we set the scene for this thesis. We begin with a brief background on the history of concurrent computing, and the concepts and issues involved in developing software that takes advantage of the additional computational capability. We then describe a number of mainstream approaches to constructing concurrent software, along with the transactional memory approach that is the focus of this thesis, illustrated with some simple examples. We conclude with a synopsis of each chapter and a summary of our primary contributions.

### 1.1 Background

#### 1.1.1 A Brief Note on Moore's Law

Since the invention of the integrated circuit over 50 years ago and the subsequent development of the microprocessor, the number of transistors that engineers can manufacture on a single silicon chip for the same cost has been increasing at an exponential pace, roughly doubling every two years. This growth has been remained consistent, so much so that it has been informally codified as 'Moore's Law' [Moo65]. The related

statement<sup>1</sup> that “microprocessors *performance* roughly doubles every 18 months” has also held true, once we factor in the improved performance of individual transistors.

The popular understanding of Moore’s Law tends to be simplified to “computer speed doubles roughly every 18 months.” Until half a decade ago, this interpretation sufficed, because in order to pack more transistors next to each other, each one had to be made smaller. This in turn meant faster signal propagation between components, and so faster switching (or clock speeds), increasing performance. The implication of this is that one could expect the same piece of software to run twice as fast on the available hardware, 18 months down the line.

### 1.1.2 The End of the Free Lunch

Moore’s Law had become self-perpetuating as the industry assumed its truth to make projections for their technology roadmaps. By shrinking the size of individual transistors, not only were silicon manufacturers able to increase how many transistors can be economically placed on a single piece of silicon, they were also able to clock their processors at progressively higher frequencies due to reduced switching and signal propagation delays.

Sadly, miniaturisation has some undesirable side-effects: on the sub-micron scales of a modern microprocessor, current leakage due to quantum tunnelling effects across the on-chip insulation is very much detrimental to signal integrity: the smaller the features, the more power the integrated circuit needs to counteract these side-effects. This additional power must be dissipated in the form of waste heat, limiting the extent to which we can simply increase the clock speed. Indeed, some recent desktop processors expend up to a third [LS08] of their power solely to ensure accurate clock signal distribution to outlying areas of the silicon die, and expel as much as 150W of excess heat in an area less than 15mm<sup>2</sup>.

---

<sup>1</sup>Due to David House—an Intel executive at the time—as claimed by Moore:  
<http://news.cnet.com/2100-1001-984051.html>

Given the restriction that we cannot reasonably clock individual processors at increasingly higher speeds, how could we pack more computing power onto each silicon die? With the extra transistors afforded to us by Moore’s Law, the most obvious and easiest solution is to resort to *symmetric multiprocessing* (SMP), by fabricating multiple independent processors on the same die that share access to the same memory.

### 1.1.3 A Brief Look at Parallel and Concurrent Computing

The concept of SMP had been put to practice as early as the 1960s with the introduction of the Burroughs B5500 mainframe <sup>2</sup>. In the decades that followed, the entire computing industry resorted one by one to some form of parallelism in order to achieve their stated performance. First steps in this direction included the development of vector processors, where each instruction simultaneously operate on tens or sometimes hundreds of words of data. In today’s parlance, we refer to such architectures as *single instruction multiple data* (SIMD).

In contrast, a *multiple instruction multiple data* (MIMD) architecture comprises a number of independent processing units, each concurrently executing its own sequence of instructions. However, programming for multiprocessing systems is a task fraught with pitfalls, as Seymour Cray was alleged to have once quipped: “If you were ploughing a field, which would you rather use: two strong oxen or 1024 chickens?” His remark alludes to the challenge of synchronising a large number of independent processors with each one working on a small part of a larger problem while sharing the same working memory. It is much easier for people to work with a few oxen than to try and herd a large number of chickens.

Such systems with multiple, independent processors are therefore suited to domains involving very large datasets and have intrinsically parallelisable solutions that

---

<sup>2</sup>[https://wiki.cc.gatech.edu/folklore/index.php/Burroughs\\_Third-Generation\\_Computers](https://wiki.cc.gatech.edu/folklore/index.php/Burroughs_Third-Generation_Computers)

do not require much synchronisation or communication between individual processors. This correlates closely with many scientific computation and simulation problems, and the insatiable demand for computational power in these domains drove the development of massively parallel computers in the decades that followed. Even Cray eventually conceded to the fact that multiprocessing was inevitable, as the costs and resources required to breed and feed the two hypothetical oxen became prohibitive compared to breeding, feeding and herding 1024 chickens.

Meanwhile, as multiprocessor computers grew increasingly larger, it became difficult to maintain fast access to the same shared memory for all processor nodes. Cutting-edge systems therefore moved towards a more *non-uniform memory architecture* (NUMA), where each node had fast access to local memory, but access to globally shared or non-local data took correspondingly longer. The lessons learnt have strongly influenced the design today's high-performance hardware, even in the context of personal computing, as seen with the recent development of *general-purpose graphical processing units* (GPGPUs). On a more distributed and larger scale, Beowulf clusters of networked computers may be seen as a looser interpretation of the NUMA paradigm. Such systems typically form the workhorse behind many of today's large-scale scientific simulations, at least in part due to the fact that they can be and often are built from cheap and readily available commodity hardware.

#### 1.1.4 A Supercomputer on Every Desk

Moving back to the sphere of personal computing, it was not until the last decade that shared memory SMP computers managed to establish a presence. The explanation for this is twofold, yet complementary: on the one hand, the cost of motherboards that can accommodate multiple processor packages were significantly more expensive, and so were only sought after by users with specialist needs. On the other, the inability for predominantly single-threaded software—such as a game or a word-processor—to



take advantage of multiple processors, meant that the vast majority of users had no interest in resorting to SMP: simply waiting another 18 months had been sufficient.

However as raw processor speeds have plateaued in recent years, we can no longer afford to dismiss multiprocessor systems as being too difficult to program. Traditional supercomputing problems—large-scale, loosely-coupled computations such as physical simulations or graphical rendering—have been well-catered for, but they encompass only a small fraction of our day-to-day software usage. We need to figure out how to make concurrency accessible and manageable for everyday software.

## 1.2 Approaches to Concurrent Software

We resort to concurrency with the hope that the more computational power we have at our disposal, the faster our programs will run. How successfully this is in practice depends on how much concurrency we can exploit in our programs. How we model concurrency has a large influence on how we—as software engineers—think and reason about our concurrent programs, which in turn influences the ease with which we can exploit concurrency. This chapter demonstrates using a few examples why concurrent programming has such a reputation for being difficult, then review some of the basic models of concurrency.

### 1.2.1 Concurrency versus Parallelism

In general computing literature, the terms ‘concurrency’ and ‘parallelism’ are often taken as synonyms and used interchangeably by many, while others make a clear distinction between the two. We will afford a few sentences to clarify what we mean by each, in the context of this thesis.

When we say *parallelism*, we mean the extraction of better *performance* from a program by inferring computations that do not interact with each other then simul-

taneously carrying them out, and the writing of programs in such a way as to make this process easy or possible for the compiler. *Concurrency* on the other hand means the use of interdependent threads of execution as a means of *structuring* the control flow of programs. The focus of this thesis is on explicit concurrency.

### 1.2.2 Counting: easy as 0, 1, 2...

Consider the problem of incrementing a counter, represented in Haskell using a mutable reference:

```
type CounterIORef = IORef Integer
```

```
makeCounterIORef :: IO CounterIORef
```

```
makeCounterIORef = newIORef 0
```

The `incrementIORef` program could then be implemented as follows:

```
incrementIORef :: CounterIORef → IO ()
```

```
incrementIORef counter = do
```

```
  n ← readIORef counter
```

```
  writeIORef counter (n + 1)
```

When only a single instance of `incrementIORef` is executing, the above code behaves as expected. Suppose however, that two instances of `incrementIORef` were executing at the same time. This results in four possible interleavings of the two `readIORef` and `writeIORef` operations, not all of which would have the intended effect of incrementing the counter twice. For example, the following interleaving would only increment the counter by one:

Thread A	Thread B	<i>counter</i>
$n_A \leftarrow \text{read}_{\text{IORef}} \text{counter}$		0
	$n_B \leftarrow \text{read}_{\text{IORef}} \text{counter}$	0

<code>write<sub>IORef</sub> counter (n<sub>A</sub> + 1)</code>	<code>write<sub>IORef</sub> counter (n<sub>B</sub> + 1)</code>	1
		1

Typically, reading from and writing to mutable variables are relatively fast primitive operations. When they occur in immediate succession, the probability of Thread A being interleaved by Thread B in the above manner is very small, and can easily slip through seemingly thorough empirical testing. Such errors are termed *race conditions*, and can occur whenever there is the possibility of concurrent access to any shared resource.

### 1.2.3 Shared Memory and Mutual Exclusion

The most widely used approach to prevent the kind of race conditions we have seen in the previous section is to simply prevent concurrent accesses to the shared resource, via a selection of related techniques—locks, semaphores, critical sections, mutexes—all of which are based on the principle of mutual exclusion.

Without discussing implementation details, let us assume the existence of two primitive operations—*lock* and *release*—with the following behaviour: *lock* attempts to acquire exclusive access to a given mutable variable; if the variable is already locked, *lock* waits until it becomes available before proceeding. Its counterpart *release* relinquishes the exclusivity previously obtained.

We can now eliminate the earlier race condition as follows:

```

incrementlock :: Counterlock → IO ()
incrementlock counter = do
  lock counter
  incrementlock counter
  release counter

```

Even if Thread A were interleaved mid-way, Thread B cannot proceed past the **lock** primitive until Thread A releases *counter*, ruling out the earlier race condition:

Thread A	Thread B	<i>counter</i>
<b>lock</b> <i>counter</i>		0
$n_A \leftarrow \text{read}_{\text{IORef}} \text{ counter}$		0
	<b>lock</b> <i>counter</i>	0
	$\vdots$	1
<b>write</b> <sub>IORef</sub> <i>counter</i> ( $n_A + 1$ )	blocked on <i>counter</i>	1
<b>release</b> <i>counter</i>	$\vdots$	1
	$n_B \leftarrow \text{read}_{\text{IORef}} \text{ counter}$	1
	<b>write</b> <sub>IORef</sub> <i>counter</i> ( $n_B + 1$ )	2
	<b>release</b> <i>counter</i>	2

Such two-state locks can easily be generalised to  $n$  states with *counting semaphores* where some limited concurrent sharing may take place.

### 1.2.4 Example: Deadlock

Let us consider a slightly more interesting example: we are required to implement a procedure to increment two given counters in lock-step. A first attempt may be as follows:

```

increment'pair :: Counterlock → Counterlock → IO ()
increment'pair c0 c1 = do
  incrementlock c0
  incrementlock c1

```

However, this does not have the desired effect, because there is an intermediate state between the two calls to **increment**<sub>lock</sub> when  $c_0$  has been incremented but  $c_1$  has not, yet neither is locked. A better implementation might lock both counters before incrementing:

`incrementpair` :: `Counterlock` → `Counterlock` → `IO ()`

`incrementpair c0 c1 = do`

`lock c0`

`lock c1`

`incrementlock c0`

`incrementlock c1`

`release c0`

`release c1`

While this version ensures that the two counters are updated together, it however suffers from a more subtle problem. If two threads A and B both attempt to increment the same pair of counters passed to `incrementpair` in a different order, a potential deadlock situation can occur:

A: <code>increment<sub>pair</sub> c<sub>0</sub> c<sub>1</sub></code>	B: <code>increment<sub>pair</sub> c<sub>1</sub> c<sub>0</sub></code>
<code>lock c<sub>0</sub></code>	<code>lock c<sub>1</sub></code>
<code>lock c<sub>1</sub></code>	<code>lock c<sub>0</sub></code>
⋮	⋮
blocked on <code>c<sub>1</sub></code>	blocked on <code>c<sub>0</sub></code>
⋮	⋮

Neither thread can make progress, as they attempt to acquire a lock on the counter which is being held by the other. This could be solved by always acquiring locks in a specific order, but enforcing this is not always straightforward.

Correctness considerations aside, there are the issues of code reuse and scalability to consider. In terms of reuse, ideally we would not want to expose `increment`, as only `incrementlock` is safe for concurrent use. On the other hand, to build more complex operations on top of the basic ‘increment’ operation, chances are that we would need access to the unsafe `increment` implementation. Unfortunately exposing this breaks

the abstraction barrier, with nothing to enforce the safe use of *increment* other than trust and diligence on the part of the programmer.

On the issue of scalability, there is also some tension regarding lock granularity, inherent to mutual-exclusion. Suppose we have a large shared data structure, and our program makes use of as many threads as there are processors. In order to allow concurrent access to independent parts of the data structure, we would need to associate a lock with each constituent part. However acquiring a large number of locks has unacceptable overheads; particularly noticeable when there are only a small number of threads contending for access to the shared data. On the other hand, increasing lock granularity would reduced the number of locks required, and in turn the overheads associated with taking the locks, but this would also rule out some potential for concurrency.

### 1.2.5 Message Passing and Implicit Synchronisation

The message passing paradigm focuses on the sending of messages between threads in the computation as a primitive, rather than the explicit use of shared memory and mutual exclusion as a medium and protocol for communication. Conceptually, this is a higher-level notion which abstracts the act of sending a message from the how, leaving it to the run-time system to choose the appropriate means. As a result, programs written using this approach have the potential to scale from single processors to distributed networks of multiprocessor computers.

Established languages and frameworks supporting message passing concurrency include Erlang [AVWW96], the Parallel Virtual Machine (PVM) [GBD<sup>+</sup>94] and the Message Passing Interface (MPI) [GLS99]. In Haskell, we can implement our previous counter example using channels, where `Chan  $\alpha$`  is the polymorphic type of channels carrying messages of type  $\alpha$ :

```
data Action = Increment | Get (Chan Integer)
```

```
type CounterChan = Chan Action
```

Here we have defined a new datatype `Action` enumerating the operations the counter supports. A counter is then represented by a channel accepting such `Actions`, to which we can either send an `Increment` command, or another channel on which to return the current count via the `Get` command.

The `makeCounterChan` function returns a channel, to which other threads may send `Actions` to increment or query the counter:

```
makeCounterChan :: IO CounterChan
makeCounterChan = do
  counter ← newChan
  value ← newIORef 0
  forkIO ∘ forever $ do
    action ← readChan counter
    n ← readIORef value
    case action of
      Increment → writeIORef value (n + 1)
      Get result → writeChan result n
  return counter
```

Even though we make use of an `IORef` to store the current count, we have implicitly avoided the mutual exclusion problem by only allowing the forked thread access, essentially serialising access to the mutable variable. Implementing an `incrementChan` operation is now straightforward:

```
incrementChan :: CounterChan → IO ()
incrementChan counter = writeChan counter Increment
```

If concurrent threads invoke `incrementChan` on the same counter, the atomicity of the `writeChan` primitive rules out any unwanted interleavings.

Unfortunately, just like the previous mutual exclusion-based solution, it is not trivial to build upon or to reuse `incrementchan`—say—to increment two counters in lock-step.

## 1.2.6 Software Transactional Memory

The popularity of mutual exclusion could be partially attributed to the fact that its implementation is relatively easy to comprehend. On the other hand, managing and composing lock-based code is rather error-prone in practice.

Automatic garbage collection frees the programmer from having to manually manage memory allocation. Laziness in functional programming allows us to write efficient higher-level programs without having to manually schedule the order of computation. In a similar vein [Gro07], software transactional memory (STM) [ST97] allows us to write programs in a compositional style in the presence of concurrency, without requiring us to manually manage undesired interleavings of operations in a shared memory environment.

The idea of using *transactions* to tackle concurrency originated in the context of concurrent databases, which face similar issues of undesirable interleavings of basic operations when different clients attempt to access the same database at the same time. Rather than explicitly locking any requisite resources before proceeding with some sequence of operations on shared data, the client simply informs the database server that the operations are to be treated as a single transaction. From the perspective of other database clients, the server ensures that none of the intermediate states of the transaction is visible, as if the entire transaction took place as a single indivisible operation. Should it fail for whatever reason, the outside perspective would be as if the transaction hadn't taken place at all.

STM implements the same concept, but with shared memory being the 'database' and individual program threads taking the rôle of the 'clients'. STM takes an op-



timistic approach to concurrency: transactions are allowed to overlap in their execution, making the most of the available hardware. Conflicts between transactions arise only when an earlier transaction commits a change to a shared variable which a later transaction depended on. Should this happen, the later one is aborted and retried. In particular, a transaction is only aborted when another one has successfully committed, thereby ensuring overall progress and the absence of deadlocks.

Under the Haskell implementation of STM, transactional computations returning a value of type  $\alpha$  have the type `STM  $\alpha$` . We can give an almost identical implementation of `incrementSTM` as that of `incrementIORef`, but uses `TVars` (*transactional variables*) instead of `IORefs`, and results in an `STM` action rather than an arbitrary `IO` action:

```
type CounterSTM = TVar Integer

incrementSTM :: CounterSTM → STM ()

incrementSTM counter = do
  n ← readTVar counter
  writeTVar counter (n + 1)
```

To effect a transaction, we have at our disposal an `atomically` primitive, which takes an `STM  $\alpha$`  and returns a runnable `IO  $\alpha$`  action. The following program fragment increments the given counter twice.

```
do
  counter ← atomically (newTVar 0)
  forkIO (atomically (incrementSTM counter))
  forkIO (atomically (incrementSTM counter))
```

In particular, the `atomically` primitive guarantees that the two instances of `incrementSTM` do not interleave each other in any way of consequence.

STM makes it straightforward to reuse existing code: simply sequencing two transactions one after another creates a larger composite transaction that increments both

counters atomically when executed:

```
incrementBoth :: CounterSTM → CounterSTM → STM ()
```

```
incrementBoth c0 c1 = do
```

```
  incrementSTM c0
```

```
  incrementSTM c1
```

This section presented but a brief overview of STM. We will examine and discuss it in more depth in Chapter 2.

## 1.3 Thesis Overview

The remaining chapters of this thesis comprise of the following:

**Chapter 2** provides additional background on the use and implementation of transactional memory, followed by a brief primer on STM Haskell.

**Chapter 3** reviews the notions of denotational, big- and small-step operational semantics along with some reasoning techniques, illustrated using a small language. We then present a compiler for this language and its corresponding virtual machine, to show the essence of a compiler correctness proof.

**Chapter 4** implements executable semantics for the above language as a Haskell program. We demonstrate the use of QuickCheck in conjunction with the Haskell Program Coverage toolkit for randomised testing of the results established in the previous chapter.

**Chapter 5** puts the above empirical approach into practice, on a simplified subset of STM Haskell with a high-level stop-the-world semantics, linked by a compiler to a virtual machine with a log-based implementation of transactions.

**Chapter 6** gently introduces the Agda proof assistant, for the purpose of constructing formal machine-checked proofs, and culminates in a verified formalisation of the results of chapter 3.

**Chapter 7** extends the notion of compiler correctness to non-deterministic languages using our new notion of combined machines, illustrated using the simple Zap language, for which a complete compiler correctness result is produced and discussed.

**Chapter 8** scales our new technique to include explicit concurrency, by introducing a ‘fork’ primitive and threads. We replay the compiler correctness proof of the previous chapter in this new setting with the help of a few extra concurrency lemmas.

**Chapter 9** develops the concept of consistency between transaction logs and the heap, which we use to establish the correctness of a log-based implementation of software transactional memory in the presence of arbitrary interference by an external agent.

**Chapter 10** concludes with a summary of this thesis, and a list of various future research directions.

The reader is assumed to be familiar with functional programming and their type systems; knowledge of Haskell would be a bonus. Relevant ideas are introduced when appropriate, with references for further reading.

## 1.4 Contributions

The contributions of this thesis are as follows:

- Identification of a simplified subset of STM Haskell with a high-level stop-the-world semantics for transactions.

## CHAPTER 1. INTRODUCTION

- A virtual machine for this language, in which transactions are implemented following a low-level log-based approach, along with a semantics for this machine.
- A compiler linking the language to the virtual machine with a statement of compiler correctness, empirically tested using QuickCheck and HPC.
- The core idea of a combined machine and semantics, that allows us to establish a direct bisimulation between the high-level language and the low-level virtual machine.
- Putting the above technique into practice using the Agda proof assistant, giving a formal compiler correctness proof for a language with a simple notion of non-determinism.
- Showing that our technique scales to a language with explicit concurrency, complete with a formal proof.
- A formal correctness proof for a transactional language that shows the equivalence of the log-based approach and the stop-the-world semantics for transactions.

Earlier accounts of some of these have been published in the following papers,

- Liyang HU and Graham Hutton [HH08]. “Towards a Verified Implementation of Software Transactional Memory”. In *Proceedings of the Symposium on Trends in Functional Programming*. Nijmegen, The Netherlands, May 2008.
- Liyang HU and Graham Hutton [HH09]. “Compiling Concurrency Correctly: Cutting Out the Middle Man”. In *Proceedings of the Symposium on Trends in Functional Programming*. Komárno, Slovakia, June 2009.

but have since been refined and expanded upon in this thesis.

## 1.4. CONTRIBUTIONS

The complete source code to this thesis, in the form of literate Haskell and Agda documents, may be found online at <http://liyang.hu/#thesis> .

## CHAPTER 1. INTRODUCTION

# Chapter 2

## Software Transactional Memory

While mutual exclusion is the dominant paradigm for shared memory concurrent programming, it can be difficult to reason about and is error-prone in practice. Taking a cue from the distributed databases community, software transactional memory applies the concept of transactions to shared memory concurrent programming. In this chapter, we introduce the notion of transactions and transactional memory, along with high-level overviews of how transactional memory can potentially be implemented. We then give a brief history of the development of the approach up to the present day, followed by a primer to the implementation as found in the Glasgow Haskell Compiler.

### 2.1 Database Transactions

Consider a database server that supports multiple concurrent clients. Each client may need to carry out some complex sequence of operations, and it is up to the database management server to ensure that different clients do not make conflicting changes to the data store. As with any concurrent software, the clients could obtain exclusive access to part of the database by taking a lock, but without careful coordination between clients this can result in deadlock situations, as we have seen in the previous

chapter. This problem is only exacerbated by the possibility of failure on the part of the client software or hardware: a failed client could be holding a lock to critical parts of the database, thus preventing others from making progress.

The alternative is to make use of the concept of *transactions* [Dat95], which provides a higher-level approach to managing concurrency than explicit locking or mutual exclusion. A client starts a transaction by issuing a *begin* command to the server. Thereafter, all subsequent operations until the final *commit* command are considered part of the transaction. The intended behaviour of such transactions is informally captured by the following four properties, abbreviated by the term ‘ACID’ [Gra81]:

**Atomicity** The sequence of operations take place as if they were a single indivisible operation, ensuring that transactions follow a simple ‘all-or-nothing’ semantics: if any of its constituent operations fail, or if the transaction is aborted for whatever reason, the server guarantees that the resulting state of the database is as if none of the operations took place at all.

**Consistency** The sequence of operations cannot put the data store into a state that is inconsistent with pre-defined invariants of the database. Were this the case, the server would cause the commit operation to fail. Typically, such invariants would be specific to the particular database application, and serve as a safety net to catch client-side bugs.

**Isolation** As a transaction is running, other clients cannot observe any of its intermediate states. Conversely, until the transaction has completed and been committed to the data store, it cannot influence the behaviour other concurrent clients.

**Durability** Once the database server has accepted a transaction as being committed, the effects of the operations on the database store should persist, even in the event of system failure.



This approach to concurrency significantly simplifies client implementation. If a transaction fails—say because another client successfully committed a conflicting change—the original client will be notified and may simply retry the transaction at a later point. Atomicity ensures that the partially completed transaction is automatically ‘rolled back’: the client need not carefully undo the changes it had made so far to avoid affecting subsequent transactions, while isolation ensures that potentially inconsistent intermediate states of the transaction is not visible to others.

Furthermore, this approach is intended to be *optimistic* in the sense that we can always proceed with any given transaction, as there are no locks to acquire, making deadlock a non-issue. The trade-off is that the transaction may later fail or be unable to commit, should a different transaction commit a conflicting change in the meantime. Nevertheless, the system as a whole has made progress.

## 2.2 Transactional Memory

Transactional memory applies the idea of the previous section to concurrent software, with shared memory playing the rôle of the database store and individual program threads acting as the clients. Clearly there will be some differences: with shared random-access memory being volatile, we may not be able to satisfy the durability aspect of the ACID properties in the event of a power failure, for example. In addition, consistency is more an issue of sequential program correctness and hence largely orthogonal to our main concern of concurrency. The focus of transactional memory is therefore on providing atomicity and isolation in the presence of concurrency.

In this section, we give a brief overview of various developments leading up to the current state of the art in the application of transactional memory to concurrent software.

### 2.2.1 Hardware Transactional Memory

While it is possible to implement synchronisation between multiple parties using only primitive read and write operations, for example with Lamport’s bakery algorithm [Lam74], such software-only techniques do not scale well with the addition of further participants. Rather, most modern architectures feature in one form or another a *compare-and-swap* (CAS) instruction that compares a given value to a word in memory and conditionally swaps in a new value if the comparison yields true. The key property of a CAS instruction is that it does so in an atomic and isolated manner, enforced at the hardware level. This is a powerful primitive on which more sophisticated synchronisation constructs can be readily built.

In a sense, a compare-and-swap primitive already allows us to implement simple word-sized transactions on the hardware level: the CAS instruction can detect interference by comparing the current value of the word with what we had read during the transaction, and if they match, commit the new value to memory. Features resembling transactions are more readily observed on other architectures—such as the DEC Alpha [JHB87]—that provided pairs of *load-linked* and *store-conditional* (LL and SC) primitives. The load-linked instruction—as well as fetching a word from memory—additionally places a watch on the system memory bus for the address in question. The subsequent store-conditional instruction proceeds only when no writes to the address has occurred in the meantime. Of course, the programmer must still check for its success, and either manually retry the operation, or attempt an alternative route.

Herlihy and Moss [HM93] later extended this approach to explicitly support multiword transactions, building upon existing cache-coherency protocols for multiprocessor architectures. In effect, a *transactional cache* local to the processor buffers any tentative writes, which are only propagated to the rest of the system after a successfully commit. They leverage existing cache-coherency protocols to efficiently guard against potential interference. As the size of the transactional cache is limited by

hardware, this sets an upper bound on the size of the transactions that can occur. To this end, Herlihy and Moss suggest virtualisation as a potential solution, transparently falling back to a software-based handler in much the same way as virtual memory gives the illusion of a practically infinite memory space.

### 2.2.2 Software Transactional Memory

While some form of hardware support beyond the basic compare-and-swap would be desirable for the implementation of transactional memory, Shavit and Touitou [ST97] propose a software-only version of Herlihy and Moss’s approach, which can be efficiently implemented on existing architectures that support CAS or LL/SC instructions. In purely pragmatic terms, it does not require the level of initial investment required by a hardware-assisted solution.

In his thesis, Fraser [Fra03] demonstrated that non-trivial data structures based on his implementation of STM had comparable performance to other lock-based or intricately crafted lock-free designs, running on a selection of existing modern hardware. In this sense, STM could be considered practical for everyday use.

However, even though transactional algorithms can be derived from existing sequential ones by simply replacing memory accesses with calls to the STM library, the impedance mismatch of having to use library calls for mutating variables makes programming in the large somewhat impractical in many of today’s mainstream languages. Furthermore, it was not possible to prevent a programmer from directly accessing shared data and circumventing the atomicity and isolation guarantees bestowed by the transaction.

Harris and Fraser [HF03] experimented with transactional extensions to the Java language, along with an STM run-time library. Simply wrapping an existing block of code within an `atomic` construct executes it within a transaction. Upon reaching the end of such a block, the run-time system would attempt to commit its changes,

and should this fail, retries the transaction from the beginning.

Harris and Fraser’s Java bytecode translator traps in-language read and writes to the program heap, replacing them with transaction-safe alternatives. However, this does not prevent Java’s native methods from surreptitiously modifying the heap or performing irreversible side-effects such as input or output, which would be problematic given that a transaction may execute more than once before it successfully commits. Arbitrary Java code is permitted within `atomic` blocks, but calls to foreign methods within a transaction would raise run-time exceptions as these have the potential to void any atomicity or isolation guarantees of the system. At present, transactional memory has yet to be accepted by the wider Java community, although there is much pioneering work in both commercial [Goe06] and academic contexts [HLM06, KSF10].

Following on from this work, Harris et al. [HMPJH05] presented an implementation of software transactional memory for the Glasgow Haskell Compiler (GHC) as part of their paper entitled *Composable Memory Transactions*. Contributions of this work include the use of Haskell’s type system to ensure that only those operations that can be rolled back are used in transactions, along with an operator for composing pairs of alternative transactions. In later work, they introduce ‘data invariants’ [HPJ06] for enforcing consistency.

More recent work has brought transactional memory to other programming languages [Tan05, Nodir00], as well as more efficient [Enn05] or alternative low-level implementations [ATS09].

## 2.3 Implementing Transactions

The high-level view of transactions is that each one is executed atomically and in isolation from other concurrent threads, as if the entire transaction took place in a

single instant without any of its intermediate states being visible by concurrently running threads. Indeed, an implementation could simply suspend all other threads upon starting a transaction, and resume them after the transaction has ended. Pragmatically, this *stop-the-world* view can be easily achieved by ensuring that only one transaction is ever executing at any point in time, say using a global mutual-exclusion lock. While this approach would be easy to implement, it prevents transactions from proceeding concurrently, and would not be make good use of multi-core hardware.

### 2.3.1 Log-Based Transactions

A concurrent implementation of transactions might make use of the notion of a *transaction log*. During the execution of a transaction, its associated log serves to isolate it from other concurrently running threads. Rather than immediately acting on any side-effecting operations, the run-time makes a record of these in the log, applying them globally only when the transaction successfully commits.

In the case of shared mutable variables, the log essentially acts as a transaction-local buffer for read and write operations: for each variable, only the first read would come from the shared heap, and only the last value written goes to the shared heap; all intermediate reads and writes operate solely on the log. At the end of the transaction, if any undesired interference has occurred, say due to another transaction having completed in the meantime, the run-time system need only discard the log and re-run the transaction, since no globally-visible changes have been made. Therefore it would be only be appropriate to allow operations that can be buffered in some suitable manner to be recorded in the transaction log, such as changes to shared mutable variables; external side-effects—such as launching missiles [HMPJH05]—cannot be undone, and should be prevented. As long as the transaction log corresponds to a consistent state of the current shared heap on the other hand, the new values recorded in the log can then be applied to the global heap in an atomic manner.

A simple implementation of the commit operation need only ensure that globally, only one commit is taking place at any time, say using a global lock. Even then, we still make gains on multicore systems, as the bodies of transactions still run concurrently. A more sophisticated implementation might allow transactions to commit concurrently, for example by making use of specialised lock-free data structures. While concurrent commits would be trickier to implement, additional complexities are restricted to the run-time system implementer rather than the proletarian programmer

### 2.3.2 Alternative Approaches to Atomicity

Atomicity is undoubtedly a useful tool for writing programs in a modular, reusable manner in the face of the challenges posed by concurrency. As well as using the notion of a transaction log, there are a number of alternative approaches for creating an illusion of atomicity:

*Compensating transactions* [Gra81, KLS90] may be thought of as a conflict-detection framework, in which the programmer manually specifies how the whole system can be returned to its initial state, should a transaction need to be retried. A transactions could abort for a variety of reasons, and the roll-back code must be able to deal with any failure mode. The as the roll-back operation is manually defined on a case-by-case basis, there is the additional challenge of ascertaining its correctness. The advantage on the other hand is that we may perform arbitrary I/O, provided we can undo them in an undetectable way.

*Lock inference* [FFL05, CGE08] on the other hand attempts to automatically insert the fine-grained locks that a programmer might have used, via various code analysis techniques. For reasons of safety, the inference of such locks must necessarily be conservative, and does not always allow for optimal concurrency. Additionally, as code analysis techniques are generally ideally performed on whole programs, we might

lose modularity on the level of object files and/or require the use of sophisticated type-systems [CPN98]. Since roll-back is no longer necessary in this approach, we can allow arbitrary I/O side-effects, but isolation would only be afforded to mutable variables.

## 2.4 Haskell and Sequential Computation

In this section we will revisit some basic aspects of Haskell required for the understanding of the implementation of STM given by Harris et al. [HMPJH05]. The material should be accessible to the reader with a general understanding of functional programming; no working knowledge of Haskell in particular is required. To aid readability, we will also make use of the following colour scheme for different syntactic classes:

Syntactic Class	Examples
Keywords	<b>type</b> , <b>data</b> , <b>let</b> , <b>do</b> ...
Types	<b>()</b> , <b>Bool</b> , <b>Integer</b> , <b>IO</b> ...
Constructors	<b>False</b> , <b>True</b> , <b>Just</b> , <b>Nothing</b> ...
Functions	<b>return</b> , <b>getChar</b> , <b>readTVar</b> ...
Literals	<b>0</b> , <b>1</b> , <b>"hello world"</b> ...

### 2.4.1 Monads for Sequential Computation

The Haskell programming language [PJ03a, Mar10]—named after the logician Haskell B. Curry—can be characterised by its three key attributes: functional, pure, and lazy. Functional programming languages are rooted in Church’s  $\lambda$ -calculus [Chu36, Bar84], and emphasise the evaluation of mathematical functions rather than the manipulation of state. The core  $\lambda$ -calculus ideas of abstraction and application are typically given prominent status in such languages. In turn, purity means functions depend only on their arguments, eschewing state or mutable variables, akin to the mathematical

notion of functions. The same program expression will always evaluate to the same result regardless of its context, and replacing an expression with its value leaves the meaning of the program unchanged. In other words, the language is *referentially transparent* [Sab98]. The laziness aspect of Haskell means that expressions are only evaluated when their values are required, thus the evaluation order—or even whether something is evaluated at all—is not necessarily immediately apparent from the program text. Together, these properties meant that for some time, it was not clear how to write programs that are more naturally expressed in an imperative, sequential style, or to deal with input and output.

A solution was found in the form of *monads*, which Moggi [Mog89] and Wadler [Wad92] adopted from category theory [ML98]. In the context of computer science, a monad could be viewed as a ‘container’ for some general notion of computation, together with an operation for combining such computations. As it turns out, sequential computation is just one instance of a monad, as we shall see in the following section.

### 2.4.2 Modelling Mutable State

Since Haskell is referentially transparent, we cannot directly work with mutable variables, but we can model them. Let us consider the case of a single mutable variable: the basic approach involves passing around the current value of the mutable variable—say, of type  $\sigma$ —as an extra argument. Thus, rather than implementing a function of type  $\alpha \rightarrow \beta$  which cannot access any mutable state, we instead write one of type  $\alpha \rightarrow \sigma \rightarrow (\sigma, \beta)$ . This takes the current value of the mutable variable as an extra argument and returns a new value of the variable, together with the original result of type  $\beta$ .

As we will frequently make use of similar types that mutate some given state, it is convenient to define the following `State` synonym:

```
type State  $\sigma$   $\alpha = \sigma \rightarrow (\sigma, \alpha)$ 
```



## 2.4. HASKELL AND SEQUENTIAL COMPUTATION

A value of type `State  $\sigma$   $\alpha$`  can be regarded as a computation involving some mutable  $\sigma$  state that delivers an  $\alpha$  result. Thus we can write the following definitions of `read` and `write`, corresponding to our intuition of a mutable variable:

```
read :: State  $\sigma$   $\sigma$ 
read =  $\lambda s \rightarrow (s, s)$ 
write ::  $\sigma \rightarrow$  State  $\sigma$  ()
write  $s' = \lambda s \rightarrow (s', ())$ 
```

The `read` computation results in a value that is the current state, without changing it in any way. On the other hand, the `write` computation replaces the current state  $s$  with the supplied  $s'$ , giving an information-free result of the singleton `()` type.

With these two primitives, we can implement an `increment` operation as follows:

```
increment :: State Integer ()
increment =  $\lambda s \rightarrow$ 
  let  $(s', n) =$  read  $s$ 
  in write  $(n + 1)$   $s'$ 
```

The result of `read` is bound to the name  $n$ , then the state is updated with  $n + 1$  by the subsequent `write`. The initial state  $s$  we have been given is passed to `read`—potentially resulting in a different state  $s'$ —which is then passed along to `write`.

In the above instance we know that `read` does not change the state, but in general any `State  $\sigma$   $\alpha$`  computation could, therefore we must carefully thread it through each computation in order to maintain the illusion of mutable state. The following definition increments the counter twice:

```
twice :: State Integer ()
twice =  $\lambda s \rightarrow$ 
  let  $(s', -) =$  increment  $s$ 
  in increment  $s'$ 
```

Were we to inadvertently pass the initial  $s$  to the second invocation of `increment` as well, we would have made a copy of the initial state, having discarded the updated  $s'$ . The resulting `twice` would only appear to increment the counter once, despite its two invocations of `increment`. Explicit state threading is not only rather tedious, small errors can also silently lead to very unexpected results.

Fortunately, the act of threading state around is sufficiently regular that we can implement a general purpose operator that hides away the details of the plumbing:

$$\begin{aligned} (\gg=) &:: \text{State } \sigma \alpha \rightarrow (\alpha \rightarrow \text{State } \sigma \beta) \rightarrow \text{State } \sigma \beta \\ ma \gg= fmb &= \lambda s \rightarrow \mathbf{let} (s', a) = ma \ s \ \mathbf{in} \ fmb \ a \ s' \end{aligned}$$

The  $\gg=$  operator—pronounced ‘bind’—takes on its left a computation delivering an  $\alpha$ ; its right argument is a function from  $\alpha$  to a second computation delivering a  $\beta$ . Bind then passes the result of its left argument—along with the modified state  $s'$ —and threads both through its right argument, delivering a stateful computation of type  $\beta$ . Using this operation, we can rewrite `increment` and `twice` as follows:

```
increment' :: State Integer ()
increment' =
  read >>= \n ->
  write (n + 1)

twice' :: State Integer ()
twice' =
  increment' >>= \_ ->
  increment'
```

In the above definitions, we no longer need to explicitly thread state around as this is handled automatically by the  $\gg=$  operation, and the resulting code has a much more imperative appearance. In fact, Haskell provides a few helper functions as well

as some lightweight syntactic sugar to support exactly this style of programming, allowing the following succinct definition:

```
increment'' :: State Integer ()
increment'' = do
  n ← read
  write (n + 1)

twice'' :: State Integer ()
twice'' = do
  increment''
  increment''
```

Here, we may think of the expression  $n \leftarrow \text{read}$  in the definition of `increment''` as binding the result of `read` to the name  $n$ , and in fact desugars to the same code as that of `increment'`. Should we merely want to run a computation for its side-effects, as we do in the definition of `twice''`, we simply omit both the  $\leftarrow$  operator and the resulting name.

To prevent direct manipulation or duplication of the threaded state, we can make `State` an opaque data type to its users, hiding the above implementations, and offer only `read` and `write` as primitives for accessing the state.

### 2.4.3 Monadic Properties

So far in this section we have deliberately avoided using the mathematical term ‘monad’. In fact, some members of the Haskell community have jokingly remarked that they would rather have used the phrase ‘warm fuzzy thing’ [PJ03b] instead. The `>>=` operator above already constitutes the primary component of the definition of the `State`  $\sigma$  monad, and we need only one further function to complete it.

**return** ::  $\alpha \rightarrow \text{State } \sigma \alpha$

**return**  $a = \lambda s \rightarrow (s, a)$

Here, the **return** function produces a trivial computation that results in the value of its given argument. Were they to agree with the mathematical definition of a monad, our definitions of **bind** and **return** for the **State**  $\sigma$  monad must satisfy certain properties, which are as follows:

$$\text{return } a \gg= fmb \quad \equiv \quad fmb a \quad \text{(ident-left)}$$

$$ma \gg= \text{return} \quad \equiv \quad ma \quad \text{(ident-right)}$$

$$(ma \gg= fmb) \gg= fmc \quad \equiv \quad ma \gg= (\lambda a \rightarrow fmb a \gg= fmc) \quad \text{(assoc)}$$

The first two specify that **return** is a left as well as right-identity for  $\gg=$ , while the third says that the  $\gg=$  operator is associative, modulo the binding operation inherent in the use of  $\gg=$ . Using the power of equational reasoning afforded to us in this pure functional setting, we can show that our definition of  $\gg=$  and **return** for the **State**  $\sigma$  monad satisfies the above laws by simply expanding the relevant definitions. For example, the **(ident-left)** property can be shown as follows:

$$\begin{aligned} & \text{return } a \gg= fmb \\ \equiv & \{ \text{definition of } \gg= \} \\ & \lambda s \rightarrow \text{let } (s', a') = \text{return } a \text{ s in } fmb a' s' \\ \equiv & \{ \text{definition of } \text{return} \} \\ & \lambda s \rightarrow \text{let } (s', a') = (s, a) \text{ in } fmb a' s' \\ \equiv & \{ \text{substitute for } a \text{ and } s \} \\ & \lambda s \rightarrow fmb a s \\ \equiv & \{ \eta\text{-contract} \} \\ & fmb a \end{aligned}$$

### 2.4.4 Input, Output and Control Structures

Now that we have shown how we can sequence operations on mutable state, what about input and output? In a sense, we can conceptually think of I/O as mutating the outside world, and indeed this is the approach used in Haskell. By threading a token representing the state of the real world through a program in a similar way to the `State`  $\sigma$  monad, we ensure that real-world side-effects occur in a deterministic order. For example, the `IO` type in Haskell could be defined as follows,

```
type IO  $\alpha$  = State RealWorld  $\alpha$ 
```

where `RealWorld` is the opaque type of the token, inaccessible to the end-programmer. Assuming two primitives `getChar :: IO Char` and `putChar :: Char → IO ()` for interacting with the user, we can implement an `echo` procedure as follows:

```
echo :: IO ()
echo = do
  c ← getChar
  putChar c
```

In Haskell, monadic actions such as `getChar` or `echo` are first-order, and when we write a program, we are in fact just composing values—evaluating `echo` for example does not prompt the user for a character nor print one out, it merely results in a value of type `IO ()` corresponding to the composition of `getChar` and `putChar`. The only way to make `echo` actually perform input and output is to incorporate it into the definition of the system-invoked `main :: IO ()` action.

Being able to manipulate monadic actions is a very powerful concept, and allows us to create high-level control structures within the language itself. For example, there's no need for Haskell to have a built-in *for-loop* construct, because we can implement it ourselves:

```

for :: Integer → Integer → (Integer → State σ ()) → State σ ()
for m n body = case m < n of
  False → return ()
  True → do
    body m
    for (m + 1) n body

```

The `for` function invokes `body` with successive integers arguments, starting at `m` and stopping before `n`. While the type of `for` explicitly mentions the `State σ` monad, `IO` is a particular instance of this, so the expression `for 0 10 (λ_ → echo)` corresponds to an `IO` action that echoes 10 characters entered by the user.

Haskell’s typeclasses permits a form of ad-hoc polymorphism, which allows us to give type-specific instances of `>>=` and `return`, so the above definition of `for` works in any monad we care to define. However a discussion of the topic [WB89] is beyond the scope of—and not required for—this thesis.

## 2.5 Haskell and Concurrent Computation

While the Haskell language is pure and lazy, occasionally we still need to make use of certain imperative features [PJ01]. By keeping such features within the `IO` monad—where a token of the external world state is implicitly threaded through each `IO` action—not only can we then guarantee a particular execution order, we also preserve the purity of the rest of the language.

For example, in those cases where the only known efficient solution to a problem is explicitly imperative, Haskell’s standard library provides true mutable variables in the form of the `IORef` datatype, where `IORef α` is a reference to values of type `α`. Its basic interface is given below:

```

newIORef :: α → IO (IORef α)

```

```

readIORef :: IORef α → IO α
writeIORef :: IORef α → α → IO ()

```

For multiprocessor programming, Parallel Haskell [THLPJ98] provides the `par :: α → β → β` combinator, which instructs to the run-time system that it may be worth evaluating its first argument in parallel (cf. section 1.2.1), and otherwise acting as the identity function on its second argument. As is evident from its type, the `par` combinator is pure and cannot perform any side-effects, nor can there be any interaction between its arguments even if they are evaluated in parallel. In fact, it would be perfectly sound for an implementation of `par` to simply ignore its first argument.

However, explicit concurrency is a necessity as well as a convenience when used as a mechanism for structuring many real-world programs. Concurrent Haskell [PJGF96] introduced the `forkIO :: IO () → IO ()` primitive, which provides a mechanism analogous to the Unix `fork()` system call, sparking a separate thread to run its argument `IO ()` action. Forking is considered impure as threads can interact with each other via a variety of mechanisms, and this fact is correspondingly reflected in the return type of `forkIO`. With the mutability provided by `IORefs`, we can create concurrent applications in the same imperative manner as other lower-level programming languages. For example, the following program launches a secondary thread to repeatedly print the letter ‘y’, while `mainChar` carries on to print ‘n’s:

```

mainChar :: IO ()
mainChar = do
  forkIO (forever (putChar 'y'))
  forever (putChar 'n')

```

The user would observe an unending stream of ‘y’s and ‘n’s, interleaved in an unspecified manner.

To demonstrate concurrency with interaction, the following program launches two

threads, both repeatedly incrementing a shared counter, as well as an individual one. The `mainIORef` function meanwhile checks that the shared counter indeed equals the sum of the two thread-local ones.

```

type CounterIORef = IORef Integer

incrementIORef :: CounterIORef → IO ()

incrementIORef c = do
  n ← readIORef c
  writeIORef c (n + 1)

incBothIORef :: CounterIORef → CounterIORef → IO ()

incBothIORef sum local = do
  incrementIORef sum
  incrementIORef local

mainIORef :: IO ()

mainIORef = do
  sum ← newIORef 0
  a ← newIORef 0
  b ← newIORef 0

  forkIO (forever (incBothIORef sum a))
  forkIO (forever (incBothIORef sum b))

  forever (do
    nsum ← readIORef sum
    na ← readIORef a
    nb ← readIORef b
    when (nsum ≠ na + nb) (do
      putStrLn "oh dear."))

```

Such a program, while seemingly straightforward in intent, can leave the programmer



with exponential number of possibilities to consider as it scales; it would simply be impractical to apply sequential reasoning to each potential interleaving. Worse still is the fact that the unwanted interleavings are often the least likely to occur, and can easily slip through otherwise thorough empirical testing.

The above program has a number of potentially rare and unexpected behaviours. Firstly, the two forked-off children both increment the *sum* counter, and it is quite possible for one thread’s execution of `incrementIORef` to interleave the `readIORef` and `writeIORef` of the other thread—as we have witnessed in section 1.2.2—losing counts in the process. Requiring our implementation of `increment` to follow a locking discipline for each counter in question would eliminate this particular race condition. Even with this fix in place, another issue remains: each thread first increments *sum*, followed by its own specific counter; meanwhile, the main thread may interleave either child in-between the two aforementioned steps, and observe a state in which the value of *sum* disagrees with the sum of the values of *a* and *b*.

As a concurrent program increases in size, race conditions and deadlock can become much more subtle and difficult to debug. Transactional memory—amongst other high-level approaches—aims to avoid such bugs, while retaining the speed benefits of concurrency.

## 2.6 Haskell and Software Transactional Memory

The previous section outlined the standard approach to concurrency in Haskell, which makes use of explicit threading and mutable variables via `forkIO` and `IORefs` within the `IO` monad. In an analogous fashion, STM Haskell provides mutable *transactional variables* of type `TVar α`, with the following interface:

```
newTVar  :: α → STM (TVar α)
readTVar :: TVar α → STM α
```

$$\text{write}_{\text{TVar}} :: \text{TVar } \alpha \rightarrow \alpha \rightarrow \text{STM } ()$$

The `newTVar` function creates a transactional variable initialised to some given value, while `readTVar` and `writeTVar` inspect and mutate `TVars`. The type `STM  $\alpha$`  may be read as *a transaction which delivers a result of type  $\alpha$* . We may combine the above three primitives to form more elaborate transactions, using the following monadic sequencing operators:

$$(\gg=) :: \text{STM } \alpha \rightarrow (\alpha \rightarrow \text{STM } \beta) \rightarrow \text{STM } \beta$$

$$\text{return} :: \alpha \rightarrow \text{STM } \alpha$$

The definition of `bind` for the `STM` monad composes transactions in a sequential manner, while `return` takes a given value to a trivial transaction resulting in the same. Transactions are converted to runnable `IO` actions via the `atomically` primitive,

$$\text{atomically} :: \text{STM } \alpha \rightarrow \text{IO } \alpha$$

which when run, performs the while transaction as if it were a single indivisible step. The intention is that when implementing some data structure for example, we need only expose the basic operations as `STM  $\alpha$`  actions, without the need to anticipate all the potential ways in which a user may wish to compose said operations in the future. The end-programmer may compose these primitives together in any desired combination, wrapped in an outer call to `atomically`. Concurrent transactions are achieved through explicit threading, using `forkIO` as before, while STM run-time takes care of the book-keeping necessary to guarantee that each composite transaction takes place in an atomic and isolated manner.

STM Haskell makes use of the notion of a transaction log (as we mentioned previously in section 2.3.1) and may automatically re-run transactions when conflicts are detected. Therefore it is important that `STM` actions only make changes to transactional variables—which can be encapsulated within its corresponding log—rather than arbitrary and possibly irrevocable `IO` actions. This an easy guarantee

because the Haskell type system strictly and statically differentiates between `IO`  $\alpha$  and `STM`  $\alpha$ , and there is no facility for actually performing an `IO` action while inside the `STM` monad. Of course, a transaction can always manipulate and return `IO` actions as first-order values, to be performed post-commit. In any case, as we idiomatically perform the bulk of computations in Haskell using only pure functions, these are necessarily free from side-effects. Thus they do not need to be kept track of by the transaction implementation and may simply be discarded in the event of a conflict. The ability to statically make this three-fold distinction between irrevocable (namely `IO`) and revocable (or `STM`) side-effecting computations—used relatively infrequently in practice—alongside pure ones, makes Haskell an ideal environment for an implementation of STM.

Let us now revisit the example of the previous section, with two threads competing to incrementing a shared counter. Using STM, we can make the previous program behave in the intended manner as follows, with only minor changes to its structure:

```

type CounterTVar = TVar Integer

incrementTVar :: CounterTVar → STM ()

incrementTVar c = do
    n ← readTVar c
    writeTVar c (n + 1)

incBothTVar :: CounterTVar → CounterTVar → STM ()

incBothTVar sum local = do
    incrementTVar sum
    incrementTVar local

mainTVar :: IO ()

mainTVar = do
    sum ← atomically (newTVar 0)

```

```

a ← atomically (new_TVar 0)
b ← atomically (new_TVar 0)

fork_IO (forever (atomically (incBoth_TVar sum a)))
fork_IO (forever (atomically (incBoth_TVar sum b)))

forever (do
  (n_sum, n_a, n_b) ← atomically (do
    n_sum ← read_TVar sum
    n_a ← read_TVar a
    n_b ← read_TVar b
    return (n_sum, n_a, n_b))
  when (n_sum ≠ n_a + n_b) (do
    putStrLn "oh dear."))

```

That is, the counter is now represented as an integer `TVar` rather than an `IORef`. Correspondingly, the `increment_TVar` primitive and the `incBoth_TVar` function now result in `STM` rather than `IO` actions. Finally, `main_TVar` atomically samples the three counters inside a single transaction to avoid potential race conditions.

While the sequencing of transactions provides a convenient and composable way to access shared data structures, a concurrency framework ought to also provide efficient ways to perform coordination between threads, say to wait on some collection of resources to become available before proceeding. With mutual exclusion, waiting on a number of objects could be implemented by waiting on each one in turn, taking care to avoid deadlocks. However, there are often cases where we might want to proceed whenever *any* one of some collection of objects becomes ready. For example, Haskell’s standard concurrency library offers generalised counting semaphores, which could be used for coordination between multiple threads. Similarly, most flavours of Unix provides a `select(2)` system call, which takes a set of file descriptors and

blocks until at least one is ready to be read from, or written to. Unfortunately, these techniques do not scale: for example in the latter case, all the file descriptors being waited upon must be collated up to a single top-level `select()`, which runs contrary to the idea of modular software development.

STM Haskell answers this problem with a pair of primitives for blocking and composing alternative transactions. The first primitive,

```
retry :: STM  $\alpha$ 
```

forces the current transaction to fail and retry. This gives a flexible, programmatic way to signal that the transaction is not yet ready to proceed, unlike traditional approaches in which the requisite conditions must be specified upfront using only a restricted subset of the language, such as e.g. Hoare’s *conditional critical regions* [Hoa02].

Armed with the `retry` primitive, we can demonstrate how a `CounterTVar` could be used as a counting semaphore [Dij65]. The `decrementTVar` function below behaves as the `wait` primitive, decrementing the counter only when its value is strictly positive, and blocking otherwise. Correspondingly the earlier `incrementTVar` defined above behaves as `signal`, incrementing the count.

```
decrementTVar :: CounterTVar  $\rightarrow$  STM ()
```

```
decrementTVar c = do
```

```
   $n_c \leftarrow$  readTVar c
```

```
  unless ( $n_c > 0$ )
```

```
    retry
```

```
  writeTVar c ( $n_c - 1$ )
```

The `retry` statement conceptually discards any side-effects performed so far and restarts the transaction from the beginning. However, the control flow within the transaction is influenced only by the `TVar`s read up until the `retry`, so if none of these have

been modified by another concurrent thread, the transaction will only end up at the same `retry` statement, ending up in a busy-waiting situation and wasting processor cycles. The STM run-time can instead suspend the current thread, rescheduling it only when one or more of the `TVars` read has changed, thus preserving the semantics of `retry`; the `TVars` involved in the decision to `retry` are conveniently recorded within the transaction log.

Suppose we now wish to implement a function to decrement the `sum` variable of the earlier example. In order to maintain the invariant  $a + b = \text{sum}$ , we must also decrement either one of `a` or `b`. Knowing that `decrementTVAR` blocks when the counter is zero, we may conclude that if `decrementTVAR sum` succeeds, then `a` and `b` cannot both be zero, and we ought to be able to decrement one of the two without blocking. But how do we choose? View `CounterTVAR` as a counting semaphore, it is not possible to wait on multiple semaphores unless such a primitive is provided by the system. STM Haskell provides a second primitive,

```
orElse :: STM α → STM α → STM α
```

for composing alternative transactions. With this we may implement the function described above:

```
decEitherTVAR :: CounterTVAR → CounterTVAR → CounterTVAR → STM ()
decEitherTVAR sum a b = do
  decrementTVAR sum
  decrementTVAR a 'orElse' decrementTVAR b
```

The `orElse` combinator—written above using infix notation—allows us to choose between alternative transactions: the expression `t 'orElse' u` corresponds to a transaction that runs one of `t` or `u`. It is left-biased, in the sense that `t` is run first: if it retries, any changes due to `t` is rolled back, and `u` is attempted instead. Only when both `t` and `u` cannot proceed, would the transaction as a whole retry. The final line of the

above fragment would decrement  $a$  preferentially over  $b$ , and blocking when neither can proceed. (In practice, the latter case can never arise in the above program.) Note that `orElse` need not be explicitly told which variables the transactions depend on—this is inferred from their respective transaction logs by the run-time system.

Using `orElse` for composing alternative transactions also allow us to elegantly turn blocking operations into non-blocking ones, for example:

```
decrement'_TVar :: Counter_TVar → STM Bool
decrement'_TVar c = (do decrement_TVar c; return True) 'orElse' return False
```

This non-blocking `decrement'_TVar` operation attempts to decrement the given counter using the original `decrement_TVar` and return a boolean `True` to indicate success. Should that retry or fail to commit, `orElse` immediately attempts the alternative transaction, which returns `False` instead.

By design, `retry` and `orElse` satisfy the following rather elegant properties:

$$\begin{aligned} \text{retry 'orElse' } u &\equiv u && \text{(ident-left)} \\ t \text{ 'orElse' } \text{retry} &\equiv t && \text{(ident-right)} \\ (t \text{ 'orElse' } u) \text{ 'orElse' } v &\equiv t \text{ 'orElse' } (u \text{ 'orElse' } v) && \text{(assoc)} \end{aligned}$$

In other words, the type `STM`  $\alpha$  of transactions forms a monoid, with `orElse` as the associative binary operation and `retry` as the unit.

## 2.7 Conclusion

In this chapter, we have reviewed the concept of transactions in the context of databases, followed by an overview of the development of transactional memory in both hardware and software, together with how transactions can be used as a high-level concurrency primitive. In section 2.3, we examined a log-based approach to

implementing transactions, contrasted with some alternatives. Section 2.4 introduced the Haskell language, in particular how monads are used to model mutable state in a purely functional context. The penultimate section (§2.5) presented primitives for mutable state and concurrency in Haskell, and we finished with a primer on STM Haskell—in particular a novel form of composing alternative transactions—in order to motivate the study of STM.



# Chapter 3

## Semantics for Compiler

### Correctness

In the context of computer science, the primary focus of *semantics* is the study of the meaning of programming languages. Having a mathematically rigorous definition of a language allows us to reason about programs written in the language in a precise manner. In this chapter, we begin by reviewing different ways of giving a formal semantics to a language, and various techniques for proving properties of these semantics. We conclude by presenting a compiler for a simple expression language, exploring what it means for this compiler to be correct, and how this may be proved.

### 3.1 Semantics

#### 3.1.1 Natural Numbers and Addition

To unambiguously reason about what any given program means, we need to give a mathematically rigorous definition of the language in which it is expressed. To this end, let us consider the elementary language of natural numbers and addition [HW04,

HW06, HW07].

$$\begin{aligned} \text{Expression} &::= \mathbb{N} && (\text{Exp-}\mathbb{N}) \\ &| \text{Expression} \oplus \text{Expression} && (\text{Exp-}\oplus) \end{aligned}$$

That is, an `Expression` is either simply a natural number, or a pair of `Expressions`, punctuated with the  $\oplus$  symbol to represent the operation of addition. We will adhere to a naming convention of  $m, n \in \mathbb{N}$  and  $a, b, e \in \text{Expression}$ .

Although seemingly simplistic, this language has sufficient structure to illustrate two fundamental aspects of computation, namely that of sequencing computations and combining their results. We shall shortly expand on this in section 3.1.5.

### 3.1.2 Denotational Semantics

Denotational semantics attempts to give an interpretation of a source language in some suitable existing formalism that we already understand. More specifically, the denotation of a program is a representation of what the program means in the vocabulary of the chosen formalism, which could be the language of sets and functions, the  $\lambda$ -calculus, or perhaps one of the many process calculi. Thus, to formally give a denotational semantics for a language is to define a mapping from the source language into some underlying semantic domain. For example, we can give the following semantics for our earlier `Expression` language, denoted as a natural number:

$$\begin{aligned} \llbracket \_ \rrbracket &: \text{Expression} \rightarrow \mathbb{N} \\ \llbracket m \rrbracket &= m && (\text{denote-val}) \\ \llbracket a \oplus b \rrbracket &= \llbracket a \rrbracket + \llbracket b \rrbracket && (\text{denote-plus}) \end{aligned}$$

Here, a numeric **Expression** is interpreted as just the number itself. The expression  $a \oplus b$  is denoted by the sum of the denotations of its sub-expressions  $a$  and  $b$ ; alternatively, we could say that the denotation of the  $\oplus$  operator is the familiar  $+$  on natural numbers. This illustrates the essential compositional aspect of denotational semantics, that the meaning of an expression is given in terms of the meaning of its parts. The expression  $\llbracket (1 \oplus 2) \oplus (4 \oplus 8) \rrbracket$  say, has the denotation 15 by repeatedly applying the above definition:

$$\begin{aligned} \llbracket (1 \oplus 2) \oplus (4 \oplus 8) \rrbracket &= \llbracket 1 \oplus 2 \rrbracket + \llbracket 4 \oplus 8 \rrbracket \\ &= (\llbracket 1 \rrbracket + \llbracket 2 \rrbracket) + (\llbracket 4 \rrbracket + \llbracket 8 \rrbracket) \\ &= (1 + 2) + (4 + 8) \\ &= 15 \end{aligned}$$

### 3.1.3 Big-Step Operational Semantics

The notion of big-step operational semantics is concerned with the overall result of a computation. Formally, we define a relation  $\Downarrow \subseteq \mathbf{Expression} \times \mathbb{N}$  between **Expressions** and their final values, given below in a natural deduction style:

$$\frac{}{m \Downarrow m} \quad \text{(big-val)}$$

$$\frac{a \Downarrow m \quad b \Downarrow n}{a \oplus b \Downarrow m + n} \quad \text{(big-plus)}$$

The first (**big-val**) rule says that a simple numeric **Expression** evaluates to the number itself. The second (**big-plus**) rule states that, if  $a$  evaluates to  $m$  and  $b$  evaluates to  $n$ , then  $a \oplus b$  evaluates to the sum  $m + n$ . Thus according to this semantics, we can

show that  $(1 \oplus 2) \oplus (4 \oplus 8) \Downarrow 15$  by the following derivation:

$$\begin{array}{c}
 \text{(big-val)} \quad \frac{\text{---}}{1 \Downarrow 1} \quad \frac{\text{---}}{2 \Downarrow 2} \quad \frac{\text{---}}{4 \Downarrow 4} \quad \frac{\text{---}}{8 \Downarrow 8} \\
 \text{(big-plus)} \quad \frac{\text{---}}{1 \oplus 2 \Downarrow 3} \quad \frac{\text{---}}{4 \oplus 8 \Downarrow 12} \\
 \text{(big-plus)} \quad \frac{\text{---}}{(1 \oplus 2) \oplus (4 \oplus 8) \Downarrow 15}
 \end{array}$$

For this simple language, the big-step operational semantics happens to be essentially the same as the denotational semantics, expressed in a different way. However, one advantage of a relational operational semantics is that the behaviour can be non-deterministic, in the sense that each expression could potentially evaluate to multiple distinct values. In contrast, a denotational semantics deals with non-determinism in the source language by mapping it to a potentially different notion of non-determinism in the underlying formalism. For example, should we require our expression language to be non-deterministic, we would need to switch the semantic domain of the previous semantics to the power set of natural numbers, rather than just the set of natural numbers.

### 3.1.4 Small-Step Operational Semantics

Small-step semantics on the other hand is concerned with how a computation proceeds as a sequence of steps. Both big-step and small-step semantics are ‘operational’ in the sense that the meaning of a program is understood through how it operates to arrive at the result. However, in this case each reduction step is made explicit, which is particularly apt when we wish to consider computations that produce side-effects. Again we formally define a relation  $\mapsto \subseteq \text{Expression} \times \text{Expression}$ , but between pairs

of Expressions, rather than between expressions and their values:

$$\frac{}{m \oplus n \mapsto m + n} \quad \text{(small-plus)}$$

$$\frac{b \mapsto b'}{m \oplus b \mapsto m \oplus b'} \quad \text{(small-right)}$$

$$\frac{a \mapsto a'}{a \oplus b \mapsto a' \oplus b} \quad \text{(small-left)}$$

The first rule (**small-plus**) deals with the case where the expressions on both sides of  $\oplus$  are numerals: in a single step, it reduces to the sum  $m + n$ . The second (**small-right**) rule applies when the left argument of  $\oplus$  is a numeral, in which case the right argument can make a single reduction, while (**small-left**) reduces the left argument of  $\oplus$  if this is possible. There is no rule corresponding to a lone numeric Expression as no further reductions are possible in this case.

As each  $\mapsto$  step corresponds to a primitive computation, it will often be more convenient to refer to it via its reflexive, transitive closure, defined as follows:

$$\frac{}{a \mapsto^* a} \quad \text{(small-nil)} \qquad \frac{a \mapsto a' \quad a' \mapsto^* b}{a \mapsto^* b} \quad \text{(small-cons)}$$

For example, the full reduction sequence of  $(1 \oplus 2) \oplus (4 \oplus 8) \mapsto^* 15$  would begin by evaluating the  $1 \oplus 2$  sub-expression,

$$\begin{array}{c} \text{(small-plus)} \frac{}{1 \oplus 2 \mapsto 3} \\ \text{(small-left)} \frac{}{(1 \oplus 2) \oplus (4 \oplus 8) \mapsto 3 \oplus (4 \oplus 8)} \end{array}$$

followed by  $4 \oplus 8$ ,

$$\begin{array}{c} \text{(small-plus)} \frac{\text{—————}}{4 \oplus 8 \mapsto 12} \\ \text{(small-right)} \frac{\text{—————}}{3 \oplus (4 \oplus 8) \mapsto 3 \oplus 12} \end{array}$$

before delivering the final result:

$$\text{(small-plus)} \frac{\text{—————}}{3 \oplus 12 \mapsto 15}$$

### 3.1.5 Modelling Sequential Computation with Monoids

It would be perfectly reasonable to give a right-to-left, or even a non-deterministic interleaved reduction strategy for the small-step semantics of our `Expression` language. However, we enforce a left-to-right order in order to model the sequential style of computation as found in the definition of the `State`  $\alpha$  monad from §2.4.2.

In the case where the result type of monadic computations form a monoid, such computations themselves can also be viewed as a monoid. Concretely, suppose we are working in some monad  $\mathbf{M}$  computing values of type  $\mathbf{N}$ . Using the monoid of sums  $(\mathbf{N}, +, 0)$ , the following definition of  $\otimes$ :

$$\begin{aligned} \otimes & : \mathbf{M}\mathbf{N} \rightarrow \mathbf{M}\mathbf{N} \rightarrow \mathbf{M}\mathbf{N} \\ a \otimes b & = a \gg= \lambda m \rightarrow \\ & \quad b \gg= \lambda n \rightarrow \\ & \quad \text{return } (m + n) \end{aligned}$$

gives the monoid  $(\mathbf{M}\mathbf{N}, \otimes, \text{return } 0)$ . We can easily verify that the identity and associativity laws hold for this monoid via simple equational reasoning proofs, as we had done in §2.4.3. Therefore, we can view monoids as a degenerate model of monads.

## 3.2. EQUIVALENCE PROOFS AND TECHNIQUES

The expression languages of this thesis only computes values of natural numbers, so rather than work with monadic computations of type  $M \mathbb{N}$ , we may work directly with the underlying  $(\mathbb{N}, +, 0)$  monoid, since it shares the same monoidal structure. This simplification allows us to avoid the orthogonal issues of variable binding and substitution. By enforcing a left-to-right evaluation order for  $\oplus$  in our expression language to mirror that of the  $\gg=$  operator, we are able to maintain a sequential order for computations, which is the key aspect of the monads that we are interested in.

### 3.2 Equivalence Proofs and Techniques

Now that we have provided precise definitions for the semantics of the language, we can proceed to show various properties of the `Expression` language in a rigorous manner. One obvious questions arises, on the matter of whether the semantics we have given in the previous section—denotational, big-step and small-step—agree in some manner. This section reviews the main techniques involved.

#### 3.2.1 Rule Induction

The main proof tool at our disposal is that of *well-founded induction*, which can be applied to any well-founded structure. For example, we can show that the syntax of the `Expression` language satisfies the condition of well-foundedness when paired with the following sub-expression ordering:

$$a \sqsubset a \oplus b \quad b \sqsubset a \oplus b \quad (\text{Exp-}\sqsubset)$$

The partial order given by the transitive closure of  $\sqsubset$  is well-founded, since any  $\sqsubset$ -descending chain of expressions must eventually end in a numeral at the leaves of the finite expression tree. This particular ordering arises naturally from the induc-

tive definition of **Expression**: the inductive case (**Exp- $\oplus$** ) allows us to build a larger expression  $a \oplus b$  given two existing expressions  $a$  and  $b$ , while the base case (**Exp- $\mathbb{N}$** ) constructs primitive expressions out of any natural number. In this particular case, to give a proof that some property  $P(e)$  holds for all  $e \in \mathbf{Expression}$ , it suffices by the well-founded induction principle to show instead that:

$$\forall b \in \mathbf{Expression}. (\forall a \in \mathbf{Expression}. a \sqsubset b \rightarrow P(a)) \rightarrow P(b)$$

More explicitly, we are provided with the hypothesis that  $P(a)$  already holds for all sub-expressions  $a \sqsubset b$  when proving  $P(b)$ ; in those cases when  $b$  has no sub-expressions, we must show that  $P(b)$  holds directly.

The application of well-founded induction to the structure of an inductive definition is called *structural induction*: to prove that a property  $P(x)$  holds for all members  $x$  of an inductively defined structure  $X$ , it suffices to initially show that  $P(x)$  holds in all the base cases in the definition of  $X$ , and that  $P(x)$  holds in the inductive cases assuming that  $P(x')$  holds for any immediate substructure  $x'$  of  $x$ .

Our earlier reduction rules  $\Downarrow$  along with  $\mapsto$  and its transitive closure  $\mapsto^*$  are similarly inductively defined, and therefore admits the same notion of structural induction. These instances will be referred to as *rule induction*.

### 3.2.2 Proofs of Semantic Equivalence

We shall illustrate the above technique with some examples.

**Theorem 3.1.** *Denotational semantics and big-step operational semantics coincide:*

$$\forall e \in \mathbf{Expression}, m \in \mathbb{N}. \llbracket e \rrbracket \equiv m \leftrightarrow e \Downarrow m$$

*Proof.* We consider each direction of the  $\leftrightarrow$  equivalence separately. To show  $\llbracket e \rrbracket \equiv$



### 3.2. EQUIVALENCE PROOFS AND TECHNIQUES

$m \rightarrow e \Downarrow m$ , we ought to proceed by induction on the definition of the  $\llbracket - \rrbracket$  function. As it happens to be structurally recursive on its argument, we may equivalently proceed by structural induction on  $e$ , giving us two cases to consider:

**Case  $e \equiv n$ :** Substituting  $e$ , this base case specialises to showing that:

$$\llbracket n \rrbracket \equiv m \rightarrow n \Downarrow m$$

By (**denote-val**) in the definition of  $\llbracket - \rrbracket$ , the hypothesis evaluates to  $n \equiv m$ . This allows us to substitute  $m$  for  $n$  in the conclusion, which is trivially satisfied by instantiating (**big-val**) with  $m$  in the definition of  $\Downarrow$ .

**Case  $e \equiv a \oplus b$ :** Substituting  $e$  as before, we need to show that:

$$\llbracket a \oplus b \rrbracket \equiv m \rightarrow a \oplus b \Downarrow m$$

Applying (**denote-plus**) once to the hypothesis, we obtain that  $\llbracket a \rrbracket + \llbracket b \rrbracket \equiv m$ . Substituting for  $m$ , the conclusion becomes  $a \oplus b \Downarrow \llbracket a \rrbracket + \llbracket b \rrbracket$ . Instantiate the induction hypothesis twice with the trivial equalities  $\llbracket a \rrbracket \equiv \llbracket a \rrbracket$  and  $\llbracket b \rrbracket \equiv \llbracket b \rrbracket$  to yield proofs of  $a \Downarrow \llbracket a \rrbracket$  and  $b \Downarrow \llbracket b \rrbracket$ , which are precisely the two antecedents required by (**big-plus**) to obtain  $a \oplus b \Downarrow \llbracket a \rrbracket + \llbracket b \rrbracket$ .

The second half of the proof requires us to show that  $\llbracket e \rrbracket \equiv m \leftarrow e \Downarrow m$ . We may proceed by rule induction directly on our assumed hypothesis of  $e \Downarrow m$ , which must match either (**big-val**) or (**big-plus**) in the definition of  $\Downarrow$ :

**Rule (**big-val**):** Matching  $e \Downarrow m$  with the consequent of (**big-val**), we may conclude that there exists an  $n$  such that  $e \equiv n$  and  $m \equiv n$ . Substituting  $n$  for  $e$  and  $m$  in  $\llbracket e \rrbracket \equiv m$  and applying (**denote-val**) once, the desired conclusion becomes  $n \equiv n$ , which is trivially satisfied by the reflexivity of  $\equiv$ .

**Rule (big-plus):** Again by matching  $e \Downarrow m$  with the consequent of **(big-plus)**, there exists  $a, b, n_a$  and  $n_b$  where  $e \equiv a \oplus b$  and  $m \equiv n_a + n_b$ , such that  $a \Downarrow n_a$  and  $b \Downarrow n_b$ . Substituting for  $e$  and  $m$ , the conclusion becomes  $\llbracket a \oplus b \rrbracket \equiv n_a + n_b$ , which reduces to:

$$\llbracket a \rrbracket + \llbracket b \rrbracket \equiv n_a + n_b$$

by applying **(denote-plus)** once. Instantiating the induction hypothesis twice with  $a \Downarrow n_a$  and  $b \Downarrow n_b$  yields the equalities  $\llbracket a \rrbracket \equiv n_a$  and  $\llbracket b \rrbracket \equiv n_b$  respectively, which allows us to rewrite the conclusion as  $\llbracket a \rrbracket + \llbracket b \rrbracket \equiv \llbracket a \rrbracket + \llbracket b \rrbracket$  by substituting  $n_a$  and  $n_b$ . The desired result is now trivially true by reflexivity of  $\equiv$ .

Thus we have shown both directions of the theorem. □

**Theorem 3.2.** *Big-step and small-step operational semantics coincide. That is,*

$$\forall e \in \text{Expression}, m \in \mathbb{N}. \quad e \Downarrow m \leftrightarrow e \mapsto^* m$$

*Proof.* We shall consider each direction separately as before. To show the forward implication, we proceed by rule induction on the assumed  $e \Downarrow m$  hypothesis:

**Rule (big-val):** There exists an  $n$  such that  $e \equiv n$  and  $m \equiv n$ , by matching  $e \Downarrow m$  with the consequent of **(big-val)**. Substituting  $n$  for both  $e$  and  $m$ , we can readily conclude that  $n \mapsto^* n$  via **(small-nil)**.

**Rule (big-plus):** There exists  $a, b, n_a$  and  $n_b$  where  $e \equiv a \oplus b$  and  $m = n_a + n_b$ , such that  $a \Downarrow n_a$  and  $b \Downarrow n_b$ . After substituting for  $e$  and  $m$ , the desired conclusion becomes:

$$a \oplus b \mapsto^* n_a + n_b$$

Instantiating the induction hypothesis with  $a \Downarrow n_a$  and  $b \Downarrow n_b$  gives us evidence of  $a \mapsto^* n_a$  and  $b \mapsto^* n_b$  respectively. With the former, we can apply  $- \oplus b$  to

### 3.2. EQUIVALENCE PROOFS AND TECHNIQUES

each of the terms and **(small-left)** to obtain a proof of  $a \oplus b \mapsto^* n_a \oplus b$ , while with the latter, we obtain  $n_a \oplus b \mapsto^* n_a \oplus n_b$  by applying  $n_a \oplus \_$  and **(small-right)**.

By the transitivity of  $\mapsto^*$ , these two small-step reduction sequences combine to give  $a \oplus b \mapsto^* n_a \oplus n_b$ , to which we need only append an instance of **(small-plus)** to arrive at the conclusion.

We proceed by induction over the definition of  $\mapsto^*$  and using an additional lemma that we state and prove afterwards. Given  $e \mapsto^* m$ ,

**Rule (small-nil):** If the reduction sequence is empty, then it follows that  $e \equiv m$ . In this case, we can trivially satisfy the conclusion of  $m \Downarrow m$  with **(big-val)**.

**Rule (small-cons):** For non-empty reduction sequences, there exists an  $e'$  such that  $e \mapsto e'$  and  $e' \mapsto^* m$ . Invoke lemma 3.3 below with  $e \mapsto e'$  and  $e' \Downarrow m$ , where the latter is given by the induction hypothesis, to obtain proof of  $e \Downarrow m$ .

Pending the proof of lemma 3.3 below, we have thus shown the equivalence of big- and small-step semantics for the **Expression** language. □

**Lemma 3.3.** *A single small-step reduction preserves the value of expressions with respect to the big-step semantics:*

$$\forall e, e' \in \text{Expression}, m \in \mathbb{N}. \quad e \mapsto e' \rightarrow e' \Downarrow m \rightarrow e \Downarrow m$$

*Proof.* Assume the two premises  $e \mapsto e'$  and  $e' \Downarrow m$ , and proceed by induction on the structure of the first:

**Rule (small-plus):** There are  $n_a$  and  $n_b$  such that  $e \equiv n_a \oplus n_b$  and  $e' \equiv n_a + n_b$ . As  $e'$  is a numeric expression, the only applicable rule for  $e' \Downarrow m$  is **(big-val)**, which implies  $m \equiv n_a + n_b$ . Thus the desired conclusion of  $e \Downarrow m$ —after substituting

for  $e$  and  $m$ —may be satisfied as follows:

$$\begin{array}{c}
 \text{(big-val)} \quad \frac{\quad}{n_a \Downarrow n_a} \quad \frac{\quad}{n_b \Downarrow n_b} \\
 \text{(big-plus)} \quad \frac{\quad}{n_a \oplus n_b \Downarrow n_a + n_b}
 \end{array}$$

**Rule (small-right):** There exists  $n_a$ ,  $b$  and  $b'$  such that  $b \mapsto b'$  with  $e \equiv n_a \oplus b$  and  $e' \equiv n_a \oplus b'$ . Substituting for  $e'$ , the second assumption becomes  $n_a \oplus b' \Downarrow m$ , with (big-plus) as the only matching rule. This implies the existence of the premises  $n_a \Downarrow n_a$  and  $b' \Downarrow n_b$ ,

$$\frac{\frac{\quad}{n_a \Downarrow n_a} \quad \frac{\quad}{b' \Downarrow n_b}}{n_a \oplus b' \Downarrow n_a + n_b} \quad \vdots$$

for some  $n_b$  such that  $m \equiv n_a + n_b$ . Invoking the induction hypothesis with  $b \mapsto b'$  and the above derivation of  $b' \Downarrow n_b$ , we obtain a proof of  $b \Downarrow n_b$ . The conclusion is satisfied by the following derivation:

$$\frac{\frac{\quad}{n_a \Downarrow n_a} \quad \frac{\quad}{b \Downarrow n_b}}{n_a \oplus b \Downarrow n_a + n_b} \quad \vdots \quad \text{(IH)}$$

**Rule (small-left):** This case proceeds in a similar manner to the previous rule, but with  $a$ ,  $a'$  and  $b$  such that  $a \mapsto a'$ , where  $e \equiv a \oplus b$  and  $e' \equiv a' \oplus b$ . Substituting for  $e$  and  $e'$  in the second assumption and inspecting its premises, we observe

that  $a' \Downarrow n_a$  and  $b \Downarrow n_b$  for some  $n_a$  and  $n_b$  where  $m \equiv n_a + n_b$ :

$$\frac{\frac{\vdots}{a' \Downarrow n_a} \quad \frac{\vdots}{b \Downarrow n_b}}{a' \oplus b \Downarrow n_a + n_b}$$

Instantiating the induction hypothesis with  $a \mapsto a'$  and  $a' \Downarrow n_a$  delivers evidence of  $a \Downarrow n_a$ . Reusing the second premise of  $b \Downarrow n_b$  verbatim, we can then derive the conclusion of  $a \oplus b \Downarrow n_a + n_b$ :

$$\text{(IH)} \frac{\frac{\vdots}{a \Downarrow n_a} \quad \frac{\vdots}{b \Downarrow n_b}}{a \oplus b \Downarrow n_a + n_b}$$

This completes the proof of  $e \mapsto e' \rightarrow e' \Downarrow m \rightarrow e \Downarrow m$ . □

### 3.3 Compiler Correctness

Now that we have established the equivalence of our three semantics for the expression language, we consider how this language may be compiled for a simple stack-based machine, what it means for such a compiler to be correct, and how this may be proved.

#### 3.3.1 A Stack Machine and Its Semantics

Unlike the previously defined high-level semantics—which operate directly on **Expressions** themselves—real processors generally execute a *linear* sequences of instructions, each mutating the state of the machine in some primitive way. In order to give such a low-level implementation of our **Expression** language, we will make use of a stack-based

virtual machine.

Our stack machine has a stack of natural numbers as its sole form of storage, and the state of the **Machine** at any point may be conveniently represented as the pair of the currently executing **Code**, along with the current **Stack**,

$$\begin{aligned} \text{Machine} &::= \langle \text{Code}, \text{Stack} \rangle \\ \text{Code} &::= [] \mid \text{Instruction} :: \text{Code} \\ \text{Stack} &::= [] \mid \mathbb{N} :: \text{Stack} \end{aligned}$$

where **Code** comprises a sequence of **Instructions**, and **Stack** a sequence of values.

Due to the simple nature of the **Expression** language, the virtual machine only requires two **Instructions**, both of which operate on the top of the stack:

$$\text{Instruction} ::= \text{PUSH } \mathbb{N} \mid \text{ADD}$$

The **PUSH**  $m$  instruction places the number  $m$  on top of the current stack, while **ADD** replaces the top two values with their sum. Formally, the semantics of the virtual machine is defined by the  $\mapsto$  reduction relation, given below:

$$\begin{aligned} \langle \text{PUSH } m :: c, \sigma \rangle &\mapsto \langle c, m :: \sigma \rangle && \text{(vm-push)} \\ \langle \text{ADD} :: c, n :: m :: \sigma \rangle &\mapsto \langle c, m + n :: \sigma \rangle && \text{(vm-add)} \end{aligned}$$

As with the previous definition of  $\mapsto^*$ , we shall write  $\mapsto^*$  for the transitive, reflexive closure of  $\mapsto$ :

$$\begin{array}{c} \frac{}{t \mapsto^* t} \text{ (vm-nil)} \qquad \frac{a \mapsto a' \quad a' \mapsto^* b}{a \mapsto^* b} \text{ (vm-cons)} \end{array}$$

Informally, the difference between the semantics of a virtual machine versus a small-step operational semantics is that the reduction rules for the former is simply a collection of transition rules between pairs of states, and does not make use of any premises.

### 3.3.2 Compiler

Given an **Expression**, a compiler in this context produces some **Code** that when executed according to the semantics of the virtual machine just defined, computes the value of the **Expression**, leaving the result on top of the current stack. To avoid the need to define concatenation on instruction sequences and the consequent need to prove various distributive properties, the definition of *compile* below accepts an extra *code continuation* argument, to which the code for the expression being compiled is prepended. To compile a top-level expression, we simply pass in the empty sequence []. This both simplifies reasoning and results in more efficient compilers [Hut07]. A numeric expression  $m$  is compiled to a **PUSH**  $m$  instruction, while  $a \oplus b$  involves compiling the sub-expressions  $a$  and  $b$  in turn, followed by an **ADD** instruction:

$$\text{compile} : \text{Expression} \rightarrow \text{Code} \rightarrow \text{Code}$$

$$\text{compile } m \ c = \text{PUSH } m :: c \quad (\text{compile-val})$$

$$\text{compile } (a \oplus b) \ c = \text{compile } a \ (\text{compile } b \ (\text{ADD} :: c)) \quad (\text{compile-add})$$

For example,  $\text{compile } ((1 \oplus 2) \oplus (4 \oplus 8)) \ []$  produces the code below,

$$\text{PUSH } 1 :: \text{PUSH } 2 :: \text{ADD} :: \text{PUSH } 4 :: \text{PUSH } 8 :: \text{ADD} :: \text{ADD} :: []$$

which when executed with an empty initial stack, reduces as follows:

$$\begin{aligned}
& \langle \text{PUSH } 1 :: \text{PUSH } 2 :: \dots, [] \rangle \\
& \mapsto \langle \text{PUSH } 2 :: \text{ADD} :: \dots, 1 :: [] \rangle \\
& \mapsto \langle \text{ADD} :: \text{PUSH } 4 :: \dots, 2 :: 1 :: [] \rangle \\
& \mapsto \langle \text{PUSH } 4 :: \text{PUSH } 8 :: \dots, 3 :: [] \rangle \\
& \mapsto \langle \text{PUSH } 8 :: \text{ADD} :: \dots, 4 :: 3 :: [] \rangle \\
& \mapsto \langle \text{ADD} :: \text{ADD} :: [], 8 :: 4 :: 3 :: [] \rangle \\
& \mapsto \langle \text{ADD} :: [], 12 :: 3 :: [] \rangle \\
& \mapsto \langle [], 15 :: [] \rangle
\end{aligned}$$

### 3.3.3 Compiler Correctness

By compiler correctness, we mean that given an expression  $e$  which evaluates to  $m$  according to a high-level semantics, compiling  $e$  and executing the resultant code on the corresponding virtual machine must compute the same  $m$ . Earlier in the chapter, we had shown the equivalence of our denotational, big-step, and small-step semantics. While we may freely choose any of these as our high-level semantics, we shall adopt the big-step semantics, as it makes our later proofs much shorter.

Using these ideas, the correctness of our compiler can now be formalised by the following equivalence:

$$e \Downarrow m \quad \leftrightarrow \quad \langle \text{compile } e \quad [], \sigma \rangle \mapsto^* \langle [], m :: \sigma \rangle$$

The  $\rightarrow$  direction corresponds to a notion of *completeness*, and states that the machine must be able to compute any  $m$  that the big-step semantics permits. Conversely, the  $\leftarrow$  direction corresponds to *soundness*, and states that the machine can only produce



values permitted by the big-step semantics. For this proof, we will need to generalise the virtual machine on the right hand side to an arbitrary code continuation and stack.

**Theorem 3.4** (Compiler Correctness).

$$e \Downarrow m \quad \leftrightarrow \quad \forall c, \sigma. \langle \text{compile } e \ c, \ \sigma \rangle \mapsto^* \langle c, \ m :: \sigma \rangle$$

*Proof.* We shall consider each direction of the double implication separately. In the forward direction, we assume  $e \Downarrow m$  and proceed on its structure:

**Rule (big-val):** There exists an  $n$  such that  $e \equiv n$  and  $m \equiv n$ . Substituting  $n$  for both  $e$  and  $m$ , the conclusion becomes:

$$\begin{aligned} & \langle \text{compile } n \ c, \ \sigma \rangle \mapsto^* \langle c, \ n :: \sigma \rangle \quad , \\ \text{or} \quad & \langle \text{PUSH } n :: c, \ \sigma \rangle \mapsto^* \langle c, \ n :: \sigma \rangle \end{aligned}$$

by (compile-val) in the definition of *compile*. The conclusion is satisfied by simply applying (vm-cons) to (vm-push) and (vm-nil):

$$\begin{array}{c} \text{(vm-push)} \frac{\quad}{\langle \text{PUSH } n :: c, \ \sigma \rangle \mapsto \langle c, \ n :: \sigma \rangle} \qquad \frac{\quad}{\langle c, \ n :: \sigma \rangle \mapsto^* \langle c, \ n :: \sigma \rangle} \text{(vm-nil)} \\ \text{(vm-cons)} \frac{\quad}{\langle \text{PUSH } n :: c, \ \sigma \rangle \mapsto^* \langle c, \ n :: \sigma \rangle} \end{array}$$

**Rule (big-plus):** By matching the assumed  $e \Downarrow m$  with the consequent of (big-plus), we see that there exists  $a, b, n_a$  and  $n_b$  where  $e \equiv a \oplus b$  and  $m \equiv n_a + n_b$ , such that  $a \Downarrow n_a$  and  $b \Downarrow n_b$ . Substituting for  $e$  and  $m$ , the conclusion becomes

$$\begin{aligned} & \langle \text{compile } (a \oplus b) \ c, \ \sigma \rangle \mapsto^* \langle c, \ n_a + n_b :: \sigma \rangle \quad , \text{ or} \\ & \langle \text{compile } a \ (\text{compile } b \ (\text{ADD} :: c)), \ \sigma \rangle \mapsto^* \langle c, \ n_a + n_b :: \sigma \rangle \end{aligned}$$

by expanding *compile*. Instantiating the induction hypothesis with  $a \Downarrow n_a$  and  $b \Downarrow n_b$  yields proofs of

$$\begin{aligned} \forall c_a, \sigma_a. \langle \text{compile } a \ c_a, \ \sigma_a \rangle \rightsquigarrow^* \langle c_a, \ n_a :: \sigma_a \rangle, \text{ and} \\ \forall c_b, \sigma_b. \langle \text{compile } b \ c_b, \ \sigma_b \rangle \rightsquigarrow^* \langle c_b, \ n_b :: \sigma_b \rangle \end{aligned}$$

respectively. Note that crucially, these two hypotheses are universally quantified over  $c$  and  $\sigma$ , written with distinct subscripts above to avoid ambiguity. Now substitute  $c_b = \text{ADD} :: c$ ,  $c_a = \text{compile } b \ c_b$ ,  $\sigma_a = \sigma$ ,  $\sigma_b = n_a :: \sigma_a$  and we obtain via the transitivity of  $\rightsquigarrow^*$ :

$$\begin{aligned} \forall c, \sigma. \langle \text{compile } a \ (\text{compile } b \ (\text{ADD} :: c)), \ \sigma \rangle \\ \rightsquigarrow^* \langle (\text{compile } b \ (\text{ADD} :: c), \ n_a :: \sigma \rangle \\ \rightsquigarrow^* \langle \text{ADD} :: c, \ n_b :: n_a :: \sigma \rangle \end{aligned}$$

A second application of transitivity to (vm-add) instantiated as follows,

$$\langle \text{ADD} :: c, \ n_b :: n_a :: \sigma \rangle \rightsquigarrow \langle c, \ n_a + n_b :: \sigma \rangle$$

gives the required conclusion of:

$$\forall c, \sigma. \langle \text{compile } a \ (\text{compile } b \ (\text{ADD} :: c)), \ \sigma \rangle \rightsquigarrow \langle c, \ n_a + n_b :: \sigma \rangle$$

For the backward direction, we proceed on the structure of  $e$ :

**Case  $e \equiv n$ :** Substituting  $e$  with  $n$ , the base case becomes:

$$\begin{aligned} \forall c, \sigma. \langle \text{compile } n \ c, \ \sigma \rangle \rightsquigarrow^* \langle c, \ m :: \sigma \rangle \rightarrow n \Downarrow m, \text{ or} \\ \forall c, \sigma. \langle \text{PUSH } n :: c, \ \sigma \rangle \rightsquigarrow^* \langle c, \ m :: \sigma \rangle \rightarrow n \Downarrow m \end{aligned}$$

Assume the hypothesis and set both  $c$  and  $\sigma$  to  $[]$  to obtain  $\langle \text{PUSH } n :: [], [] \rangle \mapsto^* \langle [], m :: [] \rangle$ , which must be a single reduction corresponding to **(vm-push)**. Therefore  $m$  and  $n$  must be one and the same, and the conclusion of  $n \Downarrow n$  is trivially satisfied by **(big-val)**.

**Case  $e \equiv a \oplus b$ :** Substituting  $e$  with  $a \oplus b$  and expanding the definition of *compile*, we need to show that:

$$\forall c, \sigma. \langle \text{compile } a (\text{compile } b (\text{ADD} :: c)), \sigma \rangle \mapsto^* \langle c, m :: \sigma \rangle \rightarrow a \oplus b \Downarrow m$$

Now, for both  $a$  and  $b$ , we know that there exists  $n_a$  and  $n_b$  such that:

$$\begin{aligned} \forall c_a, \sigma_a. \langle \text{compile } a c_a, \sigma_a \rangle \mapsto^* \langle c_a, n_a :: \sigma_a \rangle, \text{ and} \\ \forall c_b, \sigma_b. \langle \text{compile } b c_b, \sigma_b \rangle \mapsto^* \langle c_b, n_b :: \sigma_b \rangle \end{aligned}$$

Substituting for the subscripted  $c_a, c_b, \sigma_a$  and  $\sigma_b$  as we had done in the **(big-plus)** case of the first half of this proof, we obtain:

$$\forall c, \sigma. \langle \text{compile } a (\text{compile } b (\text{ADD} :: c)), \sigma \rangle \mapsto^* \langle c, n_a + n_b :: \sigma \rangle$$

Contrast this with the hypothesis:

$$\forall c, \sigma. \langle \text{compile } a (\text{compile } b (\text{ADD} :: c)), \sigma \rangle \mapsto^* \langle c, m :: \sigma \rangle$$

Since the  $\mapsto$  reduction relation is deterministic, it must be the case that  $m$  and  $n_a + n_b$  are the same. Substituting in  $n_a + n_b$  for  $m$ , the conclusion becomes  $a \oplus b \Downarrow n_a + n_b$ —an instance of **(big-plus)**—whose premises of  $a \Downarrow n_a$  and  $b \Downarrow n_b$  are in turn satisfied by instantiating the induction hypothesis with  $\forall c_a, \sigma_a. \langle \text{compile } a c_a, \sigma_a \rangle \mapsto^* \langle c_a, n_a :: \sigma_a \rangle$  and  $\forall c_b, \sigma_b. \langle \text{compile } b c_b, \sigma_b \rangle \mapsto^*$

$$\langle c_b, n_b :: \sigma_b \rangle.$$

This completes the proof of the compiler correctness theorem. □

## 3.4 Conclusion

In this chapter, we have shown by example what it means to give the semantics of simple language in denotational, big-step and small-step styles. We justified the use of a monoidal model of natural numbers and addition—with left-to-right evaluation order—as simplification of monadic sequencing. We then proved the three given semantics to be equivalent, and demonstrate the use of well-founded induction on the structure of the reduction rules (that is, rule induction) and of the syntax. Finally, we defined a stack-based virtual machine and a compiler for running programs of the Expression language, and presented a proof of compiler correctness.

# Chapter 4

## Randomised Testing in Haskell

During the initial development of a language, it is useful to be able to check that its behaviour agrees with our intuitive expectations. Formal proofs require a significant investment of effort, and are not always straightforward to revise in light of any underlying changes. Using Haskell, we can implement our semantics as an executable program: its high-level expressiveness leads to a much narrower semantic gap between the mathematical definitions and the implementation, giving us much greater confidence in the fidelity of the latter, in contrast to more traditional programming languages.

In turn, if we state the expected properties as Boolean functions in Haskell, these can then be subjected to randomly generated inputs using the QuickCheck tool [CH00], which displays any counterexamples found. While successful runs of hundreds and thousands of such tests do not comprise a formal proof of the corresponding property, they do however corroborate the existence of one. Additionally, the use of the Haskell Program Coverage (HPC) toolkit offers confidence in the validity of these tests by highlighting any unexercised fragments of the implementation.

## 4.1 Executable Semantics

Let us begin by implementing the syntax of the expression language of the previous chapter:

```
data Expression = Val Integer | Expression  $\oplus$  Expression
```

The `Expression` algebraic data type defined above has two constructors `Val` and `( $\oplus$ )`, corresponding to numerals and addition.

### 4.1.1 Denotational Semantics

The denotational semantics for our language given in the previous chapter map expressions to its underlying domain of numbers:

$$\begin{aligned} \llbracket \_ \rrbracket &: \text{Expression} \rightarrow \mathbb{N} \\ \llbracket m \rrbracket &= m && \text{(denote-val)} \\ \llbracket a \oplus b \rrbracket &= \llbracket a \rrbracket + \llbracket b \rrbracket && \text{(denote-plus)} \end{aligned}$$

This can be directly implemented as the following `denot` function, mapping any given `Expression` to an `Integer`:

```
denot :: Expression  $\rightarrow$  Integer
denot (Val m) = m
denot (a  $\oplus$  b) = denot a + denot b
```

### 4.1.2 Big-Step Operational Semantics

While denotations can be implemented directly as functions, there is in general no corresponding notion in Haskell for the transition relations given for a definition of

operational semantics, since transitions need not necessarily be deterministic. Nevertheless, we can accommodate non-deterministic transitions by implementing them as set-valued functions, that return a set of possible reducts for each expression.

Let us first define two type synonyms `REL` and `Rel` corresponding to heterogeneous and homogeneous relations respectively:

```
type REL  $\alpha$   $\beta$  =  $\alpha \rightarrow$  Set  $\beta$ 
```

```
type Rel  $\alpha$  = REL  $\alpha$   $\alpha$ 
```

That is, a heterogeneous relation between  $\alpha$  and  $\beta$  can be realised as a function from  $\alpha$  to a set of possible  $\beta$ s.

For convenience, we also define some auxiliary functions not found in the Haskell standard library. Firstly `joinSet` flattens nested sets,

```
joinSet :: Set (Set  $\alpha$ )  $\rightarrow$  Set  $\alpha$ 
```

```
joinSet = Set.fold ( $\cup$ )  $\{\}$ 
```

by folding the binary set union operation ( $\cup$ ) over the outer set. Note that we have typeset `Set.union` and `Set.empty` as their usual mathematical notations.

The `productWith` function computes a generalised Cartesian product of two sets, combining the elements from each set with the function  $f$ :

```
productWith :: ( $\alpha \rightarrow \beta \rightarrow \gamma$ )  $\rightarrow$  Set  $\alpha \rightarrow$  Set  $\beta \rightarrow$  Set  $\gamma$ 
```

```
productWith  $f$  as bs = joinSet (( $\lambda a \rightarrow f$  a 'Set.map' bs) 'Set.map' as)
```

Let us remind ourselves of the big-step  $\Downarrow$  relation given in the previous chapter:

$$\frac{}{m \Downarrow m} \quad \text{(big-val)}$$

$$\frac{a \Downarrow m \quad b \Downarrow n}{a \oplus b \Downarrow m + n} \quad \text{(big-plus)}$$

Using our previously defined helpers, we can now implement the  $\Downarrow$  relation as the following `bigStep` function, where `Set.singleton m` is rendered as  $\{m\}$ :

```
bigStep :: REL Expression Integer
bigStep (Val m) = {m}
bigStep (a ⊕ b) = productWith (+) (bigStep a) (bigStep b)
```

The first case corresponds to the (`big-val`) rule, with `Val m` reducing to  $m$ . The second case of  $a \oplus b$  recursively computes the possible reducts for  $a$  and  $b$ , then combines the two sets using `productWith (+)` to obtain the possible values corresponding to  $m + n$ .

### 4.1.3 Small-Step Operational Semantics

In the previous chapter, we had defined the small-step reduction relation  $\mapsto$  as follows:

$$\frac{}{m \oplus n \mapsto m + n} \quad \text{(small-plus)}$$

$$\frac{b \mapsto b'}{m \oplus b \mapsto m \oplus b'} \quad \text{(small-right)}$$

$$\frac{a \mapsto a'}{a \oplus b \mapsto a' \oplus b} \quad \text{(small-left)}$$

The above rules are implemented as the three cases of the `smallStep` function below, corresponding to (`small-plus`), (`small-right`) and (`small-left`) in that order:

```
smallStep :: Rel Expression
smallStep (Val m ⊕ Val n) = {Val (m + n)}
smallStep (Val m ⊕ b     ) = (λb' → Val m ⊕ b') 'Set.map' smallStep b
smallStep (a     ⊕ b     ) = (λa' → a'     ⊕ b) 'Set.map' smallStep a
```



The three patterns for `smallStep` above are not exhaustive, as we are missing a case for `Val m`. Since such expressions cannot reduce any further, we return an empty set to indicate the lack of such transitions:

$$\text{smallStep } (\text{Val } m) = \{\}$$

In a small-step semantics, the eventual result of a computation is given by repeated application of the small-step reduction relation to the initial expression. The following `reduceUntil`—parametrised on the reduction relation—performs this task:

$$\text{reduceUntil} :: (\alpha \rightarrow \text{Maybe } \beta) \rightarrow \text{Rel } \alpha \rightarrow \text{REL } \alpha \beta$$

$$\text{reduceUntil } p \text{ reduce } \text{init} = \text{step } (\{\text{init}\}, \{\}) \text{ where}$$

$$\text{step} :: (\text{Set } \alpha, \text{Set } \beta) \rightarrow \text{Set } \beta$$

$$\text{step } (\text{running}, \text{finished}) = \text{case } \text{Set.null } \text{running} \text{ of}$$

$$\text{True} \rightarrow \text{finished}$$

$$\text{False} \rightarrow \text{step } (\text{first } (\text{join}_{\text{Set}} \circ \text{Set.map } \text{reduce})$$

$$(\text{Set.fold } \text{partition } (\{\}, \text{finished}) \text{ running}))$$

$$\text{partition} :: \alpha \rightarrow (\text{Set } \alpha, \text{Set } \beta) \rightarrow (\text{Set } \alpha, \text{Set } \beta)$$

$$\text{partition } e = \text{case } p \ e \ \text{of}$$

$$\text{Nothing} \rightarrow \text{first } (\text{Set.insert } e)$$

$$\text{Just } n \rightarrow \text{second } (\text{Set.insert } n)$$

The `step` helper takes a pair of *running* and *finished* sets of states, accumulating those that satisfy *p* into the finished set for the next iteration with the aid of `partition`, and repeatedly applies *reduce* to the set of remaining running states until it becomes exhausted.

Together with the auxiliary `isVal` function,

$$\text{isVal} :: \text{Expression} \rightarrow \text{Maybe Integer}$$

$$\text{isVal } (\text{Val } n) = \text{Just } n$$

$$\text{isVal } \_ = \text{Nothing}$$

we obtain an executable implementation of  $\mapsto^*$  as the following Haskell function:

```
smallStepStar :: REL Expression Integer
smallStepStar = reduceUntil isVal smallStep
```

#### 4.1.4 Virtual Machine

The virtual machine presented previously is implemented in a similarly straightforward manner. We begin with a trio of type synonyms defining what comprises a virtual machine:

```
type Machine = (Code, Stack)
type Code = [Instruction]
type Stack = [Integer]
```

We make use of Haskell’s built-in list type for `Code` and `Stack` rather than giving our own definitions. Instructions are defined as the following algebraic data type, with one constructor for each corresponding instruction:

```
data Instruction = PUSH Integer | ADD
```

Intuitively, a `PUSH`  $m$  instruction places the number  $m$  onto the top of the stack, while the `ADD` instruction replaces the top two numbers on the stack with their sum:

$$\langle \text{PUSH } m :: c, \sigma \rangle \mapsto \langle c, m :: \sigma \rangle \quad (\text{vm-push})$$

$$\langle \text{ADD} :: c, n :: m :: \sigma \rangle \mapsto \langle c, m + n :: \sigma \rangle \quad (\text{vm-add})$$

The (vm-push) and (vm-add) rules are directly implemented as the first two rules of the following `stepVM` function:

```
stepVM :: Rel Machine
stepVM (PUSH m : c,      \sigma) = {(c, m : \sigma)}
```

```

stepVM (ADD      : c, n : m : σ) = {(c, m + n : σ)}
stepVM (c        ,      σ) = {}

```

Since no other transitions are possible for the virtual machine, we return the empty set in the final catch-all case.

The virtual machine is considered halted when its sequence of instructions is exhausted, and the stack consists of only a single number corresponding to the result of the computation:

```

isHalted :: Machine → Maybe Integer
isHalted ([], n : []) = Just n
isHalted _           = Nothing

```

In the same way as small-step semantics, we may make use of our earlier `reduceUntil` function to repeatedly iterate `stepVM` until the virtual machine halts:

```

stepVMStar :: REL Machine Integer
stepVMStar = reduceUntil isHalted stepVM

```

Finally, the compiler is more-or-less a direct transliteration of our previous definition:

```

compile :: Expression → Code → Code
compile (Val n) c = PUSH n : c
compile (a ⊕ b) c = compile a (compile b (ADD : c))

```

## 4.2 Randomised Testing with QuickCheck and HPC

QuickCheck is a system for testing properties of Haskell programs with randomly-generated inputs. Informally, properties are specified as Haskell functions that return

a boolean result. For example, to assert that  $(2 \times n) \div 2 \equiv n$  holds in virtually all cases for floating-point numbers, we may interactively invoke the following call to `quickCheck`:

```
*Main> quickCheck (\n -> (2 * n) / 2 == n)
+++ OK, passed 100 tests.
```

It is important to highlight the fact that unlike certain model-checking tools, `QuickCheck` does not attempt to exhaustively generate all possible inputs. Thus even given many successful repeated invocations of `quickCheck`, some rare corner cases may remain unprobed. For example, even discounting the possibility of *NaN* and *infinity*, the following expression evaluates to `False` for 32-bit IEEE-754 floating-point numbers, due to its limited range and finite precision:

```
*Main> (\n -> (2 * n) / 2 == n) (3e38 :: Float)
False
```

`QuickCheck` does not obviate the need for formal proofs. However, it is nevertheless very helpful while the implementation is still in a state of flux, allowing us to detect many flaws ahead of committing to more rigorous and laborious analyses.

### 4.2.1 Generating Arbitrary Expressions

While `QuickCheck` comes with existing generators for many built-in Haskell data types, custom generators can also be seamlessly added for new data types defined in our own programs. This is achieved using Haskell's *type class* mechanism—a form of ad-hoc polymorphism based on type-directed dispatch that in essence allows us to overload function names for more than one type.

The following code defines the generator for an arbitrary expression, by instantiating the `Arbitrary` type class for `Expressions`:

**instance Arbitrary Expression where**

**arbitrary** :: Gen Expression

**arbitrary** = oneof

[Val <\$> arbitrary  
, (⊕) <\$> arbitrary <\*> arbitrary]

The `oneof :: [(Gen α)] → Gen α` combinator above picks one of the listed generators with equal probability. In the first case, `Val` is applied to an `Integer` generated by invoking a different instance of `arbitrary`, while the latter recursively generate the two sub-expression arguments to `(⊕)`. We can test this with the following incantation:

```
*Main> sample (arbitrary :: Gen Expression)
```

```
Val 5 ⊕ (Val 8 ⊕ Val 4)
```

```
Val 1
```

```
...
```

However, we soon find that this generates some unacceptably large expressions. Writing  $\mathbb{E}[|e(n)|]$  for the *expected size* of the generated expression, we see why this is the case:

$$\mathbb{E}[|e(n)|] = \frac{1}{2} + \frac{1}{2}(\mathbb{E}[|e(n)|] + \mathbb{E}[|e(n)|]) = \infty$$

To bring the size of the generated expressions under control, we can use the `sized` combinator as follows, to allow QuickCheck some influence over the size of the resulting `Expression`:

**instance Arbitrary Expression where**

**arbitrary** :: Gen Expression

**arbitrary** = sized ( $\lambda n \rightarrow$  frequency

[( $n + 1$ , Val <\$> arbitrary)  
, ( $n$ , (⊕) <\$> arbitrary <\*> arbitrary)])

The **frequency** combinator behaves analogously to **oneof**, but chooses each alternative with a probability proportional to the accompanying weight. That is, the above definition of **arbitrary** produces a **Val** constructor with probability  $(n + 1)/(2n + 1)$ , and an  $(\oplus)$  constructor with probability  $n/(2n + 1)$ . Applying the same analysis as above, we expect the generated expressions to be much more manageable:

$$\mathbb{E}[|e(n)|] = \frac{(n + 1) + 2n\mathbb{E}[|e(n)|]}{2n + 1} = n + 1$$

When QuickCheck finds a counterexample to the proposition in question, we often find that it is rarely the smallest such case, which makes it difficult to understand exactly where and how the problem arises. QuickCheck provides a mechanism to automate the search for smaller counterexamples, via the **shrink** method of the **Arbitrary** type class:

```
shrink :: Expression → [Expression]
shrink e = case e of
  Val m → Val ($) shrinkIntegral m
  a ⊕ b → [a, b]
```

QuickCheck expects our definition of **shrink** to return a list of similar values that are ‘smaller’ in some sense. This is implemented in a straightforward manner for **Expressions** by simply returning a list of direct sub-expressions for the  $(\oplus)$  constructor. For values, we use the built-in **shrinkIntegral** to obtain a list of ‘smaller’ candidates.

## 4.2.2 Semantic Equivalence and Compiler Correctness

While we have already given a formal proof of the equivalence between denotational, big-step and small-step semantics in the previous chapter, we shall illustrate how we could have informally asserted our theorems using QuickCheck with a much smaller investment of effort.

## 4.2. RANDOMISED TESTING WITH QUICKCHECK AND HPC

Let us recall theorem 3.1, which states that our denotational and big-step semantics are equivalent:

$$\forall e \in \text{Expression}, m \in \mathbb{N}. \llbracket e \rrbracket \equiv m \leftrightarrow e \Downarrow m$$

That is, for all expressions whose denotation is  $m$ , the same expression also evaluates under our big-step semantics to  $m$ . Written literally, this corresponds to the following Haskell predicate,

```
prop_DenotBig' :: Expression -> Integer -> Bool
prop_DenotBig' e m = (denot e == m) == (bigStep e == m)
```

where  $\ni$  is the flipped `Set` membership operator. This can be checked as follows:

```
*Main> quickCheck prop_DenotBig'
+++ OK, passed 100 tests.
```

There are some subtleties involved: in the above test, QuickCheck generates a random `Integer` as well as an `Expression`, and checks that the implication holds in both directions. However, given the unlikelihood of some unrelated  $m$  coinciding with the value of  $e$  according to either semantics, both sides of the outer `==` will be `False` a majority of the time. This is clearly not a fruitful exploration of the test space.

We can write the same test much more efficiently by using only a single random `Expression`, by rephrasing the property as below:

```
prop_DenotBig :: Expression -> Bool
prop_DenotBig e = {denot e} == bigStep e
```

That is to say, `denot e` is the unique value that  $e$  can evaluate to under the big-step semantics. When fed to QuickCheck, the response is as we would expect:

```
*Main> quickCheck prop_DenotBig
+++ OK, passed 100 tests.
```

## CHAPTER 4. RANDOMISED TESTING IN HASKELL

Theorem 3.2—that is, the correspondence between big-step and small-step semantics—can be approached in the same way:

$$\forall e \in \text{Expression}, m \in \mathbb{N}. \quad e \Downarrow m \leftrightarrow e \mapsto^* m$$

Transliterated as a QuickCheck property, this corresponds to the following definition:

```
prop_BigSmall' :: Expression → Integer → Bool
prop_BigSmall' e m = bigStep e ∋ m ≡ smallStepStar e ∋ m
```

However, the above is just an instance of what it means for two sets to be equal. Thus we may elide the  $m$  argument and just use `Set`'s built-in notion of equality:

```
prop_BigSmall :: Expression → Bool
prop_BigSmall e = bigStep e ≡ smallStepStar e
```

Again, QuickCheck does not find any counterexamples:

```
*Main> quickCheck prop_BigSmall
+++ OK, passed 100 tests.
```

Finally, we revisit our previous statement of compiler correctness:

$$\forall e \in \text{Expression}, m \in \mathbb{N}. \quad e \Downarrow m \leftrightarrow \langle \text{compile } e \ [], [] \rangle \mapsto^* \langle [], m :: [] \rangle$$

Taking into account what we discussed in the previous paragraphs, the above can be defined as the following property:

```
prop_Correctness :: Expression → Bool
prop_Correctness e = smallStepStar e ≡ stepVMStar (compile e [], [])
```

Feeding this to QuickCheck yields no surprises:



```
*Main> quickCheck prop_Correctness
+++ OK, passed 100 tests.
```

### 4.2.3 Coverage Checking with HPC

With any testing process, it is important to ensure that all relevant parts of the program has been exercised during testing. The Haskell Program Coverage (HPC) toolkit [GR07] supports just this kind of analysis. It instruments every (sub-)expression in the given Haskell source file to record if it has been evaluated, and accumulates this information to a `.tix` file each time the program is executed. Analysis of the resulting `.tix` file enables us to quickly identify unevaluated fragments of code.

Using HPC is a straightforward process, as it is included with all recent releases of the Glasgow Haskell Compiler. We do however need to turn our implementation so far into a fully-fledged program, namely by implementing a `main` function:

```
main :: IO ()
main = do
    quickCheck prop_DenotBig
    quickCheck prop_BigSmall
    quickCheck prop_Correctness
```

This *literate Haskell* source file is then compiled with HPC instrumentation—then executed—as follows:

```
$ ghc -fhpc testing.lhs --make
[1 of 1] Compiling Main ( testing.lhs, testing.o )
Linking testing ...
$ ./testing
+++ OK, passed 100 tests.
```

```
+++ OK, passed 100 tests.
```

```
+++ OK, passed 100 tests.
```

Running the `hpc report` command on the resulting `testing.tix` file displays the following report,

```
94% expressions used (181/192)
82% alternatives used (19/23)
50% local declarations used (2/4)
89% top-level declarations used (26/29)
```

which is not quite the 100% coverage we might have expected. To see where, we may use the `hpc draft` command to produce a set of missing ‘ticks’, corresponding to unevaluated expressions in our code. (Alternatively, we could have visualised the results using the `hpc markup` command, which generates a highlighted copy of the source code for our perusal.) The salient<sup>1</sup> ones are listed below:

```
module "Main" {
  inside "smallStep" { tick "Set.empty" on line 201; }
  inside "stepVM"    { tick "Set.empty" on line 295; }
  tick function "shrink" on line 433;
}
```

The first two missing ticks refer to the use of `Set.empty`—otherwise typeset as `{}`—in the cases where `smallStep` and `stepVM` are invoked on expressions or virtual machine states that cannot reduce any further, which would imply the existence of a ‘stuck’ state. The third tick refers to the `shrink` function, which QuickCheck only invokes when it finds a counterexample for any of the properties under test. Thus incomplete coverage in these cases is not unexpected.

---

<sup>1</sup>The omitted ticks correspond to compiler-generated helper definitions.

## 4.3 Conclusion

The combination of QuickCheck for randomised testing and HPC to confirm complete code coverage was first pioneered in-the-large by the XMonad project [SJ07], providing a high level of assurance of implementation correctness with minimal effort. Leveraging these same tools and techniques as a means of verifying semantic properties, we have conferred similar levels of confidence that our implementation satisfies the compiler correctness property. Given the small semantic gap between our implementation and the mathematical definitions of the previous chapter, it would be reasonable to interpret this as evidence towards the correctness of theorem 3.4 for our expression language.



# Chapter 5

## A Model of STM

In this chapter, we identify a simplified subset of STM Haskell that is suitable for exploring design and implementation issues, and define a high-level stop-the-world semantics for this language. We then define a low-level virtual machine for this language in which transactions are made explicit, along with a semantics for this machine. Finally, we relate the two semantics using a compiler correctness theorem, and test the validity of this theorem using QuickCheck and HPC on an executable implementation this language.

### 5.1 A Simple Transactional Language

In chapter §2, we introduced STM Haskell, which provides a small set of primitives for working with transactional variables,

`newTVar` ::  $\alpha \rightarrow \text{STM } (\text{TVar } \alpha)$

`readTVar` :: `TVar`  $\alpha \rightarrow \text{STM } \alpha$

`writeTVar` :: `TVar`  $\alpha \rightarrow \alpha \rightarrow \text{STM } ()$

along with a primitive for running transactions, and another for explicit concurrency:

`atomically` :: `STM`  $\alpha \rightarrow$  `IO`  $\alpha$

`forkIO` :: `IO` ()  $\rightarrow$  `IO` ()

The `STM` and `IO` types are both monads (§2.4), so we may use `>>=` and `return` on both levels for sequencing effectful computations.

### 5.1.1 Syntax

As a first step towards a verified implementation of STM, let us consider a simplified model of the above, to allow us to focus on the key issues. The language we consider has a two-level syntax—mirroring that of the STM Haskell primitives—which can be represented as the following `Tran` and `Proc` data types in Haskell:

`data Tran = ValT Integer | Tran  $\oplus_T$  Tran | Read Var | Write Var Tran`

`data Proc = ValP Integer | Proc  $\oplus_P$  Proc | Atomic Tran | Fork Proc`

The two types correspond to actions in the `STM` and `IO` monads respectively. The language is intentionally minimal, because issues such as name binding and performing general-purpose computations are largely orthogonal to our goal of verifying an implementation of STM. Thus, we replace the `>>=` and `return` of both monads with the monoid of addition and integers, as motivated in section §3.1.5 of the previous chapter. In our language, `ValT` and `ValP` correspond to `return` for the `STM` and `IO` monads, while  `$\oplus_T$`  and  `$\oplus_P$`  combine `Tran` and `Proc` computations in an analogous way to `bind`. By enforcing left-to-right reduction semantics for  `$\oplus_T$`  and  `$\oplus_P$` , we nevertheless retain the fundamental idea of using monads to sequence computations and combine their results.

The remaining constructs emulate the `STM` and `IO` primitives provided by STM Haskell: `Read` and `Write` correspond to `readTVar` and `writeTVar`, where `Var` represents a finite collection of transactional variables. Due to the lack of name binding, we omit an analogue of `newTVar` from our language, and assume all variables are initialised to

zero. **Atomic** runs a transaction to completion, delivering a value, while **Fork** spawns off its argument as a concurrent process in the style of `forkIO`.

For simplicity, we do not consider **orElse** or **retry**, as they are not required to illustrate the basic implementation of a log-based transactional memory system.

## Example

Let us revisit the transactional counter example from section 2.6:

```

type CounterTVar = TVar Integer
incrementTVar :: CounterTVar → STM ()
incrementTVar c = do
  n ← readTVar c
  writeTVar c (n + 1)

```

In our STM model, the corresponding **increment** function would be written as follows:

```

increment :: Var → Tran
increment c = Write c (Read c ⊕T ValT 1)

```

To increment the same counter twice using concurrent threads, we would write:

```

incTwice :: Var → Proc
incTwice c = Fork (Atomic (increment c)) ⊕P Fork (Atomic (increment c))

```

### 5.1.2 Transaction Semantics

We specify the meaning of transactions in this language using a small-step operational semantics, following the approach of [HMPJH05]. Formally, we give a reduction relation  $\mapsto_{\top}$  on pairs  $\langle h, e \rangle$  consisting of a heap  $h :: \text{Heap}$  and a transaction  $e :: \text{Tran}$ . In this section we explain each of the rules defining  $\mapsto_{\top}$ , and simultaneously describe its implementation in order to highlight their similarity.

First, we model the heap as a map from variable names to their values—initialised to zero—and write  $h ? v$  to denote the value of variable  $v$  in the heap  $h$ . This may be implemented in Haskell using the standard `Map` datatype:

```
type Heap = Map Var Integer
(?) :: Heap → Var → Integer
h ? v = Map.findWithDefault 0 v h
```

While the  $\mapsto_{\tau}$  relation cannot be implemented directly, we may nevertheless model it as a set-valued function where each state reduces to a *set of possible results*:

```
type REL  $\alpha$   $\beta$  =  $\alpha \rightarrow$  Set  $\beta$  -- Heterogeneous binary relations
type Rel  $\alpha$  = REL  $\alpha$   $\alpha$  -- Homogeneous binary relations
reduceTran :: Rel (Heap, Tran)
```

Reading a variable  $v$  looks up its value in the heap,

$$\langle h, \text{Read } v \rangle \mapsto_{\tau} \langle h, \text{Val}_{\tau} (h ? v) \rangle \quad (\text{Read})$$

which is implemented by the following code, where we have written `Set.singleton`  $x$  as  $\{x\}$  for clarity of presentation:

```
reduceTran (h, Read v) = {h, Valτ (h ? v)}
```

Writing to a variable is taken care of by two rules: (`WriteZ`) updates the heap with the new value for a variable in the same manner as the published semantics of STM Haskell [HMPJH05], while (`WriteT`) allows its argument expression to be repeatedly



reduced until it becomes a value,

$$\langle h, \mathbf{Write} \ v \ (\mathbf{Val}_\tau \ n) \rangle \mapsto_\tau \langle h \uplus \{v \mapsto n\}, \mathbf{Val}_\tau \ n \rangle \quad (\mathbf{WriteZ})$$

$$\frac{\langle h, e \rangle \mapsto_\tau \langle h', e' \rangle}{\langle h, \mathbf{Write} \ v \ e \rangle \mapsto_\tau \langle h', \mathbf{Write} \ v \ e' \rangle} \quad (\mathbf{WriteT})$$

writing  $h \uplus \{v \mapsto n\}$  to denote the heap  $h$  with the variable  $v$  updated to  $n$ .

We implement these two rules by inspecting the subexpression  $e$  whose value we wish to update the heap with. In the the former case,  $e$  is just a plain number—corresponding to  $(\mathbf{WriteZ})$ —and we update the heap with the new value of  $v$  accordingly. The latter case implements  $(\mathbf{WriteT})$  by recursively reducing the subexpression  $e$ , then reconstructing  $\mathbf{Write} \ v \ e'$  by mapping  $\mathbf{second} \ (\mathbf{Write} \ v)$  over the resulting set of  $(h', e')$ :

$$\begin{aligned} \mathbf{reduce}_{\mathbf{Tran}} (h, \mathbf{Write} \ v \ e) &= \mathbf{case} \ e \ \mathbf{of} \\ \mathbf{Val}_\tau \ n &\rightarrow \{h \uplus \{v \mapsto n\}, \mathbf{Val}_\tau \ n\} \\ - &\rightarrow \mathbf{second} \ (\mathbf{Write} \ v) \ \mathbf{'Set.map'} \ \mathbf{reduce}_{\mathbf{Tran}} (h, e) \end{aligned}$$

As we replace the act of sequencing computations with addition in our language, it is therefore important to enforce a sequential evaluation order. The final group of three rules define reduction for  $\oplus_\tau$ , and ensure the left argument is evaluated to completion, before starting on the right hand side:

$$\langle h, \mathbf{Val}_\tau \ m \oplus_\tau \ \mathbf{Val}_\tau \ n \rangle \mapsto_\tau \langle h, \mathbf{Val}_\tau \ (m + n) \rangle \quad (\mathbf{AddZ}_\tau)$$

$$\frac{\langle h, b \rangle \mapsto_\tau \langle h', b' \rangle}{\langle h, \mathbf{Val}_\tau \ m \oplus_\tau \ b \rangle \mapsto_\tau \langle h', \mathbf{Val}_\tau \ m \oplus_\tau \ b' \rangle} \quad (\mathbf{AddR}_\tau)$$

$$\frac{\langle h, a \rangle \mapsto_\tau \langle h', a' \rangle}{\langle h, a \oplus_\tau \ b \rangle \mapsto_\tau \langle h', a' \oplus_\tau \ b \rangle} \quad (\mathbf{AddL}_\tau)$$

Our implementation of  $\oplus_{\tau}$  mirrors the above rules, as below:

```

reduceTran (h, a  $\oplus_{\tau}$  b) = case (a, b) of
  (Val $\tau$  m, Val $\tau$  n)  $\rightarrow$  {h, Val $\tau$  (m + n)}
  (Val $\tau$  m, -      )  $\rightarrow$  second (Val $\tau$  m  $\oplus_{\tau}$  -) ‘Set.map’ reduceTran (h, b)
  (-      , -      )  $\rightarrow$  second (       $\oplus_{\tau}$  b) ‘Set.map’ reduceTran (h, a)

```

To complete the definition of `reduceTran`, we require a further case,

```

reduceTran (h, Val $\tau$  m) = {}

```

where we return the empty set for `Val $\tau$  m`, as it has no associated reduction rules.

Because  $\mapsto_{\tau}$  only describes a single reduction step, we also need to implement a helper function to run a given initial expression to completion for our executable model. Let us first define `joinSet`, which flattens nested `Sets`:

```

joinSet :: Set (Set  $\alpha$ )  $\rightarrow$  Set  $\alpha$ 
joinSet = Set.fold (U) {}

```

Here, `Set.union` and `Set.empty` are written as `(U)` and `{}` respectively.

The following definition of `reduceUntil`—parameterised over a relation `reduce`—reduces the given `init` state to completion, according to the predicate `p`:

```

reduceUntil :: ( $\alpha \rightarrow$  Maybe  $\beta$ )  $\rightarrow$  Rel  $\alpha \rightarrow$  REL  $\alpha \beta$ 
reduceUntil p reduce init = step ({init}, {}) where
  step :: (Set  $\alpha$ , Set  $\beta$ )  $\rightarrow$  Set  $\beta$ 
  step (running, finished) = case Set.null running of
    True  $\rightarrow$  finished
    False  $\rightarrow$  step (first (joinSet  $\circ$  Set.map reduce)
                    (Set.fold partition ({}, finished) running))
  partition ::  $\alpha \rightarrow$  (Set  $\alpha$ , Set  $\beta$ )  $\rightarrow$  (Set  $\alpha$ , Set  $\beta$ )

```

```

partition e = case p e of
  Nothing → first (Set.insert e)
  Just n   → second (Set.insert n)

```

The `step` helper takes a pair of *running* and *finished* sets of states, accumulating those that satisfy  $p$  into the finished set for the next iteration with the aid of `partition`, and repeatedly applies *reduce* to the set of remaining running states until it becomes exhausted.

Finally, given the following `isValT` predicate,

```

isValT :: (Heap, Tran) → Maybe (Heap, Integer)
isValT (h, ValT n) = Just (h, n)
isValT (h, _       ) = Nothing

```

the expression `reduceUntil isValT reduceTran` then corresponds to an implementation of  $\mapsto_{\top}^*$ , which produces a set of `(Heap, Integer)` pairs from an initial `(Heap, Tran)`.

### 5.1.3 Process Soup Semantics

The reduction relation  $\mapsto_p$  for processes acts on pairs  $\langle h, s \rangle$  consisting of a heap  $h$  as before, and a ‘soup’  $s$  of running processes [PJ01]. While the soup itself is to be regarded as a multi-set, here we use a more concrete representation, namely a list of `Procs`.

The reduction rules for process are in general defined by matching on the first process in the soup. However, we begin by giving the `(Preempt)` rule, which allows the rest of the soup to make progress, giving rise to non-determinism in the language:

$$\frac{\langle h, s \rangle \mapsto_p \langle h', s' \rangle}{\langle h, p : s \rangle \mapsto_p \langle h', p : s' \rangle} \quad (\text{Preempt})$$

Our implementation of  $\mapsto_p$  comprise a pair of mutually-recursive definitions,

$$\text{reduce}_p :: \text{Proc} \rightarrow \text{Rel} (\text{Heap}, [\text{Proc}])$$

$$\text{reduce}_s :: \text{Rel} (\text{Heap}, [\text{Proc}])$$

where  $\text{reduce}_p$  performs a single-step reduction of a particular  $\text{Proc}$  in the context of the given heap and soup, and  $\text{reduce}_s$  corresponds to the general case. We begin with the definition of  $\text{reduce}_s$ :

$$\text{reduce}_s (h, [] ) = \{ \}$$

$$\text{reduce}_s (h, p : s) = (\text{second } (p : ) \text{ 'Set.map' } \text{reduce}_s (h, s))$$

$$\cup \text{reduce}_p p (h, s)$$

That is, when the soup is empty, no reduction is possible, so we return an empty set.

When the soup is not empty, we can either apply ( $\text{Preempt}$ ) to reduce the rest of the soup  $s$ , or reduce only the first process  $p$  using  $\text{reduce}_p$ . These two sets of reducts are combined using ( $\cup$ ).

In turn, for the definition of  $\text{reduce}_p$ , it is not possible for values to reduce any further in our semantics, so we return an empty set when  $p$  is a  $\text{Val}_p$ .

$$\text{reduce}_p (\text{Val}_p n) (h, s) = \{ \}$$

Executing  $\text{Fork } p$  adds  $p$  to the process soup, and evaluates to  $\text{Val}_p 0$  (which corresponds to  $\text{return } ()$  in Haskell) as the result of this action:

$$\langle h, \text{Fork } p : s \rangle \mapsto_p \langle h, \text{Val}_p 0 : p : s \rangle \quad (\text{Fork})$$

This is handled by the following case in the definition of  $\text{reduce}_p$ :

$$\text{reduce}_p (\text{Fork } p) (h, s) = \{ h, \text{Val}_p 0 : p : s \}$$

Next, the ( $\text{Atomic}$ ) rule has a premise which evaluates the given expression until it reaches a value (where  $\mapsto_\tau^*$  denotes the reflexive, transitive closure of  $\mapsto_\tau$ ), and a

conclusion which encapsulates this as a single transition on the process level:

$$\frac{\langle h, e \rangle \mapsto_{\top}^* \langle h', \mathbf{Val}_{\top} n \rangle}{\langle h, \mathbf{Atomic} e : s \rangle \mapsto_{\mathbf{P}} \langle h', \mathbf{Val}_{\mathbf{P}} n : s \rangle} \quad (\mathbf{Atomic})$$

In this manner we obtain a *stop-the-world* semantics for atomic transactions, preventing interference from other concurrently executing processes. Note that while the use of  $\mapsto_{\top}^*$  may seem odd in a small-step semantics, it expresses the intended meaning in a clear and concise way [HMPJH05], namely that the transaction executes as if it were a single atomic step.

Our model of the **(Atomic)** rule implements the same stop-the-world semantics using **reduceUntil** defined in the previous section. The values resulting from the execution of  $t$  are then placed back into the soup:

$$\begin{aligned} \mathbf{reduce}_{\mathbf{P}} (\mathbf{Atomic} t) (h, s) &= \mathbf{second} (\lambda n \rightarrow \mathbf{Val}_{\mathbf{P}} n : s) \\ &\quad \text{‘Set.map’ } \mathbf{reduceUntil} \text{ isVal}_{\top} \mathbf{reduce}_{\mathbf{Tran}} (h, t) \end{aligned}$$

Finally, it is straightforward to handle  $\oplus_{\mathbf{P}}$  on the process level using three rules, in an analogous manner to  $\oplus_{\top}$  on the transaction level:

$$\langle h, \mathbf{Val}_{\mathbf{P}} m \oplus_{\mathbf{P}} \mathbf{Val}_{\mathbf{P}} n : s \rangle \mapsto_{\mathbf{P}} \langle h, \mathbf{Val}_{\mathbf{P}} (m + n) : s \rangle \quad (\mathbf{AddZ}_{\mathbf{P}})$$

$$\frac{\langle h, b : s \rangle \mapsto_{\mathbf{P}} \langle h', b' : s' \rangle}{\langle h, \mathbf{Val}_{\mathbf{P}} m \oplus_{\mathbf{P}} b : s \rangle \mapsto_{\mathbf{P}} \langle h', \mathbf{Val}_{\mathbf{P}} m \oplus_{\mathbf{P}} b' : s' \rangle} \quad (\mathbf{AddR}_{\mathbf{P}})$$

$$\frac{\langle h, a : s \rangle \mapsto_{\mathbf{P}} \langle h', a' : s' \rangle}{\langle h, a \oplus_{\mathbf{P}} b : s \rangle \mapsto_{\mathbf{P}} \langle h', a' \oplus_{\mathbf{P}} b : s' \rangle} \quad (\mathbf{AddL}_{\mathbf{P}})$$

The corresponding implementation mirrors that of  $\oplus_{\top}$ , evaluating expressions in a left-to-right order:

$$\mathbf{reduce}_{\mathbf{P}} (a \oplus_{\mathbf{P}} b) (h, s) = \mathbf{case} (a, b) \mathbf{of}$$

$$\begin{aligned}
& (\text{Val}_p\ m, \text{Val}_p\ n) \rightarrow \{h, \text{Val}_p\ (m + n) : s\} \\
& (\text{Val}_p\ m, b) \rightarrow \text{second}\ (\text{mapHead}\ (\text{Val}_p\ m \oplus_p\ ))\ \text{'Set.map'}\ \text{reduces}_s\ (h, b : s) \\
& (a, \quad b) \rightarrow \text{second}\ (\text{mapHead}\ (\quad \oplus_p\ b))\ \text{'Set.map'}\ \text{reduces}_s\ (h, a : s) \\
& \text{where mapHead}\ f\ (p : s) = f\ p : s
\end{aligned}$$

In a similar way to our earlier definition for  $\text{isVal}_T$ , we define an  $\text{isVals}_S$  predicate to determine when an entire soup has finished reducing.

$$\begin{aligned}
& \text{isVals}_S :: (\text{Heap}, [\text{Proc}]) \rightarrow \text{Maybe}\ (\text{Heap}, [\text{Integer}]) \\
& \text{isVals}_S\ (h, s) = \text{case}\ \text{traverse}\ \text{isVal}_p\ s\ \text{of} \\
& \quad \text{Nothing} \rightarrow \text{Nothing} \\
& \quad \text{Just}\ ns \rightarrow \text{Just}\ (h, ns) \\
& \text{where} \\
& \quad \text{isVal}_p :: \text{Proc} \rightarrow \text{Maybe}\ \text{Integer} \\
& \quad \text{isVal}_p\ (\text{Val}_p\ n) = \text{Just}\ n \\
& \quad \text{isVal}_p\ \_ = \text{Nothing}
\end{aligned}$$

This completes our executable model of the high-level semantics: in particular, the term  $\text{reduceUntil}\ \text{isVals}_S\ \text{reduces}_S$  then corresponds to an implementation of  $\mapsto_p$ , which produces a set of  $(\text{Heap}, [\text{Integer}])$  pairs from an initial  $(\text{Heap}, [\text{Proc}])$ .

In summary, the above semantics for transactions and processes mirror those for STM Haskell, but for a simplified language. Moreover, while the original semantics uses evaluation contexts to identify the point at which transition rules such as  $(\text{AddZ}_p)$  can be applied, our language is sufficiently simple to allow the use of explicit structural rules such as  $(\text{AddL}_p)$  and  $(\text{AddR}_p)$ , which for our purposes have the advantage of being directly implementable.

## 5.2 A Simple Transactional Machine

The (Atomic) rule of the previous section simply states that the evaluation sequence for a transaction may be seen as a single indivisible transition with respect to other concurrent processes. However, to better exploit the available multi-core hardware, an actual implementation of this rule would have to allow multiple transactions to run concurrently, while still maintaining the illusion of atomicity. In this section we consider how this notion of concurrent transactions can be implemented, and present a compiler and virtual machine for our language.

### 5.2.1 Instruction Set

Let us consider compiling expressions into code for execution on a stack machine, in which `Code` comprises a sequence of `Instructions`:

```

type Code      = [Instruction]
data Instruction = PUSH Integer | ADD | READ Var | WRITE Var
                  | BEGIN | COMMIT | FORK Code

```

The `PUSH` instruction leaves its argument on top of the stack, while `ADD` replaces the top two numbers with their sum. The behaviour of the remaining instructions is more complex in order to maintain atomicity, but conceptually, `READ` pushes the value of the named variable onto the stack, while `WRITE` updates the variable with the topmost value. In turn, `BEGIN` and `COMMIT` mark the start and finish of a transaction, and `FORK` executes the given code concurrently.

### 5.2.2 Compiler

We define the `compileT` and `compileP` functions to provide translations from `Tran` and `Proc` to `Code`, both functions taking an additional `Code` argument to be appended

to the instructions produced by the compilation process, as in chapter 3. In both cases, integers and addition are compiled into **PUSH** and **ADD** instructions, while the remaining language constructs map directly to their analogous machine instructions. The intention is that executing a compiled transaction or process always leaves a single result value on top of the stack.

$\text{compile}_T :: \text{Tran} \rightarrow \text{Code} \rightarrow \text{Code}$

$\text{compile}_T e c = \text{case } e \text{ of}$

$\text{Val}_T i \quad \rightarrow \text{PUSH } i : c$

$x \oplus_T y \quad \rightarrow \text{compile}_T x (\text{compile}_T y (\text{ADD} : c))$

$\text{Read } v \quad \rightarrow \text{READ } v : c$

$\text{Write } v e' \rightarrow \text{compile}_T e' (\text{WRITE } v : c)$

$\text{compile}_P :: \text{Proc} \rightarrow \text{Code} \rightarrow \text{Code}$

$\text{compile}_P e c = \text{case } e \text{ of}$

$\text{Val}_P i \quad \rightarrow \text{PUSH } i : c$

$x \oplus_P y \quad \rightarrow \text{compile}_P x (\text{compile}_P y (\text{ADD} : c))$

$\text{Atomic } e' \rightarrow \text{BEGIN} : \text{compile}_T e' (\text{COMMIT} : c)$

$\text{Fork } x \quad \rightarrow \text{FORK } (\text{compile}_P x []) : c$

For example, invoking  $\text{compile}_P (\text{incTwice } \text{counter}) []$  delivers the following code:

```
[FORK [BEGIN, READ counter, PUSH 1, ADD, WRITE counter, COMMIT]
, FORK [BEGIN, READ counter, PUSH 1, ADD, WRITE counter, COMMIT]
, ADD]
```

### 5.2.3 Implementing Transactions

The simplest method of implementing transactions would be to suspend execution of all other concurrent processes on encountering a **BEGIN**, and carry on with the current process until we reach the following **COMMIT**. In essence, this is the approach



used in the high-level semantics presented in the previous section. Unfortunately, this does not allow transactions to execute concurrently, one of the key aspects of transactional memory. This section introduces the log-based approach to implementing transactions, and discusses a number of design issues.

### Transaction Logs

In order to allow transactions to execute concurrently, we utilise the notion of a *transaction log*. Informally such a log behaves as a cache for read and write operations on transactional variables. Only the first read from any given variable accesses the heap, and only the last value written can potentially modify the heap; all intermediate reads and writes operate solely on the log. Upon reaching the end of the transaction, and provided that that no other concurrent process has ‘interfered’ with the current transaction, the modified variables in the log can then be committed to the heap. Otherwise, the log is discarded and the transaction is restarted afresh.

Note that restarting a transaction relies on the fact that it executes in complete isolation, in the sense that all its side-effects are encapsulated within the log, and hence can be revoked by simply discarding the log. For example, it would not be appropriate to ‘launch missiles’ [HMPJH05] during a transaction.

### Interference

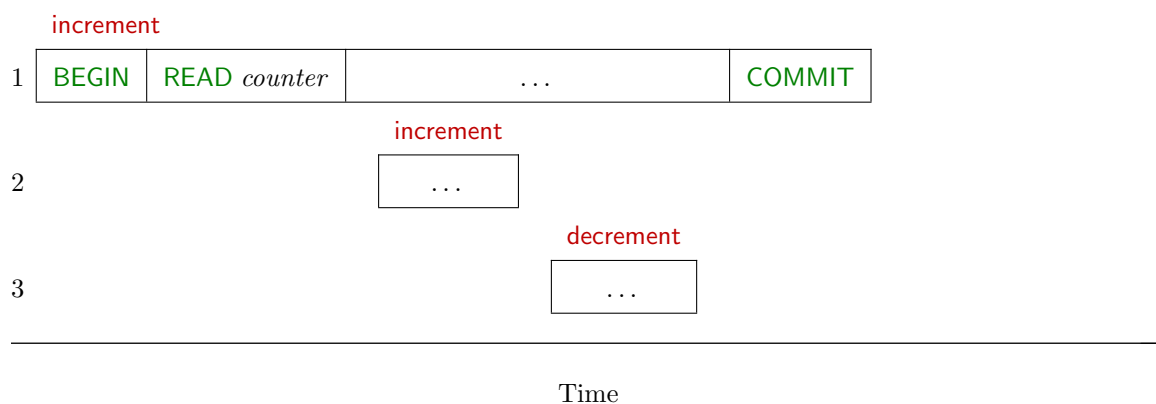
But what constitutes *interference*? When a transaction succeeds and commits its log to the heap, all of its side-effects are then made visible in a single atomic step, as if it had been executed in its entirety at that point with a stop-the-world semantics. Thus when a variable is read for the first time and its value logged, the transaction is essentially making the following bet: at the end of the transaction, the value of the variable in the heap will still be the same as that in the log.

In this manner, interference arises when any such bet fails, as the result of other

concurrent processes changing the heap in a way that invalidates the assumptions about the values of variables made in the log. In this case, the transaction fails and is restarted. Conversely, the transaction succeeds if the logged values of all the variables read are ‘equal’ to their values in the heap at the end of the transaction.

## Equality

But what constitutes *equality*? To see why this is an important question, and what the design choices are, let us return to our earlier example of a transaction that increments a given counter. Consider the following timeline:



Suppose the *counter* starts at zero, which is read by the first transaction and logged. Prior to its final **COMMIT**, a second concurrent transaction successfully increments the *counter*, which is subsequently decremented by a third transaction. When the first finally attempts to commit, the count is back to zero as originally logged, even though it has changed in the interim. Is this acceptable? That is, are the two zeros ‘equal’? We can consider a hierarchy of notions of equality, in increasing order of permissiveness:

- The most conservative choice is to increment a global counter every time the heap is updated. Under this scheme, a transaction fails if the heap is modified at any point during its execution, reflected by a change in the counter, even if this does not actually interfere with the transaction itself.

## 5.2. A SIMPLE TRANSACTIONAL MACHINE

- A more refined approach is provided by the notion of *version equality*, where a separate counter is associated with each variable, and is incremented each time the variable is updated. In this case, our example transaction would still fail to commit, since the two zeros would have different version numbers, and hence considered different.
- For a pure language such as Haskell, in which values are represented as pointers to immutable structures, *pointer equality* can be used as an efficient but weaker form of version equality. In this case, whether the two zeros are considered equal or not depends on whether the implementation created a new instance of zero, or reused the old zero by sharing.
- We can also consider *value equality*, in which two values are considered the same if they have the same representation. In this case, the two zeros are equal and the transaction succeeds.
- The most permissive choice would be a *user-defined equality*, beyond that built-in to the programming language itself, in order to handle abstract data structures in which a single value may have several representations, e.g. sets encoded as lists. Haskell provides this capability via the `Eq` typeclass.

Which of the above is the appropriate notion of equality when committing transactions? Recall that under a stop-the-world semantics, a transaction can be considered to be executed in its entirety at the point when it successfully commits, and any prior reads are effectively bets on the state of the heap at the commit point. Any intermediate writes that may have been committed by other transactions do not matter, as long as the final heap is consistent with the bets made in the log. Hence in our instance, there is no need at commit time to distinguish between the two zeroes in our example, as they are equal in the high-level expression language.

From a semantics point of view, therefore, value or user-defined equality are the best choices. Practical implementations may wish to adopt a more efficient notion of equality (e.g. STM Haskell utilises pointer equality), but for the purposes of this thesis, we will use value equality.

### 5.2.4 Virtual Machine

The state of the virtual machine is given by a pair  $\langle h, s \rangle$ , comprising a heap  $h$  mapping variables to integers, and a soup  $s$  of concurrent *threads*. A **Thread** is represented as a tuple of the form  $(c, \sigma, f, r, w)$ , where  $c$  is the code to be executed,  $\sigma$  is the thread-local stack,  $f$  gives the code to be rerun if a transaction fails to commit, and finally,  $r$  and  $w$  are two logs (partial maps from variables to integers) acting as read and write caches between a transaction and the heap.

```
type Thread = (Code, Stack, Code, Log, Log)
```

```
type Stack  = [Integer]
```

```
type Log    = Map Var Integer
```

We specify the behaviour of the machine using a transition relation  $\mapsto_M$  between machine states, defined via a collection of transition rules that proceed by case analysis on the first thread in the soup. As with the previous semantics, we begin by defining a (PREEMPT) rule to allow the rest of the soup to make progress, giving rise to non-determinism in the machine:

$$\frac{\langle h, s \rangle \mapsto_M \langle h', s' \rangle}{\langle h, t : s \rangle \mapsto_M \langle h', t : s' \rangle} \quad (\text{PREEMPT})$$

This rule corresponds to an idealised scheduler that permits context switching at every instruction, as our focus is on the implementation of transactions rather than scheduling policies. We return to this issue when we consider the correctness of our

compiler later on in this chapter.

We implement  $\mapsto_M$  using a pair of mutually recursive functions `stepM` and `stepT` in a similar fashion to that of  $\mapsto_P$  earlier. The former implements reduction between arbitrary soups of threads:

```

stepM :: Rel (Heap, [Thread])
stepM (h, [] ) = {}
stepM (h, t : s) = (second (t :) 'Set.map' stepM (h, s)) ∪ stepT t (h, s)

```

The first case handles empty soups, returning an empty set of resulting states. The second case takes the first thread  $t$  in the soup, and implements the (PREEMPT) rule by reducing the rest of the soup  $s$  before placing  $t$  back at the head. These are then combined with the states resulting from a single step of  $t$ , implemented by `stepT`:

```

stepT :: Thread → Rel (Heap, [Thread])
stepT ([], σ, f, r, w) (h, s) = {}
stepT (i : c, σ, f, r, w) (h, s) = stepi i where
  ~ (n : σ1@~(m : σ2)) = σ
stepi :: Instruction → Set (Heap, [Thread])
... defined below

```

A thread with an empty list of instructions cannot make any transitions, so we return an empty set. When there is at least one instruction remaining, we use the `stepi` helper function to handle each particular instruction. The above code also brings into scope the names  $c$ ,  $\sigma$ ,  $f$ ,  $r$  and  $w$  as detailed previously, as well as  $\sigma_1$  and  $\sigma_2$  for the current stack with one or two values popped.

Let us detail the semantics of each instruction in turn. Firstly, executing `FORK` adds a new thread  $t$  to the soup, comprising the given code  $c'$  with an initially empty

stack, restart code and read and write logs:

$$\langle h, (\text{FORK } c' : c, \sigma, f, r, w) : s \rangle \mapsto_M \langle h, (c, 0 : \sigma, f, r, w) : t : s \rangle \quad (\text{FORK})$$

$$\text{where } t = (c', [], [], \emptyset, \emptyset)$$

The above transition may be implemented directly, as follows:

$$\begin{aligned} \text{step}_I (\text{FORK } c') &= \{h, (c, \quad 0 : \sigma, f, \quad r, w) : t : s\} \\ \text{where } t &= (c', [], [], \{\}, \{\}) \end{aligned}$$

The **PUSH** instruction places its argument  $n$  on top of the stack, while **ADD** takes the top two integer from the stack and replaces them with their sum:

$$\begin{aligned} \langle h, (\text{PUSH } n, c, \sigma, f, r, w) : s \rangle &\mapsto_M \langle h, (c, n : \sigma, f, r, w) : s \rangle && (\text{PUSH}) \\ \langle h, (\text{ADD}, c, n : m : \sigma, f, r, w) : s \rangle &\mapsto_M \langle h, (c, m + n : \sigma, f, r, w) : s \rangle && (\text{ADD}) \end{aligned}$$

The corresponding cases in the definition of  $\text{step}_I$  are almost identical:

$$\begin{aligned} \text{step}_I (\text{PUSH } n) &= \{h, (c, \quad n : \sigma, f, \quad r, w) : s\} \\ \text{step}_I \text{ ADD} &= \{h, (c, m + n : \sigma_2, f, \quad r, w) : s\} \end{aligned}$$

Executing **BEGIN** starts a transaction, which involves clearing the read and write logs, while making a note of the code to be executed if the transaction fails:

$$\langle h, (\text{BEGIN} : c, \sigma, f, r, w) : s \rangle \mapsto_M \langle h, (c, \sigma, \text{BEGIN} : c, \emptyset, \emptyset) : s \rangle \quad (\text{BEGIN})$$

Accordingly,  $\text{step}_I$  sets up the retry code and initialises both read and write logs:

$$\text{step}_I \text{ BEGIN} = \{h, (c, \quad \sigma, \text{BEGIN} : c, \{\}, \{\}) : s\}$$

Next, **READ** places the appropriate value for the variable  $v$  on top of the stack. The instruction first consults the write log. If the variable has not been written to, the

read log is then consulted. Otherwise, if the variable has not been read from either, its value is looked up from the heap and the read log updated accordingly:

$$\langle h, (\text{READ } v : c, \sigma, f, r, w) : s \rangle \mapsto_M \langle h, (c, n : \sigma, f, r', w) : s \rangle \quad (\text{READ})$$

$$\text{where } \langle n, r' \rangle = \begin{cases} \langle w(v), r \rangle & \text{if } v \in \text{dom}(w) \\ \langle r(v), r \rangle & \text{if } v \in \text{dom}(r) \\ \langle h(v), r \uplus \{v \mapsto h ? v\} \rangle & \text{otherwise} \end{cases}$$

The transliteration of the (READ) rule to our implementation is as follows:

$$\begin{aligned} \text{step}_1 (\text{READ } v) &= \{h, (c, \quad n : \sigma, f, \quad r', w) \quad : s\} \\ \text{where } (n, r') &= \text{case } (\text{Map.lookup } v \ w, \text{Map.lookup } v \ r, h ? v) \text{ of} \\ &(\text{Just } n', \quad -, \quad -) \rightarrow (n', r) \\ &(\text{Nothing}, \text{Just } n', \quad -) \rightarrow (n', r) \\ &(\text{Nothing}, \text{Nothing}, n') \rightarrow (n', r \uplus \{v \mapsto n'\}) \end{aligned}$$

In turn, **WRITE** simply updates the write log for the variable  $v$  with the value on the top of the stack, without changing the heap or the stack:

$$\langle h, (\text{WRITE } v : c, n : \sigma, f, r, w) : s \rangle \mapsto_M \langle h, (c, n : \sigma, f, r, w') : s \rangle \quad (\text{WRITE})$$

$$\text{where } w' = w \uplus \{v \mapsto n\}$$

The (WRITE) rule has a similarly straightforward implementation:

$$\begin{aligned} \text{step}_1 (\text{WRITE } v) &= \{h, (c, \quad n : \sigma_1, f, \quad r, w') \quad : s\} \text{ where} \\ &w' = w \uplus \{v \mapsto n\} \end{aligned}$$

Finally, **COMMIT** checks the read log  $r$  for consistency with the current heap  $h$ , namely that the logged value for each variable read is equal to its value in the heap. According to the above definitions of (READ) and (WRITE), variables written to

before being read from during the same transaction will not result in a read log entry, since the corresponding value in the global heap cannot influence the results of the transaction. Therefore, we do not need to perform consistency checks for variables that occur only in the write log.

Using our representation of logs and heaps, the consistency condition can be concisely stated as  $r \subseteq h$ . That is, if they are consistent, then the transaction has succeeded, so we may commit its write log  $w$  to the heap. This update is expressed in terms of the overriding operator on maps as  $h \uplus w$ . Otherwise the transaction has failed, in which case the heap is not changed, the result on the top of the stack is discarded, and the transaction is restarted at  $f$ :

$$\langle h, (\text{COMMIT} : c, n : \sigma, f, r, w) : s \rangle \mapsto_{\mathbb{M}} \langle h', (c', \sigma', f, r, w) : s \rangle$$

$$\text{where } \langle h', c', \sigma' \rangle = \begin{cases} \langle h \uplus w, c, n : \sigma \rangle & \text{if } r \subseteq h \\ \langle h, f, \sigma \rangle & \text{otherwise} \end{cases} \quad (\text{COMMIT})$$

There is no need to explicitly clear the logs in the above rule, since this is already taken care of by the fact that the first instruction of  $f$  is always a **BEGIN**.

$$\begin{aligned} \text{step}_1 \text{ COMMIT} &= \{h', (c', \sigma', f, r, w) : s\} \\ \text{where } (h', c', \sigma') &= \text{case } (r \cap h) \subseteq h \text{ of} \\ \text{True} &\rightarrow (h \uplus w, c, n : \sigma_1) \\ \text{False} &\rightarrow (h, f, \sigma) \end{aligned}$$

Finally we define a **halted<sub>M</sub>** function on virtual machines to discriminate between completed and running threads.

$$\begin{aligned} \text{halted}_{\mathbb{M}} &:: (\text{Heap}, [\text{Thread}]) \rightarrow \text{Maybe} (\text{Heap}, [\text{Integer}]) \\ \text{halted}_{\mathbb{M}} (h, s) &= \text{case traverse halted}_{\top} s \text{ of} \\ \text{Just } ns &\rightarrow \text{Just } (h, ns) \end{aligned}$$



Nothing  $\rightarrow$  Nothing

where

halted $_{\top}$  :: Thread  $\rightarrow$  Maybe Integer

halted $_{\top}$  ([], n : [], -, -, -) = Just n

halted $_{\top}$  \_ = Nothing

## 5.3 Correctness of the Implementation

As we have seen, the high-level semantics of atomicity is both clear and concise, comprising a single inference rule (**Atomic**) that wraps up a complete evaluation sequence as a single transition. On the other hand, the low-level implementation of atomicity using transactions is rather more complex and subtle, involving the management of read and write logs, and careful consideration of the conditions that are necessary in order for a transaction to commit. How can we be sure that these two different views of atomicity are consistent? Our approach to establishing the correctness of the low-level implementation is to formally relate it to the high-level semantics via a compiler correctness theorem.

### 5.3.1 Statement of Correctness

In order to formulate our correctness result, we utilise a number of auxiliary definitions. First of all, because our semantics is non-deterministic, we define a relation  $\Downarrow_{\mathcal{P}}$  that encapsulates the idea of completely evaluating a process using our high-level semantics:

$$\langle h, ps \rangle \Downarrow_{\mathcal{P}} \langle h', ps' \rangle \quad \leftrightarrow \quad \langle h, ps \rangle \mapsto_{\mathcal{P}}^* \langle h', ps' \rangle \not\mapsto_{\mathcal{P}}$$

That is, a process soup  $ps :: [\text{Proc}]$  with the initial heap  $h$  can evaluate to any heap  $h'$  and soup  $ps'$  that results from completely reducing  $ps$  using our high-level semantics,

where  $\nrightarrow_p$  expresses that no further transitions are possible. We may implement the  $\Downarrow_p$  relation as the following **eval** function:

```
eval :: REL (Heap, [Proc]) (Heap, [Integer])
eval = reduceUntil isValsS reducesS
```

Similarly, we define a relation  $\Downarrow_M$  that encapsulates complete execution of a thread soup  $ts :: [\text{Thread}]$  with the initial heap  $h$  using our virtual machine, resulting in a heap  $h'$  and a thread soup  $ts'$ :

$$\langle h, ts \rangle \Downarrow_M \langle h', ts' \rangle \leftrightarrow \langle h, ts \rangle \mapsto_M^* \langle h', ts' \rangle \nrightarrow_M$$

Likewise, we may implement  $\Downarrow_M$  as the following **exec** function:

```
exec :: REL (Heap, [Thread]) (Heap, [Integer])
exec = reduceUntil haltedM stepM
```

Next, we define a function **load** that converts a process into a corresponding thread for execution, which comprises the compiled code for the process, together with an empty stack, restart code and read and write logs:

```
load :: [Proc] → [Thread]
load = map (λp → (compilep p [], [], [], {}, {}))
```

Using these definitions, the correctness of our compiler can now be expressed by the following property:

$$\forall p \in \text{Proc}, h \in \text{Heap}, s \in [\text{Integer}]. \\ \langle \{\}, p : [] \rangle \Downarrow_p \langle h, s \rangle \leftrightarrow \langle \{\}, \text{load } (p : []) \rangle \Downarrow_M \langle h, s \rangle$$

That is, evaluating a process  $p$  starting with an initial heap using our high-level stop-the-world process semantics is equivalent to compiling and loading the process, and

executing the resulting thread using the interleaved virtual machine semantics. For the purposes of proving the result, we generalise the above over a process soup rather than a single process, as well as an arbitrary initial heap:

**Theorem 5.1** (Compiler Correctness).

$$\forall ps \in [\text{Proc}], h, h' \in \text{Heap}, s \in [\text{Integer}].$$

$$\langle h, ps \rangle \Downarrow_P \langle h', s \rangle \leftrightarrow \langle h, \text{load } ps \rangle \Downarrow_M \langle h', s \rangle$$

The above  $\leftrightarrow$  equivalence can also be considered separately, where the  $\rightarrow$  direction corresponds to soundness, and states that the compiled code will always produce a result that is permitted by the semantics. Dually, the  $\leftarrow$  direction corresponds to completeness, and states that the compiled code can indeed produce every result permitted by the semantics.

In practice, some language implementations are not complete with respect to the semantics for the language by design, because implementing every behaviour that is permitted by the semantics may not be practical or efficient. For example, a real implementation may utilise a scheduler that only permits a context switch between threads at fixed intervals, rather than after every transition as in our semantics, because doing so would be prohibitively expensive.

### 5.3.2 Validation of Correctness

Proving the correctness of programs in the presence of concurrency is notoriously difficult. Ultimately we would like to have a formal proof, but the randomised testing approach—using QuickCheck and HPC, as described in Chapter 4—can provide a high level of assurance with relatively minimal effort. In the case of theorem 5.1, we can transcribe it as the following property:

```
prop_Correctness :: Heap → [Proc] → Bool
prop_Correctness h ps = eval (h, ps) ≡ exec (h, load ps)
```

In other words, from any initial heap  $h$  and process soup  $ps$ , the stop-the-world semantics produces the same set of possible outcomes as that from executing the compiled thread soup using a log-based implementation of transactions.

Given suitable **Arbitrary** instances for *Prop* and **Heap**, we can use QuickCheck to generate a large number of random test cases, and check that the theorem holds in each and every one:

```
*Main> quickCheck prop_Correctness
OK, passed 100 tests.
```

Having performed many thousands of tests in this manner, we can be highly confident in the validity of our compiler correctness theorem. However, as with any testing process, it is important to ensure that all the relevant parts of the program have been exercised in the process. Repeating the coverage checking procedure described in §4.2.3, we obtain the following report of unevaluated expressions<sup>1</sup>:

```
module "Main" {
  inside "reduceTran" {
    tick "Set.empty" on line 303;
  }
}
```

The `{` corresponds to the **Val<sub>τ</sub>** case of the `reduceTran` function, which remains unused simply because we never ask it for the reduct of **Val<sub>τ</sub>**  $m$  expressions.

---

<sup>1</sup>As before, unused expressions corresponding to compiler-generated instances have been omitted.

## 5.4 Conclusion

In this chapter we have shown how to implement software transactional memory correctly, for a simplified language inspired by STM Haskell. Using QuickCheck and HPC, we tested a low-level, log-based implementation of transactions with respect to a high-level, stop-the-world semantics, by means of a compiler and its correctness theorem. This appears to be the first time that the correctness of a compiler for a language with transactions has been mechanically tested.

The lightweight approach provided by QuickCheck and HPC was indispensable in allowing us to experiment with the design of the language and its implementation, and to quickly check any changes. Our basic definitions were refined many times during the development of this work, both as a result of correcting errors, and streamlining the presentation. Ensuring that our changes were sound was simply a matter of re-running QuickCheck and HPC.

On the other hand, it is important to recognise the limitations of this approach. First of all, randomised testing does not constitute a formal proof, and the reliability of QuickCheck depends heavily on the quality of the test-case generators. Secondly, achieving 100% code coverage with HPC does not guarantee that all possible interactions between parts of the program have been tested. Nonetheless, we have found the use of these tools to be invaluable in our work.



# Chapter 6

## Machine-Assisted Proofs in Agda

To give a formal proof of the correctness property posited in the previous chapter, we may make use of a mechanised proof assistant. Agda [Nor07, The10] is a dependently-typed functional programming language based on Martin-Löf intuitionistic type theory [ML80, NPS90]. Via the Curry-Howard correspondence—that is, viewing types as propositions and programs as proofs—it is also used as a proof assistant for constructive mathematics. In this chapter, we shall provide a gentle introduction to the language, and demonstrate how we can formalise statements of compiler correctness by means of machine-checked proofs, culminating in a verified formalisation of the proofs of chapter 3.

### 6.1 Introduction to Agda

The current incarnation of Agda has a syntax similar to that of Haskell, and should look familiar to readers versed in the latter. As in previous chapters, we will adopt a colouring convention for ease of readability:

Syntactic Class	Examples
Keywords	<b>data</b> , <b>where</b> , <b>with</b> ...

Types	$\mathbb{N}$ , List, Set...
Constructors	zero, suc, tt, []...
Functions	id, _+_, Star.gmap...

Semantically, Agda is distinguished by its foundation on *dependent types*, and is closely related to systems such as Epigram [M<sup>+</sup>08, McB05] and Coq [The08]. Dependent types systems are so-called because they allow for types to depend on values, in addition to the usual parametrisation by other types as seen in languages such as Haskell. This provides us with a much richer vocabulary of discourse for not only stating the properties of our programs, but also to be able to prove such properties within the same system. We will begin to explore how this is facilitated by dependent types from section 6.1.2 onwards.

### 6.1.1 Data and Functions

We start our introduction to Agda with some simple data and function definitions. The language itself does not specify any primitive data types, and it serves as a good introduction to see how some of these may be defined in its standard library [Dan10b]. For example, we may define the Peano numbers as follows:

```
data  $\mathbb{N}$  : Set where
  zero :  $\mathbb{N}$ 
  suc  :  $\mathbb{N} \rightarrow \mathbb{N}$ 
```

This is syntactically similar to Haskell’s generalised abstract data type (GADT) declarations [PJWW04] with a few minor differences. Firstly, arbitrary Unicode characters may be used in identifiers, and we do not use upper and lower case letters to distinguish between values and constructors<sup>1</sup>. Secondly, we write `:` to mean *has-type-of*,

<sup>1</sup>The implication here is that the processes of syntax highlighting and type-checking are inextricably linked, and that syntax colours provides more information for the reader.



and write `Set` for the *type of types*<sup>2</sup>.

Thus, the above defines `N` as a new data type inhabiting `Set`, with a nullary constructor `zero` as the base case and an unary constructor `suc` as the inductive case. These correspond to two of the Peano axioms that define the natural numbers: `zero` is a natural number, and every number  $n$  has a successor `suc n`.

We may define addition on the natural numbers as follows, by induction on its first argument:

```

_+_ : N → N → N
zero + n = n
suc m + n = suc (m + n)

```

Agda allows the use of arbitrary infix operators, where underscores ‘`_`’ denote the positions of arguments. Another difference is that all functions in Agda must be total. For example, omitting the `zero` case of `_+_` would lead to an error during the typechecking process, rather than a warning as in Haskell. Additionally, only structural recursion is permitted. In the definition of `_+_` above, we recurse on the definition of `N`: the first argument is strictly smaller on the right hand side, i.e.  $m$  rather than `suc m`. While more general forms of recursion are possible, Agda requires us to explicitly prove that the resulting definitions are total.

### 6.1.2 Programs as Proofs and Types as Predicates

The Curry-Howard correspondence refers to the initial observation by Curry that types—in the sense familiar to functional programmers—correspond to axiom-schemes for intuitionistic logic, while Howard later noted that proofs in formal systems such as natural deduction can be directly interpreted as terms in a model of computation such as the typed lambda calculus.

---

<sup>2</sup>Agda in fact has countably infinite levels of `Sets`, with `Set : Set1 : Set2 : ...`. This stratification prevents the formation of paradoxes that would lead to inconsistencies in the system.

The intuitionistic approach to logic only relies on constructive methods, disallowing notions from classical logic such as the law of the excluded middle ( $P \vee \neg P$ ) or double-negation elimination ( $\neg\neg P \rightarrow P$ ). For example, intuitionists reject  $P \vee \neg P$  because there exists a statement  $P$  in any sufficiently powerful logic that can neither be proved nor disproved within the system, by Gödel’s incompleteness theorems. In other words, intuitionism equates the truth of a statement  $P$  with the possibility of constructing a proof object that satisfies  $P$ , therefore a proof of  $\neg\neg P$ , refuting the non-existence of  $P$ , does not imply  $P$  itself.

What does this mean for the proletarian programmer? Under the Curry-Howard correspondence, the type  $A \rightarrow B$  is interpreted as the logical statement ‘ $A$  implies  $B$ ’, and vice-versa. Accordingly, a program  $p : A \rightarrow B$  corresponds to a proof of ‘ $A$  implies  $B$ ’, in that executing  $p$  constructs a witness of  $B$  as output, given a witness of  $A$  as input. Thus in a suitable type system, programming is the same as constructing proofs in a very concrete sense.

In a traditional strongly typed programming language such as Haskell, the type system exists to segregate values of different types. On the other hand, distinct values of the same type all look the same to the type-checker, which means we are unable to form types corresponding to propositions about particular values. Haskell’s GADTs break down this barrier in a limited sense, by allowing the constructors of a parametric type to target particular instantiations of the return type. While this allows us to exploit a Haskell type checker that supports GADTs as a proof-checker in some very simple cases, it comes at the cost of requiring ‘counterfeit type-level copies of data’ [McB02].

### 6.1.3 Dependent Types

Dependent type systems follow a more principled approach, being founded on Martin-Löf intuitionistic type theory [NPS90]. These have been studied over several decades,

and the current incarnation of Agda [Nor07, The10] is one example of such a system.

Coming from a Hindley-Milner background, the key distinction of dependent types is that values can influence the types of other, subsequent values. Let us introduce the notion by considering an example of a dependent type:

```
data Fin : ℕ → Set where
  fz : {n : ℕ}          → Fin (suc n)
  fs : {n : ℕ} → Fin n → Fin (suc n)
```

This defines a data type `Fin`—similar to `ℕ` above—that is additionally *indexed* by a natural number. Its two constructors are analogues of the `zero` and `suc` of `ℕ`. The `fz` constructor takes an argument of type `ℕ` named `n`, that is referred to in its resultant type of `Fin (suc n)`. The braces `{` and `}` indicate that `n` is an implicit parameter, which we may omit at occurrences of `fz`, provided that the value of `n` can be automatically inferred. We can see how this might be possible in the case of the `fs` constructor: its explicit argument has type `Fin n`, from which we may deduce `n`.

The above `Fin` represents a family of types, where each type `Fin n` has exactly `n` distinct values. This is apparent in the resultant types of the constructors, for example: neither `fz` and `fs` can inhabit `Fin zero`, as they both target `Fin (suc n)` for some `n`. The only inhabitant of `Fin (suc zero)` is `fz`, since `fs` requires an argument of type `Fin zero`, which would correspond to a proof that `Fin zero` is inhabited.

A routine application of the `Fin` family is in the safe lookup of lists or vectors, where a static guarantee on the bounds of the given position is required. The following definition defines the type of vectors, indexed by their length:

```
data Vec (X : Set) : ℕ → Set where
  []      :                               Vec X zero
  _::__ : {n : ℕ} → X → Vec X n → Vec X (suc n)
```

Unsurprisingly the empty list `[]` corresponds to a `Vec X` of length `zero`, while the `_::__`

constructor prepends an element of  $X$  to an existing `Vec`  $X$  of length  $n$ , to give a `Vec`  $X$  of length `suc`  $n$ .

A `lookup` function can then be defined as follows:

```
lookup : {X : Set} {n : ℕ} → Fin n → Vec X n → X
```

```
lookup fz (x :: xs) = x
```

```
lookup (fs i) (x :: xs) = lookup i xs
```

The first argument gives the position in the vector where we wish to extract an element, with the second argument being the vector itself.

Earlier we mentioned that all definitions in Agda must be total, yet the `lookup` function seemingly does not consider the case of the empty vector `[]`. This is because in a dependently-typed language, pattern matching on one argument can potentially influence the types of other arguments: matching on either `fz` or `fs i` forces  $n$  to `suc`  $n'$  for some  $n'$ , and in both cases the type of the vector is refined to `Vec`  $X$  (`suc`  $n'$ ). As `[]` can only inhabit `Vec`  $X$  `zero`, we needn't explicitly list this case. Had we pattern matched the vector with `[]` first on the other hand, the type of the position then becomes `Fin zero`, which is uninhabited. In this case, we may use the 'impossible' pattern `()` for the first argument, to indicate that the case is unreachable:

```
lookup () []
```

In such cases, the right hand side of the definition is simply omitted.

This is only an elementary demonstration of the power of dependent types. Being able to form any proposition in intuitionistic logic as a type gives us a powerful vocabulary with which to state and verify the properties of our programs. Conversely, by interpreting a type as its corresponding proposition, we have mapped the activity of constructing mathematical proofs to 'just' programming, albeit in a more rigorous fashion than usual.

### 6.1.4 Equality and its Properties

The previous section introduced dependent types from a programming and data types point-of-view. We shall now take a look at how we can use dependent types to state logical propositions and to construct their proofs.

A simple and commonplace construct is that of an equality type, corresponding to the proposition that two elements of the same type are definitionally equal. The following definition encodes the Martin-Löf equality relation:

```
data _≡_ {X : Set} : X → X → Set where
  refl : {x : X} → x ≡ x
```

Agda does not provide the kind of Hindley-Milner polymorphism as seen in Haskell, although we can simply take an additional parameter of the type we wish to be polymorphic over. In the above definition of equality, the variable  $X$  corresponds to the type of the underlying values, which can often be inferred from the surrounding context. By marking this parameter as implicit, we effectively achieve the same end result in its usage.

The above definition of `_≡_` is indexed by two explicit parameters of type  $X$ . Its sole constructor is `refl`, that inhabits the type  $x \equiv x$  given an argument  $x : X$ . Logically, `refl` corresponds to the axiom of reflexivity  $\forall x. x \equiv x$ . In cases where the type of the argument—such as  $x$  here—can be inferred from the surrounding context, the `∀` keyword allows us to write the type in a way that better resembles the corresponding logical notation, for example:

```
refl : ∀ {x} → x ≡ x
```

In general, Agda syntax allows us to write ‘`_`’ in place of expressions—including types—that may be automatically inferred. The above is in fact syntactic sugar for the following:

```
refl : {x : _} → x ≡ x
```

Agda includes an interactive user interface for the Emacs [Sta10] operating system that supports incremental development by the placement of ‘holes’ where arbitrary expressions are expected. Incomplete programs with holes can be passed to the type checker, which then informs the user of the expected type. Thus, writing proofs in Agda typically involves a two-way dialogue between the user and the type checker.

Given the above definition of reflexivity as an axiom, we may go on to prove that  $\equiv$  is also symmetric and transitive. The former is the proposition that given any  $x$  and  $y$ , a proof of  $x \equiv y$  implies that  $y \equiv x$ . This is implemented as the following `sym` function,

```
sym : {X : Set} (x y : X) → x ≡ y → y ≡ x
sym x y x≡y = ?
```

where the ‘?’ mark indicates an as-yet unspecified expression. Multiple arguments of the same type are simply separated with spaces, while arrows between pairs of named arguments can be omitted, since there cannot be any ambiguity.

Successfully type-checking the above turns the ‘?’ into a *hole* `{ }`, inside which we may further refine the proof. The expected type of the hole is displayed, and we can ask the type-checker to enumerate any local names that are in scope. In this case, the hole is expected to be of type  $y \equiv x$ , with the arguments  $x y : X$  and  $x \equiv y : x \equiv y$  in scope.

At this point, we can ask the system to perform case-analysis on  $x \equiv y$ . Being an equality proof, this argument must be the `refl` constructor. But something magical also happens during the case-split operation:

```
sym x .x refl = { }
```

Since `refl` only inhabits the type  $x \equiv x$ , the type checker concludes that the first two arguments  $x$  and  $y$  must be the same, and rewrites the second as `.x` to reflect

this fact. Whereas pattern matching in Haskell is an isolated affair, in a dependently typed context it can potentially cause interactions with other arguments, revealing more information about the them that can be checked and enforced by the system.

Accordingly,  $y$  is no longer in-scope, and the goal type becomes  $x \equiv x$ , which is satisfied by `refl`:

```
sym x .x refl = refl
```

As we do not explicitly refer to  $x$  on the right hand side, we could have made the  $x$  and  $y$  arguments implicit too, leading to a more succinct definition:

```
sym : {X : Set} {x y : X} → x ≡ y → y ≡ x
sym refl = refl
```

We prove transitivity in a similar fashion,

```
trans : ∀ {X : Set} {x y z : X} → x ≡ y → y ≡ z → x ≡ z
trans refl refl = refl
```

where pattern-matching the first explicit argument with `refl` unifies  $y$  with  $x$ , refining the type of the second argument to  $x \equiv z$ ; in turn, matching this with `refl` then unifies  $z$  with  $x$ . The resulting goal of  $x \equiv x$  is met on the right hand side with simply `refl`.

### 6.1.5 Existentials and Dependent Pairs

In the previous section, we had already surreptitiously modelled the universal quantifier  $\forall$  as dependent functions—that is, functions where values of earlier arguments may influence later types. Dually, we can model the existential quantifier  $\exists$  using *dependent pairs*. This is typically defined in terms of the  $\Sigma$  type<sup>3</sup>:

---

<sup>3</sup>In this thesis, I deviate from the Agda standard library by writing `∧` instead of `∧`, to avoid excessive visual overloading.

**data**  $\Sigma (X : \text{Set}) (P : X \rightarrow \text{Set}) : \text{Set}$  **where**

$\_ \wedge \_ : (x : X) \rightarrow (p : P x) \rightarrow \Sigma X P$

We can interpret the type  $\Sigma X (\lambda x \rightarrow P x)$  as the existentially quantified statement that  $\exists x \in X. P(x)$ . Correspondingly, a proof of this comprises a pair  $x \wedge p$ , where the latter is a proof of the proposition  $P x$ . Unlike classical proofs of existence which may not necessarily be constructive, a proof of the existence of some  $x$  satisfying  $P$  necessarily requires us to supply such an  $x$ . Conversely, we can always extract an  $x$  given such a proof.

As  $X$  can often be inferred from the type of the predicate  $P$ , we may define a shorthand  $\exists$  that accepts  $X$  as an implicit argument:

$\exists : \{X : \text{Set}\} (P : X \rightarrow \text{Set}) \rightarrow \text{Set}$

$\exists \{X\} = \Sigma X$

The above definition of  $\exists$  allows us to write  $\exists \lambda x \rightarrow P x$ , which better resembles the corresponding logical proposition of  $\exists x. P(x)$ .

Of course, the predicate  $P$  need not necessarily depend on the first element of a pair. In such cases, the resulting type is a non-dependent pair, which corresponds to a logical conjunction. This can be recovered as a degenerate case of the above  $\Sigma$  type, in which we simply ignore the value of the first type:

$\_ \times \_ : \text{Set} \rightarrow \text{Set} \rightarrow \text{Set}$

$X \times Y = \Sigma X (\lambda \_ \rightarrow Y)$

Putting the above into practice, we present below a definition of `splitAt` that splits a vector of length  $m + n$  at position  $m$  into left and right vectors of length  $m$  and  $n$  respectively. Unlike the eponymous Haskell function, we can also witness that the concatenation of the resulting vectors coincides with the input:

`splitAt` :  $(m : \mathbb{N}) \{n : \mathbb{N}\} \{X : \text{Set}\} (xs : \text{Vec } X (m + n)) \rightarrow$



$$\Sigma (\text{Vec } X \ m) \lambda \ ys \rightarrow \Sigma (\text{Vec } X \ n) \lambda \ zs \rightarrow xs \equiv ys ++ zs$$

$$\text{splitAt zero } xs = [] \wedge xs \wedge \text{refl}$$

For the base case where the position is `zero`, we simply return the empty list as the left part and the entirety of `xs` as the right. Since `[] ++ xs` reduces to just `xs`, a simple appeal to reflexivity completes this clause.

In the inductive case of `suc m`, the input vector must contain at least one element, followed by `xs`. We wish to split `xs` recursively at `m` and prepend `x` to the left result. In Agda, we can pattern match on intermediate results using the magic ‘`with`’ keyword, which could be thought of as the dependently-typed analogue to Haskell’s `case`:

$$\text{splitAt (suc } m) (x :: xs) \quad \text{with splitAt } m \ xs$$

$$\text{splitAt (suc } m) (x :: \cdot (ys ++ zs)) \mid ys \wedge zs \wedge \text{refl} = x :: ys \wedge zs \wedge \text{refl}$$

When we case-split on the proof of `xs ≡ ys ++ zs`, Agda sees that `≡.refl` is the only possible constructor, and correspondingly rewrites the tail of the input vector to the dotted pattern `\cdot (ys ++ zs)`. This is simply a more sophisticated instance of what we saw when case-splitting `sym` in the previous section.

### 6.1.6 Reflexive Transitive Closure

Before we conclude this brief introduction to Agda, we shall introduce the *reflexive transitive closure* of McBride, Norell and Jansson [McB07], which generalises the notion of sequences in a dependently-typed context. This general construct will prove useful later when working with sequences of small-step reductions.

We begin by defining binary relations parametrised on their underlying types:

$$\text{Rel} : \text{Set} \rightarrow \text{Set}_1$$

$$\text{Rel } X = X \rightarrow X \rightarrow \text{Set}$$

`Rel` has `Set1` as its codomain in order to avoid the `Set : Set` inconsistency. We use the `Rel X` shorthand to denote the type of binary relations on  $X$ . In fact, our earlier definition of propositional equality could equivalently have been written as follows:

```
data  $\equiv$  {X : Set} : Rel X where
  refl : {x : X} → x ≡ x
```

The definition of `Star` is parametrised on a set of indices  $I$  and a binary relation  $R$  on  $I$ . The type `Star {I} R` is itself another binary relation, indexed on the same  $I$ :

```
data Star {I : Set} (R : Rel I) : Rel I where
  ε      : {i : I} → Star R i i
  _<_<_ : {i j k : I} → (x : R i j) → (xs : Star R j k) → Star R i k
```

Here,  $\varepsilon$  gives a trivial witness of reflexivity for any  $i$ . The `_<_<_` constructor provides a heterogeneous definition of transitivity, accepting a proof of  $R\ i\ j$  and an  $R$ -chain relating  $j$  and  $k$  as its two arguments. Thus `Star R` defines the type of the reflexive transitive closure of  $R$ . Being data, we may take apart an element of `Star R i k` and inspect each step of the  $R$ -chain.

Alternatively the constructors  $\varepsilon$  and `_<_<_` could be thought of as generalisations of the `nil` and `cons` constructors of the list data type for proofs of the form  $R\ i\ k$ , with the constraint that two adjacent elements  $x : R\ i\ j$  and  $y : R\ j\ k$  must share a common index  $j$ .

In the degenerate case of a constant relation  $(\lambda\ \_ \_ \rightarrow X) : Rel\ \top$  whose indices provide no extra information, we recover the usual definition of lists of elements of type  $X$ :

```
List : Set → Set
List X = Star (λ _ _ → X) tt tt
```

Here `tt` is the unique constructor of the unit type  $\top$ , with  $\varepsilon$  and `_<_<_` taking the rôles

of *nil* and *cons* respectively. For example, we could write the two-element list of 0 and 1 as follows:

```
two : List ℕ
two = zero < suc zero < ε
```

As it is a generalisation of lists, **Star** admits many of the usual list-like functions. For example, while the `_<<<_` function below provides a proof of transitivity for **Star** *R*, it could equally be considered the generalisation of *append* on lists, as suggested by the structure of its definition:

```
_<<<_ : {I : Set} {R : Rel I} {i j k : I} →
      Star R i j → Star R j k → Star R i k
ε      <<< ys = ys
(x < xs) <<< ys = x < (xs <<< ys)
```

Similarly, we can define an analogue of *map* on lists:

```
gmap : {I I' : Set} {R : Rel I} {S : Rel I'} (f : I → I') →
      ({x y : I} → R x y → S (f x) (f y)) →
      {i j : I} → Star R i j → Star S (f i) (f j)
gmap f g ε = ε
gmap f g (x < xs) = g x < gmap f g xs
```

As well as a function *g* that is mapped over the individual elements of the sequence, **gmap** also takes a function *f* that allows for the indices to change.

To conclude, let us consider a simple use case for **Star**. Take the *predecessor* relation on natural numbers, defined as `_<₁_` below:

```
data <₁_ : Rel ℕ where
  lt₁ : (n : ℕ) → n <₁ suc n
```

Here, `lt1 n` witnesses  $n$  as the predecessor of `suc n`. We can then conveniently define the reflexive transitive closure of `<1` as follows:

```
--<_ : Rel ℕ
--<_ = Star <_1_
```

Of course, this is just the familiar *less than or equal* relation, for which we can easily prove some familiar properties, such as the following:

```
0<=n : (n : ℕ) → zero ≤ n
0<=n zero    = ε
0<=n (suc n) = 0<=n n <<< (lt1 n < ε)
```

Note that a proof of `zero ≤ n` as produced by the above `0<=n` function actually consists of a *sequence* of proofs in the style of “*0 is succeeded by 1, which is succeeded by 2, ... which is succeeded by n*” that can be taken apart and inspected one by one. We will be doing exactly this when considering reduction sequences in our later proofs.

## 6.2 Agda for Compiler Correctness

In this section, we shall revisit the language of numbers and addition from chapter 3, and demonstrate how the previous compiler correctness result can be formalised using Agda.

### 6.2.1 Syntax and Semantics

As in chapter 5, we can encode the syntax of our language as a simple algebraic data type:

```
data Expression : Set where
  #_   : (m : ℕ)      → Expression
  _⊕_  : (a b : Expression) → Expression
```

The two constructors  $\#_-$  and  $-\oplus-$  correspond to the (Exp-N) and (Exp- $\oplus$ ) rules in our original definition of the `Expression` language in chapter 3. Its denotational semantics—being a mapping from the syntax to the underlying domain of  $\mathbb{N}$ —can be simply implemented as a function:

```

[[_]] : Expression → ℕ
[[ # m ]] = m
[[ a ⊕ b ]] = [[ a ]] + [[ b ]]

```

The two cases in the definition of  $[[\_]]$  correspond to (denote-val) and (denote-plus) respectively. A numeric expression is interpreted as just its value, while the  $\oplus$  operator translates to the familiar  $+$  on natural numbers.

We had previously modelled the big-step semantics of our language as a binary relation  $\Downarrow$  between expressions and numbers. In Agda, such a relation can be implemented as a dependent data type, indexed on `Expression` and  $\mathbb{N}$ :

```

data _⇓_ : REL Expression ℕ where
  ⇓-ℕ : ∀ {m} → # m ⇓ m
  ⇓-⊕ : ∀ {a b m n} → a ⇓ m → b ⇓ n → a ⊕ b ⇓ m + n

```

The  $\Downarrow\text{-}\mathbb{N}$  constructor corresponds to the (big-val) rule: an expression  $\# m$  evaluates to just the value  $m$ . The  $\Downarrow\text{-}\oplus$  constructor takes two arguments of type  $a \Downarrow m$  and  $b \Downarrow n$ , corresponding to the two premises of (big-plus).

The small-step semantics for `Expressions` is implemented in the same fashion,

```

data _⇨_ : Rel Expression where
  ⇨-ℕ : ∀ {m n} → # m ⊕ # n ⇨ # (m + n)
  ⇨-L : ∀ {a a' b} → a ⇨ a' → a ⊕ b ⇨ a' ⊕ b
  ⇨-R : ∀ {m b b'} → b ⇨ b' → # m ⊕ b ⇨ # m ⊕ b'

```

with (small-val), (small-left) and (small-right) represented as the  $\Rightarrow\text{-}\mathbb{N}$ ,  $\Rightarrow\text{-}L$  and  $\Rightarrow\text{-}R$

constructors. Using `Star` defined earlier in this chapter, we obtain the reflexive transitive closure of  $\_ \mapsto \_$ , along with proofs of the usual properties, for free:

`\_ \mapsto* \_ : Rel Expression`

`\_ \mapsto* \_ = Star \_ \mapsto \_`

## 6.2.2 Semantic Equivalence

Let us move swiftly on to the semantic equivalence theorems of section 3.2.2. Our previous technique of rule induction essentially boils down to induction on the structure of inductively-defined data types.

The proof for the forward direction of theorem 3.1—which captures the equivalence of the denotational and big-step semantics—is implemented as `denote→big`, proceeding by case analysis on its argument  $e : \text{Expression}$ :

`denote→big : ∀ {e m} → [ e ] ≡ m → e ↓ m`

`denote→big {# n} ≡.refl = ↓-ℕ`

`denote→big {a ⊕ b} ≡.refl = ↓-⊕ (denote→big ≡.refl) (denote→big ≡.refl)`

For the  $e \equiv \# n$  case, matching the proof of  $[ \# n ] \equiv m$  with `≡.refl` convinces the typechecker that  $m$  and  $n$  are equal, so  $e \downarrow n$  is witnessed by `↓-ℕ`. Agda considers terms equal up to  $\beta$ -reduction, so the  $[ a \oplus b ] \equiv m$  argument is equivalently a proof of  $[ a ] + [ b ] \equiv m$ , by the definition of  $[ \_ ]$ . The goal type of  $a \oplus b \downarrow [ a ] + [ b ]$  is therefore met with `↓-⊕`, whose premises of  $a \downarrow [ a ]$  and  $b \downarrow [ b ]$  are obtained by recursively invoking `denote→big` with two trivial witnesses to  $[ a ] \equiv [ a ]$  and  $[ b ] \equiv [ b ]$ .

The backwards direction of theorem 3.1 is given by the following `big→denote`, which uses structural induction on the definition of  $\_ \Downarrow \_$ :

`big→denote : ∀ {e m} → e ↓ m → [ e ] ≡ m`

`big→denote ↓-ℕ = ≡.refl`

**big→denote** ( $\Downarrow\oplus a\Downarrow m b\Downarrow n$ ) **with** **big→denote**  $a\Downarrow m$  | **big→denote**  $b\Downarrow n$   
**big→denote** ( $\Downarrow\oplus a\Downarrow m b\Downarrow n$ ) |  $\equiv.\text{refl}$  |  $\equiv.\text{refl}$  =  $\equiv.\text{refl}$

For the base case of  $\Downarrow\mathbb{N}$ , it must follow that  $e \equiv \# m$ , so the goal type after  $\beta$ -reduction of  $\llbracket \_ \rrbracket$  becomes  $m \equiv m$ , which is satisfied trivially. The  $\Downarrow\oplus$  inductive case brings with it two proofs of  $a \Downarrow m$  and  $b \Downarrow n$ , which can be used to obtain proofs of  $\llbracket a \rrbracket \equiv m$  and  $\llbracket b \rrbracket \equiv n$  via the induction hypothesis. The goal of  $\llbracket a \rrbracket + \llbracket b \rrbracket \equiv m + n$  after  $\beta$ -reduction is then trivially satisfied, thus completing the proof of theorem 3.1.

The **with** keyword provides an analogue of Haskell’s **case ... of** construct that allows us to pattern match on intermediate results. Agda requires us to repeat the left hand side of the function definition when using a **with**-clause, since dependent pattern matching may affect the other arguments, as noted earlier in this chapter.

We go on to prove theorem 3.2—the equivalence between the big-step and small-step semantics—in a similar manner:

**big→small** :  $\forall \{e m\} \rightarrow e \Downarrow m \rightarrow e \mapsto^* \# m$   
**big→small**  $\Downarrow\mathbb{N} = \varepsilon$   
**big→small** ( $\Downarrow\oplus \{a\} \{b\} \{m\} \{n\} a\Downarrow m b\Downarrow n$ ) =  
**Star.gmap** ( $\lambda a' \rightarrow a' \oplus b$ )  $\mapsto\text{-L}$  (**big→small**  $a\Downarrow m$ )  $\triangleleft\triangleleft$   
**Star.gmap** ( $\lambda b' \rightarrow \# m \oplus b'$ )  $\mapsto\text{-R}$  (**big→small**  $b\Downarrow n$ )  $\triangleleft\triangleleft$   
 $\mapsto\mathbb{N} \triangleleft \varepsilon$

The goal in the case of  $\Downarrow\mathbb{N}$  is trivially satisfied by the empty reduction sequence, since  $e \equiv \# m$ . In the  $\Downarrow\oplus$  case, recursively invoking **big→small** with  $a\Downarrow m$  and  $b\Downarrow n$  gives reduction sequences for  $a \mapsto^* \# m$  and  $b \mapsto^* \# n$  respectively. We can map over the former using  $\mapsto\text{-L}$  over the witnesses and  $\lambda a' \rightarrow a' \oplus b$  over the indices to obtain the reduction sequence  $a \oplus b \mapsto^* \# m \oplus b$ , and likewise for the second to obtain  $\# m \oplus b \mapsto^* \# m \oplus \# n$ . By appending these two resulting sequences, followed by

a final application of the  $\mapsto\mathbb{N}$  rule—or equivalently invoking transitivity—we obtain the desired goal of  $a \oplus b \mapsto^* \# m + n$ .

The proof for the reverse direction of 3.2 proceeds by ‘folding’<sup>4</sup> lemma 3.3—implemented as the **sound** helper function below—over the  $e \mapsto^* \# m$  reduction sequence:

$$\begin{aligned}
 \text{small}\rightarrow\text{big} & : \forall \{e\ m\} \rightarrow e \mapsto^* \# m \rightarrow e \Downarrow m \\
 \text{small}\rightarrow\text{big} \ \varepsilon & = \Downarrow\mathbb{N} \\
 \text{small}\rightarrow\text{big} \ (e \mapsto e' \triangleleft e' \mapsto^* m) & = \text{sound} \ e \mapsto e' \ (\text{small}\rightarrow\text{big} \ e' \mapsto^* m) \ \text{where} \\
 \text{sound} & : \forall \{e\ e'\ m\} \rightarrow e \mapsto e' \rightarrow e' \Downarrow m \rightarrow e \Downarrow m \\
 \text{sound} \ \mapsto\mathbb{N} \quad \Downarrow\mathbb{N} & = \Downarrow\oplus \ \Downarrow\mathbb{N} \ \Downarrow\mathbb{N} \\
 \text{sound} \ (\mapsto\text{L} \ a \mapsto a') \ (\Downarrow\oplus \ a' \Downarrow m \ b \Downarrow n) & = \Downarrow\oplus \ (\text{sound} \ a \mapsto a' \ a' \Downarrow m) \ b \Downarrow n \\
 \text{sound} \ (\mapsto\text{R} \ b \mapsto b') \ (\Downarrow\oplus \ m \Downarrow m \ b' \Downarrow n) & = \Downarrow\oplus \ m \Downarrow m \ (\text{sound} \ b \mapsto b' \ b' \Downarrow n)
 \end{aligned}$$

The **sound** helper performs case analysis on the structure of the  $e \mapsto e'$  small-step reduction. In the case of  $\mapsto\mathbb{N}$ , it must be the case that  $e \equiv \# m \oplus \# n$  and  $e' \equiv \# m + n$  respectively, which in turn forces the  $e' \Downarrow m$  argument to be  $\Downarrow\mathbb{N}$ . The goal type of  $\# m \oplus \# n \Downarrow m + n$  is accordingly witnessed by  $\Downarrow\oplus \ \Downarrow\mathbb{N} \ \Downarrow\mathbb{N}$ .

For the two remaining  $\mapsto\text{L}$  and  $\mapsto\text{R}$  rules, it must necessarily be the case that  $e' \equiv a' \oplus b$  or  $e' \equiv \# m \oplus b'$  respectively. Thus the  $e' \Downarrow m$  argument contains a witness of either  $a' \Downarrow m$  or  $b' \Downarrow n$ , which we use to recursively obtain a proof of  $a \Downarrow m$  or  $b \Downarrow n$ .

### 6.2.3 Stack Machine, Compiler, and Correctness

In section 3.3.1 we used a stack machine as our low-level semantics, with the machine modelled as a pair of a list of instructions and a stack of values. This translates to the following Agda data declaration:

<sup>4</sup>Unfortunately we cannot use **Star.gfold** from Agda’s standard library, as  $\_ \Downarrow \_ : \text{REL Expression } \mathbb{N}$  is not a homogeneous relation.



**data** Machine : Set **where**

$\langle \_ , \_ \rangle : (c : \text{List Instruction}) \rightarrow (\sigma : \text{List } \mathbb{N}) \rightarrow \text{Machine}$

The `List` type from in Agda’s standard library implements the familiar *nil* and *cons* lists as found in Haskell. In turn, the virtual machine’s instruction set comprises the standard `PUSH` and `ADD` for stack machines:

**data** Instruction : Set **where**

`PUSH` :  $\mathbb{N} \rightarrow \text{Instruction}$

`ADD` : `Instruction`

The two reduction rules (`vm-push`) (`vm-add`) for the virtual machine are realised as the  $\rightarrow\text{-PUSH}$  and  $\rightarrow\text{-ADD}$  constructors of the  $\_ \rightarrow \_$  relation:

**data**  $\_ \rightarrow \_$  : Rel Machine **where**

$\rightarrow\text{-PUSH} : \forall \{m \ c \ \sigma\} \rightarrow \langle \text{PUSH } m :: c , \ \sigma \rangle \rightarrow \langle c , m \ \sigma \rangle$

$\rightarrow\text{-ADD} : \forall \{m \ n \ c \ \sigma\} \rightarrow \langle \text{ADD} :: c , n :: m :: \sigma \rangle \rightarrow \langle c , m + n :: \sigma \rangle$

$\_ \rightarrow^* \_$  : Rel Machine

$\_ \rightarrow^* \_ = \text{Star } \_ \rightarrow \_$

Again, we define  $\_ \rightarrow^* \_$  as `Star`  $\_ \rightarrow \_$ , and receive the usual properties of a reflexive transitive closure absolutely free.

The compiler is defined identically to that of section 3.3.2,

`compile` : `Expression`  $\rightarrow$  `List Instruction`  $\rightarrow$  `List Instruction`

`compile` (`#` *m*) *c* = `PUSH` *m* :: *c*

`compile` (*a*  $\oplus$  *b*) *c* = `compile` *a* (`compile` *b* (`ADD` :: *c*))

which compiles a number `#` *m* to `PUSH` *m*, and a sum *a*  $\oplus$  *b* to the concatenation of code that computes the value of *a* and *b*, followed by an `ADD` instruction.

The following definition of  $\_ \rightsquigarrow^* \# \_$  provides a convenient synonym for what it means when we say that executing the compiled code for an expression  $e$  computes the result  $m$ :

$\_ \rightsquigarrow^* \# \_ : \text{REL Expression } \mathbb{N}$

$$e \rightsquigarrow^* \# m = \forall \{c \sigma\} \rightarrow \langle \text{compile } e \ c, \sigma \rangle \rightsquigarrow^* \langle c, m :: \sigma \rangle$$

Note that the above proposition is quantified over any code continuation  $c$  and initial stack  $\sigma$ , and a proof of compiler correctness (theorem 3.4) amounts to functions in both directions between  $e \Downarrow m$  and  $e \rightsquigarrow^* \# m$ . In the forward direction—that is, from the high-level/big-step semantics to the low-level/virtual machine semantics—this is implemented by induction on the structure of  $e \Downarrow m$ :

$$\text{big} \rightarrow \text{machine} : \forall \{e \ m\} \rightarrow e \Downarrow m \rightarrow e \rightsquigarrow^* \# m$$

$$\text{big} \rightarrow \text{machine} \Downarrow \mathbb{N} = \rightsquigarrow \text{-PUSH} \triangleleft \varepsilon$$

$$\text{big} \rightarrow \text{machine} (\Downarrow \oplus a \Downarrow m \ b \Downarrow n) =$$

$$\text{big} \rightarrow \text{machine} \ a \Downarrow m \triangleleft \triangleleft \text{big} \rightarrow \text{machine} \ b \Downarrow n \triangleleft \triangleleft \rightsquigarrow \text{-ADD} \triangleleft \varepsilon$$

In the case of  $\Downarrow \mathbb{N}$  we have  $e \equiv \# m$ , and so  $\rightsquigarrow \text{-PUSH} \triangleleft \varepsilon$  witnesses the reduction sequence  $\forall \{c \sigma\} \rightarrow \langle \text{compile } (\# m) \ c, \sigma \rangle \rightsquigarrow^* \langle c, m :: \sigma \rangle$ . In the second case, the recursive terms  $\text{big} \rightarrow \text{machine} \ a \Downarrow m$  and  $\text{big} \rightarrow \text{machine} \ b \Downarrow n$  are of the following types:

$$\text{big} \rightarrow \text{machine} \ a \Downarrow m : \forall \{c_a \ \sigma_a\} \rightarrow \langle \text{compile } a \ c_a, \sigma_a \rangle \rightsquigarrow^* \langle c_a, m :: \sigma_a \rangle$$

$$\text{big} \rightarrow \text{machine} \ b \Downarrow n : \forall \{c_b \ \sigma_b\} \rightarrow \langle \text{compile } b \ c_b, \sigma_b \rangle \rightsquigarrow^* \langle c_b, n :: \sigma_b \rangle$$

The right hand side requires a proof of:

$$\forall \{c \sigma\} \rightarrow \langle \text{compile } (a \oplus b) \ c, \sigma \rangle \rightsquigarrow^* \langle c, m + n :: \sigma \rangle$$

which can be obtained by instantiating  $c_a = \text{compile } b \ (\text{ADD} :: c)$ ,  $c_b = \text{ADD} :: c$ ,  $\sigma_a = \sigma$ ,  $\sigma_b = m :: \sigma$ , and concatenating the resulting reduction sequences,

followed by a final application of the  $\rightarrow$ -ADD rule. As these values can be automatically inferred by Agda, we do not need to make them explicit in the definition of  $\text{big}\rightarrow\text{machine}$ .

The backwards direction of the compiler correctness proof requires us to compute a witness of  $e \Downarrow m$  from one of  $e \rightarrow^* \# m$ . We first need to implement a pair of lemmas, however. The  $\text{exec}$  lemma simply states that that for any expression  $e$ , there exists a number  $m$  for which executing  $\text{compile } e$  computes  $m$ :

$$\begin{aligned}
\text{exec} &: \forall e \rightarrow \exists \lambda m \rightarrow e \rightarrow^* \# m \\
\text{exec } (\# m) &= m \wedge \lambda \{c\} \{\sigma\} \rightarrow \rightarrow\text{-PUSH} \triangleleft \varepsilon \\
\text{exec } (a \oplus b) &\text{ with } \text{exec } a \mid \text{exec } b \\
\text{exec } (a \oplus b) \mid n_a \wedge a \rightarrow^* \# n_a \mid n_b \wedge b \rightarrow^* \# n_b &= n_a + n_b \wedge \\
&\lambda \{c\} \{\sigma\} \rightarrow a \rightarrow^* \# n_a \triangleleft\triangleleft b \rightarrow^* \# n_b \triangleleft\triangleleft \rightarrow\text{-ADD} \triangleleft \varepsilon
\end{aligned}$$

When  $e \equiv \# m$ , this number clearly has to be  $m$ , with  $\rightarrow\text{-PUSH} \triangleleft \varepsilon$  serving as the witness. Note that due to Agda's handling of implicit arguments, we must explicitly generalise over  $c$  and  $\sigma$ . In the  $e \equiv a \oplus b$  case, we simply recurse on the  $a$  and  $b$  subexpressions and concatenate the resulting reduction sequences with an  $\rightarrow\text{-ADD}$ .

The  $\text{unique}$  lemma states that given two reduction sequences from the same initial state, the resulting stacks must coincide:

$$\begin{aligned}
\text{unique} &: \forall \{c \sigma \sigma' \sigma''\} \rightarrow \langle c, \sigma \rangle \rightarrow^* \langle [], \sigma' \rangle \rightarrow \\
&\quad \langle c, \sigma \rangle \rightarrow^* \langle [], \sigma'' \rangle \rightarrow \sigma' \equiv \sigma'' \\
\text{unique } \{[]\} &\quad (()) \triangleleft c' \rightarrow^* \sigma' \mid c \rightarrow^* \sigma'' \\
\text{unique } \{[]\} &\quad \varepsilon \quad (()) \triangleleft c' \rightarrow^* \sigma'' \\
\text{unique } \{[]\} &\quad \varepsilon \quad \varepsilon \quad = \equiv \text{.refl}
\end{aligned}$$

We proceed by recursion on the code  $c$ : the first group of cases above deal with an empty  $c$ . Both sequences must be  $\varepsilon$ , since there no reduction is possible from the machine state  $\langle [], \sigma \rangle$ . In Agda, we identify such cases with the ‘*impossible*’

pattern—written as  $()$ —and accordingly the first two cases do not have a corresponding right hand side. In the third case, the proof of  $\sigma' \equiv \sigma''$  is trivial, since matching on the two  $\varepsilon$  constructors for reflexivity has already unified the  $\sigma$ ,  $\sigma'$  and  $\sigma''$  variables.

The next two cases deal with the **PUSH** and **ADD** instructions. In the first instance, we obtain the proof of  $\sigma' \equiv \sigma''$  by recursion on  $c' \mapsto^* \sigma'$  and  $c' \mapsto^* \sigma''$ , both starting from the machine state  $\langle c', m :: \sigma \rangle$ :

$$\begin{aligned}
 \text{unique } \{\text{PUSH } m :: c'\} & \quad (\mapsto\text{-PUSH} \triangleleft c' \mapsto^* \sigma') \\
 & \quad (\mapsto\text{-PUSH} \triangleleft c' \mapsto^* \sigma'') = \text{unique } c' \mapsto^* \sigma' \ c' \mapsto^* \sigma'' \\
 \text{unique } \{\text{ADD} :: c'\} \{ & \quad [] \} () \quad \triangleleft c' \mapsto^* \sigma' \ c' \mapsto^* \sigma'' \\
 \text{unique } \{\text{ADD} :: c'\} \{ & \quad m :: [] \} () \quad \triangleleft c' \mapsto^* \sigma' \ c' \mapsto^* \sigma'' \\
 \text{unique } \{\text{ADD} :: c'\} \{n :: & \quad m :: \sigma\} (\mapsto\text{-ADD} \triangleleft c' \mapsto^* \sigma') \\
 & \quad (\mapsto\text{-ADD} \triangleleft c' \mapsto^* \sigma'') = \text{unique } c' \mapsto^* \sigma' \ c' \mapsto^* \sigma''
 \end{aligned}$$

The reduction rule for **ADD** requires at least two numbers on top of the stack, which is ruled out by the first two cases in the latter group. The final case proceeds by recursion on the remainder of the reduction sequence, both of which start from the same  $\langle c', m + n :: \sigma \rangle$  machine state.

Returning to the second half of our compiler correctness theorem, it remains for us to show that  $e \mapsto^* \# m$  implies  $e \Downarrow m$ . The following definition of **machine $\rightarrow$ big** provides the proof, which proceeds by case analysis on the expression  $e$ :

$$\begin{aligned}
 \text{machine}\rightarrow\text{big} & : \forall \{e\ m\} \rightarrow e \mapsto^* \# m \rightarrow e \Downarrow m \\
 \text{machine}\rightarrow\text{big} \{ \# & \quad n \} n \mapsto^* \# m \ \text{with } n \mapsto^* \# m \ \{ [] \} \{ [] \} \\
 \text{machine}\rightarrow\text{big} \{ \# & \quad n \} n \mapsto^* \# m \mid \mapsto\text{-PUSH} \triangleleft \varepsilon = \Downarrow\text{-}\mathbb{N} \\
 \text{machine}\rightarrow\text{big} \{ \# & \quad n \} n \mapsto^* \# m \mid \mapsto\text{-PUSH} \triangleleft () \triangleleft n' \mapsto^* \# m
 \end{aligned}$$

The  $n \mapsto^* \# m$  argument is in fact a function that takes two implicit arguments, having

the type:

$$\forall \{c \sigma\} \rightarrow \langle \text{compile} (\# n) c, \sigma \rangle \rightsquigarrow^* \langle c, m :: \sigma \rangle$$

To this we apply an empty code continuation and stack, then pattern match the resulting reduction sequence against  $\rightsquigarrow\text{-PUSH} \triangleleft \varepsilon$ . This unifies  $n$  and the implicit argument  $m$ , allowing  $\Downarrow\text{-}\mathbb{N}$  to witness the goal type of  $\# m \Downarrow m$ . The second case handles the case of longer reduction sequences, which is impossible, since  $\text{compile} (\# n)$  outputs only a single  $\text{PUSH}$  instruction.

For expressions of the form  $e \equiv a \oplus b$ , we first use the  $\text{exec}$  helper to obtain the two reduction sequences  $a \rightsquigarrow^* \# n_a$  and  $b \rightsquigarrow^* \# n_b$ , making use of a  $\text{with}$  clause:

$$\begin{aligned} & \text{machine} \rightarrow \text{big} \{a \oplus b\} a \oplus b \rightsquigarrow^* \# m \text{ with } \text{exec } a \quad | \quad \text{exec } b \\ & \text{machine} \rightarrow \text{big} \{a \oplus b\} a \oplus b \rightsquigarrow^* \# m \quad | \quad n_a \wedge a \rightsquigarrow^* \# n_a \quad | \quad n_b \wedge b \rightsquigarrow^* \# n_b \\ & \quad \text{with unique } \{\sigma = []\} a \oplus b \rightsquigarrow^* \# m (a \rightsquigarrow^* \# n_a \lll b \rightsquigarrow^* \# n_b \lll \rightsquigarrow\text{-ADD} \triangleleft \varepsilon) \\ & \text{machine} \rightarrow \text{big} \{a \oplus b\} a \oplus b \rightsquigarrow^* \# m \quad | \quad n_a \wedge a \rightsquigarrow^* \# n_a \quad | \quad n_b \wedge b \rightsquigarrow^* \# n_b \\ & \quad | \equiv .\text{refl} = \Downarrow\text{-}\oplus (\text{machine} \rightarrow \text{big } a \rightsquigarrow^* \# n_a) (\text{machine} \rightarrow \text{big } b \rightsquigarrow^* \# n_b) \end{aligned}$$

The concatenation  $a \rightsquigarrow^* \# n_a \lll b \rightsquigarrow^* \# n_b \lll \rightsquigarrow\text{-ADD} \triangleleft \varepsilon$  produces a reduction sequence of  $a \oplus b \rightsquigarrow^* \# n_a + n_b$ . Using the previously defined  $\text{unique}$  lemma, we may match each step of the sequence with those of  $a \oplus b \rightsquigarrow^* \# m$  to yield a proof of  $n_a + n_b \equiv m$ . Examining this proof using a second  $\text{with}$  clause, the goal becomes  $a \oplus b \Downarrow n_a + n_b$ . This calls for an instance of  $\Downarrow\text{-}\oplus$ , whose two premises of  $a \Downarrow n_a$  and  $b \Downarrow n_b$  are obtained by recursion.

## 6.3 Conclusion

In this chapter we have given a tutorial of some basic Agda suitable for our needs, and have produced a fully machine-verified version of the results of chapter 3 as an

example. Finally we concluded with a brief to using coinductive data types in Agda.

On reflection, formalising such results in Agda has provided a number of advantages: firstly, it clarifies thinking by forcing us to be very explicit with regards to numerous low-level details that are often elided in pen-and-paper proofs. Testing out new hypotheses also becomes easier, using the type checker to guide proof development; in case of any invalid assumptions, Agda helps to pinpoint precisely where these occur.

A final benefit of using a mechanised proof system such as Agda is that it aids the incremental evolution of this work: often we would prove a simplified form of a theorem to gain some insight. At this point, we may then change, extend, or generalise some core definitions towards our target. Strictly-speaking, all the proofs leading up to the final result would need to be proved again, but in the majority of cases only minor tweaks to the simplified proofs are required. Agda helps out by pointing out where these changes need to be made. This process would not be as straightforward using a script-based theorem prover, or even possible in a pen-and-paper-based approach.

# Chapter 7


## Compiling Non-Determinism Correctly

The standard approach [Wan95] to proving compiler correctness for concurrent languages requires the use of multiple translations into an underlying process calculus. In this chapter, we present a simpler approach that avoids the need for such an underlying language, using a new method that allows us to directly establish a bisimulation between the source and target languages. We illustrate this technique on a small non-deterministic language, using the Agda system to present and formally verify our compiler correctness proofs.

### 7.1 Existing Approach

The standard approach [Wan95] to proving compiler correctness for concurrent languages requires the use of multiple translations into an intermediate process calculus.

This methodology is captured by the following diagram:



That is, for some compiler from a source language to a target language, we define separate denotational semantics  $s[-]$  and  $t[-]$  for both languages in terms of an underlying process calculus with a suitable notion of bisimulation, or observational equivalence. The compiler is said to be correct when  $s[p]$  and  $t[\text{compile } p]$  are bisimilar for all programs  $p$ . The advantage of using a traditional process calculus is that we may reuse existing theories and techniques, and perform our reasoning in a single, homogeneous framework.

However, there are two drawbacks to this method: firstly, the source languages is defined by a map  $s[-]$  into an underlying process calculus, which adds an extra level of indirection when reasoning about the operational behaviour of source programs. Secondly, the target language and process calculus are given separate *operational* semantics — represented by the two circular arrows in the above diagram — with a semantic function  $t[-]$  mapping the former to the latter. Thus we need to further show that the operational semantics of the target language is *adequate* with respect to that of the process calculus via the translation  $t[-]$ .

## 7.2 Related Work

The formal notion of *compiler correctness* dates back to 1967, when McCarthy and Painter [MP67] proved the correctness of a compiler from arithmetic expressions to a register machine. Their stated goal was “to make it possible to use a computer to



*check proofs that compilers are correct*”; we aim to do precisely this for a concurrent language.

In the four decades since, a large body of pen-and-paper correctness proofs have been produced for various experimental languages. (See [Dav03] for a detailed bibliography.) However it is only recently—by making essential use of a formal theorem prover—that it has become possible to fully verify a realistic compiler. In particular, Leroy [Ler06] has produced a certified compiler for a C-like core language in the Coq [The08] framework, relating a series of intermediate languages, eventually targeting PowerPC assembly code.

While compiler correctness for sequential languages has been well explored by many researchers, the issue of concurrency has received relatively little attention, with most of the progress being made by Wand and his collaborators. In the early 1980s, Wand [WS95] initially suggested a methodology for sequential languages: by giving the *denotational* semantics of both source and target languages in a common domain, the correctness proofs relating the *operational* behaviour of the source and target languages may then be carried out within the same domain. Wand then adapted this paradigm to the concurrent setting [Wan95] as summarised in our introduction, which is further elaborated by Gladstein [GW96, Gla94].

Our work in this chapter follows on from Hutton and Wright [HW07], who recently considered the issue of compiler correctness for a simple non-deterministic language, by relating a denotational source semantics to an operational target semantics, based on the extensional notion of comparing final results. As noted in [HW07], the addition of effects and concurrency requires an intensional notion of comparing intermediate actions via a suitable notion of bisimulation. The purpose of this chapter is to explore this idea, while retaining the approach of directly relating the source and target without the need for an intermediate language.

### 7.3 A Non-Deterministic Language

In order to focus on the essence of this problem, we abstract from the details of a real language and keep to a simple expression language consisting of integers and addition [HW04, HW06, HW07] as we had done in previous chapters. We give its operational semantics using a labelled transition system, together with an extra ‘*zap*’ rule to introduce a form of non-determinism. A virtual machine, and a compiler linking the two completes the definition. We present and justify a novel technique for proving compiler correctness in the presence of non-determinism.

We begin by recalling the syntax of our expression language—comprising natural numbers and addition—from the previous chapter:

```
data Expression : Set where
  #_   : (m : ℕ)      → Expression
  _⊕_  : (a b : Expression) → Expression
```

As we had previously noted in §3.1.5, the monoid  $(\mathbb{N}, 0, \oplus)$  may be seen as a degenerate monad, allowing us to avoid orthogonal issues such as binding and substitution. The sequencing aspect of monads is retained by giving our language a left-to-right evaluation order.

Rather than defining a reduction relation between pairs of states as in chapter 6, we generalise to a labelled transition system. For this we define a shorthand **LTS**, parametrised on the type of labels and the underlying state:

```
LTS : Set → Set → Set1
LTS L X = X → L → X → Set
```

We use **Labels** to indicate the nature of the transition:

```
data Action : Set where
  ⊞   : Action
```

$\zeta$  : Action  
 $\square_{-}$  :  $(m : \mathbb{N}) \rightarrow \text{Action}$

**data** Label : Set **where**

$\tau$  : Label  
 $!_{-}$  :  $(\alpha : \text{Action}) \rightarrow \text{Label}$

Each transition either emits (denoted by ‘!’) one of the  $\boxplus$ ,  $\zeta$  or  $\square$  actions (pronounced ‘add’, ‘zap’ and ‘result’ respectively), or has the silent label  $\tau$ . We make a two-level distinction between **Labels** and **Actions** in this language so that potentially silent and non-silent transitions may be identified by their types in the Agda proofs.

Transitions on expressions are defined as the following data type:

**data**  $\mapsto\langle\_\rangle$  : LTS Label Expression **where**

$\mapsto\boxplus$  :  $\forall \{m\ n\} \rightarrow \# m \oplus \# n \mapsto\langle !\boxplus \rangle \# m + n$   
 $\mapsto\zeta$  :  $\forall \{m\ n\} \rightarrow \# m \oplus \# n \mapsto\langle !\zeta \rangle \# 0$   
 $\mapsto R$  :  $\forall \{m\ b\ b'\ \Lambda\} (b \mapsto b' : b \mapsto\langle \Lambda \rangle b') \rightarrow$   
 $\# m \oplus b \mapsto\langle \Lambda \rangle \# m \oplus b'$   
 $\mapsto L$  :  $\forall \{a\ a'\ b\ \Lambda\} (a \mapsto a' : a \mapsto\langle \Lambda \rangle a') \rightarrow$   
 $a \oplus b \mapsto\langle \Lambda \rangle a' \oplus b$

By convention, we will use the letters  $m$  and  $n$  for natural numbers,  $a$ ,  $b$  and  $e$  for expressions,  $\alpha$  for actions and  $\Lambda$  for labels. Using Agda’s facility for infix operators, the proposition that  $e$  reduces in a single  $\Lambda$ -labelled step to  $e'$  is written quite naturally as follows:  $e \mapsto\langle \Lambda \rangle e'$ .

Each constructor of the above definition corresponds to a transition rule. Let us consider the two base rules, covering expression of the form  $\# m \oplus \# n$ . When reducing  $\# m \oplus \# n$ , one of two things can happen: either the two numbers are summed as usual by the  $\mapsto\boxplus$  rule, or they are ‘zapped’ to zero by  $\mapsto\zeta$ ; the two possible transitions are labelled accordingly. The inclusion of the  $\mapsto\zeta$  rule introduces

a simple form of non-determinism, as a first step towards moving from a sequential, deterministic language, to a concurrent, non-deterministic one. The two remaining rules  $\mapsto\text{-R}$  and  $\mapsto\text{-L}$  ensure a left-biased reduction order, as mentioned previously.

### 7.3.1 Choice of Action Set

How was the set of actions chosen? As we shall see later in §7.5, we wish to distinguish between different choices in the reduction path a given expression can take. In the instance of this Zap language, we need to know which of the  $\mapsto\text{-}\boxplus$  or  $\mapsto\text{-}\zeta$  rules were applied, hence the use of the distinct actions  $\boxplus$  and  $\zeta$  respectively. Where there is only one unique transition from some given state, we label it with a silent  $\tau$ . Later in §7.6 we also wish to distinguish between different final results for an expression, which are revealed using the  $\square$  action.

## 7.4 Compiler, Virtual Machine and its Semantics

The virtual machine for this language has a simple stack-based design, with the same two instructions as we had in the previous chapter:

```
data Instruction : Set where
  PUSH : (m : ℕ) → Instruction
  ADD  : Instruction
```

A program comprises a list of such instructions. The compiler for our expression language is the same as the previous chapter, and is repeated below. In order to make our proofs more straightforward, we take a code continuation as an additional argument [Hut07], which corresponds to writing the compiler in an accumulator-passing style:

**compile** : Expression  $\rightarrow$  List Instruction  $\rightarrow$  List Instruction

**compile** (#  $m$ )  $c = \text{PUSH } m :: c$

**compile** ( $a \oplus b$ )  $c = \text{compile } a (\text{compile } b (\text{ADD} :: c))$

To execute a program  $c : \text{List Instruction}$ , we pair it with a stack  $\sigma : \text{List } \mathbb{N}$ . This is precisely how we represent the state of a virtual machine:

**data** Machine : Set **where**

$\langle \_ , \_ \rangle : (c : \text{List Instruction}) (\sigma : \text{List } \mathbb{N}) \rightarrow \text{Machine}$

Finally, we specify the operational semantics of the virtual machine through the  $\_ \mapsto \langle \_ \rangle \_$  relation:

**data**  $\_ \mapsto \langle \_ \rangle \_ : \text{LTS Label Machine}$  **where**

$\mapsto\text{-PUSH} : \forall \{c \sigma m\} \rightarrow \langle \text{PUSH } m :: c , \sigma \rangle \mapsto \langle \tau \rangle \langle c , m :: \sigma \rangle$

$\mapsto\text{-ADD} : \forall \{c \sigma m n\} \rightarrow$

$\langle \text{ADD} :: c , n :: m :: \sigma \rangle \mapsto \langle ! \boxplus \rangle \langle c , m + n :: \sigma \rangle$

$\mapsto\text{-ZAP} : \forall \{c \sigma m n\} \rightarrow$

$\langle \text{ADD} :: c , n :: m :: \sigma \rangle \mapsto \langle ! \boxminus \rangle \langle c , 0 :: \sigma \rangle$


That is, the **PUSH** instruction takes a numeric argument  $m$  and pushes it onto the stack  $\sigma$ , with a silent label  $\tau$ . In turn, the **ADD** instruction replaces the top two numbers on the stack with either their sum, or zero—labelled respectively with  $\boxplus$  or  $\boxminus$ —in correspondence with the  $\mapsto\text{-ADD}$  and  $\mapsto\text{-ZAP}$  rules.

## 7.5 Non-Deterministic Compiler Correctness

In general, a compiler correctness theorem asserts that for any source program, the result of executing the corresponding compiled target code on its virtual machine will

## CHAPTER 7. COMPILING NON-DETERMINISM CORRECTLY

coincide with that of evaluating the source using its high-level semantics:



With a deterministic language and virtual machine—such as our Zap language without the two ‘zap’ rules—it would be natural to use a high-level denotational or big-step semantics for the expression language, which we can realise as an interpreter

$$\text{eval} : \text{Expression} \rightarrow \mathbb{N}$$

In turn, the low-level operational or small-step semantics for the virtual machine can be realised as a function


$$\text{exec} : \text{List } \mathbb{N} \rightarrow \text{List Instruction} \rightarrow \text{List } \mathbb{N}$$

that takes an initial stack along with a list of instructions and returns the final stack.

Compiler correctness is then the statement that:

$$\text{det-correct} : \forall c \sigma e \rightarrow \text{exec } \sigma (\text{compile } e \ c) \equiv \text{exec } (\text{eval } e \ :: \ \sigma) \ c$$

Equivalently, we can visualise this as the following commuting diagram:




That is to say, compiling an expression  $a$  and then executing the resulting code together with a code continuation  $c$  gives the same result—up to definitional equality—as executing  $c$  with the value of  $a$  pushed onto the original stack  $\sigma$ .

The presence of non-determinism requires a more refined approach, due to the possibility that different runs of the same program may give different results. One approach is to realise the interpreter and virtual machine as set-valued functions [HW07], restating the above equality on final values in terms of *sets* of final values. A more natural approach however, is to define the high-level semantics as a relation rather than a function, using a small-step operational semantics. Moreover, the small-step approach also allows us to consider the *intensional* (or local) notion of what choices are made in the reduction paths, in contrast to the *extensional* (or global) notion of comparing final results. In our Zap language, the available choices are reflected in our selection of transition labels, and we weaken the above definitional equality to a suitable notion of branching equivalence on intermediate states. This is just the familiar notion of bisimilarity [Mil89], which we shall make concrete in §7.7. As we shall see, the local reasoning afforded by this approach also leads to simpler and more natural proofs.

## 7.6 Combined Machine and its Semantics

In this section, we introduce our key idea of a ‘combined machine’, which we arrive at by considering the small-step analogue of the compiler correctness statement for big-step deterministic languages. The advantage of the combined machine is that it lifts source expressions and target virtual machine states into the same domain, which avoids the need for an intermediate process calculus [Wan95] and allows us to directly establish a bisimulation between the source and target languages.

Our approach to non-deterministic compiler correctness is illustrated below,



making use of such a combined machine, where the lifting by  $\text{inj}_S$  and  $\text{inj}_T$  are trivial enough to be ‘obviously correct’. Rather than considering final results, we consider combined machines up to a suitable notion of bisimilarity.

In the case of our Zap language, a combined machine  $x : \text{Combined}$  has three distinct phases of execution,

**data**  $\text{Combined} : \text{Set}$  **where**

$\langle \_ , \_ \rangle : (e : \text{Expression}) (t : \text{Machine}) \rightarrow \text{Combined}$

$\langle \_ \rangle : (t : \text{Machine}) \rightarrow \text{Combined}$

$\langle \rangle : \text{Combined}$

whose semantics is defined by the following transition relation:

**data**  $\_ \twoheadrightarrow \langle \_ \rangle \_ : \text{LTS Label Combined}$  **where**

$\twoheadrightarrow \mapsto : \forall \{e e' t \Lambda\} (e \mapsto e' : e \mapsto \langle \Lambda \rangle e') \rightarrow \langle e , t \rangle \twoheadrightarrow \langle \Lambda \rangle \langle e' , t \rangle$

$\twoheadrightarrow \twoheadrightarrow : \forall \{t t' \Lambda\} (t \twoheadrightarrow t' : t \twoheadrightarrow \langle \Lambda \rangle t') \rightarrow \langle t \rangle \twoheadrightarrow \langle \Lambda \rangle \langle t' \rangle$

$\twoheadrightarrow \text{-switch} : \forall \{m c \sigma\} \rightarrow \langle \# m , \langle c , \sigma \rangle \rangle \twoheadrightarrow \langle \tau \rangle \langle \langle c , m :: \sigma \rangle \rangle$

$\twoheadrightarrow \text{-done} : \forall \{m\} \rightarrow \langle \langle [] , m :: [] \rangle \rangle \twoheadrightarrow \langle ! \square m \rangle \langle \rangle$

The first constructor  $\langle \_ , \_ \rangle$  of  $\text{Combined}$  pairs an expression with a virtual machine continuation. In this initial phase, a combined machine  $\langle e , \langle c , \sigma \rangle \rangle$  can be understood as the small-step analogue of the right side of the **det-correct** statement—**exec** (**eval**  $a :: \sigma$ )  $c$ —which begins by effecting the reduction of  $a$ . The applicable



reductions are exactly those of the **Expression** language, inherited via the  $\rightarrow\text{-}\mapsto$  rule above.

When the expression  $a$  eventually reduces to a value  $m$ , the  $\rightarrow\text{-}\text{switch}$  rule pushes  $m$  onto the stack  $\sigma$ , switching the combined machine to its second phase of execution, corresponding to the  $\langle\_\rangle$  constructor. This can be thought of as the small-step analogue of pushing the result  $m \equiv \text{eval } a$  onto the stack, again following the right side of **det-correct**, namely  $\text{exec } (m :: \sigma) c$ .

The second **Combined** constructor  $\langle\_\rangle$  lifts a virtual machine into a combined machine, which then effects the reduction of the former via the  $\rightarrow\text{-}\mapsto$  rule. This corresponds to the small-step analogue of  $\text{exec } \sigma c$ , which matches the left side of **det-correct**, and also the right side after the evaluation of the embedded expression has completed.

Lastly, the  $\rightarrow\text{-}\text{done}$  rule embeds the computed result in a  $\square$  action, and terminates with the empty  $\langle\_\rangle$  state. This construction allows us to compare final result values using the existing bisimulation machinery.

## 7.7 Weak Bisimilarity

Now we can give a concrete definition to our notion of bisimilarity. More specifically, we shall be defining ‘weak bisimilarity’, as we are not concerned with silent transitions. First of all, it is convenient to define a ‘visible transition’  $\_ \Rightarrow \langle\_\rangle \_$  where only **Actions** are exposed, in terms of the  $\_ \rightarrow \langle\_\rangle \_$  relation from the previous section,

**data**  $\_ \Rightarrow \langle\_\rangle \_ : \text{LTS Action Combined where}$

$$\begin{aligned} \Rightarrow\text{-}\mapsto & : \forall \{x \ x_0 \ x_1 \ x' \ \alpha\} (x \rightarrow\tau^* x_0 : x \rightarrow\langle\tau\rangle^* x_0) \\ & (x_0 \rightarrow x_1 : x_0 \Rightarrow\langle!\alpha\rangle x_1) (x_1 \rightarrow\tau^* x' : x_1 \rightarrow\langle\tau\rangle^* x') \rightarrow \\ & x \Rightarrow\langle\alpha\rangle x' \end{aligned}$$

where we write  $_{-}\rightarrow\langle\tau\rangle^*_{-}$  as a shorthand for the reflexive transitive closure of  $\rightarrow\langle\tau\rangle$ , defined as follows:

$_{-}\rightarrow\langle\tau\rangle_{-}$  : Rel Combined

$$x \rightarrow\langle\tau\rangle y = x \rightarrow\langle\tau\rangle y$$

$_{-}\rightarrow\langle\tau\rangle^*_{-}$  : Rel Combined

$$_{-}\rightarrow\langle\tau\rangle^*_{-} = \text{Star } _{\rightarrow\langle\tau\rangle}_{-}$$

Two states  $x$  and  $y$  are now defined to be ‘weakly bisimilar’ if and only if whatever visible transition  $x$  can make,  $y$  is able to follow with the same action, resulting in states  $x'$  and  $y'$  that are also bisimilar, and vice versa. Since this is symmetric in both directions, we will define a helper  $_{-}\preceq_{-}$  for ease of readability:

$_{-}\preceq_{-}$  : Rel Combined

$$x \preceq y = \forall x' \{ \alpha \} \rightarrow x \Rightarrow\langle\alpha\rangle x' \rightarrow \exists \lambda y' \rightarrow y \Rightarrow\langle\alpha\rangle y' \times x' \approx y'$$

That is, we write  $x \preceq y$  iff for whatever state  $x'$  can be reached from  $x$  while emitting the action  $\alpha$ ,  $y$  can reach a corresponding state  $y'$  such that  $x'$  and  $y'$  are bisimilar. In turn, a relation  $x \approx y$  is simply a conjunction of  $x \preceq y$  and  $y \preceq x$ :

**data**  $_{-}\approx_{-}$  : Rel Combined **where**

$$_{-}\&_{-} : \forall \{x y\} (x \preceq y : \infty (x \preceq y)) (y \preceq x : \infty (y \preceq x)) \rightarrow x \approx y$$

The types of both  $x \preceq y$  and  $y \preceq x$  arguments are coinductive, as indicated by the use of the  $\infty$  type, since bisimilarity is a coinductive notion [Mil89, San09].

Note that in the above definition of  $_{-}\preceq_{-}$ , we demand a proof of  $x' \approx y'$  rather than a direct proof of  $x' \approx y'$ , in order to ‘beat Agda’s productivity checker’ [Dan10a]. The  $_{-}\approx_{-}$  data type can be seen as the syntax for an embedded language, where each constructor corresponds to an operation that we wish to perform on bisimilarity proofs. In this instance we only require symmetry and transitivity, along with a constructor  $\approx'\approx$  that embeds  $_{-}\approx_{-}$  proofs:

**data**  $\approx'_-$  : Rel Combined **where**

$\approx'-\approx$  :  $\approx_- \Rightarrow \approx'_-$

$\approx'$ -sym : Symmetric  $\approx'_-$

$\approx'$ -trans : Transitive  $\approx'_-$

We write  $\Rightarrow_-$  as a synonym for relational implication, which is—along with **Symmetric** and **Transitive**—defined in the standard library in the obvious manner. With this technique, we are obliged to show that  $\approx'_-$  implies  $\approx_-$ , i.e. provide an interpreter from this embedded language to an actual  $\approx_-$  proof. However, as far as comprehension is concerned, the reader can simply treat the use of  $\approx'_-$  in the definition of  $\preceq_-$  as if it were  $\approx_-$ .

Given the above, it is straightforward to prove that  $\approx_-$  is an equivalence relation on **Combined**. Reflexivity is shown by the following two mutual definitions:

**mutual**

$\preceq$ -refl : Reflexive  $\preceq_-$

$\preceq$ -refl  $x' x \Rightarrow x' = x' \wedge x \Rightarrow x' \wedge \approx'-\approx \approx$ -refl

$\approx$ -refl : Reflexive  $\approx_-$

$\approx$ -refl =  $\# \preceq$ -refl &  $\# \preceq$ -refl

The type of  $\preceq$ -refl is synonymous to  $\forall \{x\} \rightarrow x \preceq x$ , which in turn expands by the definition of  $\preceq_-$  to the following:

$$\forall \{x\} x' \{\alpha\} \rightarrow x \Rightarrow \langle \alpha \rangle x' \rightarrow \exists \lambda x'' \rightarrow x \Rightarrow \langle \alpha \rangle x'' \times x' \approx' x''$$

We prove the existence of an  $x''$  such that  $x \Rightarrow \langle \alpha \rangle x''$  and  $x' \approx' x''$  by returning the same  $x'$  and the witness of  $x \Rightarrow \langle \alpha \rangle x'$  as we were given. Proof of  $x' \approx' x'$  is obtained by embedding the result of a corecursive call to  $\approx$ -refl :  $x' \approx x'$ .

The proof of Reflexive  $\approx_-$  involves two symmetric invocations to  $\preceq$ -refl. Since both corecursive instances of  $\approx$ -refl are guarded by an unbroken chain of  $\&-$ ,  $\#-$ ,

$\_ \wedge \_$  and  $\_ \approx \_$  constructors with no intervening function invocations, Agda accepts the above definition as productive.

Symmetry on the other hand is much more straightforward,

```

≈-sym : Symmetric ≈
≈-sym (x ≈ y & y ≈ x) = y ≈ x & x ≈ y
    
```

as we need only swap the two halves of the bisimilarity proof. Finally to show transitivity of  $\_ \approx \_$ , we again make use of the symmetric nature of bisimilarity, with the help of a mutually corecursive definition **≈-trans**:

**mutual**

```

≈-trans : Transitive ≈
≈-trans x ≈ y y ≈ z x' x ⇨ x' with x ≈ y x' x ⇨ x'
≈-trans x ≈ y y ≈ z x' x ⇨ x' | y' ∧ y ⇨ y' ∧ x' ≈ y' with y ≈ z y' y ⇨ y'
≈-trans x ≈ y y ≈ z x' x ⇨ x' | y' ∧ y ⇨ y' ∧ x' ≈ y' | z' ∧ z ⇨ z' ∧ y' ≈ z'
    = z' ∧ z ⇨ z' ∧ ≈'-trans x' ≈ y' y' ≈ z'

≈-trans : Transitive ≈
≈-trans (x ≈ y & y ≈ x) (y ≈ z & z ≈ y)
    = # ≈-trans (b x ≈ y) (b y ≈ z) & # ≈-trans (b z ≈ y) (b y ≈ x)
    
```

The **≈-trans** proof—given  $x \approx y$ ,  $y \approx z$ ,  $x'$  and  $x \mapsto x'$ —has a goal of:

```

∃ λ z' → z ⇨ α > z' × x' ≈ z'
    
```

We can use  $x \approx y$  and  $y \approx z$  in two steps to construct evidence of  $y'$  and  $z'$  such that,

```

y ⇨ α > y' × x' ≈ y'   and   z ⇨ α > z' × y' ≈ z'
    
```

with the witness to  $x' \approx z'$  obtained by **≈'-trans**. The proof for **Transitive  $\_ \approx \_$**  proceeds in the same manner as **≈-refl**, with two corecursive calls to **≈-trans** for each half of the property.

Earlier we stated that we are obliged to show that  $\approx'_-$  implies  $\approx_-$ . Having now implemented all of the functions corresponding to the terms of the embedded  $\approx'_-$  language, we can proceed quite simply as follows:

$$\begin{aligned} \approx' \rightarrow \approx & : \approx'_- \Rightarrow \approx_- \\ \approx' \rightarrow \approx (\approx' \text{-}\approx x \approx y) & = x \approx y \\ \approx' \rightarrow \approx (\approx' \text{-}\text{sym } y \approx' x) & = \approx \text{-}\text{sym } (\approx' \rightarrow \approx y \approx' x) \\ \approx' \rightarrow \approx (\approx' \text{-}\text{trans } x \approx' z \ z \approx' y) & = \approx \text{-}\text{trans } (\approx' \rightarrow \approx x \approx' z) (\approx' \rightarrow \approx z \approx' y) \end{aligned}$$

Here it is clear that  $\approx' \rightarrow \approx$  is structurally recursive on its  $x \approx' y$  argument, and therefore must be total.

We conclude this section by noting that  $\approx_-$  is an equivalence relation on `Combined` machines, which enables us to use the equational (actually pre-order) reasoning combinators from the standard library module `Relation.Binary.PreorderReasoning`. The plumbing details have been omitted in this presentation for brevity, however.

## 7.8 The elide- $\tau$ Lemma

A key lemma used throughout our correctness proofs states that a silent transition between two states  $x$  and  $y$  implies that they are bisimilar:

$$\begin{aligned} \text{elide-}\tau & : \_ \rightarrow \langle \tau \rangle \_ \Rightarrow \approx \_ \\ \text{elide-}\tau \{x\} \{y\} \ x \rightarrow \tau y & = \# x \preceq y \ \& \ \# y \preceq x \ \text{where} \\ y \preceq x & : y \preceq x \\ y \preceq x \ y' (\text{H}\rightarrow\rightarrow y \rightarrow \tau^* y_0 \ y_0 \rightarrow y_1 \ y_1 \rightarrow \tau^* y') & \\ = y' \ \wedge \ \text{H}\rightarrow\rightarrow (x \rightarrow \tau y \ \triangleleft \ y \rightarrow \tau^* y_0) \ y_0 \rightarrow y_1 \ y_1 \rightarrow \tau^* y' \ \wedge \ \approx' \text{-}\text{refl} & \end{aligned}$$

In the  $y \preceq x$  direction, the proof is trivial: whatever  $y$  does,  $x$  can always match it by first making the given  $x \rightarrow \tau y$  transition, after which it can follow  $y$  exactly.

In the  $x \rightsquigarrow y$  direction, the proof relies on the fact that wherever there is a choice in the reduction of any give state, each possible transition is identified with a distinct and non-silent label, which we mentioned in §7.3.1. Conversely given a silent transition  $x \rightarrow \tau y : x \rightarrow \langle \tau \rangle y$ , it must in fact be the *unique* transition from  $x$ , which we can show by the following **unique** lemma:

$$\begin{aligned}
 \text{unique} & : \forall \{x \ y \ \Lambda \ y'\} \rightarrow x \rightarrow \langle \tau \rangle y \rightarrow x \rightarrow \langle \Lambda \rangle y' \rightarrow \Lambda \equiv \tau \times y' \equiv y \\
 \text{unique} & (\rightarrow \dashv \dashv e \dashv \dashv e') \ x \rightarrow y' = \perp\text{-elim} (\dashv \dashv \langle \tau \rangle e \dashv \dashv e') \\
 \text{unique} & (\rightarrow \dashv \dashv \dashv \dashv \text{PUSH}) (\rightarrow \dashv \dashv \dashv \dashv \text{PUSH}) = \equiv.\text{refl} \wedge \equiv.\text{refl} \\
 \text{unique} & \rightarrow\text{-switch} (\rightarrow \dashv \dashv ()) \\
 \text{unique} & \rightarrow\text{-switch} \rightarrow\text{-switch} = \equiv.\text{refl} \wedge \equiv.\text{refl}
 \end{aligned}$$

Using **unique**, the  $x \rightsquigarrow y$  direction of **elide- $\tau$**  is shown in two parts; the first shows that  $x$  cannot make a non-silent transition,

$$\begin{aligned}
 x \rightsquigarrow y & : x \rightsquigarrow y \\
 x \rightsquigarrow y \ x' & (\dashv \dashv \dashv \dashv \varepsilon \ x \rightarrow x_0 \ x_0 \rightarrow \tau^* x') \ \text{with unique } x \rightarrow \tau y \ x \rightarrow x_0 \\
 x \rightsquigarrow y \ x' & (\dashv \dashv \dashv \dashv \varepsilon \ x \rightarrow x_0 \ x_0 \rightarrow \tau^* x') \mid () \wedge x_0 \equiv y
 \end{aligned}$$

while the second shows that the first silent transition  $x$  makes must coincide with  $x \rightarrow \tau y$ , in which case  $y$  can transition to  $x'$  by following the subsequent transitions:

$$\begin{aligned}
 x \rightsquigarrow y \ x' & (\dashv \dashv \dashv \dashv (x \rightarrow x_0 \triangleleft x_0 \rightarrow \tau^* x_1) \ x_1 \rightarrow x_2 \ x_2 \rightarrow \tau^* x') \ \text{with unique } x \rightarrow \tau y \ x \rightarrow x_0 \\
 x \rightsquigarrow y \ x' & (\dashv \dashv \dashv \dashv (x \rightarrow \tau y \triangleleft y \rightarrow \tau^* x_1) \ x_1 \rightarrow x_2 \ x_2 \rightarrow \tau^* x') \mid \equiv.\text{refl} \wedge \equiv.\text{refl} \\
 & = x' \wedge \dashv \dashv \dashv \dashv y \rightarrow \tau^* x_1 \ x_1 \rightarrow x_2 \ x_2 \rightarrow \tau^* x' \wedge \approx'\text{-refl}
 \end{aligned}$$

Since  $\approx$  is transitive and reflexive, we can generalise the above result to handle silent transition sequences of arbitrary length, as well as a symmetric variant:

$$\begin{aligned}
 \text{elide-}\tau^* & : \_ \rightarrow \langle \tau \rangle^* \_ \Rightarrow \_ \approx \_ \\
 \text{elide-}\tau^* & = \text{Star.fold } \_ \approx \_ \approx\text{-trans } \approx\text{-refl} \circ \text{Star.map elide-}\tau
 \end{aligned}$$

$$\text{elide-}\tau^{*'} : \text{Sym } \_ \rightarrow \langle \tau \rangle^* \_ \approx \_$$

$$\text{elide-}\tau^{*'} = \approx\text{-sym} \circ \text{elide-}\tau^*$$

## 7.9 Compiler Correctness

Now we have enough machinery to formulate the compiler correctness theorem, which states that given a code continuation  $c$  and an initial stack  $\sigma$ , execution of the compiled code for an expression  $e$  followed by  $c$  is weakly bisimilar to the reduction of the expression  $e$  followed by the machine continuation  $\langle c, \sigma \rangle$ ,

$$\text{correctness} : \forall e c \sigma \rightarrow \langle e, \langle c, \sigma \rangle \rangle \approx \langle \langle \text{compile } e c, \sigma \rangle \rangle$$

or equivalently as the following commuting diagram:

$$\begin{array}{ccc} \text{correctness} : \forall c \sigma \rightarrow \text{Expression} & \xrightarrow{\text{compile } \_ c} & \text{List Instruction} \\ & \searrow \langle \_, \langle c, \sigma \rangle \rangle & \swarrow \langle \langle \_, \sigma \rangle \rangle \\ & \text{Combined} & / \approx \end{array}$$

In particular, instantiating  $c$  and  $\sigma$  to empty lists results in the corollary that for any expressions  $e$ , the bisimilarity  $\langle e, \langle [] , [] \rangle \rangle \approx \langle \langle \text{compile } e [] , [] \rangle \rangle$  holds.

### 7.9.1 Proof of correctness

We proceed to prove **correctness** by structural induction on the expression  $e$ . Two additional lemmas corresponding to the following propositions are required, which we will prove along the way:

$$\text{eval-left} : \langle a \oplus b, \langle c, \sigma \rangle \rangle \approx \langle a, \langle \text{compile } b (\text{ADD} :: c), \sigma \rangle \rangle$$

$$\text{eval-right} : \langle \# m \oplus b, \langle c, \sigma \rangle \rangle \approx \langle b, \langle \text{ADD} :: c, m :: \sigma \rangle \rangle$$

First, let us consider the base case of **correctness**, where  $e \equiv \# m$ :

```

correctness (# m) c σ =
  begin
    ⟨ # m , ⟨ c , σ ⟩ ⟩
  ≈⟨ elide-τ →-switch ⟩
    ⟨ ⟨ c , m :: σ ⟩ ⟩
  ≈⟨ elide-τ (→-→ →-PUSH)-1 ⟩
    ⟨ ⟨ PUSH m :: c , σ ⟩ ⟩
  ≡⟨ ≡.refl ⟩
    ⟨ ⟨ compile (# m) c , σ ⟩ ⟩
  □
    
```

As we mentioned in the conclusion of chapter 6, the equational reasoning combinators defined in the standard library [Dan10b] allow us to present the proof in a simple calculational style. In the code above, proofs of bisimilarity (or definitional equality) are supplied as the second argument of the  $\approx\langle\_ \rangle$  operator; the  $\approx\langle\_ \rangle^{-1}$  operator is a symmetric variant, while  $\equiv\langle\_ \rangle$  takes a proof of definitional equality instead. The proofs are combined using the appropriate reflexivity, symmetry and transitivity properties. That is, the above could have been equivalently written:

```

correctness (# m) c σ = ≈-trans (elide-τ →-switch)
  (≈-trans (≈-sym (elide-τ (→-→ →-PUSH)))) (≡→≈ ≡.refl))
    
```

However, the extra annotations in the equational reasoning version makes the proof easier to read and understand.

Moving on to the inductive case, where  $e \equiv a \oplus b$ :

```

correctness (a ⊕ b) c σ =
  begin
    
```



$$\begin{aligned}
& \langle a \oplus b , \langle c , \sigma \rangle \rangle \\
& \approx \langle \text{eval-left } a b c \sigma \rangle \\
& \langle a , \langle \text{compile } b (\text{ADD} :: c) , \sigma \rangle \rangle \\
& \approx \langle \text{correctness } a (\text{compile } b (\text{ADD} :: c)) \sigma \rangle \\
& \langle \langle \text{compile } (a \oplus b) c , \sigma \rangle \rangle \\
& \square
\end{aligned}$$

The first bisimilarity is given by the **eval-left** lemma, while the second uses the inductive hypothesis for the subexpression  $a$ , instantiating the code continuation with **compile**  $b$  (**ADD** ::  $c$ ).

### 7.9.2 The **eval-left** Lemma

The lemma **eval-left** has essentially a coinductive proof on the possible transitions  $\exists \lambda \alpha \rightarrow \exists \lambda a' \rightarrow a \Rightarrow \langle \alpha \rangle a'$  starting from  $a$ , with a base case—when no transitions are possible—that is mutually inductively defined with **correctness**. In the instance of this Zap language however, it suffices to proceed by case analysis on  $a$ , as the two alternatives happen to coincide with whether  $a$  has any possible transitions. For the base case where  $a$  is a numeral  $\# m$ , no further transitions are possible, and we can reason equationally as follows:

$$\begin{aligned}
& \text{eval-left} : \forall a b c \sigma \rightarrow \\
& \langle a \oplus b , \langle c , \sigma \rangle \rangle \approx \langle a , \langle \text{compile } b (\text{ADD} :: c) , \sigma \rangle \rangle \\
& \text{eval-left } (\# m) b c \sigma = \\
& \text{begin} \\
& \langle \# m \oplus b , \langle c , \sigma \rangle \rangle \\
& \approx \langle \text{eval-right } m b c \sigma \rangle \\
& \langle b , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle \\
& \approx \langle \text{correctness } b (\text{ADD} :: c) (m :: \sigma) \rangle
\end{aligned}$$

$$\begin{aligned}
 & \langle \langle \text{compile } b \text{ (ADD :: } c \text{)} , m :: \sigma \rangle \rangle \\
 & \approx \langle \text{elide-}\tau \rightarrow \text{-switch}^{-1} \rangle \\
 & \langle \# m , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle \\
 & \square
 \end{aligned}$$

The proof above makes use of the inductive hypothesis for the subexpression  $b$  (of  $e \equiv a \oplus b$ ) in the mutual definition of **correctness**, as well as the **eval-right** lemma. In the coinductive case, we consider the  $x \preccurlyeq y$  and  $y \preccurlyeq x$  halves of the bisimilarity separately,

$$\begin{aligned}
 \text{eval-left } (a^l \oplus a^r) b c \sigma &= \# x \preccurlyeq y \ \& \ \# y \preccurlyeq x \ \text{where} \\
 a &= a^l \oplus a^r
 \end{aligned}$$

As we mentioned earlier, this part of the proof is coinductive on the possible reductions rather than structural on the expression  $a$ , so the fact that  $a \equiv a^l \oplus a^r$  is besides the point. We define  $a$  as a synonym for  $a^l \oplus a^r$  for clarity as we do not need to refer to  $a^l$  or  $a^r$  in the proof. We also adopt the convention of writing  $x$  for the left, and  $y$  for the right hand sides of the  $\approx$  bisimilarity.

In the forward direction, we are given a witness to  $x \Rightarrow \alpha > x'$ , and must show that  $y$  can follow with the same action to some  $y'$ , such that the resulting  $x'$  and  $y'$  are bisimilar<sup>1</sup>. The coinduction hypothesis tells us that  $a$  can make some visible transition to an  $a'$  while emitting an action  $\alpha$ . Since the evaluation of  $a \oplus b$  is left-biased, it must be the case that  $a \oplus b$  reduces under the  $\mapsto\text{-L}$  rule. Therefore, we can extract the witness  $a \mapsto a'$  and repack it to witness the following:

$$\langle a , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle \Rightarrow \alpha > \langle a' , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle$$

<sup>1</sup>Recall that the type synonym  $\preccurlyeq$  was defined in §7.7 as follows:

$$\begin{aligned}
 \preccurlyeq &: \text{Rel Combined} \\
 x \preccurlyeq y &= \forall x' \{ \alpha \} \rightarrow x \Rightarrow \alpha > x' \rightarrow \exists \lambda y' \rightarrow y \Rightarrow \alpha > y' \times x' \approx y'
 \end{aligned}$$

That is,  $x \preccurlyeq y$  is a function that given an  $x'$ , an implicit  $\alpha$  and a witness of  $x \Rightarrow \alpha > x'$ , must return a triple comprising  $y'$ , a witness of  $y \Rightarrow \alpha > y'$ , along with a proof of  $x' \approx y'$ .

The proof below does not explicitly mention the action  $\alpha$  (among others identifiers) as this is already implied by its type, and is automatically inferred.

$$\begin{aligned}
x \rightsquigarrow y & : \langle a \oplus b , \langle c , \sigma \rangle \rangle \rightsquigarrow \langle a , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle \\
x \rightsquigarrow y & \_ . (\text{H} \dashrightarrow \varepsilon (\text{H} \dashrightarrow (\text{L} \{a' = a'\} a \mapsto a')) \varepsilon) \\
& = \langle a' , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle \\
& \wedge \text{H} \dashrightarrow \varepsilon (\text{H} \dashrightarrow a \mapsto a') \varepsilon \\
& \wedge \approx' \text{-} \approx (\text{eval-left } a' b c \sigma)
\end{aligned}$$

The above clause only matches visible transitions with empty sequences of silent combined transitions (i.e.  $\varepsilon$ ). Alternative cases are not possible since  $\_ \dashrightarrow \langle \_ \rangle \_$  transitions cannot be silent, and are eliminated using the  $\dashrightarrow \langle \tau \rangle$  lemma:

$$\begin{aligned}
x \rightsquigarrow y \ x' & (\text{H} \dashrightarrow \varepsilon (\text{H} \dashrightarrow (\text{L} a \mapsto a_0)) (\text{H} \dashrightarrow a_0 \oplus b \mapsto a_1 \oplus b \triangleleft x_1 \mapsto \tau^* x')) \\
& = \perp\text{-elim } (\dashrightarrow \langle \tau \rangle a_0 \oplus b \mapsto a_1 \oplus b) \\
x \rightsquigarrow y \ x' & (\text{H} \dashrightarrow (\text{H} \dashrightarrow a \oplus b \mapsto a_0 \oplus b \triangleleft x_0 \mapsto \tau^* x_1) x_1 \mapsto x_2 x_2 \mapsto \tau^* x') \\
& = \perp\text{-elim } (\dashrightarrow \langle \tau \rangle a \oplus b \mapsto a_0 \oplus b)
\end{aligned}$$

In the opposite direction, the combined machine  $y$  makes a visible transition to  $\langle a_0 , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle$  (henceforth denoted by  $y_0$ ), from which we can extract a witness  $a \mapsto a_0$ . This is followed by some sequence of silent transitions  $y_0 \mapsto \tau^* y'$ . In response,  $x$  can make a transition to  $\langle a_0 \oplus b , \langle c , \sigma \rangle \rangle$  ( $x'$  from here on) emitting the same action, that is bisimilar to  $y_0$  via the coinductive hypothesis on  $a_0$ . Finally,  $\text{elide-}\tau^*$   $y_0 \mapsto \tau^* y'$  provides a proof of  $y' \approx' y_0$ , with which we can obtain  $y' \approx' x'$  by transitivity:

$$\begin{aligned}
y \rightsquigarrow x & : \langle a , \langle \text{compile } b \text{ (ADD :: } c \text{)} , \sigma \rangle \rangle \rightsquigarrow \langle a \oplus b , \langle c , \sigma \rangle \rangle \\
y \rightsquigarrow x \ y' & (\text{H} \dashrightarrow \varepsilon (\text{H} \dashrightarrow \{e' = a_0\} a \mapsto a_0) y_0 \mapsto \tau^* y') \\
& = \langle a_0 \oplus b , \langle c , \sigma \rangle \rangle \\
& \wedge \text{H} \dashrightarrow \varepsilon (\text{H} \dashrightarrow (\text{L} a \mapsto a_0)) \varepsilon
\end{aligned}$$

$$\begin{aligned}
 & \wedge \approx\text{-trans } (\approx\text{-}\approx (\text{elide-}\tau^* y_0 \rightarrow \tau^* y')) (\approx\text{-sym } (\approx\text{-}\approx (\text{eval-left } a_0 \ b \ c \ \sigma))) \\
 & y \lesssim x \ y' \ (\text{E}\rightarrow\rightarrow (\rightarrow\text{-}\rightarrow a \rightarrow a_0 \ \triangleleft \ y_0 \rightarrow \tau^* y_1) \ y_1 \rightarrow y_2 \ y_2 \rightarrow \tau^* y_1) \\
 & = \perp\text{-elim } (\neg\rightarrow\triangleleft\tau \ a \rightarrow a_0)
 \end{aligned}$$

The second clause eliminates the impossible case of  $a$  making a  $\tau$  transition, which completes the proof of **eval-left**.

### 7.9.3 The **eval-right** Lemma

The **eval-right** lemma proceeds in a similar manner by coinduction on the possible  $\rightarrow\triangleleft\rightarrow$  transitions from  $b$ . In the base case,  $b$  cannot make any further such transitions (i.e. is some number  $\# n$ ), and we need to show:

$$\langle \# m \oplus \# n , \langle c , \sigma \rangle \rangle \approx \langle \# n , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle$$

As the right hand side can make a silent  $\rightarrow\text{-switch}$  transition, we can factor this part out to make the proof a little neater:

$$\text{eval-right} : \forall m \ b \ c \ \sigma \rightarrow \langle \# m \oplus b , \langle c , \sigma \rangle \rangle \approx \langle b , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle$$

$$\text{eval-right } m \ (\# n) \ c \ \sigma =$$

**begin**

$$\langle \# m \oplus \# n , \langle c , \sigma \rangle \rangle$$

$$\approx \langle \oplus \approx \text{ADD} \rangle$$

$$\langle \langle \text{ADD} :: c , n :: m :: \sigma \rangle \rangle$$

$$\approx \langle \text{elide-}\tau \rightarrow\text{-switch}^{-1} \rangle$$

$$\langle \# n , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle$$

□ **where**

$$\oplus \approx \text{ADD} : \langle \# m \oplus \# n , \langle c , \sigma \rangle \rangle \approx \langle \langle \text{ADD} :: c , n :: m :: \sigma \rangle \rangle$$

$$\oplus \approx \text{ADD} = \# x \lesssim y \ \& \ \# y \lesssim x \ \text{where}$$

The  $\oplus \approx \text{ADD}$  lemma is where we encounter the non-determinism in our Zap language. The proof for the two halves  $x \lesssim y$  and  $y \lesssim x$  are as follows: in the first instance,

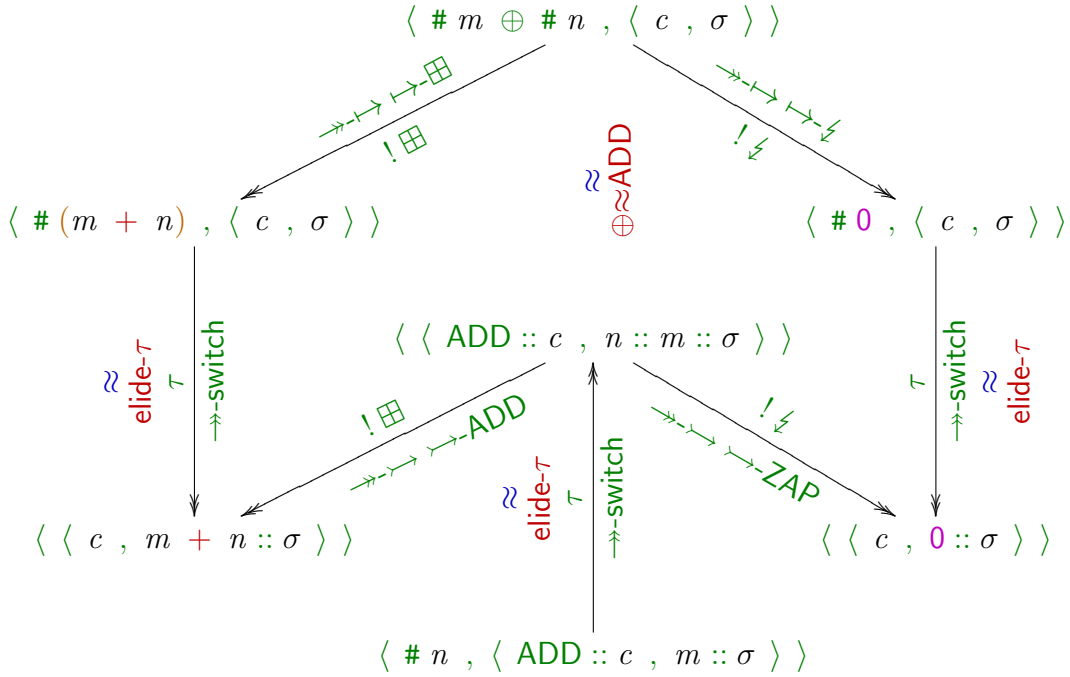
$\# m \oplus \# n$  can transition with either the  $\mapsto\text{-}\boxplus$  or  $\mapsto\text{-}\zeta$  rule, and we must show that  $\langle \text{ADD} :: c , n :: m :: \sigma \rangle$  can follow with the same action:

$$\begin{aligned}
& x \rightsquigarrow y : \langle \# m \oplus \# n , \langle c , \sigma \rangle \rangle \rightsquigarrow \langle \langle \text{ADD} :: c , n :: m :: \sigma \rangle \rangle \\
& x \rightsquigarrow y' (\mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\mapsto \mapsto\text{-}\boxplus) x_0 \rightarrow \tau^* x') \\
& \quad = \langle \langle c , m + n :: \sigma \rangle \rangle \\
& \quad \wedge \mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\rightarrow \mapsto\text{-}\text{ADD}) \varepsilon \\
& \quad \wedge \approx'\text{-}\approx (\text{elide-}\tau^{*'} x_0 \rightarrow \tau^* x') (\approx'\text{-}\approx (\text{elide-}\tau \rightarrow\text{-}\text{switch})) \\
& x \rightsquigarrow y' (\mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\mapsto \mapsto\text{-}\zeta) x_0 \rightarrow \tau^* x') \\
& \quad = \langle \langle c , 0 :: \sigma \rangle \rangle \\
& \quad \wedge \mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\rightarrow \mapsto\text{-}\text{ZAP}) \varepsilon \\
& \quad \wedge \approx'\text{-}\approx (\text{elide-}\tau^{*'} x_0 \rightarrow \tau^* x') (\approx'\text{-}\approx (\text{elide-}\tau \rightarrow\text{-}\text{switch})) \\
& x \rightsquigarrow y' (\mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\mapsto (\mapsto\text{-}\text{R} ())) x_0 \rightarrow \tau^* x') \\
& x \rightsquigarrow y' (\mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\mapsto (\mapsto\text{-}\text{L} ())) x_0 \rightarrow \tau^* x') \\
& x \rightsquigarrow y' (\mapsto\text{-}\rightarrow (\mapsto\text{-}\mapsto e \mapsto e_0 \triangleleft x_0 \rightarrow \tau^* x_1) x_1 \rightarrow x_2 x_2 \rightarrow \tau^* x') \\
& \quad = \perp\text{-}\text{elim} (\neg \mapsto\text{-}\langle \tau \rangle e \mapsto e_0)
\end{aligned}$$

This is sketched in figure 7.1—read from the top down—where the left and right branches correspond to the first two clauses above. The third and fourth clauses handle the fact that neither  $\# m$  nor  $\# n$  by themselves can reduce any further, while the last deals with silent  $\mapsto\text{-}\langle \_ \rangle\text{-}$  transitions.

The other direction of  $\oplus \approx \text{ADD}$  is illustrated by reading figure 7.1 from the bottom up, and proceeds in the same manner:

$$\begin{aligned}
& y \rightsquigarrow x : \langle \langle \text{ADD} :: c , n :: m :: \sigma \rangle \rangle \rightsquigarrow \langle \# m \oplus \# n , \langle c , \sigma \rangle \rangle \\
& y \rightsquigarrow x' (\mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\rightarrow \mapsto\text{-}\text{ADD}) y_0 \rightarrow \tau^* y') \\
& \quad = \langle \langle c , m + n :: \sigma \rangle \rangle \\
& \quad \wedge \mapsto\text{-}\rightarrow \varepsilon (\mapsto\text{-}\mapsto \mapsto\text{-}\boxplus) (\mapsto\text{-}\text{switch} \triangleleft \varepsilon) \\
& \quad \wedge \approx'\text{-}\approx (\text{elide-}\tau^{*'} y_0 \rightarrow \tau^* y')
\end{aligned}$$


 Figure 7.1: Proof sketch for the base case of the **eval-right** lemma.

$$\begin{aligned}
 & y \preceq x y' \ (\text{H}\rightarrow\rightarrow\ \varepsilon \ (\rightarrow\rightarrow\rightarrow\ \text{ZAP}) \ y_0 \rightarrow \tau^* y') \\
 & = \langle \langle c, 0 :: \sigma \rangle \rangle \\
 & \wedge \text{H}\rightarrow\rightarrow\ \varepsilon \ (\rightarrow\rightarrow\rightarrow\ \downarrow) \ (\rightarrow\rightarrow\text{switch} \triangleleft \varepsilon) \\
 & \wedge \approx'\text{-}\approx \ (\text{elide-}\tau^* y_0 \rightarrow \tau^* y') \\
 & y \preceq x y' \ (\text{H}\rightarrow\rightarrow\ (\rightarrow\rightarrow\rightarrow\ ()) \triangleleft y_0 \rightarrow \tau^* y_1) \ y_1 \rightarrow y_2 \ y_2 \rightarrow \tau^* y_1)
 \end{aligned}$$

The impossible pattern  $()$  in the final clause corresponds to the fact that there is no transition in which an **ADD** instruction emits a silent action.

The coinductive case of the **eval-right** lemma follows a similar structure to that of **eval-left**, on the possible transitions from  $b$ :

$$\begin{aligned}
 \text{eval-right } m (b^l \oplus b^r) c \sigma & = \# x \preceq y \ \& \ \# y \preceq x \ \text{where} \\
 b & = b^l \oplus b^r
 \end{aligned}$$

In the  $x \preceq y$  direction,  $\langle \# m \oplus b, \langle c, \sigma \rangle \rangle$  must reduce according to  $b \mapsto b_0$ , therefore  $y$  can follow by transitioning to  $\langle b_0, \langle \text{ADD} :: c, m :: \sigma \rangle \rangle$ :

$$\begin{aligned}
x \rightsquigarrow y & : \langle \# m \oplus b , \langle c , \sigma \rangle \rangle \rightsquigarrow \langle b , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle \\
x \rightsquigarrow y \ x' & (\mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow (\mapsto\text{-R} \{b' = b_0\} b \mapsto b_0)) x_0 \twoheadrightarrow \tau^* x') \\
& = \langle b_0 , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle \\
& \wedge \mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow b \mapsto b_0) \varepsilon \\
& \wedge \approx'\text{-trans} (\approx'\text{-}\approx (\text{elide-}\tau^* x_0 \twoheadrightarrow \tau^* x')) (\approx'\text{-}\approx (\text{eval-right } m \ b_0 \ c \ \sigma)) \\
x \rightsquigarrow y \ x' & (\mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow (\mapsto\text{-L} ())) x_0 \twoheadrightarrow \tau^* x') \\
x \rightsquigarrow y \ x' & (\mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow m \oplus b \mapsto m \oplus b_0 \triangleleft x_0 \twoheadrightarrow \tau^* x_1) x_1 \twoheadrightarrow x_2 \ x_2 \twoheadrightarrow \tau^* x') \\
& = \perp\text{-elim} (\neg \mapsto \langle \tau \rangle m \oplus b \mapsto m \oplus b_0)
\end{aligned}$$

The second clause shows that  $\# m \mapsto \langle \Lambda \rangle a'$  is uninhabited, while the third uses the fact that  $\# m \oplus b$  cannot make a silent  $\_ \mapsto \langle \_ \rangle \_$  transition, to show that both cases are impossible.

For the  $y \rightsquigarrow x$  direction, the proof simply runs in reverse, with  $x$  following  $y$  by transitioning to  $\langle \# m \oplus b_0 , \langle c , \sigma \rangle \rangle$ :

$$\begin{aligned}
y \rightsquigarrow x & : \langle b , \langle \text{ADD} :: c , m :: \sigma \rangle \rangle \rightsquigarrow \langle \# m \oplus b , \langle c , \sigma \rangle \rangle \\
y \rightsquigarrow x \ y' & (\mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow \{e' = b_0\} b \mapsto b_0) y_0 \twoheadrightarrow \tau^* y') \\
& = \langle \# m \oplus b_0 , \langle c , \sigma \rangle \rangle \\
& \wedge \mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow (\mapsto\text{-R } b \mapsto b_0)) \varepsilon \\
& \wedge \approx'\text{-trans} (\approx'\text{-}\approx (\text{elide-}\tau^* y_0 \twoheadrightarrow \tau^* y')) (\approx'\text{-sym} (\approx'\text{-}\approx (\text{eval-right } m \ b_0 \ c \ \sigma))) \\
y \rightsquigarrow x \ y' & (\mathbb{T} \dashrightarrow \varepsilon (\twoheadrightarrow \dashrightarrow b \mapsto b_0 \triangleleft y_0 \twoheadrightarrow \tau^* y_1) y_1 \twoheadrightarrow y_2 \ y_2 \twoheadrightarrow \tau^* y_1) \\
& = \perp\text{-elim} (\neg \mapsto \langle \tau \rangle b \mapsto b_0)
\end{aligned}$$

The final case deals with the impossibility of silent  $b \mapsto b_0$  transitions. This completes the proof of the **eval-right** and **eval-left** lemmas, and in turn the **correctness** theorem for our Zap language.

## 7.10 Conclusion

In this chapter we introduced a new technique for handling non-determinism in the context of compiler correctness proofs, which we illustrated using the Zap language. By carefully choosing silent and visible actions to distinguish between non-deterministic choices in the reduction of expressions and virtual machines, we were able to show the crucial `elide- $\tau$`  lemma used in the compiler correctness proof that follows. Finally, our notion of a combined machine allowed us to directly establish a bisimulation between our source and target languages without the need for an underlying process calculus.



# Chapter 8

## Compiling Concurrency Correctly

In the previous chapter, we introduced our methodology of using the notion of bisimilarity on combined machines for tacking compiler correctness for a simple non-deterministic language. In this chapter, we shall demonstrate that this technique is scalable to a concurrent setting, by extending the language with a simple `fork` primitive that introduces explicit concurrency into our system.

### 8.1 The Fork Language

#### 8.1.1 Syntax and Virtual Machine

As with the Zap language of the previous chapter, we base this Fork language on that of natural numbers and addition. The inclusion of an extra `fork` primitive introduces a simple and familiar approach to explicit concurrency:

```
data Expression : Set where  
  #_      : (m : ℕ) → Expression  
  _⊕_    : (a b : Expression) → Expression  
  fork  : (e : Expression) → Expression
```

An expression `fork e` will begin evaluation of  $e$  in a new thread, immediately returning `# 0`, in a manner reminiscent of Haskell’s `forkIO` primitive. The collection of concurrent threads in the system is modelled as a ‘thread soup’ [PJ01], defined later in §8.2. Similarly, we extend the virtual machine with a `FORK` instruction, which spawns a given sequence of instructions in a new thread:

```
data Instruction : Set where
  PUSH : (m : ℕ) → Instruction
  ADD   : Instruction
  FORK  : (c : List Instruction) → Instruction
```

The compiler includes an extra case for `fork`, but remains otherwise unchanged from the definition in §7.4:

```
compile : Expression → List Instruction → List Instruction
compile (# m) c = PUSH m :: c
compile (a ⊕ b) c = compile a (compile b (ADD :: c))
compile (fork e) c = FORK (compile e []) :: c
```

As before, each virtual machine thread comprises a list of `Instructions` along with a stack of natural numbers:

```
data Machine : Set where
  ⟨→,→⟩ : (c : List Instruction) (σ : List ℕ) → Machine
```

### 8.1.2 Actions

We extend the set of actions by  $^+ e$ , to indicate the spawning of a new thread  $e$ , and the action  $\dots \alpha$  to indicate preemption of the foremost thread:

```
data Action (L : Set) : Set where
  τ : Action L
```

$$\begin{aligned}
\boxplus & : \text{Action } L \\
\Box\_ & : (m : \mathbb{N}) \rightarrow \text{Action } L \\
^+ \_ & : (x : L) \rightarrow \text{Action } L \\
\dots\_ & : (\alpha : \text{Action } L) \rightarrow \text{Action } L
\end{aligned}$$

The above definition of **Action** is parametrised over the type of spawned threads, either **Expressions** or **Machines**. As we now have explicit concurrency in the language, we no longer require the ‘zap’ action or its associated semantics.

With the Zap language, a  $\tau$  label sufficed to identify silent actions, because its semantics does not diverge at points where silent transitions occurred. With the Fork language, we have a ‘soup’ of concurrent threads, of which more than one may be able to make a silent transition at any given point. Previously we mandated that distinct choices in the reduction path must be labelled with distinct actions: in this case, we have folded the  $\tau$  label into the definition of **Action**, such that e.g. both  $\tau$  and  $\dots \tau$  are considered to be silent, yet they remain distinct.

This choice does complicate matters somewhat: in the two-level definition of labels and actions in the Zap language, we could simply pattern match a ‘label’ with  $\tau$  to determine if a transition was silent; in the same way, we know *a priori* that ‘actions’ cannot be silent. With the Fork language, we must use a more elaborate scheme:

$$\begin{aligned}
\mathbf{data} \_ \simeq_{\tau} \{l\} & : \text{Action } l \rightarrow \text{Set} \mathbf{where} \\
\mathbf{is-}\tau & : \tau \simeq_{\tau} \\
\mathbf{is-}\dots & : \forall \{\alpha\} \rightarrow \alpha \simeq_{\tau} \rightarrow (\dots \alpha) \simeq_{\tau}
\end{aligned}$$

The above type functions as a predicate on **Actions**:  $\alpha \simeq_{\tau}$  is inhabited precisely when  $\alpha$  is considered silent. Conversely, the negation of  $\_ \simeq_{\tau}$  serves the same rôle for non-silent actions, defined as follows:

$$\begin{aligned}
\mathbf{not-}\tau & : \forall \{l\} \rightarrow \text{Action } l \rightarrow \text{Set} \\
\alpha \mathbf{not-}\tau & = \neg \alpha \simeq_{\tau}
\end{aligned}$$

### 8.1.3 Semantics

It remains for us to give the semantics for expressions and virtual machines. As per §7.3, expressions follow a left-biased reduction semantics given by the  $\mapsto\text{-}\boxplus$ ,  $\mapsto\text{-}\mathbf{R}$  and  $\mapsto\text{-}\mathbf{L}$  rules:

**data**  $\mapsto\langle\_ \rangle\_ : \text{LTS (Action Expression) Expression where}$

$$\mapsto\text{-}\boxplus : \forall \{m\ n\} \rightarrow \# m \oplus \# n \mapsto\langle \boxplus \rangle \# (m + n)$$

$$\mapsto\text{-}\mathbf{R} : \forall \{m\ b\ b'\ \alpha\} (b \mapsto\langle \alpha \rangle b') \rightarrow \# m \oplus b \mapsto\langle \alpha \rangle \# m \oplus b'$$

$$\mapsto\text{-}\mathbf{L} : \forall \{a\ a'\ b\ \alpha\} (a \mapsto\langle \alpha \rangle a') \rightarrow a \oplus b \mapsto\langle \alpha \rangle a' \oplus b$$

$$\mapsto\text{-}\text{fork} : \forall \{e\} \rightarrow \text{fork } e \mapsto\langle \text{+ } e \rangle \# 0$$

Spawning of new threads is handled by the  $\mapsto\text{-}\text{fork}$  rule, which embeds the expression in an  $\text{+ } e$  action. The expression  $\text{fork } e$  immediately reduces to  $\# 0$ , in a manner reminiscent of Haskell's  $\text{forkIO} :: \text{IO } () \rightarrow \text{IO } ()$  primitive.

In turn, the virtual machine for our Fork language inherits the  $\mapsto\text{-}\text{PUSH}$  and  $\mapsto\text{-}\text{ADD}$  rules from that of the Zap language, given in §7.4:

**data**  $\mapsto\langle\_ \rangle\_ : \text{LTS (Action Machine) Machine where}$

$$\mapsto\text{-}\text{PUSH} : \forall \{c\ \sigma\ m\} \rightarrow$$

$$\langle \text{PUSH } m :: c, \sigma \rangle \mapsto\langle \tau \rangle \langle c, m :: \sigma \rangle$$

$$\mapsto\text{-}\text{ADD} : \forall \{c\ \sigma\ m\ n\} \rightarrow$$

$$\langle \text{ADD} :: c, n :: m :: \sigma \rangle \mapsto\langle \boxplus \rangle \langle c, m + n :: \sigma \rangle$$

$$\mapsto\text{-}\text{FORK} : \forall \{c\ c'\ \sigma\} \rightarrow$$

$$\langle \text{FORK } c' :: c, \sigma \rangle \mapsto\langle \text{+ } \langle c', [] \rangle \rangle \langle c, 0 :: \sigma \rangle$$

In this instance, we have added a  $\mapsto\text{-}\text{FORK}$  rule that handles the case of a  $\text{FORK } c'$  instruction: given a sequence of instructions  $c'$ , we emit a newly initialised virtual machine embedded in an  $\text{+ } \langle c', [] \rangle$  action, leaving  $0$  on top of the stack.

## 8.2 Combined Machines and Thread Soups

Our definition of a combined machine remains unchanged from the Zap language, with the constructors  $\langle -, - \rangle$ ,  $\langle - \rangle$  and  $\langle \rangle$  corresponding to the three phases of executions:

**data** Combined : Set **where**

$\langle -, - \rangle : (e : \text{Expression}) (t : \text{Machine}) \rightarrow \text{Combined}$

$\langle - \rangle : (t : \text{Machine}) \rightarrow \text{Combined}$

$\langle \rangle : \text{Combined}$

So far, the semantics of the Fork language have been given in terms of individual expression or virtual machine threads. Since the notion of a ‘thread soup’ is common to both cases, we simply matters by introducing concurrency at the level of combined machines. It suffices to model our thread soups as **Lists** of combined machines, and we define labelled transitions between them as follows:

**data**  $\_ \rightarrow \langle \_ \rangle \_ : \text{LTS} (\text{Action Combined}) (\text{List Combined})$  **where**

$\rightarrow \mapsto : \forall \{e e' t s \alpha\} \rightarrow$

$(e \mapsto e' : e \mapsto \langle \alpha \rangle e') \rightarrow \text{let } \alpha' = \mathbf{E}^+ \langle \alpha \rangle \text{ in}$

$\langle e, t \rangle :: s \rightarrow \langle \alpha' \rangle \langle e', t \rangle :: \alpha' \vdash s$

$\rightarrow \mapsto : \forall \{t t' s \alpha\} \rightarrow$

$(t \mapsto t' : t \mapsto \langle \alpha \rangle t') \rightarrow \text{let } \alpha' = \mathbf{M}^+ \langle \alpha \rangle \text{ in}$

$\langle t \rangle :: s \rightarrow \langle \alpha' \rangle \langle t' \rangle :: \alpha' \vdash s$

$\rightarrow \text{done} : \forall \{m s\} \rightarrow$

$\langle \langle \square, m :: \square \rangle \rangle :: s \rightarrow \langle \square m \rangle \langle \rangle :: s$

$\rightarrow \text{switch} : \forall \{m c \sigma s\} \rightarrow$

$\langle \# m, \langle c, \sigma \rangle \rangle :: s \rightarrow \langle \tau \rangle \langle \langle c, m :: \sigma \rangle \rangle :: s$

$\rightarrow \text{preempt} : \forall \{x s s' \alpha\} \rightarrow (s \mapsto s' : s \mapsto \langle \alpha \rangle s') \rightarrow$

$x :: s \rightarrow \langle \dots \alpha \rangle x :: s'$

As with the Zap language, the  $\rightarrow\text{-}\mapsto$  and  $\rightarrow\text{-}\rightsquigarrow$  rules lift transitions on expression and virtual machine threads to soups of combined machines. The trivial  $E^+$  and  $M^+$  helpers lift `Expression` and `Machine` into `Combined`, given as follows,

$E^+ : \text{Expression} \rightarrow \text{Combined}$

$E^+ e = \langle e, \langle [], [] \rangle \rangle$

$M^+ : \text{Machine} \rightarrow \text{Combined}$

$M^+ = \langle \_ \rangle$

while  $\_+::\_$  inserts any spawned threads into the thread soup, defined below:

$\_+::\_ : \text{Action Combined} \rightarrow \text{List Combined} \rightarrow \text{List Combined}$

$\tau \quad +:: s = s$

$\boxplus \quad +:: s = s$

$\square m \quad +:: s = s$

$+ x \quad +:: s = x :: s$

$\dots \alpha \quad +:: s = s$

Aside from the generalisation to thread soups, the  $\rightarrow\text{-done}$  and  $\rightarrow\text{-switch}$  rules are otherwise identical to those defined for the Zap language.

Finally, we allow arbitrary thread interleaving via the  $\rightarrow\text{-preempt}$  rule. As our focus is not on the subtleties of different scheduling algorithms, we therefore do not place any restrictions on what thread the ‘scheduler’ may choose to execute next.

### 8.3 Silent and Visible Transitions

For our later proofs, it will be convenient to have a canonical definition of silent and non-silent transitions. We regard a silent transition between  $r$  and  $s$  as a triple comprising an action  $\alpha$ , a proof of  $\alpha \simeq \tau$ , along with the transition  $r \rightarrow\langle \alpha \rangle s$ :

$$\_ \rightarrow_{\mathcal{T}} \_ : \forall r s \rightarrow \text{Set}$$

$$r \rightarrow_{\mathcal{T}} s = \exists \lambda \alpha \rightarrow \alpha \simeq_{\mathcal{T}} \times r \rightarrow_{\mathcal{T}} \langle \alpha \rangle s$$

Conversely, a non-silent transition carries a proof of  $\alpha \not\rightarrow_{\mathcal{T}}$  instead:

$$\_ \rightarrow_{\not\rightarrow_{\mathcal{T}}} \_ : \forall r s \rightarrow \text{Set}$$

$$r \rightarrow_{\not\rightarrow_{\mathcal{T}}} s = \exists \lambda \alpha \rightarrow \alpha \not\rightarrow_{\mathcal{T}} \times r \rightarrow_{\mathcal{T}} \langle \alpha \rangle s$$

Finally we may write  $\_ \rightarrow_{\mathcal{T}^*} \_$  for the reflexive, transitive closure of  $\_ \rightarrow_{\mathcal{T}} \_$ , using the following definition:

$$\_ \rightarrow_{\mathcal{T}^*} \_ : \forall r s \rightarrow \text{Set}$$

$$\_ \rightarrow_{\mathcal{T}^*} \_ = \text{Star } \_ \rightarrow_{\mathcal{T}} \_$$

When we are only interested in the transitions of a single thread, the following synonyms are helpful for stating any relevant properties:

$$\_ \rightarrow_{\mathcal{T}_1} \_ : \forall x y \rightarrow \text{Set}$$

$$x \rightarrow_{\mathcal{T}_1} y = \forall s \rightarrow x :: s \rightarrow_{\mathcal{T}} y :: s$$

$$\_ \rightarrow_{\mathcal{T}_1^*} \_ : \forall x y \rightarrow \text{Set}$$

$$x \rightarrow_{\mathcal{T}_1^*} y = \forall s \rightarrow x :: s \rightarrow_{\mathcal{T}^*} y :: s$$

We subscript the above transitions with ‘<sub>1</sub>’ as a reminder that the propositions are  $\forall$ -quantified over the rest of the thread soup. For  $\_ \rightarrow_{\not\rightarrow_{\mathcal{T}_1}}$ , we must concatenate the resulting  $x'^+$  : **List Combined** to the rest of the soup,

$$\_ \rightarrow_{\not\rightarrow_{\mathcal{T}_1}} \_ : \forall x x'^+ \rightarrow \text{Set}$$

$$x \rightarrow_{\not\rightarrow_{\mathcal{T}_1}} x'^+ = \forall s \rightarrow x :: s \rightarrow_{\not\rightarrow_{\mathcal{T}}} x'^+ ++ s$$

as non-silent transitions may potentially spawn new threads. Finally, the  $\_ \triangleleft_1 \_$  function allows us to conveniently combine  $\_ \rightarrow_{\mathcal{T}_1}$  sequences, in the same manner as the  $\_ \triangleleft \_$  constructor of the **Star** type:

$$\begin{aligned} \_ \triangleleft_1 \_ & : \forall \{x\ x'\ y\} \rightarrow x \rightarrow_{\tau_1} x' \rightarrow x' \rightarrow_{\tau_1^*} y \rightarrow x \rightarrow_{\tau_1^*} y \\ x \rightarrow_{\tau_1} x' \triangleleft_1 x' \rightarrow_{\tau_1^*} y & = \lambda s \rightarrow x \rightarrow_{\tau_1} x' s \triangleleft x' \rightarrow_{\tau_1^*} y s \end{aligned}$$

Given the above definitions of silent and non-silent transitions, our notion of a visible transition is identical in essence to that of the Zap language, given back in §7.7:

**data**  $\_ \Rightarrow \_ < \_ > \_ : \text{LTS (Action } \top \text{) (List Combined) where}$

$$\begin{aligned} \_ \Rightarrow \_ \rightarrow \_ & : \forall \{s\ s_0\ s_1\ s'\} \\ (s \rightarrow_{\tau^*} s_0 & : s \rightarrow_{\tau^*} s_0) \\ (s_0 \rightarrow_{\not\tau} s_1 & : s_0 \rightarrow_{\not\tau} s_1) \\ (s_1 \rightarrow_{\tau^*} s' & : s_1 \rightarrow_{\tau^*} s') \rightarrow \\ s \Rightarrow & \langle [ s_0 \rightarrow_{\not\tau} s_1 ] \rangle s' \end{aligned}$$

As we do not have direct access to the action emitted by the non-silent  $s_0 \rightarrow_{\not\tau} s_1$  transition, we require a helper  $\llbracket \_ \rrbracket$  to extract the visible action:

$$\begin{aligned} \text{visible} & : \forall \{\alpha : \text{Action Combined}\} \rightarrow \alpha \rightarrow_{\not\tau} \rightarrow \text{Action } \top \\ \text{visible} \{\tau\} \ \alpha \rightarrow_{\not\tau} & = \perp\text{-elim } (\alpha \rightarrow_{\not\tau} \text{ is-}\tau) \\ \text{visible} \{\boxplus\} \ \alpha \rightarrow_{\not\tau} & = \boxplus \\ \text{visible} \{\square m\} \ \alpha \rightarrow_{\not\tau} & = \square m \\ \text{visible} \{+ x\} \ \alpha \rightarrow_{\not\tau} & = + \text{tt} \\ \text{visible} \{\dots \alpha\} \ \dots \alpha \rightarrow_{\not\tau} & = \text{visible } (\dots \alpha \rightarrow_{\not\tau} \circ \text{is-}\dots) \\ \llbracket \_ \rrbracket & : \forall \{s\ s'\} \rightarrow s \rightarrow_{\not\tau} s' \rightarrow \text{Action } \top \\ \llbracket \_ \rrbracket (\alpha \wedge \alpha \rightarrow_{\not\tau} \wedge s \rightarrow s') & = \text{visible } \alpha \rightarrow_{\not\tau} \end{aligned}$$

By this point, any spawned threads will have been inserted into the thread soup already, so we are no longer interested in its particulars, other than that a fork has taken place. Correspondingly,  $\llbracket \_ \rrbracket$  returns an **Action**  $\top$ —where  $\top$  is the singleton ‘unit’ type—rather than an **Action Combined**.



## 8.4 Bisimilarity

The definition of bisimilarity remains identical to that of the Zap language given in §7.7, save for a change in the carrier set from `Combined` to `List Combined`:

```

 $\_ \preceq \_ : \text{Rel } (\text{List Combined})$ 
 $x \preceq y = \forall x' \{ \alpha \} \rightarrow x \mapsto \langle \alpha \rangle x' \rightarrow \exists \lambda y' \rightarrow y \mapsto \langle \alpha \rangle y' \times x' \approx' y'$ 

data  $\_ \approx \_ (x : \text{List Combined}) : \text{List Combined} \rightarrow \text{Set}$  where
   $\_ \& \_ : \forall \{y\} \rightarrow (x \preceq y : \infty (x \preceq y)) \rightarrow (y \preceq x : \infty (y \preceq x)) \rightarrow x \approx y$ 

```

The embedded language of  $\_ \approx \_$  gains an extra symbol  $\approx'$ -`cong2` that allows us to combine two pairs of bisimilar thread soups. Formally, it corresponds to the proposition that  $\_ \approx \_$  is a congruence relation on the monoid  $(\text{List Combined}, ++, [])$ ,

```

data  $\_ \approx' \_ : \text{Rel } (\text{List Combined})$  where
   $\approx' \text{-}\approx : \_ \approx \_ \Rightarrow \_ \approx' \_$ 
   $\approx' \text{-sym} : \text{Symmetric } \_ \approx' \_$ 
   $\approx' \text{-trans} : \text{Transitive } \_ \approx' \_$ 
   $\approx' \text{-cong}_2 : \_ ++ \_ \text{ Preserves}_2 \_ \approx' \_ \rightarrow \_ \approx' \_ \rightarrow \_ \approx' \_$ 

```

where the type of the  $\approx'$ -`cong2` constructor expands to:

```

 $\approx' \text{-cong}_2 : \forall \{r^l s^l r^r s^r\} \rightarrow r^l \approx' s^l \rightarrow r^r \approx' s^r \rightarrow r^l ++ r^r \approx' s^l ++ s^r$ 

```

The same proofs from §7.7 suffices to show that  $\_ \approx \_$  forms an equivalence relation on thread soups. The obligation to show that  $\_ \approx' \_$  implies  $\_ \approx \_$  will be fulfilled at the end of §8.7, as it depends on the proof that  $\_ \approx \_$  is also a congruence relation on  $(\text{List Combined}, ++, [])$ . Before we can do that however, we have a few more lemmas to establish.

## 8.5 Properties of Thread Soups

In this section, we will highlight various lemmas concerning thread soups that are used towards the final correctness theorem for this Fork language. For brevity, we will omit the proofs, and instead hint at the proof method; the full Agda source code may be found on my website.

### 8.5.1 Soups Concatenation Preserves Silent Transitions

Concatenation of thread soups preserve silent transition sequences,

$$\rightarrow_{\mathcal{T}^*} \text{-} \text{++} \text{-} : \text{-} \text{++} \text{-} \text{ Preserves}_2 \text{-} \rightarrow_{\mathcal{T}^*} \text{-} \longrightarrow \text{-} \rightarrow_{\mathcal{T}^*} \text{-} \longrightarrow \text{-} \rightarrow_{\mathcal{T}^*} \text{-}$$

or equivalently, given  $r \rightarrow_{\mathcal{T}^*} r'$  and  $s \rightarrow_{\mathcal{T}^*} s'$ , we can produce a silent transition sequence  $r \text{ ++ } s \rightarrow_{\mathcal{T}^*} r' \text{ ++ } s'$ :

$$\rightarrow_{\mathcal{T}^*} \text{-} \text{++} \text{-} : \forall \{r \ r' \ s \ s'\} \rightarrow r \rightarrow_{\mathcal{T}^*} r' \rightarrow s \rightarrow_{\mathcal{T}^*} s' \rightarrow r \text{ ++ } s \rightarrow_{\mathcal{T}^*} r' \text{ ++ } s'$$

We can proceed by structural induction on the first thread soup argument, which is  $x :: r$  in the inductive case. Using the fact that forking cannot be silent, we can decapitate the  $x$  from the thread soup to obtain a pair of transitions  $x \rightarrow_{\mathcal{T}_1^*} x'$  and  $r \rightarrow_{\mathcal{T}^*} r'$ . The first of these can be instantiated to one half of the goal, that is:

$$x :: r \text{ ++ } s \rightarrow_{\mathcal{T}^*} x' :: r \text{ ++ } s$$

The induction hypothesis on the other hand uses the second  $r \rightarrow_{\mathcal{T}^*} r'$  to give the sequence  $r \text{ ++ } s \rightarrow_{\mathcal{T}^*} r' \text{ ++ } s'$ , which we can map  $\rightarrow\text{-preempt}$  over to arrive at the other half of the goal:

$$x' :: r \text{ ++ } s \rightarrow_{\mathcal{T}^*} x' :: r' \text{ ++ } s'$$

Concatenation of the two halves completes the proof.

### 8.5.2 Partitioning Silent Transitions

Conversely, the thread soups of a silent transition sequence  $r \dashv\vdash s \rightarrow_{\mathcal{T}^*} r's'$  can be partitioned into  $r \rightarrow_{\mathcal{T}^*} r'$  and  $s \rightarrow_{\mathcal{T}^*} s'$ :

$$\begin{aligned} \rightarrow_{\mathcal{T}^*}\text{-split} &: \forall r s \{r's'\} \rightarrow r \dashv\vdash s \rightarrow_{\mathcal{T}^*} r's' \rightarrow \\ &\exists_2 \lambda r' s' \rightarrow r' \dashv\vdash s' \equiv r's' \times r \rightarrow_{\mathcal{T}^*} r' \times s \rightarrow_{\mathcal{T}^*} s' \end{aligned}$$

Again, the proof uses structural induction on the first thread soup argument; decapitating  $x$  from the  $x :: r$  in the inductive case—as per the proof for  $\rightarrow_{\mathcal{T}^*}\text{-}\dashv\vdash$ —produces a pair of transitions  $x \rightarrow_{\mathcal{T}_1^*} x'$  and  $r \rightarrow_{\mathcal{T}^*} r'$ . The induction hypothesis delivers the  $r \rightarrow_{\mathcal{T}^*} r'$  needed to construct the sequence  $x :: r \rightarrow_{\mathcal{T}^*} x' :: r \rightarrow_{\mathcal{T}^*} x' :: r'$ , which completes the proof.

A useful corollary of  $\rightarrow_{\mathcal{T}^*}\text{-split}$  allows us to focus our attention on a single thread, by dissecting its transitions out of a transition sequence on the entire thread soup:

$$\begin{aligned} \rightarrow_{\mathcal{T}^*}\text{-dissect} &: \forall r^l r^r \{x r'\} \rightarrow r^l \dashv\vdash x :: r^r \rightarrow_{\mathcal{T}^*} r' \rightarrow \\ &\exists_2 \lambda r^{l'} r^{r'} \rightarrow \exists \lambda x' \rightarrow r' \equiv r^{l'} \dashv\vdash x' :: r^{r'} \times \\ &r^l \rightarrow_{\mathcal{T}^*} r^{l'} \times r^r \rightarrow_{\mathcal{T}^*} r^{r'} \times x \rightarrow_{\mathcal{T}_1^*} x' \end{aligned}$$

In other words, given a silent transition sequence starting from the thread soup  $r^l \dashv\vdash x :: r^r$ , there exists  $r^{l'}$ ,  $r^{r'}$  and  $x'$  satisfying  $r^l \rightarrow_{\mathcal{T}^*} r^{l'}$ ,  $r^r \rightarrow_{\mathcal{T}^*} r^{r'}$  and  $x \rightarrow_{\mathcal{T}_1^*} x'$  respectively.

### 8.5.3 Partitioning a Non-Silent Transition

We can also partition the thread soup for a non-silent transition, although the situation is a little more involved:

$$\begin{aligned} \rightarrow_{\not\mathcal{T}}\text{-split} &: \forall r s \{r's'\} \rightarrow (rs \rightarrow r's' : r \dashv\vdash s \rightarrow_{\not\mathcal{T}} r's') \rightarrow \\ &(\exists \lambda r' \rightarrow r' \dashv\vdash s \equiv r's' \times \Sigma (r \rightarrow_{\not\mathcal{T}} r')) \\ &\lambda r \rightarrow r' \rightarrow \llbracket r \rightarrow r' \rrbracket \equiv \llbracket rs \rightarrow r's' \rrbracket \uplus \end{aligned}$$

$$\begin{aligned}
 & (\exists \lambda s' \rightarrow r \text{ ++ } s' \equiv r's' \times \Sigma (s \rightarrow_{\not\tau} s')) \\
 & \lambda s \rightarrow s' \rightarrow \llbracket s \rightarrow s' \rrbracket \equiv \llbracket rs \rightarrow r's' \rrbracket
 \end{aligned}$$

Partitioning the thread soup in a non-silent transition has two possible outcomes, as the active thread responsible for the transition could be in either one of  $r$  or  $s$ . The proof proceeds by structural induction on the soup  $r$  as before, by inspecting each of its threads. If found, we construct and return the transition  $r \rightarrow_{\not\tau} r'$ ; otherwise none of the  $r$  threads are responsible for the non-silent transition, so by elimination it must be in  $s$ , and we can respond with the transition  $s \rightarrow_{\not\tau} s'$ . In both cases, we construct a proof that the action emitted by the original  $rs \rightarrow r's'$  transition is the same as either  $\llbracket r \rightarrow r' \rrbracket$  or  $\llbracket s \rightarrow s' \rrbracket$ , as appropriate.

In the same way that  $\rightarrow_{\tau^*}$ -split has its  $\rightarrow_{\tau^*}$ -dissect corollary, we can show the following  $\rightarrow_{\not\tau}$ -dissect corollary for  $\rightarrow_{\not\tau}$ -split:

$$\begin{aligned}
 \rightarrow_{\not\tau}\text{-dissect} & : \forall r^l r^r \{x x' r'\} \rightarrow \\
 & x \rightarrow_{\tau_1} x' \rightarrow (r^l x r^r \rightarrow r' : r^l \text{ ++ } x :: r^r \rightarrow_{\not\tau} r') \rightarrow \\
 & (\exists \lambda r^{l'} \rightarrow r' \equiv r^{l'} \text{ ++ } x :: r^r \times \Sigma (r^l \rightarrow_{\not\tau} r^{l'})) \\
 & \lambda r^l \rightarrow r^{l'} \rightarrow \llbracket r^l \rightarrow r^{l'} \rrbracket \equiv \llbracket r^l x r^r \rightarrow r' \rrbracket \uplus \\
 & (\exists \lambda r^{r'} \rightarrow r' \equiv r^l \text{ ++ } x :: r^{r'} \times \Sigma (r^r \rightarrow_{\not\tau} r^{r'})) \\
 & \lambda r^r \rightarrow r^{r'} \rightarrow \llbracket r^r \rightarrow r^{r'} \rrbracket \equiv \llbracket r^l x r^r \rightarrow r' \rrbracket
 \end{aligned}$$

Here we are given an additional hypothesis that the thread  $x$  makes a silent initial transition, which is unique by our choice of actions. Therefore the thread responsible for the non-silent transition  $r^l x r^r \rightarrow r'$  must reside in either one of  $r^l$  or  $r^r$ , giving rise to the two alternatives of either  $r^l \rightarrow_{\not\tau} r^{l'}$  or  $r^r \rightarrow_{\not\tau} r^{r'}$ .

### 8.5.4 Dissecting a Visible Transition

Recall that a visible transition comprises of two sequences of silent transitions either side of a single non-silent transition. Therefore combining the previous results allows

us to dissect a visible transition in much the same way:

$$\begin{aligned}
\text{\color{red}\(\Rightarrow\)-dissect} & : \forall r^l r^r \{x x' r' \alpha\} \rightarrow \\
& x \rightarrow_{\tau_1} x' \rightarrow r^l \text{\color{red}\(+\)} x :: r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r' \rightarrow \\
& r^l \text{\color{red}\(+\)} x' :: r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r' \text{\color{blue}\(\Updownarrow\)} \\
& \exists_2 \lambda r^{l'} r^{r'} \rightarrow r' \equiv r^{l'} \text{\color{red}\(+\)} x :: r^{r'} \times \\
& ((r^l \text{\color{blue}\(\Rightarrow\)} \alpha > r^{l'} \times r^r \rightarrow_{\tau^*} r^{r'}) \text{\color{blue}\(\Updownarrow\)} (r^l \rightarrow_{\tau^*} r^{l'} \times r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r^{r'}))
\end{aligned}$$

Given a witness of  $x \rightarrow_{\tau_1} x'$  and the visible transition  $r^l \text{\color{red}\(+\)} x :: r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r'$ , there are two possibilities regarding the thread  $x$ : either  $x \rightarrow_{\tau_1} x'$  takes place somewhere within the visible transition, so that removing it results in a witness of  $r^l \text{\color{red}\(+\)} x' :: r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r'$ ; or that  $x$  remains inactive throughout while  $r^l$  and  $r^r$  make transitions to some  $r^{l'}$  and  $r^{r'}$  respectively. Depending on which of  $r^l$  or  $r^r$  the thread responsible for the non-silent action is found in, we provide witnesses of either  $r^l \text{\color{blue}\(\Rightarrow\)} \alpha > r^{l'}$  and  $r^r \rightarrow_{\tau^*} r^{r'}$ , or  $r^l \rightarrow_{\tau^*} r^{l'}$  and  $r^r \text{\color{blue}\(\Rightarrow\)} \alpha > r^{r'}$  respectively.

The proof—totalling approximately 60 wrapped lines of code—follows a straightforward method, namely using  $\rightarrow_{\tau^*}$ -dissect and  $\rightarrow_{\cancel{\tau}}\text{-dissect}$  to tease apart the thread soup, then putting the pieces back together in the right way, depending on what we find inside.

### 8.5.5 Extracting the Active Thread

A number of the previous results were concerned with transitions from thread soups of the form  $r^l \text{\color{red}\(+\)} x :: r^r$ , where  $x$  makes an initial silent transition. This final lemma shows that every silent transition  $r \rightarrow_{\tau} r'$  is in fact of this form. In other words, we can extract from  $r \rightarrow_{\tau} r'$  the active thread  $x x'$  and a witness of its transition  $x \rightarrow_{\tau_1} x'$ , along with evidence that other threads in  $r$  remain unchanged:

$$\begin{aligned}
\rightarrow_{\tau}\text{-extract} & : \forall \{r r'\} \rightarrow r \rightarrow_{\tau} r' \rightarrow \exists_2 \lambda r^l r^r \rightarrow \exists_2 \lambda x x' \rightarrow \\
& r \equiv r^l \text{\color{red}\(+\)} x :: r^r \times r' \equiv r^l \text{\color{red}\(+\)} x' :: r^r \times x \rightarrow_{\tau_1} x'
\end{aligned}$$

The proof proceeds simply by induction on the structure of  $r \twoheadrightarrow_{\tau} r'$ .

## 8.6 The **elide- $\tau$** Lemma

Given the arsenal assembled in the previous section, the proof of the **elide- $\tau$**  lemma is relatively straightforward:

$$\mathbf{elide-}\tau : \_ \twoheadrightarrow_{\tau} \_ \Rightarrow \_ \approx \_$$

$$\mathbf{elide-}\tau \{r\} \{s\} r \twoheadrightarrow_{\tau} s = \# r \preceq s r \twoheadrightarrow_{\tau} s \ \& \ \# s \preceq r \ \mathbf{where}$$

$$s \preceq r : s \preceq r$$

$$s \preceq r s' (\text{H} \twoheadrightarrow s \twoheadrightarrow^* s_0 s_0 \twoheadrightarrow s_1 s_1 \twoheadrightarrow^* s')$$

$$= s' \wedge \text{H} \twoheadrightarrow (r \twoheadrightarrow_{\tau} s \triangleleft s \twoheadrightarrow^* s_0) s_0 \twoheadrightarrow s_1 s_1 \twoheadrightarrow^* s' \wedge \approx'\text{-refl}$$

For  $s \preceq r$  the proof is trivial: whatever  $s$  does,  $r$  can always match it by first making the given  $r \twoheadrightarrow_{\tau} s$  transition, after which it can follow  $s$  exactly.

In the other direction, we begin by extracting (§8.5.5) the active thread  $x$  from  $r \twoheadrightarrow_{\tau} s$ . This refines  $r$  to  $r^l \# x :: r^r$ , which allows us to dissect (§8.5.4)  $r \twoheadrightarrow_{\tau} r'$  using  $x$  as the pivot:

$$r \preceq s : \forall \{r s\} \rightarrow r \twoheadrightarrow_{\tau} s \rightarrow r \preceq s$$

$$r \preceq s r \twoheadrightarrow_{\tau} s r' r \twoheadrightarrow_{\tau} r'$$

$$\mathbf{with} \twoheadrightarrow_{\tau}\text{-extract } r \twoheadrightarrow_{\tau} s$$

$$r \preceq s r \twoheadrightarrow_{\tau} s r' r \twoheadrightarrow_{\tau} r'$$

$$| r^l \wedge r^r \wedge x \wedge x' \wedge \equiv.\text{refl} \wedge \equiv.\text{refl} \wedge x \twoheadrightarrow_{\tau_1} x'$$

$$\mathbf{with} \twoheadrightarrow_{\tau_1}\text{-dissect } r^l r^r x \twoheadrightarrow_{\tau_1} x' r \twoheadrightarrow_{\tau} r'$$

In the instance where  $r \twoheadrightarrow_{\tau} r'$  happens to already include the  $x \twoheadrightarrow_{\tau_1} x'$  transition, the proof is trivial:

$$r \preceq s r \twoheadrightarrow_{\tau} s r' r \twoheadrightarrow_{\tau} r'$$

$$\begin{array}{l}
| r^l \wedge r^r \wedge x \wedge x' \wedge \equiv.\text{refl} \wedge \equiv.\text{refl} \wedge x \rightarrow_{\tau_1} x' \\
| \text{in}_L s \Rightarrow r' = r' \wedge s \Rightarrow r' \wedge \approx'\text{-refl}
\end{array}$$

Here,  $\Rightarrow\text{-dissect}$  provides the witness  $s \Rightarrow r'$  showing that  $s$  can transition to  $r'$  too, with  $r' \approx' r'$  given by reflexivity of  $\approx'_-$ .

Otherwise  $r \Rightarrow r'$  has yet to make the  $x \rightarrow_{\tau_1} x'$  transition. Two alternatives arise, as the non-silent transition could have been on either side of  $x$ . Without loss of generalisation suppose this is on the left, in which case  $\Rightarrow\text{-dissect}$  refines  $r'$  to  $r^{l'} \text{ ++ } x :: r^{r'}$ , and delivers witnesses of  $r^l \Rightarrow \langle \alpha \rangle r^{l'}$  and  $r^r \rightarrow_{\tau^*} r^{r'}$ :

$$\begin{array}{l}
\mathbf{r} \Leftarrow s \ r \rightarrow_{\tau} s \ . \_ \ r \Rightarrow r' \\
| r^l \wedge r^r \wedge x \wedge x' \wedge \equiv.\text{refl} \wedge \equiv.\text{refl} \wedge x \rightarrow_{\tau_1} x' \\
| \text{in}_R (r^{l'} \wedge r^{r'} \wedge \equiv.\text{refl} \\
\quad \wedge \text{in}_L (\Rightarrow \rightarrow r^l \rightarrow_{\tau^*} r_0^l \ r_0^l \rightarrow_{r_1^l} r_1^l \rightarrow_{\tau^*} r^{l'} \wedge r^r \rightarrow_{\tau^*} r^{r'})) \\
\mathbf{with} \ \rightarrow_{\tau} \text{-append} (x' :: r^r) \ r_0^l \rightarrow_{r_1^l} \\
\mathbf{r} \Leftarrow s \ r \rightarrow_{\tau} s \ . \_ \ r \Rightarrow r' \\
| r^l \wedge r^r \wedge x \wedge x' \wedge \equiv.\text{refl} \wedge \equiv.\text{refl} \wedge x \rightarrow_{\tau_1} x' \\
| \text{in}_R (r^{l'} \wedge r^{r'} \wedge \equiv.\text{refl} \\
\quad \wedge \text{in}_L (\Rightarrow \rightarrow r^l \rightarrow_{\tau^*} r_0^l \ r_0^l \rightarrow_{r_1^l} r_1^l \rightarrow_{\tau^*} r^{l'} \wedge r^r \rightarrow_{\tau^*} r^{r'})) \\
| r_0^l x' r^r \rightarrow_{r_1^l} x' r^r \wedge \llbracket r_0^l \rightarrow_{r_1^l} \rrbracket \equiv \llbracket r_0^l x' r^r \rightarrow_{r_1^l} x' r^r \rrbracket \\
\mathbf{rewrite} \ \llbracket r_0^l \rightarrow_{r_1^l} \rrbracket \equiv \llbracket r_0^l x' r^r \rightarrow_{r_1^l} x' r^r \rrbracket \\
= r^{l'} \text{ ++ } x' :: r^{r'} \\
\wedge \Rightarrow \rightarrow (\rightarrow_{\tau^*} \text{-++} \ r^l \rightarrow_{\tau^*} r_0^l \ (\varepsilon \{x = x' :: r^r\})) \ r_0^l x' r^r \rightarrow_{r_1^l} x' r^r \\
(\rightarrow_{\tau^*} \text{-++} \ r_1^l \rightarrow_{\tau^*} r^{l'} \ (\varepsilon \{x = x' :: r^r\}) \ \llcorner \\
\rightarrow_{\tau^*} \text{-++} \ (\varepsilon \{x = r^{l'}\}) \ (\rightarrow_{\tau^*} \text{-++} \ (\varepsilon \{x = x' :: \square\}) \ r^r \rightarrow_{\tau^*} r^{r'})) \\
\wedge \approx'\text{-}\approx \ (\text{elide-}\tau \ (\rightarrow_{\tau} \text{-prepend} \ r^{l'} \ (x \rightarrow_{\tau_1} x' \ r^{r'})))
\end{array}$$

Note that the earlier  $\rightarrow_{\tau}\text{-extract}$  had established that  $s$  is in fact equal to  $r^l \text{ ++ } x' :: r^r$ .

Therefore, we can construct a visible transition from  $s$ ,

$$r^l \text{ ++ } x' :: r^r \Rightarrow \langle \alpha \rangle r^{l'} \text{ ++ } x' :: r^{r'}$$

by reconstituting the aforementioned  $r^l \Rightarrow \langle \alpha \rangle r^{l'}$  and  $r^r \rightarrow_{\tau^*} r^{r'}$ . The final bisimilarity component of the proof is obtained by coinduction on:

$$r^{l'} \text{ ++ } x :: r^{r'} \rightarrow_{\tau} r^{l''} \text{ ++ } x' :: r^{r''}$$

For the case where the non-silent transition is to the right of  $x$ , the proof follows the same approach.

Using the transitivity and reflexivity of  $\approx$  we can generalise `elide- $\tau$`  to silent transition sequences, as well as a symmetric variant:

$$\text{elide-}\tau^* : \_ \rightarrow_{\tau^*} \_ \Rightarrow \_ \approx \_$$

$$\text{elide-}\tau^* = \text{Star.fold } \_ \approx \_ \approx\text{-trans } \approx\text{-refl} \circ \text{Star.map elide-}\tau$$

$$\text{elide-}\tau^{*'} : \_ \rightarrow_{\tau^*} \_ \Rightarrow \text{flip } \_ \approx \_$$

$$\text{elide-}\tau^{*'} = \approx\text{-sym} \circ \text{elide-}\tau^*$$

## 8.7 Soup Concatenation Preserves Bisimilarity

With the introduction of explicit concurrency in this Fork language, another important lemma used in our compiler correctness proof concerns the result of combining two pairs of bisimilar thread soups. That is, given  $r^l \approx s^l$  and  $r^r \approx s^r$ , concatenating the soups pairwise results in a pair of bisimilar soups,  $r^l \text{ ++ } r^r \approx s^l \text{ ++ } s^r$ .

Intuitively, one can appeal to the following reasoning to see why this is true; without loss of generality, we need only consider  $r^l \text{ ++ } r^r \preceq s^l \text{ ++ } s^r$  as the other direction can be obtained by symmetry. We must show that whatever visible transition  $r^l \text{ ++ } r^r$  makes,  $s^l \text{ ++ } s^r$  is able to follow with the same action. If non-silent transition is due to  $r^l$ , then the  $s^l$  half of  $s^l \text{ ++ } s^r$  can match it by  $r^l \approx s^l$ , and vice versa for  $r^r$ . Any silent transitions can be bridged using the `elide- $\tau^*$`  lemma.



## 8.7. SOUP CONCATENATION PRESERVES BISIMILARITY

Let us now formalise the above argument. We are given  $r'$  and the three transition sequences  $r \rightarrow \tau^* r_0$ ,  $r_0 \rightarrow \not\rightarrow \tau r_1$  and  $r_1 \rightarrow \tau^* r'$  comprising  $r \Rightarrow \langle \alpha \rangle r'$ . By  $r$ , we actually mean  $r^l \dashv\vdash r^r$ ; therefore we can use  $\rightarrow \tau^*$ -split to partition  $r \rightarrow \tau^* r_0$  into two sequences  $r^l \rightarrow \tau^* r_0^l$  and  $r^r \rightarrow \tau^* r_0^r$ , which refines  $r_0$  to  $r_0^l \dashv\vdash r_0^r$ :

$$\begin{aligned}
 & \Leftarrow\text{-cong}_2 : \_ \dashv\vdash \_ \text{ Preserves}_2 \_ \approx \_ \longrightarrow \_ \approx \_ \longrightarrow \_ \Leftarrow \_ \\
 & \Leftarrow\text{-cong}_2 \{r^l\} \{s^l\} \{r^r\} \{s^r\} \ r^l \approx s^l \ r^r \approx s^r \ r' \ (\Rightarrow \rightarrow r \rightarrow \tau^* r_0 \ r_0 \rightarrow \not\rightarrow \tau r_1 \ r_1 \rightarrow \tau^* r') \\
 & \quad \text{with } \rightarrow \tau^*\text{-split } r^l \ r^r \ r \rightarrow \tau^* r_0 \\
 & \Leftarrow\text{-cong}_2 \{r^l\} \{s^l\} \{r^r\} \{s^r\} \ r^l \approx s^l \ r^r \approx s^r \ r' \ (\Rightarrow \rightarrow r \rightarrow \tau^* r_0 \ r_0 \rightarrow \not\rightarrow \tau r_1 \ r_1 \rightarrow \tau^* r') \\
 & \quad | \ r_0^l \wedge r_0^r \wedge \equiv.\text{refl} \wedge r^l \rightarrow \tau^* r_0^l \wedge r^r \rightarrow \tau^* r_0^r \\
 & \quad \text{with } \rightarrow \not\rightarrow \tau\text{-split } r_0^l \ r_0^r \ r_0 \rightarrow \not\rightarrow \tau r_1
 \end{aligned}$$

Carrying on in the same vein, we use  $\rightarrow \not\rightarrow \tau$ -split on  $r_0 \rightarrow \not\rightarrow \tau r_1$  to locate which side of  $r_0^l \dashv\vdash r_0^r$  the non-silent transition comes from. The proofs for both cases are symmetrical, so let us consider just the left instance:  $\rightarrow \not\rightarrow \tau$ -split returns a witness  $r_0^l \rightarrow \not\rightarrow \tau r_1^l$ , and refines  $r_1$  to  $r_1^l \dashv\vdash r_0^r$ :

$$\begin{aligned}
 & \Leftarrow\text{-cong}_2 \{r^l\} \{s^l\} \{r^r\} \{s^r\} \ r^l \approx s^l \ r^r \approx s^r \ r' \ (\Rightarrow \rightarrow r \rightarrow \tau^* r_0 \ r_0 \rightarrow \not\rightarrow \tau r_1 \ r_1 \rightarrow \tau^* r') \\
 & \quad | \ r_0^l \wedge r_0^r \wedge \equiv.\text{refl} \wedge r^l \rightarrow \tau^* r_0^l \wedge r^r \rightarrow \tau^* r_0^r \\
 & \quad | \ \text{in}_L (r_1^l \wedge \equiv.\text{refl} \wedge r_0^l \rightarrow \not\rightarrow \tau r_1^l \wedge \llbracket r_0^l \rightarrow \tau r_1^l \rrbracket \equiv \llbracket r_0 \rightarrow \tau r_1 \rrbracket) \\
 & \quad \text{with } \rightarrow \tau^*\text{-split } r_1^l \ r_0^r \ r_1 \rightarrow \tau^* r' \\
 & \Leftarrow\text{-cong}_2 \{r^l\} \{s^l\} \{r^r\} \{s^r\} \ r^l \approx s^l \ r^r \approx s^r \ \_ \ (\Rightarrow \rightarrow r \rightarrow \tau^* r_0 \ r_0 \rightarrow \not\rightarrow \tau r_1 \ r_1 \rightarrow \tau^* r') \\
 & \quad | \ r_0^l \wedge r_0^r \wedge \equiv.\text{refl} \wedge r^l \rightarrow \tau^* r_0^l \wedge r^r \rightarrow \tau^* r_0^r \\
 & \quad | \ \text{in}_L (r_1^l \wedge \equiv.\text{refl} \wedge r_0^l \rightarrow \not\rightarrow \tau r_1^l \wedge \llbracket r_0^l \rightarrow \tau r_1^l \rrbracket \equiv \llbracket r_0 \rightarrow \tau r_1 \rrbracket) \\
 & \quad | \ r^{l'} \wedge r^{r'} \wedge \equiv.\text{refl} \wedge r_1^l \rightarrow \tau^* r^{l'} \wedge r_0^r \rightarrow \tau^* r^{r'} \\
 & \quad \text{with } \approx \rightarrow \Leftarrow \ r^l \approx s^l \ r^{l'} \ (\Rightarrow \rightarrow r^l \rightarrow \tau^* r_0^l \ r_0^l \rightarrow \not\rightarrow \tau r_1^l \ r_1^l \rightarrow \tau^* r^{l'})
 \end{aligned}$$

Partitioning  $r_1 \rightarrow \tau^* r'$  along the two sides of  $r_1$  then gives us the witnesses  $r_1^l \rightarrow \tau^* r^{l'}$  and  $r_0^r \rightarrow \tau^* r^{r'}$ ; the  $r'$  argument is refined to  $r^{l'} \dashv\vdash r^{r'}$ . The transitions  $r^l \rightarrow \tau^* r_0^l$ ,  $r_0^l \rightarrow \not\rightarrow \tau r_1^l$  and  $r_1^l \rightarrow \tau^* r^{l'}$  are just what we need to build a witness of  $r^l \Rightarrow \langle \alpha \rangle r^{l'}$ .



$$\begin{aligned}
\approx' \rightarrow \approx (\approx'\text{-trans } r \approx' s \ s \approx' t) &= \approx\text{-trans } (\approx' \rightarrow \approx r \approx' s) (\approx' \rightarrow \approx s \approx' t) \\
\approx' \rightarrow \approx (\approx'\text{-cong}_2 r^l \approx' s^l \ r^r \approx' s^r) &= \approx\text{-cong}_2 (\approx' \rightarrow \approx r^l \approx' s^l) (\approx' \rightarrow \approx r^r \approx' s^r)
\end{aligned}$$

## 8.8 Compiler Correctness

The compiler correctness property for our Fork language is essentially the same as that of the Zap language, but on singleton thread soups rather than combined machines:

$$\text{correctness} : \forall e \ c \ \sigma \rightarrow \langle e , \langle c , \sigma \rangle \rangle :: [] \approx \langle \langle \text{compile } e \ c , \sigma \rangle \rangle :: []$$

There is no need to generalise over an arbitrary thread soup, since the  $\approx\text{-cong}_2$  lemma of §8.7 allows us to concatenate as many pairs of bisimilar thread soups as is required.

The proof comprises of two parts, each showing one direction of the bisimulation. Proceeding by case analysis on the visible transition, the **fork $\preceq$ FORK** part first shows that **fork**  $e$  cannot make a non-silent transition:

$$\begin{aligned}
\text{correctness } (\text{fork } e) \ c \ \sigma &= \# \text{fork} \preceq \text{FORK} \ \& \ \# \text{FORK} \preceq \text{fork} \ \text{where} \\
\text{fork} \preceq \text{FORK} &: \langle \text{fork } e , \langle c , \sigma \rangle \rangle :: [] \preceq \langle \langle \text{FORK } (\text{compile } e \ []) \rangle :: c , \sigma \rangle :: [] \\
\text{fork} \preceq \text{FORK } s' &(\mathbb{H} \rightarrow \rightarrow ((\_ \wedge ()) \wedge \rightarrow \rightarrow \rightarrow \text{fork}) \triangleleft \_ ) \_ \_ \\
\text{fork} \preceq \text{FORK } s' &(\mathbb{H} \rightarrow \rightarrow ((\_ \wedge \alpha \approx \tau \wedge \rightarrow \rightarrow \text{preempt} ()) \triangleleft \_ ) \_ \_ ) \\
\text{fork} \preceq \text{FORK } s' &(\mathbb{H} \rightarrow \rightarrow \varepsilon (\_ \wedge \alpha \not\approx \tau \wedge \rightarrow \rightarrow \text{preempt} ()) \ s_0 \rightarrow \tau^* s')
\end{aligned}$$

The two  $\rightarrow\text{-preempt}$  clauses correspond to the fact that the empty soup  $[]$  cannot make any transitions at all. In the case of a non-silent  $\rightarrow\text{-fork}$  transition, the expression side transitions to,

$$\langle \# 0 , \langle c , \sigma \rangle \rangle :: \langle e , \langle [], [] \rangle \rangle :: []$$

while the virtual machine follows by the  $\rightarrow\text{-FORK}$  rule:

$$\begin{aligned}
\text{fork} \preceq \text{FORK } s' &(\mathbb{H} \rightarrow \rightarrow \varepsilon (\_ \wedge \alpha \not\approx \tau \wedge \rightarrow \rightarrow \rightarrow \text{fork}) \ s_0 \rightarrow \tau^* s') \\
&= \langle \langle c , 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ [], [] \rangle \rangle :: []
\end{aligned}$$

$$\begin{aligned}
 & \wedge \text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right)
 \end{aligned}$$

The two reducts of the original threads are bisimilar by the **elide- $\tau$**  lemma, while bisimilarity of the two spawned threads is obtained from the induction hypothesis on  $e$ . Finally, we use the  **$\approx\text{-cong}_2$**  lemma to combine these results, to give the overall bisimilarity of the two thread soups.

In the opposite direction of **FORK  $\preceq$  fork**, the proof follows the same steps, with the first clause showing that the **FORK** instruction cannot be silent:

$$\begin{aligned}
 \text{FORK} \preceq \text{fork} & : \langle \langle \text{FORK} (\text{compile } e \ \langle [] \rangle) :: c, \sigma \rangle \rangle :: [] \preceq \langle \text{fork } e, \langle c, \sigma \rangle \rangle :: [] \\
 \text{FORK} \preceq \text{fork } s' & (\text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right) \triangleleft -) \text{ --} \\
 \text{FORK} \preceq \text{fork } s' & (\text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right) \triangleleft -) \text{ --} \\
 \text{FORK} \preceq \text{fork } s' & (\text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right) \triangleleft -) \text{ --}
 \end{aligned}$$

When the virtual machine makes a transition by the  **$\rightarrow\text{FORK}$**  rule, the expression follows with  **$\mapsto\text{fork}$** ,

$$\begin{aligned}
 \text{FORK} \preceq \text{fork } s' & (\text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right) \triangleleft -) \text{ --} \\
 & = \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\
 & \wedge \text{FORK} \varepsilon \left( \begin{array}{l} \text{begin} \\ s' \\ \text{elide-}\tau^* s_0 \rightarrow \tau^* s'^{-1} \\ \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle [], [] \rangle \rangle :: [] \\ \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\} )) (\text{correctness } e \ \langle [] \rangle) \rangle \\ \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \langle [] \rangle, \langle [] \rangle \rangle \rangle :: [] \\ \square \end{array} \right)
 \end{aligned}$$

$$\begin{aligned}
& s' \\
& \approx \langle \text{elide-}\tau^{*'} s_0 \rightarrow \tau^* s' \rangle \\
& \quad \langle \langle c, 0 :: \sigma \rangle \rangle :: \langle \langle \text{compile } e \ \square, \ \square \rangle \rangle :: \square \\
& \approx \langle \approx\text{-cong}_2 (\text{elide-}\tau (\rightarrow\tau\text{-switch } \{\square\})) (\text{correctness } e \ \square \ \square)^{-1} \rangle \\
& \quad \langle \# 0, \langle c, \sigma \rangle \rangle :: \langle e, \langle \square, \square \rangle \rangle :: \square \\
& \square)
\end{aligned}$$

and we obtain the bisimilarity of the two pairs of threads via the  $\approx\text{-cong}_2$  lemma as before. The remaining clauses of `correctness` deal with  $\# m$  and  $a \oplus b$ , following the same steps as that for the Zap language, modulo cosmetic changes to account for the thread soup. This completes the compiler correctness proof for the Fork language.

## 8.9 Conclusion

We have demonstrated that our previously introduced technique of showing bisimilarity between combined machines does indeed scale to the explicitly concurrent Fork language, modelled as a simple ‘thread soup’ of combined machines. The `elide- $\tau$`  lemma was updated for this context using our arsenal of thread soup lemmas, while the result that soup concatenation preserves bisimilarity meant that we could phrase our compiler correctness statement on singleton thread soups. As a result, we were able to reuse most of the correctness proof for the Zap language, with only the `fork e` case requiring further attention.



# Chapter 9

## Transaction Correctness

The previous chapter scaled our proof technique to a language with explicit concurrency. In this chapter, we now consider a language with transactions. In order to reconcile the stop-the-world and log-based semantics, we make two simplifications to our approach. First of all, we replace concurrency with arbitrary interference by an external agent, and secondly we replace the compiler and virtual machine with a direct log-based transactional semantics for the source language.

### 9.1 The Atomic Language

#### 9.1.1 Syntax

For the Atomic language, we migrate to a two-level syntax in a similar manner to the `Tran` and `Proc` types of chapter 5. On the transaction level, we extend our base language of natural numbers and addition with the `read` and `write` keywords for manipulating transactional variables,

```
data Expression' : Set where  
  #_      : (m : ℕ) → Expression'  
  _⊕_    : (a b : Expression') → Expression'
```

$$\text{read} : (v : \text{Variable}) \rightarrow \text{Expression}'$$

$$\text{write} : (v : \text{Variable}) (e : \text{Expression}') \rightarrow \text{Expression}'$$

while the ‘IO’ level is extended with an **atomic** keyword that runs a transaction of type  $\text{Expression}'$ <sup>1</sup>:

**data**  $\text{Expression} : \text{Set}$  **where**

$$\#_ : (m : \mathbb{N}) \rightarrow \text{Expression}$$

$$_ \oplus _ : (a \ b : \text{Expression}) \rightarrow \text{Expression}$$

$$\text{atomic} : (e : \text{Expression}') \rightarrow \text{Expression}$$

Note that we have not included the **fork** construct that spawns additional threads as we did in chapters 5 and 8:

$$\text{fork} : (e : \text{Expression}) \rightarrow \text{Expression}$$

The reason for this is that the presence of **fork** turned out to significantly complicate the formal reasoning process, so we investigated a simpler approach. First we replace concurrency with a ‘mutate’ rule that can change the heap at any time during a transaction, which simulates the worst possible concurrent environment, in a similar manner to the worst-case interrupt rule of [HW07]. Secondly we replace the compiler and virtual machine with a direct log-based transactional semantics for the source language, which makes the proof more manageable.

### 9.1.2 Heaps and Variables

Recall that previously in chapter 5, we modelled the heap as a total map from a fixed set of variable names to their values, initialised to zero. In Agda, we can realise this using the indexed **Vec** type (§6.1.3) from the standard library. As our proof is to be

---

<sup>1</sup>Throughout this chapter, I adopt the convention of using a  $'$  to identify types that are associated with the transaction level of the Atomic language.



independent of the heap size—rather than parametrising the entire proof by it—we simply postulate a number  $|\mathbf{Heap}|$ ,

**postulate**  $|\mathbf{Heap}| : \mathbb{N}$

with the  $\mathbf{Heap}$  type defined as follows:

$\mathbf{Heap} : \mathbf{Set}$

$\mathbf{Heap} = \mathbf{Vec} \ \mathbb{N} \ |\mathbf{Heap}|$

Correspondingly, a  $\mathbf{Variable}$  is just a synonym for the finite set (§6.1.3) with  $|\mathbf{Heap}|$  distinct elements:

$\mathbf{Variable} : \mathbf{Set}$

$\mathbf{Variable} = \mathbf{Fin} \ |\mathbf{Heap}|$

### 9.1.3 Stop-the-World Semantics

Now we have enough machinery to describe the high-level stop-the-world semantics for the Atomic language. Due to the two-level stratification of the language, this involves two separate transitions for  $\mathbf{Expression}'$  and  $\mathbf{Expression}$ . The first of these is  $\_ \mapsto' \_$ , defined on the transaction level between pairs of  $\mathbf{Heaps}$  and  $\mathbf{Expression}'$ s. We begin with the familiar rules for left-biased addition:

**data**  $\_ \mapsto' \_ : \mathbf{Rel} \ (\mathbf{Heap} \times \mathbf{Expression}') \ \mathbf{where}$

$\mapsto' \text{-}\oplus \mathbb{N} : \forall \{h \ m \ n\} \rightarrow$

$h, \# \ m \ \oplus \ \# \ n \mapsto' h, \# \ (m + n)$

$\mapsto' \text{-}\oplus \mathbb{L} : \forall \{h \ h' \ a \ a'\} \ b \rightarrow$

$(a \mapsto a' : h, a \mapsto' h', a') \rightarrow$

$h, a \ \oplus \ b \mapsto' h', a' \ \oplus \ b$

$\mapsto' \text{-}\oplus \mathbb{R} : \forall \{h \ h' \ b \ b'\} \ m \rightarrow$

$$\begin{aligned}
 & (b \mapsto b' : h, \quad b \mapsto' h', \quad b') \rightarrow \\
 & \quad h, \# m \oplus b \mapsto' h', \# m \oplus b' \\
 \\
 & \mapsto'\text{-read} : \forall h v \rightarrow \\
 & \quad h, \text{read } v \mapsto' h, \# h[v] \\
 & \mapsto'\text{-write}_{\mathbb{N}} : \forall \{h v m\} \rightarrow \\
 & \quad h, \text{write } v (\# m) \mapsto' h[v] := m, \# m \\
 & \mapsto'\text{-write}_{\mathbb{E}} : \forall \{h e h' e' v\} \rightarrow \\
 & \quad (e \mapsto e' : h, \quad e \mapsto' h', \quad e') \rightarrow \\
 & \quad h, \text{write } v e \mapsto' h', \text{write } v e'
 \end{aligned}$$

Here the  $\mapsto'\text{-read}$  and  $\mapsto'\text{-write}_{\mathbb{N}}$  rules refer directly to the heap, while  $\mapsto'\text{-write}_{\mathbb{E}}$  effects the reduction of the sub-expression argument to **write**. We write  $\_ \mapsto^* \_$  for the reflexive, transitive closure of  $\_ \mapsto' \_$ , defined using the **Star** type:

$$\begin{aligned}
 \_ \mapsto^* \_ & : \text{Rel} (\text{Heap} \times \text{Expression}') \\
 \_ \mapsto^* \_ & = \text{Star } \_ \mapsto' \_
 \end{aligned}$$

On the ‘IO’ level, transitions are labelled with a choice of actions,

$$\begin{aligned}
 \text{data Action} & : \text{Set where} \\
 \tau \boxplus \clubsuit & : \text{Action}
 \end{aligned}$$

where  $\tau$  is the silent action,  $\boxplus$  corresponds to the addition operation, and  $\clubsuit$  indicates the successful completion of a transaction. These simple observable actions make it possible to define a notion of bisimilarity for the stop-the-world and log-based semantics, where there need not be a one-to-one correspondence of transition rules on each side.

The labelled transition is defined as follows, with the first three rules corresponding to the familiar left-biased addition:

**data**  $\triangleright_{\rightarrow} \rightarrow_{\rightarrow} : \text{Action} \rightarrow \text{Rel} (\text{Heap} \times \text{Expression})$  **where**

$$\rightarrow_{\rightarrow} \oplus \mathbb{N} : \forall \{h \ m \ n\} \rightarrow$$

$$\boxplus \triangleright h, \# m \oplus \# n \mapsto h, \# (m + n)$$

$$\rightarrow_{\rightarrow} \oplus \mathbb{R} : \forall \{\alpha \ h \ h' \ b \ b'\} m \rightarrow$$

$$(b \mapsto b' : \alpha \triangleright h, \quad b \mapsto h', \quad b') \rightarrow$$

$$\alpha \triangleright h, \# m \oplus b \mapsto h', \# m \oplus b'$$

$$\rightarrow_{\rightarrow} \oplus \mathbb{L} : \forall \{\alpha \ h \ h' \ a \ a'\} b \rightarrow$$

$$(a \mapsto a' : \alpha \triangleright h, \ a \quad \mapsto h', \ a') \rightarrow$$

$$\alpha \triangleright h, \ a \oplus b \mapsto h', \ a' \oplus b$$

The  $\rightarrow_{\rightarrow}$ -**atomic** rule implements a stop-the-world semantics for **atomic** blocks by taking a reduction sequence  $e \mapsto^* m$  on the transaction level, and encapsulating it in a single step:

$$\rightarrow_{\rightarrow} \text{atomic} : \forall \{h \ e \ h' \ m\} \rightarrow$$

$$(e \mapsto^* m : h, \quad e \mapsto^* h', \ \# m) \rightarrow$$

$$\blacktriangleright \triangleright h, \ \text{atomic } e \mapsto h', \ \# m$$

$$\rightarrow_{\rightarrow} \text{mutate} : \forall h' \{h \ e\} \rightarrow$$

$$\tau \triangleright h, \ \text{atomic } e \mapsto h', \ \text{atomic } e$$

Since there are no explicit threads in the Atomic language, we introduce a silent  $\rightarrow_{\rightarrow}$ -**mutate** rule to allow the heap to change at any time, which reflects the idea that a concurrent thread may modify the heap while the current thread is running. The above rule implements the worst possible case in which the heap  $h$  can be replaced by a completely different heap  $h'$ . For simplicity, note that  $\rightarrow_{\rightarrow}$ -**mutate** is limited to contexts where the expression is of the form **atomic**  $e$ , as this is the only construct that interacts with the heap. We shall later examine how a corresponding rule in the log-based semantics allows the heap to mutate during a transaction.

### 9.1.4 Transaction Logs and Consistency

Before we give the log-based semantics for *atomic* blocks, let us first define what transaction logs are. Recall from chapter 5 that we modelled them as partial maps from variables to numbers, where an entry for a variable exists only when it has been read from or written to. We take a similar approach to **Heaps**, using vectors of **Maybe**  $\mathbb{N}^2$  initialised to  $\circ$ :

```

data Logs : Set where
  constructor _&_
  field
     $\rho$   $\omega$  : Vec (Maybe  $\mathbb{N}$ ) |Heap|
   $\emptyset$  : Logs
   $\emptyset$  = Vec.replicate  $\circ$  & Vec.replicate  $\circ$ 

```

The  $\rho$  and  $\omega$  fields of a **Logs** record correspond to the read and write logs of chapter 5, and are used in an identical manner to keep track of variables during a running transaction. Let us quickly review the rules for log-based writes and reads in the the context of the current chapter.

Writing to a transaction variable is the most straightforward of the two operations, and is implemented by the following **Write** function that returns a new pair of logs with the entry for  $v$  in  $\omega$  updated to the new value  $m$ .

```

Write : Logs  $\rightarrow$  Variable  $\rightarrow$   $\mathbb{N}$   $\rightarrow$  Logs
Write ( $\rho$  &  $\omega$ )  $v$   $m$  =  $\rho$  &  $\omega$  [  $v$  ] :=  $\bullet$   $m$ 

```

The **Read** function on the other hand takes a heap, a pair of logs and a variable as arguments, and returns a potentially modified read log along with the in-transaction value of the variable:

---

<sup>2</sup>For aesthetic reasons I have renamed **nothing** and **just** of **Maybe** to  $\circ$  and  $\bullet$  respectively.

**Read** : Heap  $\rightarrow$  Logs  $\rightarrow$  Variable  $\rightarrow$  Logs  $\times$   $\mathbb{N}$

**Read**  $h (\rho \ \& \ \omega) v$  **with**  $\omega [ v ]$

... |  $\bullet m = \rho \ \& \ \omega, m$

... |  $\circ$  **with**  $\rho [ v ]$

... |  $\bullet m = \rho \ \& \ \omega, m$

... |  $\circ = \rho [ v ] := \bullet m \ \& \ \omega, m$  **where**  $m = h [ v ]$

If a variable has been written to according to  $\omega$ , we immediately return the new value. Otherwise we consult the read log  $\rho$ : if a previous value for  $v$  exists, we return that. In both cases the transaction logs remain unchanged. Only when no cached value for  $v$  exists—that is, when we are reading a variable for the first time—do we update the read log  $\rho$  with the value of  $v$  from the current heap. Note that if a variable is written to before it is read, the corresponding read log entry will never be filled.

On reaching the end of a transaction we either commit or roll back, depending on whether the values of the variables gleaned from the heap during the transaction are consistent with their corresponding values at the end. That is, all values recorded in the read log must match those currently in the heap for corresponding variables. The following predicate allows us to state this succinctly:

**Consistent** : Heap  $\rightarrow$  Logs  $\rightarrow$  Set

**Consistent**  $h (\rho \ \& \ \_ ) = \forall v m \rightarrow \rho [ v ] \equiv \bullet m \rightarrow h [ v ] \equiv m$

A read log  $\rho$  is consistent with the heap  $h$  precisely when all non-empty entries in  $\rho$  have the same values as the corresponding entries in  $h$ . Note that a transactional variable that is written to before being read from will not have a corresponding entry in  $\rho$ ; this is acceptable since its original value in the heap could not possibly have influenced the behaviour of the transaction. Naturally the empty log  $\emptyset$  is consistent with any heap:

**$\emptyset$ -Consistent** :  $\forall \{h\} \rightarrow$  **Consistent**  $h \ \emptyset$

$\emptyset$ -Consistent  $v$   $m$  **rewrite** `Vec.lookup◊replicate`  $v$  ( $\circ : \text{Maybe } \mathbb{N}$ ) =  $\lambda$  ()

The above uses the `Vec.lookup◊replicate` function to obtain a proof that the entry for  $v$  in the newly-initialised read log is  $\circ$ , in which case we can use an absurd lambda to eliminate the  $\circ \equiv \bullet m$  argument.

The `Dec P` type corresponds to the decidability of some proposition  $P$ . It has two constructors `yes` and `no`, carrying the appropriate evidence in either case:

**data** `Dec (P : Set) : Set where`

`yes : (p : P) → Dec P`

`no : (¬p : ¬ P) → Dec P`

Thus an element of `Dec P` is strictly more informative than a boolean value. Using this, we can give a decision procedure for whether a given heap and read log are indeed consistent. This is implemented below as `Consistent?`, via the `dec` helper that decides consistency for one particular variable:

`Consistent? : (h : Heap) (l : Logs) → Dec (Consistent h l)`

`Consistent? h (ρ & ω) = Dec.map' Vec.Pointwise.app Vec.Pointwise.ext`

`(Vec.Pointwise.decidable dec h ρ) where`

`dec : (hv : ℕ) (ρv : Maybe ℕ) → Dec (∀ m → ρv ≡ • m → hv ≡ m)`

`dec hv ◦ = yes (λ m ())`

`dec hv (• n) with hv  $\stackrel{?}{\equiv}_{\mathbb{N}}$  n`

`... | yes hv ≡ n rewrite hv ≡ n = yes (λ m → •-inj)`

`... | no hv ≠ n = no (λ p → hv ≠ n (p n ≡.refl))`

The library functions `Dec.map'` and `Vec.Pointwise.decidable` are used to generalise the pointwise decision procedure over all variables.

Finally when a transaction is ready to commit, we can use the `Update` function to commit the contents of the write log to the heap:

**Update-lookup** : **Heap**  $\rightarrow$  **Logs**  $\rightarrow$  **Variable**  $\rightarrow$   $\mathbb{N}$

**Update-lookup**  $h (\rho \ \& \ \omega) v = \text{maybe id } (h [ v ]) (\omega [ v ])$

**Update** : **Heap**  $\rightarrow$  **Logs**  $\rightarrow$  **Heap**

**Update**  $h l = \text{Vec.tabulate } (\text{Update-lookup } h l)$

This is implemented using the library function **Vec.tabulate** that takes a function that gives the new value for each index or variable. We have factored out **Update-lookup** in order to leverage existing proofs in the standard library.

### 9.1.5 Log-Based Semantics

The log-base semantics makes transitions between pairs of **Logs** and **Expression'** rather than operating directly on a **Heap**. We can still read from the heap, but it is never modified by the following rules. The first three rules corresponding to left-biased addition should look familiar:

**data**  $\_ \vdash \_ \rightsquigarrow' \_$  ( $h : \text{Heap}$ ) : **Rel** (**Logs**  $\times$  **Expression'**) **where**

$\rightsquigarrow' \text{-}\oplus \mathbb{N} : \forall \{l \ m \ n\} \rightarrow$

$h \vdash l, \# m \oplus \# n \rightsquigarrow' l, \# (m + n)$

$\rightsquigarrow' \text{-}\oplus \mathbb{R} : \forall \{l \ b \ l' \ b'\} m \rightarrow$

$(b \rightsquigarrow b' : h \vdash l, \quad b \rightsquigarrow' l', \quad b') \rightarrow$

$h \vdash l, \# m \oplus b \rightsquigarrow' l', \# m \oplus b'$

$\rightsquigarrow' \text{-}\oplus \mathbb{L} : \forall \{l \ a \ l' \ a'\} b \rightarrow$

$(a \rightsquigarrow a' : h \vdash l, \ a \rightsquigarrow' l', \ a') \rightarrow$

$h \vdash l, \ a \oplus b \rightsquigarrow' l', \ a' \oplus b$

The  $\rightsquigarrow' \text{-read}$  rule reduces a **read**  $v$  expression to the value of  $v$  using the **Read** function defined in the previous section, potentially also resulting in a new log:

$\rightsquigarrow' \text{-read} : \forall l \ v \rightarrow \text{let } l'm = \text{Read } h \ l \ v \text{ in}$

$h \vdash l, \ \text{read } v \rightsquigarrow' \text{fst } l'm, \ \# \text{snd } l'm$

$$\begin{aligned}
 \multimap\text{-write}_{\mathbb{N}} & : \forall \{l v m\} \rightarrow \\
 & h \vdash l, \text{write } v (\# m) \multimap\text{' Write } l v m, \# m \\
 \multimap\text{-write}_{\mathbb{E}} & : \forall \{l e l' e' v\} \rightarrow \\
 & (e \multimap e' : h \vdash l, \quad e \multimap\text{' } l', \quad e') \rightarrow \\
 & h \vdash l, \text{write } v e \multimap\text{' } l', \text{write } v e'
 \end{aligned}$$

The  $\multimap\text{-write}_{\mathbb{N}}$  rule updates the write log via the **Write** helper when the expression argument to **write** is just a number, while  $\multimap\text{-write}_{\mathbb{E}}$  effects the reduction of  $e$  in the same manner as the stop-the-world semantics.

We write  $\perp\text{-}\multimap\text{'*}$  for the reflexive, transitive closure of  $\perp\text{-}\multimap\text{'}$  under the same heap, again defined using the **Star** type:

$$\begin{aligned}
 \perp\text{-}\multimap\text{'*} & : \text{Heap} \rightarrow \text{Rel} (\text{Logs} \times \text{Expression}') \\
 \perp\text{-}\multimap\text{'*} h & = \text{Star} (\perp\text{-}\multimap\text{' } h)
 \end{aligned}$$

For the ‘IO’ level of this log-based semantics, we define a transition  $\triangleright\text{-}\multimap\text{-}$  between triples of heaps, transaction states and expressions, labelled with the same **Actions** we used earlier. During a running transaction, the state comprises of the original expression and the transaction **Logs**; otherwise it is empty:

$$\begin{aligned}
 \text{TState} & : \text{Set} \\
 \text{TState} & = \text{Maybe} (\text{Expression}' \times \text{Logs})
 \end{aligned}$$

The rules for addition are identical to those of  $\triangleright\text{-}\multimap\text{-}$ :

$$\begin{aligned}
 \text{data } \triangleright\text{-}\multimap\text{-} & : \text{Action} \rightarrow \text{Rel} (\text{Heap} \times \text{TState} \times \text{Expression}) \text{ where} \\
 \multimap\text{-}\oplus\mathbb{N} & : \forall \{h m n\} \rightarrow \\
 & \boxplus \triangleright h, \circ, \# m \oplus \# n \multimap h, \circ, \# (m + n) \\
 \multimap\text{-}\oplus\mathbb{R} & : \forall \{\alpha h t b h' t' b'\} m \rightarrow \\
 & (b \multimap b' : \alpha \triangleright h, t, \quad b \multimap h', t', \quad b') \rightarrow \\
 & \alpha \triangleright h, t, \# m \oplus b \multimap h', t', \# m \oplus b'
 \end{aligned}$$



$$\begin{aligned}
\multimap\oplus\mathbf{L} & : \forall \{\alpha \ h \ t \ a \ h' \ t' \ a'\} \ b \ \rightarrow \\
& (a \multimap a' : \alpha \triangleright h, t, a \quad \multimap h', t', a') \rightarrow \\
& \alpha \triangleright h, t, a \oplus b \multimap h', t', a' \oplus b
\end{aligned}$$

Next we move on to the transaction rules: when the expression to be reduced is of the form **atomic**  $e$  and we have yet to enter the transaction, the  $\multimap$ -**begin** rule sets up the restart expression to  $e$  and initialises the transaction logs to  $\emptyset$ :

$$\begin{aligned}
\multimap\text{-begin} & : \forall \{h \ e\} \ \rightarrow \\
& \tau \triangleright h, \circ, \text{atomic } e \multimap h, \bullet (e, \emptyset), \text{atomic } e
\end{aligned}$$

The second  $\multimap$ -**step** rule allows us to make a single  $\perp\text{-}\multimap'$  transition on the transaction level; note that the heap  $h$  does not change:

$$\begin{aligned}
\multimap\text{-step} & : \forall \{h \ r \ l \ e \ l' \ e'\} \ \rightarrow \\
& (e \multimap e' : h \perp l, e \multimap' l', e') \rightarrow \\
& \tau \triangleright h, \bullet (r, l), \text{atomic } e \multimap h, \bullet (r, l'), \text{atomic } e'
\end{aligned}$$

While the Atomic language does not contain explicit parallelism, we can model interference using a  $\multimap$ -**mutate** rule that changes to an arbitrary heap  $h'$  at any time during a transaction:

$$\begin{aligned}
\multimap\text{-mutate} & : \forall h' \{h \ t \ e\} \ \rightarrow \\
& \tau \triangleright h, \bullet t, \text{atomic } e \multimap h', \bullet t, \text{atomic } e
\end{aligned}$$

Finally we come to the  $\multimap$ -**abort** and  $\multimap$ -**commit** rules, one of which applies when the transactional expression has reduced down to a number:

$$\begin{aligned}
\multimap\text{-abort} & : \forall \{h \ r \ l \ m\} (\neg\text{cons} : \neg \text{Consistent } h \ l) \ \rightarrow \\
& \tau \triangleright h, \bullet (r, l), \text{atomic } (\# m) \multimap h, \bullet (r, \emptyset), \text{atomic } r \\
\multimap\text{-commit} & : \forall \{h \ r \ l \ m\} (\text{cons} : \text{Consistent } h \ l) \ \rightarrow \\
& \clubsuit \triangleright h, \bullet (r, l), \text{atomic } (\# m) \multimap \text{Update } h \ l, \circ, \# m
\end{aligned}$$

Both rules carry proof of the consistency or otherwise of the log  $l$  with respect to  $h$ . While this is not technically necessary and we could make do with a single rule—as consistency is decidable—having two rules labelled with distinct *Actions* makes later proofs easier to work with.

In any case if we do have consistency, we commit the transaction by applying the write log to the heap using the **Update** function, setting the transaction state to  $\circ$ , and reducing **atomic** ( $\# m$ ) to  $\# m$ . Otherwise the  $\rightarrow\text{-abort}$  rule applies, and we silently restart the transaction by resetting the transaction state and the expression.

## 9.2 Combined Semantics and Bisimilarity

### 9.2.1 Combined Semantics

In a similar spirit to the combined machines of previous chapters, let us define a **Combined** type that allows us to select which of our two semantics to use to interpret some given expression:

```
data Combined : Set where
   $\mapsto$  : Combined
   $\rightarrow\text{-}_-$  : (t : TState)  $\rightarrow$  Combined
```

The  $\mapsto$  constructor indicates that the associated expression should be interpreted according to the stop-the-world semantics, while  $\rightarrow\text{-}_-$  implies the log-based semantics. The latter is also used to carry around the transaction state. How this works can be seen in the definition of  $\triangleleft\text{-}\rightarrow\text{-}$  below:

```
data  $\triangleleft\text{-}\rightarrow\text{-}$  ( $\alpha$  : Action) : Rel (Heap  $\times$  Combined  $\times$  Expression) where
   $\rightarrow\text{-}\mapsto$  :  $\forall \{h\ e\ h'\ e'\} \rightarrow$ 
    ( $e\mapsto e'$  :  $\alpha \triangleright h$ ,  $e \mapsto h'$ ,  $e'$ )  $\rightarrow$ 
       $\alpha \triangleright h$ ,  $\mapsto$ .,  $e \rightarrow h'$ ,  $\mapsto$ .,  $e'$ 
```

$$\begin{aligned}
 \multimap \multimap & : \forall \{h \ t \ e \ h' \ t' \ e'\} \rightarrow \\
 (e \multimap e' : \alpha \triangleright h, \quad t, e \multimap h', \quad t', e') & \rightarrow \\
 \alpha \triangleright h, \multimap : t, e \multimap h', \multimap : t', e' &
 \end{aligned}$$

This combined transition is essentially a disjoint union of  $\multimap \multimap$  and  $\multimap \multimap$ . We will write  $\multimap_{\tau}^*$  for an arbitrary sequence of silent  $\tau$  transitions:

$$\begin{aligned}
 \multimap_{\tau}^* & : \text{Rel} (\text{Heap} \times \text{Combined} \times \text{Expression}) \\
 \multimap_{\tau}^* & = \text{Star} (\multimap \multimap \tau)
 \end{aligned}$$

Finally, let us define a visible transition as a sequence of silent  $\tau$  transitions followed by a single non- $\tau$  transition:

```

data  $\multimap \multimap$  ( $\alpha : \text{Action}$ ) ( $x \ x'' : \text{Heap} \times \text{Combined} \times \text{Expression}$ ) : Set where
constructor  $\multimap$ :
field
    { $h'$ } : Heap
    { $c'$ } : Combined
    { $e'$ } : Expression
     $\alpha \neq \tau$  :  $\alpha \neq \tau$ 
     $e \multimap_{\tau}^* e'$  :  $x \multimap_{\tau}^* h', c', e'$ 
     $e' \multimap e''$  :  $\alpha \triangleright h', c', e' \rightarrow x''$ 
    
```

The visible transition above is basically the same idea as that of the Fork and Zap languages of the previous two chapters, but without a second sequence of  $\tau$ -transitions after the main  $e' \multimap e''$  transition.

## 9.2.2 Bisimilarity of Semantics

The definition of bisimilarity differs in some details compared with the previous two chapters. Since we are just comparing two different semantics for the same expression, we define one half of bisimilarity as follows:

$$\begin{aligned}
 & \_ \vdash \_ \rightsquigarrow \_ : \text{Heap} \times \text{Expression} \rightarrow \text{Rel Combined} \\
 & h, e \vdash x \rightsquigarrow y = \forall \{h' x' e' \alpha\} \rightarrow \\
 & \quad (x \mapsto x' : \alpha \triangleright h, x, e \mapsto h', x', e') \rightarrow \\
 & \quad \exists \lambda y' \rightarrow \alpha \triangleright h, y, e \mapsto h', y', e' \times (h', e' \vdash x' \approx y')
 \end{aligned}$$

That is given a heap  $h$  and expression  $e$ , whenever it can make a visible transition to  $h'$  and  $e'$  under the semantics represented by  $x$ , an equivalent transition to the same  $h'$  and  $e'$  exists under  $y$ , such that  $x'$  and  $y'$  are also bisimilar for  $h'$  and  $e'$ .

Bisimilarity is then defined as a pair of coinductive  $\_ \vdash \_ \rightsquigarrow \_$  relations:

```

data  $\_ \vdash \_ \approx \_$  (he : Heap × Expression) (x y : Combined) : Set where
  constructor  $\_ \& \_$ 
  field
     $\approx \rightarrow \rightsquigarrow$  :  $\infty$  (he  $\vdash$  x  $\rightsquigarrow$  y)
     $\approx \rightarrow \rightsquigarrow$  :  $\infty$  (he  $\vdash$  y  $\rightsquigarrow$  x)
    
```

Utilising the same proofs as before, we can show that  $\_ \vdash \_ \approx \_$  is reflexive, symmetric, and transitive on **Combined**, given a heap and an expression:

```

 $\rightsquigarrow$ -refl : {he : Heap × Expression} → Reflexive ( $\_ \vdash \_ \rightsquigarrow \_$  he)
 $\rightsquigarrow$ -refl  $x \mapsto x' = \_ , x \mapsto x' , \rightsquigarrow$ -refl
 $\approx$ -refl : {he : Heap × Expression} → Reflexive ( $\_ \vdash \_ \approx \_$  he)
 $\approx$ -refl =  $\# \rightsquigarrow$ -refl &  $\# \rightsquigarrow$ -refl
 $\approx$ -sym : {he : Heap × Expression} → Symmetric ( $\_ \vdash \_ \approx \_$  he)
 $\approx$ -sym (x  $\rightsquigarrow$  y & y  $\rightsquigarrow$  x) = y  $\rightsquigarrow$  x & x  $\rightsquigarrow$  y
 $\rightsquigarrow$ -trans : {he : Heap × Expression} → Transitive ( $\_ \vdash \_ \rightsquigarrow \_$  he)
 $\rightsquigarrow$ -trans x  $\rightsquigarrow$  y y  $\rightsquigarrow$  z  $x \mapsto x'$  with x  $\rightsquigarrow$  y  $x \mapsto x'$ 
... | y', y  $\mapsto$  y', x'  $\approx$  y' with y  $\rightsquigarrow$  z y  $\mapsto$  y'
... | z', z  $\mapsto$  z', y'  $\approx$  z' = z', z  $\mapsto$  z',  $\approx$ -trans x'  $\approx$  y' y'  $\approx$  z'
    
```

$$\begin{aligned}
& \approx\text{-trans} : \{he : \text{Heap} \times \text{Expression}\} \rightarrow \text{Transitive} (\_ \vdash \_ \approx \_ he) \\
& \approx\text{-trans} (x \preceq y \ \& \ y \preceq x) (y \preceq z \ \& \ z \preceq y) = \\
& \quad \# \preceq\text{-trans} (\triangleright x \preceq y) (\triangleright y \preceq z) \ \& \ \# \preceq\text{-trans} (\triangleright z \preceq y) (\triangleright y \preceq x)
\end{aligned}$$

In this chapter, we managed to avoid the use of an embedded language (the  $\_ \approx \_$  relation of previous chapters) to convince Agda that our coinductive definitions are properly constructor-guarded, and hence productive. That said, we will need to manually inline several uses of  $\approx\text{-sym}$  in later proofs, as the guardedness checker cannot see through function definitions yet at the time of writing.

### 9.2.3 Definition of Correctness

Having accumulated enough machinery, we can now give a definition of correctness of the log-based semantics, which is simply the following:

$$\text{correct} : \forall h e \rightarrow h, e \vdash \mapsto : \approx \triangleright : \circ$$

That is for any heap  $h$  and expression  $e$ , the stop-the-world semantics (as represented by  $\mapsto$ ) and the log-based semantics with an empty transaction state (proxied by  $\triangleright : \circ$ ) are bisimilar up to visible transitions.

## 9.3 Reasoning Transactionally

In this section, we will cover some useful lemmas concerning heaps and transaction logs that are used to show that the stop-the-world and log-based transaction semantics coincide.

### 9.3.1 Consistency-Preserving Transitions

First of all, recall that when the log-based semantics needs to read a variable  $v$  and it is not present in either of the read and write logs, we update the read log with the

value of  $v$  from the heap. The following lemma shows that this operation preserves log consistency:

$$\begin{aligned}
 \text{Read-Consistent} & : \forall \{h\} l v \rightarrow \text{Consistent } h l \rightarrow \\
 & \text{Consistent } h (\text{Logs.}\rho l [v] := \bullet (h [v]) \ \& \ \text{Logs.}\omega l) \\
 \text{Read-Consistent } \{h\} (\rho \ \& \ \omega) v \text{ cons } v' m & \text{ with } v' \stackrel{?}{=}_{\text{Fin}} v \\
 \dots \mid \text{yes } v' \equiv v & \text{ rewrite } v' \equiv v \mid \text{Vec.lookup}\circ\text{update } v \rho (\bullet (h [v])) = \bullet\text{-inj} \\
 \dots \mid \text{no } v' \not\equiv v & \text{ rewrite } \text{Vec.lookup}\circ\text{update}' v' \not\equiv v \rho (\bullet (h [v])) = \text{cons } v' m
 \end{aligned}$$

We have  $\eta$ -expanded **Read-Consistent** with a second variable  $v'$  and  $m$  taken by the resulting **Consistent** type, and need to show that  $\rho [v'] \equiv \bullet m \rightarrow h [v'] \equiv m$ .

There are two cases to consider, depending on whether  $v'$  coincides with the variable  $v$  whose read log entry is being updated. If they are indeed the same, we can use **Vec.lookup** $\circ$ **update** to show that the updated read log entry is  $\bullet h_v$ , in which case the goal of  $\bullet h [v'] \equiv \bullet m \rightarrow h [v'] \equiv m$  can be satisfied by injectivity of  $\bullet$ . When  $v$  and  $v'$  correspond to different variables, **Vec.lookup** $\circ$ **update**' gives us a proof that the read log entry for  $v'$  remains unchanged, and the *cons* argument suffices.

Using the above result, we can demonstrate that any transaction transition under the log-based semantics preserves consistency:

$$\begin{aligned}
 \rightsquigarrow'\text{-Consistent} & : \forall \{h l e l' e'\} \rightarrow \text{Consistent } h l \rightarrow \\
 & h \vdash l, e \rightsquigarrow' l', e' \rightarrow \text{Consistent } h l' \\
 \rightsquigarrow'\text{-Consistent } \text{cons } \rightsquigarrow'\text{-}\oplus\mathbb{N} & = \text{cons} \\
 \rightsquigarrow'\text{-Consistent } \text{cons } (\rightsquigarrow'\text{-}\oplus\mathbb{R} m b \rightsquigarrow b') & = \rightsquigarrow'\text{-Consistent } \text{cons } b \rightsquigarrow b' \\
 \rightsquigarrow'\text{-Consistent } \text{cons } (\rightsquigarrow'\text{-}\oplus\mathbb{L} b a \rightsquigarrow a') & = \rightsquigarrow'\text{-Consistent } \text{cons } a \rightsquigarrow a' \\
 \rightsquigarrow'\text{-Consistent } \text{cons } (\rightsquigarrow'\text{-write}_{\mathbb{E}} e \rightsquigarrow e') & = \rightsquigarrow'\text{-Consistent } \text{cons } e \rightsquigarrow e' \\
 \rightsquigarrow'\text{-Consistent } \text{cons } \rightsquigarrow'\text{-write}_{\mathbb{N}} & = \text{cons} \\
 \rightsquigarrow'\text{-Consistent } \text{cons } (\rightsquigarrow'\text{-read } l v) & \text{ with } \text{Logs.}\omega l [v] \\
 \dots \mid \bullet m = \text{cons} &
 \end{aligned}$$

$\dots \mid \circ$  **with**  $\text{Logs}.\rho\ l\ [v]$   
 $\dots \mid \bullet$   $m = \text{cons}$   
 $\dots \mid \circ =$  **Read-Consistent**  $l\ v\ \text{cons}$

The proof proceeds by induction on the structure of the reduction rules, making use of the **Read-Consistent** lemma when the read log changes. Naturally, we can extend the above to an arbitrary  $\_ \vdash \_ \rightsquigarrow^* \_$  sequence by folding over it:

$\rightsquigarrow^*$ -**Consistent** :  $\forall \{h\ l\ e\ l'\ e'\} \rightarrow$  **Consistent**  $h\ l \rightarrow$   
 $h \vdash l, e \rightsquigarrow^* l', e' \rightarrow$  **Consistent**  $h\ l'$   
 $\rightsquigarrow^*$ -**Consistent**  $\{h\}\ \{l\}\ \{e\} =$   
 $\text{Star.gfoldl fst (const (Consistent } h)) \rightsquigarrow'$ -**Consistent**  $\{i = l, e\}$

In the opposite direction, we can show a pair of similar but slightly more general consistency-preservation lemmas. This extra generality in fact turns out to be crucial to our later proofs. The **Read-Consistent'** lemma shares an analogous structure to that of **Read-Consistent**, but requires an extra argument showing that the pre-transition read log entry for  $v$  is empty:

**Read-Consistent'** :  $\forall \{h\ n\} l\ v \rightarrow$   $\text{Logs}.\rho\ l\ [v] \equiv \circ \rightarrow$   
 $\text{Consistent } h\ (\text{Logs}.\rho\ l\ [v] := \bullet\ n \ \& \ \text{Logs}.\omega\ l) \rightarrow$  **Consistent**  $h\ l$   
**Read-Consistent'**  $\{h\}\ \{n\}\ (\rho \ \& \ \omega)\ v\ \rho_v \equiv \circ\ \text{cons}'\ v'\ m$  **with**  $v' \stackrel{?}{=}_{\text{Fin}} v$   
 $\dots \mid$  **yes**  $v' \equiv v$  **rewrite**  $v' \equiv v \mid \rho_v \equiv \circ = \lambda\ ()$   
 $\dots \mid$  **no**  $v' \not\equiv v$  **rewrite**  $\equiv \text{sym (Vec.lookupupdate}'\ v' \not\equiv v\ \rho\ (\bullet\ n)) = \text{cons}'\ v'\ m$

As before, there are two alternatives: when  $v'$  coincides with the variable  $v$  whose read log entry is being updated, we use the  $\rho_v \equiv \circ$  argument to rewrite the goal to  $\circ \equiv \bullet\ m \rightarrow h\ [v] \equiv m$ , which is then discharged with an absurd  $\lambda$ . This is essentially making use of the fact that each read log entry is only ever updated once, from  $\circ$  to  $\bullet$ . When  $v'$  differs, the  $\text{cons}'$  argument suffices.

In the **yes** case of **Read-Consistent**, we required that the post-transition read log entry for  $v$  be  $\bullet (h \llbracket v \rrbracket)$ . Since the corresponding case here is absurd, this is no longer necessary, and the proof can be generalised to any  $\bullet n$ . This means that the heap  $h$  under which the logs and expression make their transition need not be the same as the heap  $h'$  with which  $l$  and  $l'$  are consistent in the following lemma:

$$\begin{aligned}
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} & : \forall \{h \ h' \ l \ e \ l' \ e'\} \rightarrow \text{Consistent } h' \ l' \rightarrow \\
 & h \vdash l, e \text{\color{blue}\mathrel{\twoheadrightarrow}'} l', e' \rightarrow \text{Consistent } h' \ l \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\oplus\mathbb{N}} & = \text{cons}' \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\oplus\mathbb{R}} \ m \ b \rightarrow b' & = \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \ b \rightarrow b' \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\oplus\mathbb{L}} \ b \ a \rightarrow a' & = \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \ a \rightarrow a' \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\text{write}_{\mathbb{E}}} \ e \rightarrow e' & = \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \ e \rightarrow e' \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\text{write}_{\mathbb{N}}} & = \text{cons}' \\
 \text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent'} \text{ } \text{cons}' \text{\color{green}\mathrel{\twoheadrightarrow}'-\text{read}} \ (\rho \ \& \ \omega) \ v \ \text{with } \omega \llbracket v \rrbracket & \\
 \dots \mid \bullet \ m = \text{cons}' & \\
 \dots \mid \circ \ \text{with } \rho \llbracket v \rrbracket \mid \equiv.\text{inspect} \ (\_ \llbracket \_ \rrbracket) \ \rho \ v & \\
 \dots \mid \bullet \ m \mid \_ = \text{cons}' & \\
 \dots \mid \circ \ \mid \llbracket \rho_v \equiv \circ \rrbracket = \text{Read-Consistent}' \ (\rho \ \& \ \omega) \ v \ \rho_v \equiv \circ \ \text{cons}' &
 \end{aligned}$$

This follows an identical structure to  $\text{\color{red}\mathrel{\twoheadrightarrow}'-Consistent}$ , with the only difference being the use of the  $\equiv.\text{inspect}$  idiom to obtain a proof of  $\rho \llbracket v \rrbracket \equiv \circ$ .

### 9.3.2 Heaps and Logs Equivalence

Recall that a pair of read and write logs is used to give an local view of the heap during a running transaction. For our correctness proof, it will be convenient to define a predicate stating that the view of the heap during the transaction—that is,  $h$  overlaid with read and write logs—is equivalent to another heap  $h'$  that is accessed directly using the stop-the-world semantics:



**Equivalent** :  $\text{Heap} \rightarrow \text{Logs} \rightarrow \text{Heap} \rightarrow \text{Set}$

**Equivalent**  $h \ l \ h' = \text{snd} \circ \text{Read} \ h \ l \doteq \text{\_}[\_] \ h'$

We write  $f \doteq g$  to mean pointwise equality of  $f$  and  $g$ , and is a synonym for  $\forall x \rightarrow f \ x \equiv g \ x$ . In other words,  $\text{Read} \ h \ l \ v$  gives the same value as  $h' \ [ \ v \ ]$  for all variables.

On commencing a transaction, the logs are initialised to  $\emptyset$  by the  $\rightarrow\text{-begin}$  rule, while the heaps according to both semantics have yet to diverge. The following definition shows that every heap  $h$  is equivalent to itself overlaid with empty logs:

**$\emptyset$ -Equivalent** :  $\forall \{h\} \rightarrow \text{Equivalent} \ h \ \emptyset \ h$

**$\emptyset$ -Equivalent**  $v \ \text{rewrite} \ \text{Vec.lookup}\circ\text{replicate} \ v \ (\circ : \text{Maybe } \mathbb{N})$

|  $\text{Vec.lookup}\circ\text{replicate} \ v \ (\circ : \text{Maybe } \mathbb{N}) = \equiv \text{.refl}$

The two rewrites correspond to showing that the write and read logs are always empty, using the  $\text{Vec.lookup}\circ\text{replicate}$  lemma to obtain proofs of  $\text{Vec.replicate} \ \circ \ [ \ v \ ] \equiv \circ$ , so that the value returned by  $\text{Read}$  reduces to just  $h \ [ \ v \ ]$ . The goal is then trivially satisfied by reflexivity.

In a similar manner to **Read-Consistent**, the operation of updating the read log for a variable  $v$  when it is first read preserves heap-log equivalence.

**Read-Equivalent** :  $\forall \{h \ l \ h' \ v\} \rightarrow \text{Logs}.\rho \ l \ [ \ v \ ] \equiv \circ \rightarrow$

**Equivalent**  $h \ l \ h' \rightarrow \text{Equivalent} \ h \ (\text{Logs}.\rho \ l \ [ \ v \ ] := \bullet (h \ [ \ v \ ]) \ \& \ \text{Logs}.\omega \ l) \ h'$

**Read-Equivalent**  $\{h\} \ \{\rho \ \& \ \omega\} \ \{h'\} \ \{v\} \ \rho_v \equiv \circ \ \text{equiv} \ v' \ \text{with} \ \text{equiv} \ v'$

... |  $\text{equiv-}v' \ \text{with} \ \omega \ [ \ v' \ ]$

... |  $\bullet \ m = \text{equiv-}v'$

... |  $\circ \ \text{with} \ v' \stackrel{?}{=}_{\text{Fin}} \ v$

We start by binding the application  $\text{equiv} \ v'$  to  $\text{equiv-}v'$ , which starts off with a type of  $\text{snd} \ (\text{Read} \ h \ l \ v') \equiv h' \ [ \ v' \ ]$ . This is so that the  $\text{Read}$  function in its type can

be refined as we perform case analyses on the write and read log entries for  $v'$ . Since the write log does not change, the types of both the goal and  $equiv\text{-}v'$  reduces to  $m \equiv h' [ v' ]$  when  $\omega [ v' ]$  is  $\bullet m$ . Otherwise we must consider whether  $v'$  refers to the same variable as  $v$  whose read log entry is being updated:

$$\begin{aligned} \dots & \quad | \text{yes } v' \equiv v \text{ rewrite } v' \equiv v \quad | \rho_v \equiv \circ \\ & \quad | \text{Vec.lookupupdate } v \rho (\bullet (h [ v ])) = equiv\text{-}v' \end{aligned}$$

If  $v'$  is indeed the variable being updated, we can use the  $\rho_v \equiv \circ$  argument to refine the type of  $equiv\text{-}v'$  to  $h [ v ] \equiv h' [ v ]$ , and a final **Vec.lookupupdate** rewrites the goal to the same type. Otherwise, we use the **Vec.lookupupdate'** lemma to show that  $\rho [ v' ]$  is unaffected by the update:

$$\begin{aligned} \dots & \quad | \text{no } v' \not\equiv v \text{ rewrite } \text{Vec.lookupupdate}' v' \not\equiv v \rho (\bullet (h [ v ])) \text{ with } \rho [ v' ] \\ \dots & \quad | \bullet m = equiv\text{-}v' \\ \dots & \quad | \circ = equiv\text{-}v' \end{aligned}$$

In the two alternatives above, the types of the goals and  $equiv\text{-}v'$  reduce to  $m \equiv h' [ v' ]$  and  $h [ v' ] \equiv h' [ v' ]$ , corresponding to the cases where  $v'$  was already cached in the read log, and when it is read for the first time respectively.

Unlike the **Consistent** property which only involves the read log, **Equivalent** also depends on the write log (indirectly via **Read**). Therefore we must demonstrate that write log updates preserve some notion of heap-log equivalence. We proceed by applying  $equiv$  to  $v'$ , and checking whether  $v'$  and  $v$  are the same variable:

$$\begin{aligned} \text{Write-Equivalent} & : \forall \{h \ l \ h' \ v \ m\} \rightarrow \\ & \quad \text{Equivalent } h \ l \ h' \rightarrow \text{Equivalent } h \ (\text{Write } l \ v \ m) (h' [ v ] := m) \\ \text{Write-Equivalent } \{h\} \{ \rho \ \& \ \omega \} \{h'\} \{v\} \{m\} & equiv \ v' \text{ with } equiv \ v' \quad | \ v' \stackrel{?}{=}_{\text{Fin}} v \\ \dots & \quad | equiv\text{-}v' \quad | \text{yes } v' \equiv v \text{ rewrite } v' \equiv v \quad | \text{Vec.lookupupdate } v \ \omega (\bullet m) \\ & \quad | \text{Vec.lookupupdate } v \ h' \ m = \equiv.\text{refl} \end{aligned}$$

In the **yes** case, we use **Vec.lookupupdate** to first show that the value returned by **Read**  $h$  (**Write**  $l$   $v$   $m$ )  $v$  is in fact  $m$ , which corresponds to the left-hand side of the  $\_ \equiv \_$  goal. The next clause rewrites the right-hand side from  $(h' [ v ] := m) [ v ]$  to the same  $m$ , and  $\equiv$ .**refl** completes this half of the proof.

For the **no** half where  $v'$  is not the variable being written to, the write log entry  $\omega [ v' ]$  and the value of  $h' [ v' ]$  are not updated, which is taken care of by the two **Vec.lookupupdate'** rewrites. Thus the existing  $equiv\text{-}v'$  suffices to complete the proof, although we do have to inspect the appropriate log entries to verify that  $equiv\text{-}v'$  and the goal have the correct types in all cases:

```

... | equiv-v' | no v'≠v rewrite Vec.lookupupdate' v'≠v ω (• m)
      | Vec.lookupupdate' v'≠v h' m with ω [ v' ]
...   | • n = equiv-v'
...   | ◦ with ρ [ v' ]
...   | • n = equiv-v'
...   | ◦ = equiv-v'
    
```

### 9.3.3 Post-Commit Heap Equality

When a transaction completes successfully, we proceed to update the unmodified heap with the contents of the write log, using the **Update** function defined at the end of §9.1.4. Given an  $h'$  that is equivalent to some heap  $h$  overlaid with logs  $l$  and that  $h$  and  $l$  are mutually consistent, we can proceed to show that updating  $h$  with the contents of the write log results in a heap identical to one that is modified in-place by the stop-the-world semantics:

```

Commit : ∀ {h l h'} →
  Consistent h l → Equivalent h l h' → Update h l ≡ h'
Commit {h} {l} {h'} cons equiv =
    
```

```

Equivalence.to Vec.Pointwise-≡ ⟨$⟩ Vec.Pointwise.ext hω≐h' where
hω≐h' : ∀ v → Update h l [ v ] ≡ h' [ v ]
hω≐h' v rewrite Vec.lookup◦tabulate (Update-lookup h l) v
  with Logs.ω l [ v ] | equiv v
... | • m | equiv-v = equiv-v

```

The main  $h\omega\equiv h'$  part of the proof shows pointwise equality of `Update h l` and  $h'$ , by considering the entry for  $v$  in the write and read logs. When the write log contains  $\bullet m$ , the corresponding entry of  $h$  would be updated with  $m$ ; not coincidentally  $equiv-v$  has been refined to a proof of  $m \equiv h' [ v ]$ . Otherwise the write log contains a  $\circ$ , and the goal type reduces to  $h [ v ] \equiv h' [ v ]$ :

```

... | ◦ | equiv-v with Logs.ρ l [ v ] | ≡.inspect ([-_] (Logs.ρ l)) v
... | • m | [ [ ρ_v ≡ m ] ] = ≡.trans (cons v m ρ_v ≡ m) equiv-v
... | ◦ | _ = equiv-v

```

We then proceed to inspect the read log: if it contains  $\bullet m$  then  $equiv-v$  refines to a proof of  $m \equiv h' [ v ]$ , so we use `cons` to show that  $h [ v ]$  is also equal to  $m$ , and transitivity completes the proof. In the last case where both log entries are empty, the `Read` on the left-hand side of the type of  $equiv-v$  becomes simply  $h [ v ]$ , and so completes the proof. Finally we use the proof of pointwise/definitional equivalence for `Vec` from the Agda standard library to convert  $h\omega\equiv h'$  to a proof of definitional equality.

## 9.4 Transaction Correctness

During a transaction, the high-level stop-the-world semantics manipulates the heap directly, while the log-based semantics accumulates its reads and writes in a transaction log, eventually committing it to the heap. In this section we show that for any

transaction sequence under one semantics, there exists a matching sequence under the other semantics.

### 9.4.1 Completeness of Log-Based Transactions

We shall tackle the completeness part of transactional correctness first, as it is the simpler of the two directions. Let us begin by defining a function that extracts the  $\dashv\vdash^*$  sequence from a visible transition starting from **atomic**  $e$ :

$$\begin{aligned}
 \text{\color{red}\dashv\vdash-extract} & : \forall \{ \alpha \triangleright h \ e \ h'' \ c'' \ e'' \} \rightarrow \\
 & \alpha \triangleright h, \text{\color{blue}\dashv\vdash} : , \text{\color{green}atomic} \ e \ \text{\color{blue}\dashv\vdash} h'' , c'' , e'' \rightarrow \\
 & \exists_2 \lambda h_0 \ m \rightarrow \alpha , c'' , e'' \equiv \text{\color{green}\clubsuit} , \text{\color{blue}\dashv\vdash} : , \# \ m \times \\
 & h_0 , e \ \text{\color{blue}\dashv\vdash}^* h'' , \# \ m \\
 \text{\color{red}\dashv\vdash-extract} & (\text{\color{blue}\dashv\vdash} : \alpha \not\equiv \tau \ \varepsilon \ (\rightarrow \dashv\vdash (\text{\color{green}\dashv\vdash-mutate} \ h_1))) = \text{\color{red}\perp-elim} (\alpha \not\equiv \tau \ \equiv .\text{refl}) \\
 \text{\color{red}\dashv\vdash-extract} & (\text{\color{blue}\dashv\vdash} : \alpha \not\equiv \tau \ \varepsilon \ (\rightarrow \dashv\vdash (\text{\color{green}\dashv\vdash-atomic} \ e \dashv\vdash^* m))) = \\
 & \text{\color{green}\dashv\vdash} , \text{\color{green}\dashv\vdash} , \equiv .\text{refl} , e \dashv\vdash^* m \\
 \text{\color{red}\dashv\vdash-extract} & (\text{\color{blue}\dashv\vdash} : \alpha \not\equiv \tau \ (\rightarrow \dashv\vdash (\text{\color{green}\dashv\vdash-mutate} \ h_1) \triangleleft e \dashv\vdash^*_{\tau} e') \ e' \dashv\vdash e'') \\
 & \text{\color{red}with} \ \text{\color{blue}\dashv\vdash-extract} (\text{\color{blue}\dashv\vdash} : \alpha \not\equiv \tau \ e \dashv\vdash^*_{\tau} e' \ e' \dashv\vdash e'') \\
 \dots \mid h_0 , m , \equiv .\text{refl} , e \dashv\vdash^* m & = \\
 h_0 , m , \equiv .\text{refl} , e \dashv\vdash^* m &
 \end{aligned}$$

Under the stop-the-world semantics, the only non-silent transition **atomic**  $e$  can make is  $\dashv\vdash\text{-atomic}$ , which carries the  $e \dashv\vdash^* m$  we desire. Note that the silent sequence preceding it may contain some number of  $\dashv\vdash\text{-mutate}$  rules which we simply discard here, hence the heap at the start of the transaction may not necessarily be the same as that of the visible transition. We also give a proof that will allow us to refine  $\alpha$ ,  $c''$  and  $e''$ .

Next, we show that for each of the transaction rules under the stop-the-world semantics, there is an corresponding rule that preserves the equivalence of the heap

and log of the log-based semantics with the in-place modified heap of the stop-the-world semantics:

$$\begin{aligned}
 \mapsto' \rightarrow \mapsto' & : \forall \{l \ h_0 \ h \ e \ h' \ e'\} \rightarrow \\
 & \text{Equivalent } h_0 \ l \ h \ \rightarrow \ h, \ e \ \mapsto' \ h', \ e' \ \rightarrow \\
 & \exists \lambda \ l' \ \rightarrow \text{Equivalent } h_0 \ l' \ h' \times \ h_0 \ \vdash \ l, \ e \ \mapsto' \ l', \ e' \\
 \mapsto' \rightarrow \mapsto' \text{ equiv } \mapsto' \text{-}\oplus\mathbb{N} & = \_, \text{equiv}, \mapsto' \text{-}\oplus\mathbb{N} \\
 \mapsto' \rightarrow \mapsto' \text{ equiv } (\mapsto' \text{-}\oplus\mathbb{R} \ m \ b \mapsto b') & = \Sigma.\text{map}_3 \ (\mapsto' \text{-}\oplus\mathbb{R} \ m) \ (\mapsto' \rightarrow \mapsto' \text{ equiv } b \mapsto b') \\
 \mapsto' \rightarrow \mapsto' \text{ equiv } (\mapsto' \text{-}\oplus\mathbb{L} \ b \ a \mapsto a') & = \Sigma.\text{map}_3 \ (\mapsto' \text{-}\oplus\mathbb{L} \ b) \ (\mapsto' \rightarrow \mapsto' \text{ equiv } a \mapsto a') \\
 \mapsto' \rightarrow \mapsto' \text{ equiv } (\mapsto' \text{-}\text{write}_E \ e \mapsto e') & = \Sigma.\text{map}_3 \ \mapsto' \text{-}\text{write}_E \ (\mapsto' \rightarrow \mapsto' \text{ equiv } e \mapsto e')
 \end{aligned}$$

For  $\mapsto' \text{-}\oplus\mathbb{N}$ , this is simply  $\mapsto' \text{-}\oplus\mathbb{N}$ ; the corresponding rules for  $\mapsto' \text{-}\oplus\mathbb{R}$ ,  $\mapsto' \text{-}\oplus\mathbb{L}$  and  $\mapsto' \text{-}\text{write}_E$  are similarly named, and we can simply map them over a recursive call to fulfil the goal.

The  $\mapsto' \text{-}\text{write}_N$  rule that modifies the heap directly has a counterpart  $\mapsto' \text{-}\text{write}_N$  that updates the write log, and we use the **Write-Equivalent** lemma given previously to show that equivalence is maintained:

$$\begin{aligned}
 \mapsto' \rightarrow \mapsto' \text{ equiv } \mapsto' \text{-}\text{write}_N & = \_, \text{Write-Equivalent equiv}, \mapsto' \text{-}\text{write}_N \\
 \mapsto' \rightarrow \mapsto' \{l\} \text{ equiv } (\mapsto' \text{-}\text{read } h_0 \ v) \text{ with } \text{equiv } v \mid \mapsto' \text{-}\text{read } l \ v \\
 \dots \mid \text{equiv-v} \mid \mapsto' \text{-}\text{read-l-v} \text{ with } \text{Logs.}\omega \ l \ [ \ v \ ] \\
 \dots \mid \bullet \ m \ \text{rewrite } \text{equiv-v} & = \_, \text{equiv}, \mapsto' \text{-}\text{read-l-v} \\
 \dots \mid \circ \ \text{with } \text{Logs.}\rho \ l \ [ \ v \ ] \mid \equiv.\text{inspect } (\_ \ [ \ \_ \ ] \ (\text{Logs.}\rho \ l)) \ v \\
 \dots \mid \bullet \ m \ \mid \_ \ \text{rewrite } \text{equiv-v} & = \_, \text{equiv}, \mapsto' \text{-}\text{read-l-v} \\
 \dots \mid \circ \ \mid \llbracket \rho_v \equiv \circ \rrbracket \ \text{rewrite } \equiv.\text{sym} \ \text{equiv-v} & = \_ \\
 & , \text{Read-Equivalent } \rho_v \equiv \circ \ \text{equiv}, \mapsto' \text{-}\text{read-l-v}
 \end{aligned}$$

The matching rule for  $\mapsto' \text{-}\text{read}$  is of course  $\mapsto' \text{-}\text{read}$  in all cases, although we need to pre-apply it with  $l$  and  $v$  and inspect the write and read logs to allow its type and that of the goal to refine so that they coincide. In the last alternative when both both

logs are empty, the **Read-Equivalent** lemma lets us show that heap-log equivalence still holds with the updated read log.

Lastly we generalise the above to transition sequences of any length, proceeding in the usual manner by applying  $\mapsto' \rightarrow \mapsto'$  to each transition from left to right:

$$\begin{aligned}
 \mapsto'^* \rightarrow \mapsto'^* & : \forall \{h_0 \mid h \ e \ h' \ e'\} \rightarrow \\
 & \text{Equivalent } h_0 \mid h \rightarrow h, e \mapsto^* h', e' \rightarrow \\
 & \exists \lambda l' \rightarrow \text{Equivalent } h_0 \mid h' \times h_0 \vdash l, e \mapsto'^* l', e' \\
 \mapsto'^* \rightarrow \mapsto'^* \text{ equiv } \varepsilon = \_ , \text{equiv} , \varepsilon \\
 \mapsto'^* \rightarrow \mapsto'^* \text{ equiv } (e \mapsto e' \triangleleft e' \mapsto^* e'') \text{ with } \mapsto' \rightarrow \mapsto' \text{ equiv } e \mapsto e' \\
 \dots \mid l' , \text{equiv}' , e \mapsto e' \text{ with } \mapsto'^* \rightarrow \mapsto'^* \text{ equiv}' e' \mapsto^* e'' \\
 \dots \mid l'' , \text{equiv}'' , e' \mapsto^* e'' = l'' , \text{equiv}'' , e \mapsto e' \triangleleft e' \mapsto^* e''
 \end{aligned}$$

Using the combination of  $\mapsto$ -**extract** and  $\mapsto'^* \rightarrow \mapsto'^*$ , we can construct a transition sequence under the log-based semantics given a visible transition under the stop-the-world semantics, such that the heap and final log of the former is equivalent to the final heap of the latter.

### 9.4.2 Soundness of Log-Based Transactions

The soundness part of transactional correctness involves showing that the stop-the-world semantics can match the log-based semantics when the latter successfully commits. This is more difficult as the heap can be changed at any point during a log-based transaction by the  $\mapsto$ -**mutate** rule. Let us begin by defining a variation on  $\_ \vdash \mapsto' \_$  that encapsulates the heap used for each transition:

$$\begin{aligned}
 \text{H}\vdash \_ \mapsto' \_ & : \text{Rel} (\text{Logs} \times \text{Expression}') \\
 \text{H}\vdash x \mapsto' x' & = \Sigma \text{Heap} (\lambda h \rightarrow h \vdash x \mapsto' x') \\
 \text{H}\vdash \_ \mapsto'^* \_ & : \text{Rel} (\text{Logs} \times \text{Expression}') \\
 \text{H}\vdash \_ \mapsto'^* \_ & = \text{Star } \text{H}\vdash \_ \mapsto' \_
 \end{aligned}$$

We write  $\text{H} \vdash \_ \rightsquigarrow^* \_$  for its reflexive, transitive closure. Every step of this transition potentially uses a different heap, in contrast to the  $\_ \vdash \_ \rightsquigarrow^* \_$  relation where the entire sequence runs with the same heap.

Next we define a function that discards the steps from any aborted transactions, leaving only the final sequence of transitions leading up to a  $\rightsquigarrow\text{-commit}$ , along with the heaps used at each step:

$$\begin{aligned}
 \rightsquigarrow\text{-extract}' &: \forall \{h \ \alpha \ r \ l \ e \ h' \ c' \ e' \ h'' \ c'' \ e''\} \rightarrow \\
 &\text{H} \vdash \emptyset, r \rightsquigarrow^* l, e \rightarrow \\
 &\alpha \not\equiv \tau \rightarrow h, \rightsquigarrow: \bullet (r, l), \text{atomic } e \rightsquigarrow_{\tau}^* h', c', e' \rightarrow \\
 &\alpha \triangleright h', c', e' \rightarrow h'', c'', e'' \rightarrow \\
 &\exists_2 \lambda l' m \rightarrow \alpha, h'', c'', e'' \equiv \clubsuit, \text{Update } h' l', \rightsquigarrow: \circ, \# m \times \\
 &\text{Consistent } h' l' \times \text{H} \vdash \emptyset, r \rightsquigarrow^* l', \# m \\
 \rightsquigarrow\text{-extract}' \ r \rightsquigarrow^* e \ \alpha \not\equiv \tau \ \varepsilon \ (\rightsquigarrow\text{-}\rightsquigarrow \ (\rightsquigarrow\text{-step } e \rightsquigarrow e')) &= \perp\text{-elim } (\alpha \not\equiv \tau \equiv.\text{refl}) \\
 \rightsquigarrow\text{-extract}' \ r \rightsquigarrow^* e \ \alpha \not\equiv \tau \ \varepsilon \ (\rightsquigarrow\text{-}\rightsquigarrow \ (\rightsquigarrow\text{-mutate } h')) &= \perp\text{-elim } (\alpha \not\equiv \tau \equiv.\text{refl}) \\
 \rightsquigarrow\text{-extract}' \ r \rightsquigarrow^* e \ \alpha \not\equiv \tau \ \varepsilon \ (\rightsquigarrow\text{-}\rightsquigarrow \ (\rightsquigarrow\text{-abort } \neg\text{cons})) &= \perp\text{-elim } (\alpha \not\equiv \tau \equiv.\text{refl})
 \end{aligned}$$

The first argument to  $\rightsquigarrow\text{-extract}'$  accumulates the sequence of transactions steps starting from the initial  $r$ , while the next three correspond to the three fields of a visible transition. The three clauses above eliminate the cases where a silent transition appears in the non-silent position.

If no further transaction steps remain, the only possible rule is  $\rightsquigarrow\text{-commit}$ , in which case we return the accumulated sequence  $r \rightsquigarrow^* e$  and the proof of consistency carried by  $\rightsquigarrow\text{-commit}$ , along with an equality showing the values of  $\alpha$ ,  $h''$ ,  $c''$  and  $e''$ :

$$\begin{aligned}
 \rightsquigarrow\text{-extract}' \ r \rightsquigarrow^* e \ \alpha \not\equiv \tau \ \varepsilon \ (\rightsquigarrow\text{-}\rightsquigarrow \ (\rightsquigarrow\text{-commit } \text{cons})) &= \\
 \_, \_, \equiv.\text{refl}, \text{cons}, r \rightsquigarrow^* e
 \end{aligned}$$

When the transaction makes a single step, we append it to the end of the accumulator, taking a copy of the heap used for that step:



$$\begin{aligned}
 \multimap\text{-extract}' \{h\} r \multimap'^* e \alpha \not\equiv \tau (\multimap\text{-}\multimap (\multimap\text{-step } e \multimap e') \triangleleft e' \multimap_{\tau}^* e'') e'' \multimap e''' &= \\
 \multimap\text{-extract}' (r \multimap'^* e \triangleleft (h, e \multimap e') \triangleleft \varepsilon) \alpha \not\equiv \tau e' \multimap_{\tau}^* e'' e'' \multimap e''' &
 \end{aligned}$$

Should we encounter a  $\multimap\text{-mutate}$  rule, we simply move on to the next step, albeit with a different heap:

$$\begin{aligned}
 \multimap\text{-extract}' r \multimap'^* e \alpha \not\equiv \tau (\multimap\text{-}\multimap (\multimap\text{-mutate } h') \triangleleft e' \multimap_{\tau}^* e'') e'' \multimap e''' &= \\
 \multimap\text{-extract}' r \multimap'^* e \alpha \not\equiv \tau e' \multimap_{\tau}^* e'' e'' \multimap e''' &
 \end{aligned}$$

Lastly if  $e$  has reduced completely to a number, but the read log was not consistent with the heap at that point, the transaction aborts and retries. In this case, we reset the accumulator to  $\varepsilon$  and carry on with the rest of the sequence:

$$\begin{aligned}
 \multimap\text{-extract}' r \multimap'^* e \alpha \not\equiv \tau (\multimap\text{-}\multimap (\multimap\text{-abort } \neg\text{cons}) \triangleleft e' \multimap_{\tau}^* e'') e'' \multimap e''' &= \\
 \multimap\text{-extract}' \varepsilon \alpha \not\equiv \tau e' \multimap_{\tau}^* e'' e'' \multimap e''' &
 \end{aligned}$$

We can write a wrapper for the above that takes a visible transition, and strips off the initial  $\multimap\text{-begin}$  rule before invoking  $\multimap\text{-extract}'$  itself:

$$\begin{aligned}
 \multimap\text{-extract} &: \forall \{\alpha h r h'' c'' e''\} \rightarrow \\
 &\alpha \triangleright h, \multimap: \circ, \text{atomic } r \Rightarrow h'', c'', e'' \rightarrow \\
 &\exists_3 \lambda h' l' m \rightarrow \alpha, h'', c'', e'' \equiv \clubsuit, \text{Update } h' l', \multimap: \circ, \# m \times \\
 &\text{Consistent } h' l' \times \text{HI} \vdash \emptyset, r \multimap'^* l', \# m \\
 \multimap\text{-extract} (\Rightarrow: \alpha \not\equiv \tau \varepsilon (\multimap\text{-}\multimap \multimap\text{-begin})) &= \perp\text{-elim } (\alpha \not\equiv \tau \equiv \text{refl}) \\
 \multimap\text{-extract} (\Rightarrow: \alpha \not\equiv \tau (\multimap\text{-}\multimap \multimap\text{-begin} \triangleleft e \multimap_{\tau}^* e') e' \multimap e'') &= \\
 \_, \multimap\text{-extract}' \varepsilon \alpha \not\equiv \tau e \multimap_{\tau}^* e' e' \multimap e'' &
 \end{aligned}$$

The next lemma says that we can swap the heap used for any  $\perp\text{-}\multimap'$  transition, as long as the target heap is consistent with the original post-transition log  $l'$ :

$$\begin{aligned}
 \multimap'\text{-swap} &: \forall \{h h' l e l' e'\} \rightarrow \text{Consistent } h' l' \rightarrow \\
 &h \vdash l, e \multimap' l', e' \rightarrow h' \vdash l, e \multimap' l', e'
 \end{aligned}$$

$$\begin{aligned}
 \multimap\text{'-swap } cons' \multimap\text{'-}\oplus\mathbb{N} &= \multimap\text{'-}\oplus\mathbb{N} \\
 \multimap\text{'-swap } cons' (\multimap\text{'-}\oplus\mathbb{R} m b \multimap b') &= \multimap\text{'-}\oplus\mathbb{R} m (\multimap\text{'-swap } cons' b \multimap b') \\
 \multimap\text{'-swap } cons' (\multimap\text{'-}\oplus\mathbb{L} b a \multimap a') &= \multimap\text{'-}\oplus\mathbb{L} b (\multimap\text{'-swap } cons' a \multimap a') \\
 \multimap\text{'-swap } cons' (\multimap\text{'-write}_{\mathbb{E}} e \multimap e') &= \multimap\text{'-write}_{\mathbb{E}} (\multimap\text{'-swap } cons' e \multimap e') \\
 \multimap\text{'-swap } cons' \multimap\text{'-write}_{\mathbb{N}} &= \multimap\text{'-write}_{\mathbb{N}}
 \end{aligned}$$

The first few cases are trivial since they either don't interact with the heap, or the proof burden can be deferred to a recursive call. The last  $\multimap\text{'-read}$  case however does require our attention:

$$\begin{aligned}
 &\multimap\text{'-swap } \{h\} \{h'\} cons' (\multimap\text{'-read } l v) \text{ with } \multimap\text{'-read } \{h'\} l v \\
 &\dots \mid \multimap\text{'-read-}l\text{-}v \text{ with } \text{Logs.}\omega l [ v ] \\
 &\dots \mid \bullet m = \multimap\text{'-read-}l\text{-}v \\
 &\dots \mid \circ \text{ with } \text{Logs.}\rho l [ v ] \\
 &\dots \mid \bullet m = \multimap\text{'-read-}l\text{-}v \\
 &\dots \mid \circ \text{ rewrite } cons' v (h [ v ]) \\
 &\quad (\text{Vec.lookupupdate } v (\text{Logs.}\rho l) (\bullet (h [ v ]))) = \multimap\text{'-read-}l\text{-}v
 \end{aligned}$$

As long as one of the log entries for  $v$  is not empty, the transaction does not interact with the heap, so  $\multimap\text{'-read-}l\text{-}v$  by itself suffices. Otherwise by the time we see that both entries are empty,  $\text{Logs.}\rho l'$  has been refined to  $\text{Logs.}\rho l [ v ] := \bullet (h [ v ])$ , so the type of  $cons'$  is now:

$$cons' : \forall v' m \rightarrow (\text{Logs.}\rho l [ v ] := \bullet (h [ v ])) [ v' ] \equiv \bullet m \rightarrow h' [ v' ] \equiv m$$

Instantiating  $v'$  to  $v$  and  $m$  to  $h [ v ]$  and invoking the  $\text{Vec.lookupupdate}$  lemma leads to a witness of  $h' [ v ] \equiv h [ v ]$ , which we use in a rewrite clause to show that  $\multimap\text{'-read}$  under  $h'$  does indeed result in the same  $l'$  and  $e'$  as it did under  $h$ , completing the proof of  $\multimap\text{'-swap}$ .

Of course, we can generalise  $\multimap\text{'-swap}$  to  $\text{H}\vdash \_ \multimap\text{'*} \_$  sequences of any length:

$$\begin{aligned}
 \multimap^{/*}\text{-swap} & : \forall \{h' l e l'' e''\} \rightarrow \text{Consistent } h' l'' \rightarrow \\
 & \text{H}\vdash l, e \multimap^{/*} l'', e'' \rightarrow h' \vdash l, e \multimap^{/*} l'', e'' \\
 \multimap^{/*}\text{-swap } \{h'\} \text{ cons}'' & = \text{snd} \circ \text{Star.gfold id } \text{C}\vdash \multimap^{/*} \text{-trans } (\text{cons}'', \varepsilon) \text{ where} \\
 \text{C}\vdash \multimap^{/*} \text{-} & : \text{Logs} \times \text{Expression}' \rightarrow \text{Logs} \times \text{Expression}' \rightarrow \text{Set} \\
 \text{C}\vdash (l, e) \multimap^{/*} (l', e') & = \text{Consistent } h' l \times h' \vdash l, e \multimap^{/*} l', e' \\
 \text{trans} & : \forall \{x x' x''\} \rightarrow \text{H}\vdash x \multimap' x' \rightarrow \text{C}\vdash x' \multimap^{/*} x'' \rightarrow \text{C}\vdash x \multimap^{/*} x'' \\
 \text{trans } (h, e \multimap e') & (\text{cons}', e' \multimap^* e'') = \\
 & \multimap'\text{-Consistent}' \text{ cons}' e \multimap e', \multimap'\text{-swap } \text{cons}' e \multimap e' \triangleleft e' \multimap^* e''
 \end{aligned}$$

The auxiliary  $\text{C}\vdash \multimap^{/*} \text{-}$  relation pairs  $\text{H}\vdash \multimap^{/*} \text{-}$  with a proof of the consistency of  $h'$  with the read logs at the start of the sequence, while **trans** corresponds to the transitivity of a one-step  $\text{H}\vdash \multimap' \text{-}$  and  $\text{C}\vdash \multimap^{/*} \text{-}$ . The proof of  $\multimap^{/*}\text{-swap}$  results from folding **trans** over the  $\text{H}\vdash \multimap^{/*} \text{-}$  argument, using a final **snd** to discard the consistency part of the  $\text{C}\vdash \multimap^{/*} \text{-}$  pair.

What we have shown with  $\multimap^{/*}\text{-swap}$  is that provided the read log is consistent with the heap just before the commit, then regardless of what different heaps the transaction had originally used, re-running the transaction with just the pre-commit heap—without any intervening heap mutations—delivers the same result.

It remains for us to show that we can construct an equivalent transition under the stop-the-world semantics, given one that uses the same pre-commit heap. We start by taking a single log-based transition step and returning its corresponding stop-the-world rule, while showing that heap-log equivalence is preserved:

$$\begin{aligned}
 \multimap' \rightarrow \mapsto' & : \forall \{h l e l' e' h_0\} \rightarrow \\
 & \text{Equivalent } h_0 l h \rightarrow h_0 \vdash l, e \multimap' l', e' \rightarrow \\
 & \exists \lambda h' \rightarrow \text{Equivalent } h_0 l' h' \times h, e \mapsto' h', e' \\
 \multimap' \rightarrow \mapsto' \text{ equiv } \multimap'\text{-}\oplus\mathbb{N} & = \text{-}, \text{equiv}, \mapsto'\text{-}\oplus\mathbb{N} \\
 \multimap' \rightarrow \mapsto' \text{ equiv } (\multimap'\text{-}\oplus\mathbb{R} m b \multimap b') & = \Sigma.\text{map}_3 (\mapsto'\text{-}\oplus\mathbb{R} m) (\multimap' \rightarrow \mapsto' \text{equiv } b \multimap b')
 \end{aligned}$$

$$\begin{aligned}
 \multimap' \rightarrow \multimap' \text{ equiv } (\multimap' \oplus \mathbf{L} \ b \ a \multimap a') &= \Sigma.\text{map}_3 (\multimap' \oplus \mathbf{L} \ b) (\multimap' \rightarrow \multimap' \text{ equiv } a \multimap a') \\
 \multimap' \rightarrow \multimap' \text{ equiv } (\multimap' \text{-write}_{\mathbb{E}} \ e \multimap e') &= \Sigma.\text{map}_3 \multimap' \text{-write}_{\mathbb{E}} (\multimap' \rightarrow \multimap' \text{ equiv } e \multimap e') \\
 \multimap' \rightarrow \multimap' \text{ equiv } \multimap' \text{-write}_{\mathbb{N}} &= \_, \text{Write-Equivalent equiv}, \multimap' \text{-write}_{\mathbb{N}} \\
 \multimap' \rightarrow \multimap' \{h\} \text{ equiv } (\multimap' \text{-read } l \ v) \text{ with } \text{equiv } v \mid \multimap' \text{-read } h \ v \\
 \dots \mid \text{equiv-v} \mid \multimap' \text{-read-h-v} \text{ with } \text{Logs}.\omega \ l \ [ \ v \ ] \\
 \dots \mid \bullet \ m \ \text{rewrite } \text{equiv-v} &= \_, \text{equiv}, \multimap' \text{-read-h-v} \\
 \dots \mid \circ \ \text{with } \text{Logs}.\rho \ l \ [ \ v \ ] \mid \equiv.\text{inspect } (\_ \ [ \ \_ \ ] \ (\text{Logs}.\rho \ l)) \ v \\
 \dots \mid \bullet \ m \mid \_ \ \text{rewrite } \text{equiv-v} &= \_, \text{equiv}, \multimap' \text{-read-h-v} \\
 \dots \mid \circ \mid \llbracket \rho_v \equiv \circ \rrbracket \ \text{rewrite } \equiv.\text{sym} \ \text{equiv-v} &= \_ \\
 &, \text{Read-Equivalent } \rho_v \equiv \circ \ \text{equiv}, \multimap' \text{-read-h-v}
 \end{aligned}$$

The above definition has an identical structure to that of  $\mapsto' \rightarrow \mapsto'$  from the previous section, using the same **Write-Equivalent** and **Read-Equivalent** lemmas for the  $\multimap' \text{-write}_{\mathbb{N}}$  and  $\multimap' \text{-read}$  cases respectively, so we will let the code speak for itself.

Finally we extend  $\multimap' \rightarrow \multimap'$  to handle any  $\perp \_ \multimap'^* \_$  sequence, in the same manner as  $\mapsto'^* \rightarrow \mapsto'^*$ :

$$\begin{aligned}
 \multimap'^* \rightarrow \multimap'^* &: \forall \{h \ l \ e \ l' \ e' \ h_0\} \rightarrow \\
 &\quad \text{Equivalent } h_0 \ l \ h \ \rightarrow \ h_0 \ \vdash \ l, \ e \ \multimap'^* \ l', \ e' \ \rightarrow \\
 &\quad \exists \lambda \ h' \ \rightarrow \text{Equivalent } h_0 \ l' \ h' \times \ h, \ e \ \mapsto^* \ h', \ e' \\
 \multimap'^* \rightarrow \multimap'^* \text{ equiv } \varepsilon &= \_, \text{equiv}, \varepsilon \\
 \multimap'^* \rightarrow \multimap'^* \text{ equiv } (e \multimap e' \triangleleft e' \multimap^* e'') \text{ with } \multimap' \rightarrow \multimap' \text{ equiv } e \multimap e' \\
 \dots \mid h', \text{equiv}', e \mapsto e' \text{ with } \multimap'^* \rightarrow \multimap'^* \text{ equiv}' e' \multimap^* e'' \\
 \dots \mid h'', \text{equiv}'', e' \mapsto^* e'' = h'', \text{equiv}'', e \mapsto e' \triangleleft e' \mapsto^* e''
 \end{aligned}$$

To summarise, given a visible transition in which the log-based semantics commits a transaction, we can use  $\multimap \text{-extract}$  to obtain the final successful sequence of  $\perp \_ \multimap' \_$  transitions leading up to the commit, along with the heaps used at each step. The  $\multimap'^* \text{-swap}$  lemma then lets us swap the different heaps for the pre-commit heap, while

$\mapsto^{!*} \rightarrow \mapsto^{!*}$  maps each log-based transition to their corresponding stop-the-world ones, allowing us to construct an equivalent transaction under the stop-the-world semantics.

## 9.5 Bisimilarity of Semantics

The proof that the stop-the-world semantics for our Atomic language is bisimilar to the log-based semantics proceeds for the most part by corecursion on the applicable transition rules, as well as structural recursion in the case of  $a \oplus b$  : Expressions when either  $a$  or  $b$  can make further transitions. We will show each of the cases individually, then assemble the pieces to give the full correctness property.

We begin by showing that bisimilarity holds for numbers, where no further transitions are possible:

$$\begin{aligned} \text{correct-}\# & : \forall \{h\ m\} \rightarrow h, \# m \vdash \mapsto : \approx \mapsto : \circ \\ \text{correct-}\# & = \# (\perp\text{-elim} \circ \# \not\mapsto) \ \& \ \# (\perp\text{-elim} \circ \# \not\mapsto) \end{aligned}$$

The proof makes use of a trivial  $\# \not\mapsto$  lemma showing that no visible transitions are possible from expressions of the form  $\# m$ , under either semantics.

### 9.5.1 Addition

For the first non-trivial case, we define  $\text{correct-}\oplus\mathbb{N}$  which handles expressions of the form  $\# m \oplus \# n$ . In this case, the only applicable rules are  $\mapsto\text{-}\oplus\mathbb{N}$  and  $\mapsto\text{-}\oplus\mathbb{N}$ . We show each direction of bisimilarity separately:

$$\begin{aligned} \text{correct-}\oplus\mathbb{N} & : \forall \{h\ m\ n\} \rightarrow h, \# m \oplus \# n \vdash \mapsto : \approx \mapsto : \circ \\ \text{correct-}\oplus\mathbb{N} \{h\} \{m\} \{n\} & = \# \mapsto \rightsquigarrow \rightsquigarrow \ \& \ \# \mapsto \rightsquigarrow \rightsquigarrow \ \mathbf{where} \\ \mapsto \rightsquigarrow \rightsquigarrow & : h, \# m \oplus \# n \vdash \mapsto : \rightsquigarrow \rightsquigarrow : \circ \\ \mapsto \rightsquigarrow \rightsquigarrow (\mapsto : \alpha \neq \tau \ \varepsilon \ (\rightarrow \mapsto \mapsto \oplus\mathbb{N})) & = \\ \_, \mapsto : \alpha \neq \tau \ \varepsilon \ (\rightarrow \rightsquigarrow \rightsquigarrow \rightsquigarrow \oplus\mathbb{N}), \text{correct-}\# & \end{aligned}$$

$$\begin{aligned}
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto (\mapsto \oplus \mathbf{R} \ \_ \ b \mapsto b'))) &= \perp\text{-elim} (\# \not\mapsto b \mapsto b') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto (\mapsto \oplus \mathbf{L} \ \_ \ a \mapsto a'))) &= \perp\text{-elim} (\# \not\mapsto a \mapsto a') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ (\mapsto \mapsto (\mapsto \oplus \mathbf{R} \ \_ \ b \mapsto b') \triangleleft \_ \ \_)) &= \perp\text{-elim} (\# \not\mapsto b \mapsto b') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ (\mapsto \mapsto (\mapsto \oplus \mathbf{L} \ \_ \ a \mapsto a') \triangleleft \_ \ \_)) &= \perp\text{-elim} (\# \not\mapsto a \mapsto a')
 \end{aligned}$$

To show that the log-based semantics can simulate the stop-the-world semantics we inspect the visible transition that  $\# m \oplus \# n$  makes under the latter. As hinted above, the only applicable transition is  $\mapsto \oplus \mathbf{N}$ , for which we use  $\mapsto \oplus \mathbf{N}$  to show that the log-based semantics can follow. The resulting expression of  $\# (m + n)$  is then bisimilar by the **correct-#** lemma. The remaining clauses amount to showing that further transitions by  $\# m$  or  $\# n$  alone are impossible.

The proof for the opposite direction proceeds in exactly the same way:

$$\begin{aligned}
 \mapsto \rightsquigarrow \mapsto : h, \# m \oplus \# n \vdash \mapsto : \circ \rightsquigarrow \mapsto : \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto \mapsto \oplus \mathbf{N})) &= \\
 \_, \mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto \mapsto \oplus \mathbf{N}), \approx\text{-sym} \text{ correct-#} \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto (\mapsto \oplus \mathbf{R} \ \_ \ b \mapsto b'))) &= \perp\text{-elim} (\# \not\mapsto b \mapsto b') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ \varepsilon \ (\mapsto \mapsto (\mapsto \oplus \mathbf{L} \ \_ \ a \mapsto a'))) &= \perp\text{-elim} (\# \not\mapsto a \mapsto a') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ (\mapsto \mapsto (\mapsto \oplus \mathbf{R} \ \_ \ b \mapsto b') \triangleleft \_ \ \_)) &= \perp\text{-elim} (\# \not\mapsto b \mapsto b') \\
 \mapsto \rightsquigarrow \mapsto (\mathbb{E}: \alpha \neq \tau \ (\mapsto \mapsto (\mapsto \oplus \mathbf{L} \ \_ \ a \mapsto a') \triangleleft \_ \ \_)) &= \perp\text{-elim} (\# \not\mapsto a \mapsto a')
 \end{aligned}$$

## 9.5.2 Right Evaluation

Given an induction hypothesis of  $h, b \vdash \mapsto : \approx \mapsto : \circ$ , we can show that the two semantics are bisimilar for expressions of the form  $\# m \oplus b$ :

$$\begin{aligned}
 \text{correct-}\oplus \mathbf{R} : \forall \{h \ m \ b\} \rightarrow h, b \vdash \mapsto : \approx \mapsto : \circ \rightarrow h, \# m \oplus b \vdash \mapsto : \approx \mapsto : \circ \\
 \text{correct-}\oplus \mathbf{R} \ \{h\} \ \{m\} \ \{b\} \ b \vdash \mapsto \approx \mapsto = \# \mapsto \rightsquigarrow \mapsto \ \& \ \# \mapsto \rightsquigarrow \mapsto \ \text{where}
 \end{aligned}$$

For the completeness direction, we first use a  $\mapsto^* / \mapsto \oplus \mathbf{R}$  helper that peels off any  $\mapsto \oplus \mathbf{R}$  rules in the visible transition starting from  $\# m \oplus b$ . This is not always

possible: when  $b$  is already a number  $\# n$ , the full expression cannot make any transitions under  $\mapsto\oplus\mathbf{R}$ , so it returns a witness  $b\equiv n$  that allows us to defer the rest of the proof to one half of the **correct- $\oplus\mathbf{N}$**  lemma given earlier:

$$\begin{aligned} & \mapsto\rightsquigarrow : h, \# m \oplus b \vdash \mapsto : \rightsquigarrow : \circ \\ & \mapsto\rightsquigarrow (\mapsto : \alpha \neq \tau \ e \mapsto_{\tau}^* e' \ e' \mapsto e'') \text{ with } \mapsto^*/\mapsto\oplus\mathbf{R} \ \alpha \neq \tau \ e \mapsto_{\tau}^* e' \ e' \mapsto e'' \\ & \dots \mid \text{inl} (n, b\equiv n, \equiv.\text{refl}, \equiv.\text{refl}, \equiv.\text{refl}) \text{ rewrite } b\equiv n = \\ & \quad b (\approx\rightsquigarrow\rightsquigarrow \text{correct-}\oplus\mathbf{N}) (\mapsto : \alpha \neq \tau \ \varepsilon (\mapsto\mapsto\mapsto\oplus\mathbf{N})) \end{aligned}$$

Otherwise  $b$  must make some visible transition under  $\mapsto\oplus\mathbf{R}$ , and  $\mapsto^*/\mapsto\oplus\mathbf{R}$  returns  $b\mapsto_{\tau}^* b' : h, \mapsto : , b \mapsto_{\tau}^* h', \mapsto : , b'$  as well as  $b'\mapsto b'' : \alpha \triangleright h', \mapsto : , b' \mapsto h'', \mapsto : , b''$ , essentially constituting a visible transition made by just  $b$  itself. The latter transition is labelled with the same  $\alpha$  as the original  $e'\mapsto e''$ , which in turn has been refined to  $h', \mapsto : , \# m \oplus b' \mapsto_{\tau}^* h'', \mapsto : , \# m \oplus b''$  by the two equality proofs returned from  $\mapsto^*/\mapsto\oplus\mathbf{R}$ :

$$\begin{aligned} & \dots \mid \text{inr} (h', b', h'', b'', \equiv.\text{refl}, \equiv.\text{refl}, b\mapsto_{\tau}^* b', b'\mapsto b'') \\ & \quad \text{with } b (\approx\rightsquigarrow\rightsquigarrow b\vdash\mapsto\approx\rightsquigarrow) (\mapsto : \alpha \neq \tau \ b\mapsto_{\tau}^* b' \ b'\mapsto b'') \\ & \dots \mid c'', b\mapsto b'', b''\vdash\mapsto\approx\rightsquigarrow \text{ with } \mapsto\circ\rightsquigarrow\oplus\mathbf{R} \ m \ b\mapsto b'' \\ & \dots \mid c''\equiv\rightsquigarrow, m\oplus b\mapsto m\oplus b'' \text{ rewrite } c''\equiv\rightsquigarrow = \\ & \quad -, m\oplus b\mapsto m\oplus b'', \text{correct-}\oplus\mathbf{R} \ b''\vdash\mapsto\approx\rightsquigarrow \end{aligned}$$

Next we invoke one half of the induction hypothesis  $b\vdash\mapsto\approx\rightsquigarrow$  with the aforementioned stop-the-world visible transition of  $\mapsto : \alpha \neq \tau \ b\mapsto_{\tau}^* b' \ b'\mapsto b''$ , which returns an equivalent log-based visible transition  $b\mapsto b'' : \alpha \triangleright h, \rightsquigarrow : \circ, b \mapsto h'', \rightsquigarrow : \circ, b''$ . Another lemma  $\mapsto\circ\rightsquigarrow\oplus\mathbf{R}$  then replaces the  $\mapsto\oplus\mathbf{R}$  rules peeled off earlier with their corresponding  $\rightsquigarrow\oplus\mathbf{R}$  rules, and a corecursive call to **correct- $\oplus\mathbf{R}$**  completes this part of the proof.

The definitions of  $\mapsto^*/\mapsto\oplus\mathbf{R}$  and  $\mapsto\circ\rightsquigarrow\oplus\mathbf{R}$  are straightforward but rather tedious, and they can be found in the full source code on the author's website.





$$\begin{aligned}
\text{correct-}\oplus\text{L} &: \forall \{h \ a \ b\} \rightarrow h, a \vdash \mapsto: \approx \rightsquigarrow: \circ \rightarrow \\
& (\forall h' \rightarrow h', b \vdash \mapsto: \approx \rightsquigarrow: \circ) \rightarrow h, a \oplus b \vdash \mapsto: \approx \rightsquigarrow: \circ \\
\text{correct-}\oplus\text{L} \{h\} \{a\} \{b\} a \vdash \mapsto \approx \rightsquigarrow \forall b \vdash \mapsto \approx \rightsquigarrow &= \# \mapsto \rightsquigarrow \& \# \rightsquigarrow \rightsquigarrow \mapsto \text{ where}
\end{aligned}$$

The  $\rightarrow^*/\mapsto\oplus\text{L}$  lemma then lets us distinguish between the case when  $a$  is just a number  $\# m$ , and the case where  $a$  makes a visible transition. In the former case, we pass the proof obligation on to the  $\text{correct-}\oplus\text{R}$  lemma, instantiating  $\forall b \vdash \mapsto \approx \rightsquigarrow$  with the current heap:

$$\begin{aligned}
& \mapsto \rightsquigarrow \rightsquigarrow : h, a \oplus b \vdash \mapsto: \rightsquigarrow \rightsquigarrow: \circ \\
& \mapsto \rightsquigarrow \rightsquigarrow (\mapsto: \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e'') \text{ with } \rightarrow^*/\mapsto\oplus\text{L} \ \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e'' \\
& \dots \mid \text{inl} (m, a \equiv m) \text{ rewrite } a \equiv m = \\
& \quad b (\approx \rightarrow \rightsquigarrow (\text{correct-}\oplus\text{R} (\forall b \vdash \mapsto \approx \rightsquigarrow h))) (\mapsto: \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e'') \\
& \dots \mid \text{inr} (h', a', h'', a'', \equiv.\text{refl}, \equiv.\text{refl}, a \rightarrow_{\tau}^* a', a' \rightarrow a'') \\
& \quad \text{with } b (\approx \rightarrow \rightsquigarrow a \vdash \mapsto \approx \rightsquigarrow) (\mapsto: \alpha \neq \tau \ a \rightarrow_{\tau}^* a' \ a' \rightarrow a'') \\
& \dots \mid c'', a \mapsto a'', a'' \vdash \mapsto \approx \rightsquigarrow \text{ with } \mapsto \circ \rightarrow \oplus\text{L} \ b \ a \mapsto a'' \\
& \dots \mid c'' \equiv \rightsquigarrow, a \oplus b \mapsto a'' \oplus b \text{ rewrite } c'' \equiv \rightsquigarrow = \\
& \quad -, a \oplus b \mapsto a'' \oplus b, \text{correct-}\oplus\text{L} \ a'' \vdash \mapsto \approx \rightsquigarrow \forall b \vdash \mapsto \approx \rightsquigarrow
\end{aligned}$$

Otherwise we have a visible transition  $\mapsto: \alpha \neq \tau \ a \rightarrow_{\tau}^* a' \ a' \rightarrow a''$ , and the first inductive hypothesis allows us to obtain a corresponding visible transition  $a \mapsto a''$  under the log-based semantics. Next we replace the  $b$  on the right hand side of  $-\oplus-$  using the  $\mapsto \circ \rightarrow \oplus\text{L}$  lemma to obtain the transition  $a \oplus b \mapsto a'' \oplus b$ . Since  $b$  has not made any transitions, a corecursive call with  $a'' \vdash \mapsto \approx \rightsquigarrow$  and the original  $\forall b \vdash \mapsto \approx \rightsquigarrow$  completes the proof.

We proceed with the soundness half of  $\text{correct-}\oplus\text{L}$  in exactly the same way as that of  $\text{correct-}\oplus\text{R}$ :

$$\begin{aligned}
& \rightsquigarrow \rightsquigarrow \rightsquigarrow : h, a \oplus b \vdash \rightsquigarrow: \circ \rightsquigarrow \rightsquigarrow: \\
& \rightsquigarrow \rightsquigarrow \rightsquigarrow (\mapsto: \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e'') \text{ with } \rightarrow^*/\rightsquigarrow\oplus\text{L} \ \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e''
\end{aligned}$$

$$\begin{aligned}
 & \dots \mid \mathbf{inl} (m, a \equiv m) \mathbf{rewrite} \ a \equiv m = \\
 & \quad \mathbf{b} (\approx \rightarrow \rightsquigarrow (\mathbf{correct}\text{-}\oplus\mathbf{R} (\forall b \vdash \rightarrow \approx \rightarrow h))) (\Rightarrow: \alpha \neq \tau \ e \rightarrow_{\tau}^* e' \ e' \rightarrow e'') \\
 & \dots \mid \mathbf{inr} (h', t', a', h'', a'', \equiv.\mathbf{refl}, \equiv.\mathbf{refl}, a \rightarrow_{\tau}^* a', a' \rightarrow a'') \\
 & \quad \mathbf{with} \ \mathbf{b} (\approx \rightarrow \rightsquigarrow a \vdash \rightarrow \approx \rightarrow) (\Rightarrow: \alpha \neq \tau \ a \rightarrow_{\tau}^* a' \ a' \rightarrow a'') \\
 & \dots \mid c'', a \Rightarrow a'', a'' \vdash \rightarrow \approx \rightarrow \mathbf{with} \ \Rightarrow \circ \rightarrow \oplus \mathbf{L} \ b \ a \Rightarrow a'' \\
 & \dots \mid c'' \equiv \rightarrow, a \oplus b \Rightarrow a'' \oplus b \mathbf{rewrite} \ c'' \equiv \rightarrow = \\
 & \quad -, a \oplus b \Rightarrow a'' \oplus b, \approx \rightarrow \rightsquigarrow \rightarrow \approx \rightarrow \ \& \ \approx \rightarrow \rightsquigarrow \rightarrow \approx \rightarrow \mathbf{where} \\
 & \quad \rightarrow \approx \rightarrow = \mathbf{correct}\text{-}\oplus \mathbf{L} (\approx\text{-sym} \ a'' \vdash \rightarrow \approx \rightarrow) \forall b \vdash \rightarrow \approx \rightarrow
 \end{aligned}$$

Observe how  $\mathbf{correct}\text{-}\oplus \mathbf{L}$  shares the same overall structure as  $\mathbf{correct}\text{-}\oplus \mathbf{R}$ .

### 9.5.4 Transactions

Finally we arrive at the correctness proof for  $\mathbf{atomic}$  expressions, where we need to show that transactions run under the stop-the-world semantics coincide with those using our log-based semantics:

$$\begin{aligned}
 \mathbf{correct}\text{-}\mathbf{atomic} & : \forall \{h \ e\} \rightarrow h, \mathbf{atomic} \ e \vdash \Rightarrow: \approx \rightarrow: \circ \\
 \mathbf{correct}\text{-}\mathbf{atomic} \ \{h\} \ \{e\} & = \# \rightarrow \rightsquigarrow \ \& \ \# \rightarrow \rightsquigarrow \mathbf{where}
 \end{aligned}$$

In the completeness direction, we show that the log-based semantics can follow the stop-the-world one by simply running the entire transaction uninterrupted at the same point as the  $\rightarrow\text{-}\mathbf{atomic}$  rule. First the  $\rightarrow\text{-}\mathbf{extract}$  helper from §9.4.1 picks out the  $e \mapsto^* m$  sequence:

$$\begin{aligned}
 \rightarrow \rightsquigarrow & : h, \mathbf{atomic} \ e \vdash \Rightarrow: \rightsquigarrow \rightarrow: \circ \\
 \rightarrow \rightsquigarrow \ e \Rightarrow e'' & \mathbf{with} \ \rightarrow\text{-}\mathbf{extract} \ e \Rightarrow e'' \\
 \dots \mid h_0, m, \equiv.\mathbf{refl}, e \mapsto^* m \\
 & \mathbf{with} \ \mapsto^* \rightarrow \rightarrow^* \ \emptyset\text{-Equivalent} \ e \mapsto^* m \\
 \dots \mid l', \mathit{equiv}', e \mapsto^* m & \mathbf{with} \ \rightarrow^* \text{-Consistent} \ \emptyset\text{-Consistent} \ e \mapsto^* m
 \end{aligned}$$

```

... | cons' rewrite  $\equiv$ .sym (Commit cons' equiv') =
    -, e $\Rightarrow$ m , correct-# where

```

The  $\mapsto^*/\rightarrow^*$  function from the same section then translates each transition of  $e \mapsto^* m$  to its log-based equivalent, as well as constructing a proof  $equiv'$  of the equivalence of  $h_0$  and  $l'$  with the heap  $h'$  at the end of the  $e \mapsto^* m$  sequence. By the  $\mapsto^*$ -Consistent lemma, we show that the consistency of  $h_0$  and the logs is preserved along the  $e \mapsto^* m$  sequence, culminating in a witness  $cons'$  of Consistent  $h_0 l' h'$ . Finally, a rewrite clause using the Commit lemma convinces Agda that Update  $h_0 l'$  and  $h'$  are definitionally equal. Since running a transaction results in just a number  $\# m$ , correct-# suffices to show that both semantics are bisimilar in this case.

It remains for us to construct the visible transition  $e \Rightarrow m$  that uses the log-based semantics. Should the heap be changed just before the stop-the-world semantics runs the transaction, we need a corresponding  $\mapsto$ -mutate rule in the log-based transition sequence. The `mutate?` helper checks whether this is necessary:

```

mutate? : h ,  $\mapsto$ : • (e ,  $\emptyset$ ) , atomic e  $\rightarrow^*$  $\tau$  h0 ,  $\mapsto$ : • (e ,  $\emptyset$ ) , atomic e
mutate? with h  $\stackrel{?}{=}_{\text{Heap}}$  h0
... | yes h $\equiv$ h0 rewrite h $\equiv$ h0 =  $\varepsilon$ 
... | no h $\not\equiv$ h0 =  $\mapsto$ - $\mapsto$  ( $\mapsto$ -mutate h0)  $\triangleleft$   $\varepsilon$ 

```

Next, the auxiliary definition  $e \rightarrow^*_\tau m$  lifts each transition of  $e \mapsto^* m$  up to the  $\triangleright \_ \_ \_$  level using the  $\mapsto$ -step rule, prepending a  $\mapsto$ -mutate rules when necessary:

```

e $\rightarrow^*_\tau$ m : h ,  $\mapsto$ : • (e ,  $\emptyset$ ) , atomic e  $\rightarrow^*$  $\tau$  h0 ,  $\mapsto$ : • (e , l') , atomic (# m)
e $\rightarrow^*_\tau$ m = mutate?  $\triangleleft$  Star.gmap _ ( $\mapsto$ - $\mapsto$   $\circ$   $\mapsto$ -step) e $\mapsto^*$ m
e $\Rightarrow$ m :  $\clubsuit$   $\triangleright$  h ,  $\mapsto$ :  $\circ$  , atomic e  $\Rightarrow$  Update h0 l' ,  $\mapsto$ :  $\circ$  , # m
e $\Rightarrow$ m =  $\Rightarrow$ : (  $\lambda$  () ) ( $\mapsto$ - $\mapsto$   $\mapsto$ -begin  $\triangleleft$  e $\rightarrow^*_\tau$ m) ( $\mapsto$ - $\mapsto$  ( $\mapsto$ -commit cons'))

```

Lastly we add a  $\mapsto$ -begin to beginning of the visible transition to initialise the transaction state, followed by  $e \rightarrow^*_\tau m$  as the main body of the transaction. A final non-silent

$\mapsto$ -**commit** carrying the  $cons'$  witness produces the desired visible transition.

The proof of soundness relies on us having shown that for every transaction that commits successfully, re-running the entire transaction without any interference at the point of  $\mapsto$ -**commit** computes the same result. We can then use this uninterrupted transaction sequence to derive a corresponding visible transition under the stop-the-world semantics.

We start with a similar  $\mapsto$ -**extract** lemma defined in §9.4.2 that returns a sequence  $e \mapsto^{!*} m : \mathbb{H} \vdash \emptyset, e \mapsto^{!*} l', \# m$  where each transition potentially uses a different heap, as well as the  $cons'$  proof carried by the final  $\mapsto$ -**commit**:

$$\begin{aligned}
 & \mapsto \preccurlyeq \mapsto : h, \text{atomic } e \vdash \mapsto : \circ \preccurlyeq \mapsto : \\
 & \mapsto \preccurlyeq \mapsto e \Rightarrow e'' \text{ with } \mapsto\text{-extract } e \Rightarrow e'' \\
 & \dots \mid h_0, l', m, \equiv.\text{refl}, cons', e \mapsto^{!*} m \\
 & \quad \text{with } \mapsto^{!*} \rightarrow \mapsto^{!*} \emptyset\text{-Equivalent } (\mapsto^{!*}\text{-swap } cons' e \mapsto^{!*} m) \\
 & \dots \mid h', equiv', e \mapsto^{!*} m \text{ rewrite } \equiv.\text{sym } (\text{Commit } cons' equiv') = \\
 & \quad \_ , e \Rightarrow m, \approx\text{-sym correct-}\# \text{ where}
 \end{aligned}$$

There is one additional step involved: we must use the  $\mapsto^{!*}$ -**swap** lemma to replace the different heaps used throughout  $e \mapsto^{!*} m$  with  $h_0$ —to give a witness of  $h_0 \vdash \emptyset, e \mapsto^{!*} l', \# m$ —before we can use  $\mapsto^{!*} \rightarrow \mapsto^{!*}$  to convert this to the sequence  $e \mapsto^{!*} m : h_0, e \mapsto^{!*} h', \# m$ . This is necessary because the  $\mapsto^{!*} \rightarrow \mapsto^{!*}$  lemma requires its input log-based transitions to be under the same heap in order to show equivalence preservation.

The result of both transaction is the same expression  $\# m$ , and we use the symmetry of the earlier **correct-#** lemma to provide evidence of bisimilarity. All that is left is to wrap up the stop-the-world transactional transition sequence  $e \mapsto^{!*} m$  as the visible transition  $e \Rightarrow m$ . We define a **mutate?** sequence to correspond to any necessary pre-transaction  $\mapsto$ -**mutate** rules,

$$\begin{aligned}
& \text{mutate?} : h, \mapsto:, \text{atomic } e \xrightarrow{\tau^*} h_0, \mapsto:, \text{atomic } e \\
& \text{mutate? with } h \stackrel{?}{=}_{\text{Heap}} h_0 \\
& \dots \mid \text{yes } h \equiv h_0 \text{ rewrite } h \equiv h_0 = \varepsilon \\
& \dots \mid \text{no } h \not\equiv h_0 = \xrightarrow{\tau} \mapsto ( \mapsto\text{-mutate } h_0 ) \triangleleft \varepsilon \\
& e \Rightarrow m : \clubsuit \triangleright h, \mapsto:, \text{atomic } e \Rightarrow \text{Update } h_0 l', \mapsto:, \# m \\
& e \Rightarrow m = \Rightarrow: (\lambda () \text{ mutate? } ( \xrightarrow{\tau} \mapsto ( \mapsto\text{-atomic } e \mapsto^{\tau^*} m ) ))
\end{aligned}$$

which simply slots in as the silent transitions before the final  $\mapsto\text{-atomic}$ , to give the desired visible transition  $e \Rightarrow m$ . This completes the bisimilarity proof for transactions.

### 9.5.5 Putting It Together

Having shown for each individual case that our stop-the-world and log-based semantics are indeed bisimilar, all that remains for us is to combine them to give the proof of bisimilarity for any [Expression](#):

$$\begin{aligned}
& \text{correct} : \forall h e \rightarrow h, e \vdash \mapsto: \approx \mapsto: \circ \\
& \text{correct } h (\# m) = \text{correct-}\# \\
& \text{correct } h (a \oplus b) = \text{correct-}\oplus\text{L } (\text{correct } h a) (\lambda h' \rightarrow \text{correct } h' b) \\
& \text{correct } h (\text{atomic } e) = \text{correct-atomic}
\end{aligned}$$

In the  $a \oplus b$  case, observe how  $\text{correct-}\oplus\text{L}$  is supplied with the induction hypothesis on  $a$ , but that for  $b$  is abstracted over an arbitrary heap  $h'$ .

## 9.6 Conclusion

In this chapter we considered a language with transactions, and a worst-case ‘mutate’ rule in place of explicitly modelling interference by concurrent threads. We gave a fully formalised proof that our stop-the-world and log-based semantics for transactions do indeed coincide, without resorting to any postulated lemmas: the completeness part

simply ran the log-based transaction uninterrupted at the same time as the stop-the-world  $\mapsto$ -**atomic** rule. For soundness we used a key  $\mapsto$ -**\*-swap** lemma showing that regardless of any heap interference during the log-based transaction, as long as the pre-commit heap  $h'$  and logs  $l'$  were consistent, replaying the transaction under  $h'$  gives the same results. This underlines the importance of **Consistent**  $h' l'$  as the correct criteria for a log-based transaction to commit.

# Chapter 10

## Conclusion

To conclude, these final pages will comprise an overview of this thesis and an account of how it came to be, followed by a summary of its contributions and some directions for further work.

### 10.1 Retrospection

The quest for higher-level abstractions to manage the complexities of concurrent programming has been an especially apt topic in recent years, due to reasons outlined in the introductory chapter. With respect to software transactional memory (Chapter 2), I was fortunate enough to be in the right places at the right times to have attended two of Tim Harris’s talks on the topic: the first in Cambridge during my undergraduate years, on the JVM-based implementation; and a second time at Imperial College during my MSc course in the early part of 2005, on the composability of STM Haskell.

My work for this thesis began in 2006 under the guidance of Graham Hutton, with an initial goal of reasoning about concurrent programs, in particular those written using STM. To this end, we opted for a simple formal language based on Hutton’s Razor, extended with a minimal set of transactional primitives, described in Chapter

§5. While this language—following the reference stop-the-world semantics given by Harris et al. [HMPJH05]—had a simple implementation, it was not immediately clear how STM Haskell dealt with conflicting transactions internally, consequently drawing our attention towards the correctness of the low-level concurrent implementation.

To better understand the implementation issues behind software transactional memory, we began building a stack-based virtual machine and a compiler for our minimal language, of which the final version is given in §5.2. Using Haskell as a metalanguage, it was a straightforward task to transcribe the syntax and semantics of our model as an executable program. Combined with the use of QuickCheck and the Haskell Program Coverage toolkit (Chapter 4), this allowed us to take a ‘rapid prototyping’ approach to the design of the virtual machine. Most notably, we were able to clarify the appropriate conditions needed for a transaction to commit successfully (§5.2.3), and to realise that writing to an as-yet unread variable within a transaction does not imply a dependency on the current state of the heap.

Of course, regardless of how many times we run QuickCheck on our STM model, overwhelming evidence does not constitute a proof of compiler correctness for our interleaved semantics and log-based implementation of transactional memory. Due to the influence of numerous type-theorists at Nottingham, I had become interested in dependently-typed programming (Chapter 6) and dually, the application of intuitionistic type theory as a framework for conducting formal proofs. Therefore, the goal of a mechanised compiler correctness proof for our model in a dependently-typed language/proof-assistant seemed a natural choice.

Whereas compiler correctness theorems in a deterministic setting (Chapter 3) are concerned only with the final results, with the introduction of non-determinism (of which concurrency is one form) we can no longer afford to ignore *how* results are computed, in addition to what is being computed. Bisimilarity as a notion of behavioural equivalence is a standard tool in the field of process calculi, and



Wand [Wan95, WS95, GW96] et al. were the first to use it to tackle concurrent compiler correctness over a decade ago. Their work relied on giving denotational semantics to both source and target languages in an underlying process calculus, and showing that compilation preserved bisimilarity of denotations. In contrast, we defined our language and virtual machine in an operational manner, and sought a simpler and more direct approach.

Thus, the idea of the combined machine was born, detailed in Chapter 7. A key realisation was that certain kinds transitions preserve bisimilarity, giving rise to the **elide- $\tau$**  lemma. Having tested the waters with the non-deterministic Zap language, Chapter 8 then demonstrates that our approach can indeed scale to handle concurrency, at least for that of the Fork language. As well as updating the **elide- $\tau$**  lemma for explicit concurrency, we also showed that combining bisimilar pairs of thread soups preserves bisimilarity.

Finally, Chapter 9 considers a log-based implementation of transactions. In order to simplify the proof of its equivalence with a stop-the-world semantics, we replaced concurrency with a ‘mutate’ rule that simulates the worst possible concurrent environment, and directly defined a log-based transaction semantics on the source language. The final correctness proof makes essential use of a notion of equivalence between heaps and logs, and confirms our earlier hypothesis that consistency of the read log with the heap is a sufficient condition for a transaction to commit.

## 10.2 Summary of Contributions

This thesis addresses the familiar question of compiler correctness in a current context, with a particular emphasis on the implementation of software transactional memory. We have identified a simplified subset of STM Haskell that is amenable to formal reasoning, which has a stop-the-world transactional semantics, together with

a concurrent virtual machine for this language, using the notion of transaction logs. A compiler linking this simplified language to its virtual machine then allowed us to formulate a concrete statement of compiler correctness. We were able to implement the above semantics in a fairly direct manner using the high-level vocabulary provided by Haskell, enabling us to empirically test compiler correctness with the help of QuickCheck and HPC.

Working towards a formal proof of the above hypothesis, we stripped down to a minimal language with trivial non-determinism, and moving to a labelled transition system. The core idea of a combined machine and semantics then allowed us to establish a direct bisimulation between this language and its virtual machine. This technique was put into practice using the Agda proof assistant, giving a machine-checked compiler correctness proof for a language with a simple notion of non-determinism. We then extended the above proof and our approach in the direction of the initially identified subset of STM Haskell, in an incremental manner: first introducing explicit concurrency in the form of a `fork` primitive, before finally tackling a language with an `atomic` construct in a simplified setting, resulting in a formal proof of the equivalence of the high-level stop-the-world semantics with a low-level log-based approach.

### 10.3 Directions for Further Research

Our original aim was to formally verify the compiler correctness result from Chapter 5, whereas in Chapter 9 we verified this result with some simplifications, in particular the removal of the `fork` construct. While this establishes what we feel to be the essence of the result, we believe our original aim is still achievable with further work, since our worst-case ‘mutate’ rule subsumes any interference by other threads in a concurrent setting.

Our simplified model of STM Haskell focuses on the essence of implementing

transactions, and consequently omits many of the facilities expected of a realistic language. Namely, the lack of primitive recursion or even name binding limits the computational power of our model in a very tangible sense. We could tackle this using a lightweight approach, by borrowing said facilities from the metalanguage and defining the high-level semantics as a functional specification. For example, Gordon’s thesis [Gor92] presents such a specification of teletype IO for a subset of Haskell in terms of a low-level metalanguage, while Swierstra [Swi08] advocates the use of Agda as the metalanguage, due to its enforcement of totality.

Given the above as a basis, a machine-verified formalisation of the omitted parts of the STM Haskell specification—in particular `retry/orElse` and the interaction with exceptions—becomes a much more tractable proposition. Open questions include: How will these additions affect the design of the corresponding virtual machine? Can we maintain the simplicity of our combined machine approach? Is the outline of our reasoning for transactions still valid in this richer language? Our current virtual machine immediately retries failed transactions, rather than waiting until some relevant transactional variable has changed. How can our virtual machine more faithfully model the implementation of STM in GHC?

Going further, we could extend the set of side-effects that can be safely rolled back by the transactional machinery. One widely asserted advantage of STM Haskell over other STM implementations is that its type system restricts transactions—aside from modifying `TVars`—to pure computations, guaranteeing that rollback is always possible. During my initial work on the model of STM Haskell, the notion of running multiple nested transactions concurrently arose quite naturally, when considering their rôle in the implementation of `retry/orElse`. While the `forkIO` primitive is considered impure, forking a nested transaction need not be, as its side-effects can only escape as part of that of its enclosing transaction. It could be interesting to flesh out the precise behaviour of such a `forkSTM :: STM () -> STM ThreadId` primitive, and

## CHAPTER 10. CONCLUSION

to evaluate its utility for concurrent programming in the real world.

# Bibliography

- [ATS09] Ali-Reza Adl-Tabatabai and Tatiana Shpeisman. *Draft Specification of Transactional Language Constructs for C++*. Intel, August 2009.
- [AVWW96] Joe Armstrong, Robert Virding, Claes Wikström, and Mike Williams. *Concurrent Programming in ERLANG*. Prentice Hall, second edition edition, 1996.
- [Bar84] Hendrick Pieter Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. Elsevier, 1984.
- [CGE08] David Cunningham, Khilan Gudka, and Susan Eisenbach. Keep Off the Grass: Locking the Right Path for Atomicity. In *Compiler Construction 2008*, April 2008.
- [CH00] Koen Claessen and John Hughes. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *ICFP Proceedings*, 2000.
- [Chu36] Alonzo Church. An Unsolvable Problem of Elementary Number Theory. *American Journal of Mathematics*, 58(2):345–363, April 1936.
- [CPN98] David G. Clarke, John M. Potter, and James Noble. Ownership Types for Flexible Alias Protection. In *Proceedings of the 13th Conference on Object-Oriented Programming, Systems, Languages, and Applications*

## BIBLIOGRAPHY

- (*OOPSLA-98*), volume 33:10 of *ACM SIGPLAN Notices*, pages 48–64, New York, October 1998. ACM Press.
- [Dan10a] Nils Anders Danielsson. Beating the Productivity Checker Using Embedded Languages. In *Workshop on Partiality and Recursion in Interactive Theorem Provers*, July 2010.
- [Dan10b] Nils Anders Danielsson. The Agda Standard Library. Available from <http://cs.nott.ac.uk/~nad/listings/lib/>, September 2010.
- [Dat95] Christopher J Date. *An Introduction to Database Systems*. Addison-Wesley, 6<sup>th</sup> edition, 1995.
- [Dav03] Maulik A Dave. Compiler Verification: A Bibliography. *ACM SIGSOFT Software Engineering Notes*, 28(6):2–2, November 2003.
- [Dij65] Edsger Wybe Dijkstra. Cooperating Sequential Processes. Lecture notes, Technological University, Eindhoven, The Netherlands, September 1965.
- [Enn05] Robert Ennals. Software Transactional Memory Should Not be Obstruction-Free. Submitted to SCOOOL 2005, 2005.
- [FFL05] Cormac Flanagan, Stephen N Freund, and Marina Lifshin. Type Inference for Atomicity. In *Proceedings of Types in Language Design and Implementation*, January 2005.
- [Fra03] Keir Fraser. *Practical Lock-Freedom*. PhD thesis, University of Cambridge, September 2003.
- [GBD<sup>+</sup>94] Al Geist, Adam Beguelin, Jack Dongarra, Weicheng Jiang, Robert Manchek, and Vaidy Sunderam. *PVM: Parallel Virtual Machine—A Users’ Guide and Tutorial for Networked Parallel Computing*. MIT Press, 1994.

- [Gla94] David S Gladstein. *Compiler Correctness for Concurrent Languages*. PhD Thesis, Northeastern University, Massachusetts, December 1994.
- [GLS99] William Gropp, Ewing Lusk, and Anthony Skjellum. *Using MPI—Portable Parallel Programming with the Message Passing Interface*. MIT Press, second edition edition, 1999.
- [Goe06] Brian Goetz. Optimistic Thread Concurrency. Whitepaper, Azul Systems, Inc., January 2006.
- [Gor92] Andrew Donald Gordon. *Functional Programming and Input/Output*. PhD Thesis, University of Cambridge, 1992.
- [GR07] Andy Gill and Colin Runciman. Haskell Program Coverage. In *Haskell Workshop Proceedings*, September 2007.
- [Gra81] Jim Gray. The Transaction Concept: Virtues and Limitations. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, 1981.
- [Gro07] Dan Grossman. The Transactional Memory / Garbage Collection Analogy. In *OOPSLA Proceedings*, pages 695–706, October 2007.
- [GW96] David S Gladstein and Mitchell Wand. Compiler Correctness for Concurrent Languages. In *Proceedings of Coordination*, volume 1061 of *Lecture Notes in Computer Science*. Springer, April 1996.
- [HF03] Tim Harris and Keir Fraser. Language Support for Lightweight Transactions. In *OOPSLA Proceedings*, October 2003.
- [HH08] Liyang HU and Graham Hutton. Towards a Verified Implementation of Software Transactional Memory. In *Proceedings of Trends in Functional Programming*, May 2008.

## BIBLIOGRAPHY

- [HH09] Liyang HU and Graham Hutton. Compiling Concurrency Correctly: Cutting Out the Middle Man. In *Trends in Functional Programming*, 2009.
- [HLM06] Maurice P Herlihy, Victor Luchangco, and Mark Moir. A Flexible Framework for Implementing Software Transactional Memory. In *Proceedings of OOPSLA*, 2006.
- [HM93] Maurice P Herlihy and J Eliot B Moss. Transactional Memory: Architectural Support for Lock-Free Data Structures. In *International Symposium on Computer Architecture Proceedings*, 1993.
- [HMPJH05] Tim Harris, Simon Marlow, Simon Peyton Jones, and Maurice P Herlihy. Composable Memory Transactions. In *Proceedings of Principles and Practice of Parallel Programming*, June 2005.
- [Hoa02] Charles Antony Richard Hoare. Towards a Theory of Parallel Programming (*first published 1971*). In *The Origin of Concurrent Programming: from Semaphores to Remote Procedure Calls*, pages 231–244. Springer-Verlag, 2002.
- [HPJ06] Tim Harris and Simon Peyton Jones. Transactional Memory with Data Invariants. In *TRANSACT Proceedings*, March 2006.
- [Hut07] Graham Hutton. *Programming in Haskell*. Cambridge University Press, January 2007.
- [HW04] Graham Hutton and Joel Wright. Compiling Exceptions Correctly. In *Proceedings of International Conference on Mathematics of Program Construction*, number 3125 in Lecture Notes in Computer Science. Springer, July 2004.



- [HW06] Graham Hutton and Joel Wright. Calculating an Exceptional Machine. In *Proceedings of Trends in Functional Programming*, volume 5, February 2006.
- [HW07] Graham Hutton and Joel Wright. What is the Meaning of These Constant Interruptions? *Journal of Functional Programming*, 17(6):777–792, November 2007.
- [JHB87] Eric H Jensen, Gary W Hagensen, and Jeffrey M Broughton. A New Approach to Exclusive Data Access in Shared Memory Multiprocessors. Technical Report UCRL-97663, Lawrence Livermore National Laboratory, November 1987.
- [KLS90] Henry F Korth, Eliezer Levy, and Abraham Silberschatz. A Formal Approach to Recovery by Compensating Transactions. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, 1990.
- [KSF10] Guy Korland, Nir Shavit, and Pascal Felber. Noninvasive Concurrency with Java STM. In *Programmability Issues for Multi-Core Computers (MULTIPROG)*, January 2010.
- [Lam74] Leslie Lamport. A New Solution of Dijkstra’s Concurrent Programming Problem. *Communications of the ACM*, 17(8):453–455, August 1974.
- [Ler06] Xavier Leroy. Formal Certification of a Compiler Back-End, or: Programming a Compiler with a Proof Assistant. In *Proceedings of Principles of Programming Languages*, volume 33, pages 42–54, 2006.
- [LS08] Anand Lal Shimpi. Intel’s Atom Architecture: The Journey Begins. Retrieved from <http://www.anandtech.com/printarticle.aspx?i=3276>, page 14., April 2008.

## BIBLIOGRAPHY

- [M<sup>+</sup>08] Conor McBride et al. The Epigram System. Available from <http://www.e-pig.org/>, 2008.
- [Mar10] Simon Marlow, editor. *Haskell 2010 Language Report*. Online, 2010.
- [McB02] Conor McBride. Faking It (Simulating Dependent Types in Haskell). *Journal of Functional Programming*, 12(4–5):375–392, July 2002.
- [McB05] Conor McBride. Epigram: Practical Programming with Dependent Types. In *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [McB07] Conor McBride. R\* is the new  $[\alpha]$ . Talk given at Fun in the Afternoon, York., November 2007.
- [Mil89] Robin Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice Hall, 1989.
- [ML80] Per Martin-Löf. Intuitionistic Type Theory. Lecture notes, 1980.
- [ML98] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, second edition edition, September 1998.
- [Mog89] Eugenio Moggi. Computational Lambda-Calculus and Monads. In *Proceedings of Logic in Computer Science*, pages 14–23. IEEE Computer Society Press, June 1989.
- [Moo65] Gordon Moore. Moore’s Law. Retrieved from <http://www.intel.com/technology/mooreslaw/>, 1965.
- [MP67] John McCarthy and James Painter. Correctness of a Compiler for Arithmetic Expressions. In *Proceedings of Symposia in Applied Mathematics*, volume 19. AMS, 1967.

- [Nodir0o0] Cyprien Noël. Extensible Software Transactional Memory. In *Proceedings of the Conference on Computer Science and Software Engineering*, 2010.
- [Nor07] Ulf Norell. *Towards a Practical Programming Language Based on Dependent Type Theory*. PhD thesis, Chalmers University of Technology, September 2007.
- [NPS90] Bengt Nordström, Kent Petersson, and Jan M Smith. *Programming in Martin-Löf's Type Theory*. Oxford University Press, 1990.
- [PJ01] Simon Peyton Jones. Tackling the Awkward Squad: Monadic Input/Output, Concurrency, Exceptions, and Foreign-Language Calls in Haskell. In *Engineering Theories of Software Construction*, pages 47–96. IOS Press, 2001.
- [PJ03a] Simon Peyton Jones, editor. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [PJ03b] Simon Peyton Jones. Wearing the Hair Shirt: A Retrospective on Haskell. (Invited talk.). In *Principles of Programming Languages*, 2003.
- [PJGF96] Simon Peyton Jones, Andrew Gordon, and Sigbjorn Finne. Concurrent Haskell. In *Proceedings of Principles of Programming Languages*, pages 295–308, 1996.
- [PJWW04] Simon Peyton Jones, George Washburn, and Stephanie Weirich. Wobbly Types: Type Inference for Generalised Algebraic Data Types. Technical Report MS-CIS-05-26, University of Pennsylvania, 2004.
- [Sab98] Amr Sabry. What is a Purely Functional Programming Language? *Journal of Functional Programming*, 8:1–22, 1998.

## BIBLIOGRAPHY

- [San09] Davide Sangiorgi. On the Origins of Bisimulation and Coinduction. *ACM Transactions on Programming Languages and Systems*, 31(4), May 2009.
- [SJ07] Don Stewart and Spencer Janssen. XMonad: A Tiling Window Manager. In *Haskell Workshop*, September 2007.
- [ST97] Nir Shavit and Dan Touitou. Software Transactional Memory. *Distributed Computing*, 10(2):99–116, February 1997.
- [Sta10] Richard Matthew Stallman. GNU Emacs. Available from <http://www.gnu.org/software/emacs/>, 2010.
- [Swi08] Wouter S Swierstra. *A Functional Specification of Effects*. PhD thesis, University of Nottingham, November 2008.
- [Tan05] Audrey Tang. Pugs: an implementation of Perl 6 in Haskell, February 2005.
- [The08] The Coq Development Team. The Coq Proof Assistant. Available from <http://coq.inria.fr/>, June 2008.
- [The10] The Agda Team. The Agda Wiki. Available from <http://wiki.portal.chalmers.se/agda/>, September 2010.
- [THLPJ98] Philip W Trinder, Kevin Hammond, Hans-Wolfgang Loidl, and Simon Peyton Jones. Algorithm + Strategy = Parallelism. *Journal of Functional Programming*, 8(1):23–60, January 1998.
- [Wad92] Philip Wadler. Comprehending Monads. *Mathematical Structures in Computer Science*, 2:461–493, 1992.
- [Wan95] Mitchell Wand. Compiler Correctness for Parallel Languages. In *Proceedings of Functional Programming Languages and Computer Architecture*, pages 120–134, June 1995.

## BIBLIOGRAPHY

- [WB89] Philip Wadler and Stephen Blott. How to Make Ad-Hoc Polymorphism Less Ad-Hoc. In *Proceedings of Principles of Programming Languages*, January 1989.
- [WS95] Mitchell Wand and Gregory T Sullivan. A Little Goes a Long Way: A Simple Tool to Support Denotational Compiler-Correctness Proofs. Technical Report NU-CCS-95-19, Northeastern University College of Computer Science, October 1995.