

Imperial College London
Department of Computing

Reasoning with Time and Data Abstractions

Pedro da Rocha Pinto

September 2016

Supervised by Philippa Gardner

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Computing of Imperial College London
and the Diploma of Imperial College London

Declaration

I herewith certify that all material in this thesis which is not my own work has been properly acknowledged.

Pedro da Rocha Pinto

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

In this thesis, we address the problem of verifying the functional correctness of concurrent programs, with emphasis on fine-grained concurrent data structures. Reasoning about such programs is challenging since data can be concurrently accessed by multiple threads: the reasoning must account for the interference between threads, which is often subtle. To reason about interference, concurrent operations should either be at distinct times or on distinct data.

We present TaDA, a sound program logic for verifying clients and implementations that use abstract specifications that incorporate both *abstract atomicity*—the abstraction that operations take effect at a single, discrete instant in time—and *abstract disjointness*—the abstraction that operations act on distinct data resources. Our key contribution is the introduction of *atomic triples*, which offer an expressive approach for specifying program modules.

We also present Total-TaDA, a sound extension of TaDA with which we can verify total correctness of concurrent programs, i.e. that such programs both produce the correct result and terminate. With Total-TaDA, we can specify constraints on a thread’s concurrent environment that are necessary to guarantee termination. This allows us to verify total correctness for nonblocking algorithms and express lock- and wait-freedom. More generally, the abstract specifications can express that one operation cannot impede the progress of another, a new non-blocking property that we call *non-impedance*.

Finally, we describe how to extend TaDA for proving abstract atomicity for data structures that make use of *helping*—where one thread is performing an abstract operation on behalf of another—and *speculation*—where an abstract operation is determined by future behaviour.

To the people who are still alive.

Contents

1	Introduction	19
1.1	Contributions and Thesis Outline	19
2	Background	22
2.1	Concurrent Modules	22
2.1.1	A Spin Counter Implementation	22
2.1.2	A Ticket Lock Client	23
2.2	Proof Outlines	24
2.3	Sequential Specification	26
2.4	Auxiliary State	27
2.5	Interference Abstraction	29
2.6	Resource Ownership	30
2.7	Atomicity	33
3	Informal Development	35
3.1	Spin Counter	36
3.1.1	Atomic Specification	36
3.1.2	Implementation	38
3.2	Ticket Lock	40
3.2.1	CAP Specification	40
3.2.2	Implementation	42
4	TaDA Logic	46
4.1	Programming Language	46
4.2	Operational Semantics	47
4.3	Assertion Language	53
4.4	Program Logic	54
4.5	Model	61
4.5.1	Semantic Assertions	66
4.5.2	Semantic Judgements	68
4.6	Soundness	71
5	Using TaDA	90
5.1	Spin Lock	90
5.1.1	Atomic Specification	90
5.1.2	CAP Specification	92
5.1.3	Implementation	95

5.2	Multiple Compare-and-Set (MCAS)	96
5.2.1	Atomic Specification	96
5.2.2	Implementation	97
5.2.3	Resource Transfer	99
5.3	Deque	101
5.3.1	Atomic specification	101
5.3.2	Implementation	102
6	Related Work	107
7	Reasoning about Termination	108
7.1	Motivating Examples	110
7.1.1	Atomic Specification	111
7.1.2	Clients	112
7.1.3	Implementations	115
7.2	Logic	118
7.3	Case Study: Treiber's Stack	120
7.3.1	Atomic Specification	121
7.3.2	Implementation	121
7.4	Semantics and Soundness	123
7.5	Non-blocking Properties	127
7.5.1	Lock-freedom	127
7.5.2	Wait-freedom	129
7.5.3	Non-impedance	131
7.5.4	Related Work	132
8	Extending the logic	133
8.1	Motivating Examples	134
8.1.1	Ticket Lock	134
8.1.2	Two-step Counter	138
8.2	Case Study: Michael-Scott Queue	141
8.2.1	Atomic Specification	143
8.2.2	Implementation	143
9	Conclusions	149
9.1	Future Work	149
9.1.1	Tool Support	149
9.1.2	Helping/Speculation	149
9.1.3	Higher-order Support	150
9.1.4	Weak Memory	150
9.1.5	Liveness	150
9.1.6	Fault Tolerance	150

List of Figures

2.1	A counter module.	23
2.2	A ticket lock implemented using the counter module.	24
2.3	Reasoning about concurrent increments using auxiliary state.	28
2.4	Reasoning about concurrent increments using interference abstraction.	29
2.5	Ownership-based reasoning for concurrent increments.	33
3.1	Proof outline for the read operation.	40
3.2	Proof outline for the incr operation.	41
3.3	Proof outline for the wkincr operation.	42
3.4	A ticket lock implementation using the counter module.	42
3.5	Proof outline for the release operation.	44
3.6	Proof outline for the acquire operation.	45
4.1	Programming language proof rules of the TaDA logic.	56
4.2	Atomicity proof rules of the TaDA logic.	57
4.3	Auxiliary proof rules of the TaDA logic.	59
5.1	Derivation of the release specification.	94
5.2	Derivation of the acquire specification.	94
5.3	Spin lock operations.	94
5.4	Proof outline for the acquire operation.	96
5.5	The abstract specification for the MCAS module.	97
5.6	Multiple compare-and-set module operations.	98
5.7	Proof of the dcas implementation.	100
5.8	Snark deque operations.	103
5.9	Examples of a deque before and after performing popLeft , which uses 3cas to updated pointers <i>c</i> , <i>d</i> and <i>e</i>	104
5.10	Proof of the popLeft implementation.	106
7.1	Spin counter operations.	111
7.2	Proof of a sequential client of the counter.	113
7.3	Proof of a concurrent client of the counter.	115
7.4	Proof of total correctness of increment.	117
7.5	Backoff increment.	118
7.6	Proof of total correctness of backoff increment.	119
7.7	Stack operation specifications.	120
7.8	Treiber’s stack operations.	122

7.9	Proof of total correctness of Treiber’s stack pop operation.	124
7.10	Proof of total correctness of Treiber’s stack push operation.	125
8.1	Axioms for proxy modalities.	137
8.2	Auxiliary predicate describing the shared resources of the ticket lock.	138
8.3	Proof of correctness of the acquire operation.	139
8.4	Proof of correctness of the release operation.	140
8.5	Auxiliary predicate describing the set of configurations for the two-step counter. . . .	140
8.6	Proof of correctness of the readVal operation.	142
8.7	Proof of correctness of the incrVal operation.	142
8.8	Michael-Scott Queue implementation.	144
8.9	Auxiliary predicates to describe the shared resources of the Michael-Scott queue. . . .	146
8.10	Proxies and witnesses for the Michael-Scott queue.	146
8.11	Proof of correctness of dequeue operation.	148

List of Theorems

2.1	Remark (Conventions)	25
2.2	Remark (On disjoint resources)	31
3.1	Remark (On the abstractly-typed parameters)	37
4.1	Definition (Variable Names)	46
4.2	Definition (Expressions)	46
4.3	Definition (Boolean Expressions)	46
4.4	Definition (Function Names)	47
4.5	Definition (Commands)	47
4.6	Definition (Function Bodies)	47
4.7	Definition (Programs)	47
4.8	Definition (Program Values)	47
4.9	Definition (Variable Store)	47
4.10	Definition (Semantics of Expressions)	48
4.11	Definition (Semantics of Boolean Expressions)	48
4.12	Definition (Extended Commands)	48
4.13	Definition (Continuations)	48
4.14	Definition (Threads)	49
4.15	Definition (Labels of Atomic Commands)	49
4.16	Definition (Function Environment)	49
4.17	Definition (Thread Operational Semantics)	49
4.18	Definition (Extend Commands Operational Semantics)	49
4.19	Definition (Thread Identifiers)	51
4.20	Definition (Thread Pools)	51
4.21	Definition (Heaps)	51
4.22	Definition (Program States)	51
4.23	Definition (Thread Pool Operational Semantics)	51
4.24	Definition (Addresses)	52
4.25	Definition (Interpretation of Atomic Commands)	52
4.26	Definition (Program Operational Semantics)	53
4.27	Definition (Logical Expressions)	53
4.28	Definition (Assertions)	53
4.29	Definition (Function Specification Context)	54
4.30	Definition (Atomic Judgements)	55
4.31	Definition (Proof System)	55
4.32	Definition (Program Variables Sets)	58

4.33 Definition (Modified Sets)	58
4.34 Definition (Variable Names)	61
4.35 Definition (Guards and Guard Algebras)	61
4.36 Definition (Abstract States and Transition Systems)	62
4.37 Definition (Abstract Region Types)	62
4.38 Definition (Abstract Predicates)	62
4.39 Definition (Levels)	62
4.1 Remark (On levels)	62
4.40 Definition (Region Identifiers)	62
4.41 Definition (Region Assignments)	62
4.42 Definition (Guard Assignments)	63
4.43 Definition (Region States)	63
4.44 Definition (Worlds)	63
4.45 Definition (World Predicates)	63
4.2 Remark (Intuitionistic Interpretation)	64
4.46 Definition (Worlds with Atomic Tracking)	64
4.47 Definition (Atomicity Context)	64
4.48 Definition (Rely Relation)	64
4.49 Definition (Stable Predicates)	65
4.50 Definition (Region Interpretation)	65
4.51 Definition (Abstract Predicate Interpretation)	65
4.52 Definition (Region Collapse)	65
4.53 Definition (Abstract Predicate Collapse)	66
4.3 Remark (On the interpretation of abstract predicates)	66
4.54 Definition (Reification)	66
4.55 Definition (Guarantee Relation)	66
4.4 Remark (On reification)	66
4.56 Definition (Variable Interpretations)	67
4.57 Definition (Logical Expressions Semantics)	67
4.58 Definition (Assertion Semantics)	67
4.59 Definition (Primitive Atomic Satisfaction Judgement)	68
4.60 Definition (View Shift)	68
4.61 Definition (Semantic Judgement)	68
4.62 Definition (Function Environment and Function Context Agreement)	70
4.63 Definition (Semantic Judgement with Function Context)	70
4.1 Theorem (Soundness of Commands)	71
4.2 Lemma (SKIP Rule)	71
4.3 Lemma (SEQUENCING Rule)	71
4.4 Lemma (LOOP Rule)	74
4.5 Lemma (CONDITIONAL Rule)	75
4.6 Lemma (Stack Frame)	75
4.7 Lemma (FUNCTIONCALL Rule)	77

4.8	Lemma (FORK Rule)	78
4.9	Lemma (ALLOCATION Rule)	79
4.10	Lemma (ASSIGNMENT Rule)	79
4.11	Lemma (LOOKUP Rule)	79
4.12	Lemma (MUTATION Rule)	80
4.13	Lemma (COMPAREANDSET Rule)	80
4.14	Lemma (OPENREGION Rule)	80
4.15	Lemma (USEATOMIC Rule)	82
4.16	Lemma (UPDATEREGION Rule)	83
4.17	Lemma (Drop Context)	86
4.18	Lemma (MAKEATOMIC Rule)	86
4.19	Lemma (FRAME Rule)	88
4.20	Lemma (SUBSTITUTION Rule)	88
4.21	Lemma (CONSEQUENCE Rule)	88
4.22	Lemma (AEXISTS Rule)	88
4.23	Lemma (AWEAKENING1 Rule)	88
4.24	Lemma (AWEAKENING2 Rule)	89
4.25	Lemma (AWEAKENING3 Rule)	89
5.1	Remark (On alternative implementations)	96
7.1	Definition (Semantic Judgement)	123
7.1	Lemma (LOOP Rule)	126
7.2	Definition ((m, n) -bounded General Client)	128
7.2	Theorem (Hoffmann <i>et al.</i> [30])	128
7.3	Theorem (Lock-freedom)	128
7.4	Theorem (Wait-freedom)	129
7.5	Theorem (Wait-free Operation)	130

Acknowledgements

Foremost, I am grateful to my supervisor, Philippa Gardner, for her continued encouragement, enthusiasm and guidance during this research. I am equally thankful to Thomas Dinsdale-Young for his constant support, mentoring and patience. I am also indebted to Gian Ntzik for all the discussions throughout the years that have vastly improved my work.

I would like to thank my amazing collaborators — Kristoffer Just Andersen, Lars Birkedal, Mike Dodds, Julian Sutherland, Mark Wheelhouse, Shale Xiong — who have been wonderful to work with. As well as my colleagues at Imperial who make up the stimulating and friendly environment in which I have been privileged to work.

Finally, I would like to thank my family and my friends for the unconditional support and encouragement to pursue my interests and believing in me. In particular, my wonderful wife, Anita, whose love and support made this possible.

1 Introduction

The specification and verification of concurrent program modules is a difficult problem. When concurrent threads work with shared data, the resulting behaviour can be complex. Two abstractions provide useful simplifications: *atomicity*, stating that operations effectively act at distinct times; and *disjointness*, stating that operations effectively act on disjoint resources. While programmers routinely work with sophisticated combinations of these two abstractions, existing reasoning techniques, in contrast, tend to have limitations with respect to one or the other.

Two of the most widely used and influential techniques related to these two abstractions are *linearisability* [26], a correctness condition that specified that the operations of a concurrent module appear to behave atomically; and Concurrent Abstract Predicates (CAP) [12], an approach that offers support for reasoning about abstract disjoint resources and for specification of program modules.

With linearisability, each operation is given a sequential specification, and operations are asserted to behave atomically *with respect to each other*. Therefore, linearisability is a whole-module property: if we extend a module with an operation, we have to redo the linearisability proof to check that this new operation respects the atomicity of the others, and vice versa. Moreover, all operations are required to be atomic, so we can not specify non-atomic operations.

The CAP approach is based on concurrent separation logic, which provide concurrent reasoning based on the fundamental principle that resources are disjoint. Using CAP, it is possible to reason about shared resources as if they were disjoint, allowing for sophisticated concurrent implementations to be verified against simple specifications. However, in CAP, shared resources can only be accessed using *primitive atomic* operations, and operations provided by concurrent modules are rarely primitive atomic. Consequently, the abstract resources provided by a module are not easily shared and the nesting of modules is difficult.

Linearisability and CAP have complementary virtues and weaknesses. Linearisability gives strong, whole-module specifications based on abstract atomicity; CAP gives weaker, independent specifications based on abstract disjointness. Linearisability supports nested modules, but whole-module specifications make it difficult to extend modules; CAP supports the extension of modules, but the weak specifications make building up nested modules more difficult. Linearisability does not constrain the client, thus placing significant burden on the implementation; CAP constrains the client to use specific disjoint resources, enabling more flexibility in the implementation.

We propose a solution that combines the virtues of both approaches.

1.1 Contributions and Thesis Outline

Our main contribution is TaDA, a program logic for Time and Data Abstraction, which extends Concurrent Abstract Predicates [12] with rules for deriving and using atomic triples. We show how the logic supports vertical reasoning about modules by verifying implementations that are built on the

top of each other. We thus demonstrate that TaDA combines the benefits of abstract atomicity and abstract disjointness within a single program logic.

Moreover, we show how to extend TaDA to prove termination of non-blocking programs. We name this extension Total-TaDA and use it to prove several examples with increasing difficulty. We provide and prove relations between our specifications and lock-freedom and wait-freedom for modules and operations. We also present a new non-blocking property, called non-impedance, which enables us to describe that one operation cannot prevent the progress of another.

The last part of the thesis proposes a possible extension to the original logic that focusses on algorithms that exhibit helping and speculation with respect to an atomic specification. We propose a way to adapt TaDA to handle these algorithms and apply it to the ticket lock, which exhibits helping, a counter that requires speculation, and a Michael-Scott queue, which uses speculation.

The outline of the thesis is as follows:

- Chapter 2 provides the technical background that motivates the rest of the thesis. We start by introducing two concurrent modules — a counter and a ticket lock — and describing some of the challenges with reasoning about them. We give a brief introduction to Hoare reasoning, introduces the proof rules and how they can be used to reason about programs. We then survey a range of verification techniques for specifying concurrent modules, highlighting four key concepts in particular: auxiliary state, interference abstraction, resource ownership and atomicity.
- In chapter 3, we present an informal development of the TaDA logic using the counter and the ticket lock examples previously introduced. We show how our approach can be used to give abstract specifications, using atomic triples, which allow for modular reasoning and overcome the problems described throughout chapter 2.
- We formalise the programming language in chapter 4. We define the syntax and semantics of the TaDA logic and prove the soundness of the logic.
- In chapter 5, we apply the logic to a series of modules that are built on the top of each other. We show how the logic supports vertical reasoning about modules, by verifying an implementation of multiple-compare-and-set (MCAS) using the lock specification, and an implementation of a concurrent double-ended queue (deque) using the MCAS specification. We show that our approach is modular and that it can verify the operations of a module independently and provide abstract specifications.
- In chapter 6, we discuss related work and its advantages and disadvantages with respect to our work.
- In chapter 7, we present Total-TaDA, a program logic that extends TaDA with support for verifying the total correctness of non-blocking concurrent programs. We verify the total correctness of non-blocking algorithms, such as counters and Treiber's stack. We relate our specifications to the standard definitions of lock- and wait-freedom and propose a new non-blocking property we call non-impedance. We show that our extension is sound by adapting the semantics used for TaDA.

- In chapter 8, we propose a technique to reason about advanced concurrent algorithms that make use of helping and speculation. We apply the technique to the ticket lock, a novel two-step counter, and a Michael-Scott queue.
- Finally, in chapter 9 we summarise the contributions and raise several research questions.

Collaboration

Chapter 2 is heavily inspired by work done in collaboration with Dinsdale-Young and Gardner. Much of it was previously published in *Steps in Modular Specifications for Concurrent Modules (Invited Tutorial Paper)* [9]. Chapter 4 and chapter 5 are based on joint work with Dinsdale-Young and Gardner and part of it was published in *TaDA: A Logic for Time and Data Abstraction* [8]. The termination work in chapter 7 was undertaken with Dinsdale-Young, Gardner and Sutherland, and published in *Modular Termination Verification for Non-blocking Concurrency* [10]. Further examples of modules verified using TaDA are available in *Abstract Specifications for Concurrent Maps* [64].

2 Background

This chapter introduces the terminology and notation used in the thesis and provides a systematic overview of a range of techniques for specifying and verifying concurrent modules that introduced the concepts that this work is built upon.

The specification and verification of a concurrent program are difficult problems. When concurrent threads work with shared data, the resulting behaviour can be complicated. In order to specify such modules, we require effective abstractions for capturing such complex behaviour. We will describe the main concepts related to specifying and verifying concurrent programs that have emerged over the last few decades. In particular, we restrict our exposition to those concepts on which this work is based: auxiliary state, interference abstraction, resource ownership and atomicity.

We highlight the challenges of specifying a concurrent module by using the counter module as an illustrative example in §2.1. We require a specification to be expressive enough for verifying the intended clients of the module. We also require the specification to be opaque, in that the implementation details do not leak into the specification.

Using the counter module, we present a range of historical verification techniques for concurrency:

- Owicki-Gries [49] introduces *auxiliary state* to abstract the internal state of threads (§2.4);
- rely/guarantee [34] introduces *interference abstraction* to abstract the interactions between different threads (§2.5);
- concurrent separation logic [47] introduces *resource ownership* to encode interference abstraction as auxiliary state (§2.6);
- linearisability [26] introduces *atomicity* as a way to abstract the effects of an operation (§2.7).

Later, we show in §3 how to combine these techniques to provide expressive ways for specifying concurrent modules.

2.1 Concurrent Modules

2.1.1 A Spin Counter Implementation

We consider the implementation of a concurrent counter shown in Figure 2.1. We make use of three *atomic* operations (operations that take effect at a single, discrete instant in time) that manipulate the heap. The operation $x := [\mathbb{E}]$; reads the value of the heap position $\mathbb{E}$ and assigns it to the variable x . The operation $[\mathbb{E}_1] := \mathbb{E}_2$; stores the value $\mathbb{E}_2$ in the heap position $\mathbb{E}_1$. Finally, the compare-and-set (CAS) operation $x := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3)$; checks if the value at heap position $\mathbb{E}_1$ is equal to $\mathbb{E}_2$: if so, it replaces it with the value $\mathbb{E}_3$ and assigns 1 to x ; otherwise, x is assigned 0.

Here, the **read** operation returns the value of the counter. The **incr** operation increments the value of the counter and returns the old value, it is robust to other threads incrementing the counter.

```

function makeCounter() {
  x := alloc(1);           // Allocate a single cell
  [x] := 0;                 // Initialise the value at x with 0
  return x;
}

function read(x) {
  v := [x];                // Get value at address x
  return v;
}

function incr(x) {
  do {
    v := [x];              // Get value at x
    b := CAS(x, v, v + 1); // Compare value at x with v and
                          // set it to v + 1 if they are the same
  } while (b = 0);         // Retry if the CAS failed
  return v;
}

function wkIncr(x) {
  v := [x];                // Get value at address x
  [x] := v + 1;            // Set value at x with v + 1
  return v;
}

```

Figure 2.1: A counter module.

This is done by using the compare-and-set to atomically update the value of the counter. The `wkincr` just increments the value of the counter assuming that no other thread is performing any increment concurrently.

A specification should describe how each operation affects the value of the counter. A specification should require the counter to exist as a precondition for each operation, since operations will not work unless the memory holding the counter is allocated. A specification should also describe the permitted interference from the context of concurrent operations. Intuitively, the `read` and `incr` operations are robust with respect to concurrent operations that change the value of the counter. By contrast, the potentially faster `wkIncr` requires that there is no concurrent operation that changes the value of the counter between the read and the increment of the value in order to behave as intended.

2.1.2 A Ticket Lock Client

Let us now consider the ticket lock [42] module that uses the counter module to provide synchronisation. The code for the lock is given in Figure 2.2. The lock uses two counters, the ticket counter `next` and the serving counter `owner`, which both initially have the value 0. A thread acquires the lock by calling the `acquire` operation. This operation increments the `next` counter to obtain a notional ticket. When the value of the `owner` counter agrees with this ticket, the thread has acquired the lock. It can then use whatever resources are protected by the lock, without interference from other threads. Control of these resources is relinquished by calling the `release` operation. This increments the `owner` counter, passing the lock on to the next waiting thread. Intuitively, the use of `incr` for the `acquire` operation

```

function makeLock() {
  next := makeCounter(); // Allocate ticket counter
  owner := makeCounter(); // Allocate serving counter
  x := alloc(2); // Allocate two cells
  [x.next] := next; // Initialise first cell with ticket counter address
  [x.owner] := owner; // Initialise second cell with serving counter address
  return x;
}

function acquire(x) {
  next := [x.next]; // Get address of the ticket counter
  owner := [x.owner]; // Get address of the serving counter
  t := incr(next); // Take a ticket
  do {
    v := read(owner); // Get serving counter value
  } while (v ≠ t); // Wait for serving value to match the ticket
}

function release(x) {
  owner := [x.owner]; // Get address of the serving counter
  wkIncr(owner); // Increment the serving counter
}

```

where $\mathbb{E}.\text{next} \stackrel{\text{def}}{=} \mathbb{E}$ $\mathbb{E}.\text{owner} \stackrel{\text{def}}{=} \mathbb{E} + 1$.

Figure 2.2: A ticket lock implemented using the counter module.

is necessary, since it needs to be robust with respect to concurrent threads taking tickets. The use of `wkIncr` for the `release` operation is possible, since only the thread holding the lock should release it.

The challenge is to develop a concurrent specification of the counter module that would be strong enough to allow us to reason about the ticket lock. This mandates a precise description of how each operation affects the value of the counter, and a detailed account of interference to capture the fundamental distinction between `incr` and `wkIncr`.

The counter and its ticket lock client are realistic examples that illustrate some of the key difficulties in specifying and reasoning about concurrent modules. We are looking for a formal specification technique that would allow us to express such specifications and formally verify both their implementations as well as programs that use such specifications.

2.2 Proof Outlines

Before examining any techniques, we need to introduce some notation for formally proving properties of programs. Tony Hoare introduced Hoare triples [29] of the form $\{P\} \mathbb{C} \{Q\}$ with the main goal of specifying the behaviour of sequential programs. A triple associates the *predicates* P and Q with a program $\mathbb{C}$, where P is the precondition and Q is the postcondition of the program. The meaning of the triple is that if a program $\mathbb{C}$ is run in a state described by P then it will not fault, and if it terminates, the resulting state will be described by Q .

Hoare also introduced a formal system with a set of rules for rigorous reasoning about imperative programs, which included the following proof rules:

$$\begin{array}{c}
\text{SKIP} \\
\hline
\vdash \{P\} \text{ skip}; \{P\} \\
\\
\text{SEQUENCING} \\
\hline
\frac{\vdash \{P\} \mathbb{C}_1 \{R\} \quad \vdash \{R\} \mathbb{C}_2 \{Q\}}{\vdash \{P\} \mathbb{C}_1 \mathbb{C}_2 \{Q\}} \\
\\
\text{LOOP} \\
\hline
\frac{\vdash \{P \wedge \mathbb{B}\} \mathbb{C} \{P\}}{\vdash \{P\} \text{ while } (\mathbb{B}) \{ \mathbb{C} \} \{P \wedge \neg \mathbb{B}\}} \\
\\
\text{ASSIGNMENT} \\
\hline
\vdash \{P[\mathbb{E}/x]\} x := \mathbb{E}; \{P\} \\
\\
\text{CONDITIONAL} \\
\hline
\frac{\vdash \{P \wedge \mathbb{B}\} \mathbb{C}_1 \{Q\} \quad \vdash \{P \wedge \neg \mathbb{B}\} \mathbb{C}_2 \{Q\}}{\vdash \{P\} \text{ if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \} \{Q\}} \\
\\
\text{CONSEQUENCE} \\
\hline
\frac{P' \vdash P \quad \vdash \{P\} \mathbb{C} \{Q\} \quad Q \vdash Q'}{\vdash \{P'\} \mathbb{C} \{Q'\}}
\end{array}$$

The SKIP rule asserts that the `skip`; does not change the state of the program.

The SEQUENCING rule allows us to compose programs that are executed sequentially.

The LOOP rule states that there is a loop invariant, P , which is preserved by the loop body $\mathbb{C}$. After the loop is finished, the invariant P still holds, and moreover $\mathbb{B}$ must be caused the loop to end.

The CONDITIONAL rule states that if both branches of the conditional establish Q , then the postcondition of the whole satisfies Q .

The ASSIGNMENT rule states that, after the assignment, any precondition that was true still holds for the variable. $P[\mathbb{E}/x]$ denotes the assertion P in which each free occurrence of x has been replaced by the expression $\mathbb{E}$.

The ASSIGNMENT rule states that, after the assignment, the variable x has value $\mathbb{E}$ where each occurrence of x has been replaced by v .

The CONSEQUENCE rule allows to strengthen the precondition and weaken the postcondition.

Remark 2.1 (Conventions). When doing proofs with Hoare triples, we generally do not mention the names of the rules, such as the rules that are used to reason about the constructs of the language. We implicitly use the rule of consequence and do not expand its steps. Moreover, when we explicitly apply rules we denote them with a vertical bar, where the name of the rule is given sideways on the left of the bar.

In order to illustrate how this reasoning system works in practice, we will start by showing a program which manipulates the store. Given a program $\mathbb{C}$ with the following specification:

$$\vdash \{x = v\} \mathbb{C} \{x = 2\},$$

by the precondition, we know that, before the execution of the program, x contains some value v . During the execution, the program is allowed to change x to any value, as long as when the program terminates, the value of x is set to 2.

We can show that $\mathbb{C} \stackrel{\text{def}}{=} \text{if } (x \neq 2) \{x := 2;\} \text{ else } \{\text{skip};\}$ satisfies the specification by constructing

the following derivation:

$$\begin{array}{c}
\{x = v\} \\
\text{if } (x \neq 2) \{ \\
\quad \{x = v \wedge x \neq 2\} \\
\quad \{x = v\} \\
\quad \text{CONSEQUENCE} \left| \begin{array}{l} \text{ASSIGNMENT} \\ \{x = v\} \\ x := 2; \\ \{x = 2\} \end{array} \right. \\
\quad \{x = 2\} \\
\quad \{x = 2\} \\
\quad \} \text{ else } \{ \\
\quad \{x = v \wedge x = 2\} \\
\quad \{x = 2\} \\
\quad \text{CONSEQUENCE} \left| \begin{array}{l} \text{SKIP} \\ \{x = 2\} \\ \text{skip;} \\ \{x = 2\} \end{array} \right. \\
\quad \{x = 2\} \\
\quad \{x = 2\} \\
\quad \} \\
\} \\
\{x = 2\}
\end{array}$$

The original proof system did not account for programs that manipulate the heap, such as the counter, or concurrency. We extend the first-order assertion language to heaps, using the syntax $x \mapsto v$ to assert that the heap, at address x , holds value v .

2.3 Sequential Specification

We can give a sequential specification for the counter module using Hoare triples:

$$\begin{array}{l}
\vdash \{ \text{True} \} \text{makeCounter}() \{ \text{ret} \mapsto 0 \} \\
\vdash \{ x \mapsto n \} \text{read}(x) \{ x \mapsto n \wedge \text{ret} = n \} \\
\vdash \{ x \mapsto n \} \text{incr}(x) \{ x \mapsto n + 1 \wedge \text{ret} = n \} \\
\vdash \{ x \mapsto n \} \text{wkIncr}(x) \{ x \mapsto n + 1 \wedge \text{ret} = n \}
\end{array}$$

With standard Hoare logic, we can use this specification to verify sequential clients that call the counter operations. However, this specification gives no information about the behaviour of the operations in a concurrent setting.

2.4 Auxiliary State

Owicki and Gries [49] developed the first tractable proof technique for concurrent programs, identifying the importance of reasoning about *interference* between threads and of using *auxiliary state*. With the Owicki-Gries method, each thread is given a sequential proof. When the threads are composed, we must check that they do not interfere with each other's proofs. This is achieved by extending standard Hoare logic with the Owicki-Gries rule for parallel composition:

$$\frac{\text{OG-PARALLEL} \quad \begin{array}{c} \vdash_{\text{OG}} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad \vdash_{\text{OG}} \{P_2\} \mathbb{C}_2 \{Q_2\} \quad \text{non-interference} \end{array}}{\vdash_{\text{OG}} \{P_1 \wedge P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 \wedge Q_2\}}$$

The **non-interference** side condition constrains the proof derivations for $\mathbb{C}_1$ and $\mathbb{C}_2$. It requires that every intermediate assertion between atomic actions in the proof of $\mathbb{C}_1$ must be preserved by every atomic action in the proof of $\mathbb{C}_2$, and vice-versa.

An abstract specification for the counter needs to be robust with respect to the **non-interference** condition. However, in general, the condition will vary depending on the concurrent context. Let us assume that the client may invoke any of the counter operations concurrently, but will not directly interact with the state of the counter. That is, we will only consider interference caused by the counter operations themselves. To this end, we can use an *invariant* — an assertion that is preserved by each atomic action in the module. For the counter, the invariant $\exists n. \mathbf{x} \mapsto n$ asserts that the counter at $\mathbf{x}$ is allocated and has some value.

We can give the following specification for the counter module:

$$\begin{array}{l} \vdash_{\text{OG}} \{\text{True}\} \text{makeCounter}() \{ \exists n. \text{ret} \mapsto n \} \\ \vdash_{\text{OG}} \{ \exists n. \mathbf{x} \mapsto n \} \text{read}(\mathbf{x}) \{ \exists n, m. \mathbf{x} \mapsto n \wedge \text{ret} = m \} \\ \vdash_{\text{OG}} \{ \exists n. \mathbf{x} \mapsto n \} \text{incr}(\mathbf{x}) \{ \exists n, m. \mathbf{x} \mapsto n \wedge \text{ret} = m \} \\ \vdash_{\text{OG}} \{ \exists n. \mathbf{x} \mapsto n \} \text{wkIncr}(\mathbf{x}) \{ \exists n, m. \mathbf{x} \mapsto n \wedge \text{ret} = m \} \end{array}$$

However, these specifications are too weak to specify clients such as the ticket lock. They lose all information about the value of the counter, and give no information about how the operations change this value. In fact, the **read** operation could change the value of the counter and still satisfy the specification! Unfortunately, assertions that describe the precise value of the counter are not invariant.

The Owicki-Gries method is able to provide stronger specifications by using *auxiliary state*, which records extra information about the execution history via auxiliary variables. The code is instrumented with *auxiliary code*, which updates the auxiliary variables. Since the auxiliary code only updates auxiliary variables, it has no effect on the program behaviour, and so can be erased — it is not required when the program is run.

By way of example, consider two threads that both increment a counter, as in Figure 2.3. The auxiliary variables **y** and **z**, with initial values 0, are used to record the contribution (*i.e.* the number of increments) of each thread. For each thread, the code of the **incr** operation is instrumented with code that updates the auxiliary variables when the **CAS** operations succeed. The auxiliary variables




Figure 2.3: Reasoning about concurrent increments using auxiliary state.

must be updated at the same instant as the counter, so that the counter always holds the sum of the two contributions — our invariant. This is expressed in the angle brackets, $\langle \_ \rangle$, which indicate that the **CAS** and auxiliary code should be executed in a single atomic step.

The resulting specification of the two-increment program is strong, with precise information about the initial and final value of the counter. However, it comes at the price of modularity. Firstly, each use of the **incr** operation requires the underlying implementation to be extended with auxiliary code to increment the appropriate auxiliary variable. A modular proof would not modify the module code for each use by the client. Secondly, the **incr** operations require different specifications depending on the client's use: in our example, the assertion $x \mapsto y + z$ uses auxiliary variables y and z ; with three threads, the specification requires three auxiliary variables. A modular proof would give a specification for the module that captures all use cases. Thirdly and more subtly, the Owicki-Gries method requires the global **non-interference** condition. To meet this, we made the implicit assumption that the client only interacts with the state of the counter through the counter operations. A modular proof would be explicit about such assumptions about the behaviour of the client.

Thesis

The concept of *auxiliary state*, introduced in the Owicki-Gries method, is important in specifying concurrent modules. Auxiliary state abstracts the internal state of threads. It is more convenient to reason using auxiliary variables than to consider the program counter and local variables of each thread in describing invariants. This abstraction is a step towards compositional reasoning. As we shall see, various subsequent approaches have taken a more modular approach to auxiliary state than that provided by auxiliary variables in the Owicki-Gries method.

$$\begin{array}{c}
\{\text{True}\} \\
\mathbf{x} := \text{makeCounter}(); \\
\{\mathbf{x} \mapsto 0\} \\
// \text{ Weaken assertion} \\
\{\exists n. \mathbf{x} \mapsto n \wedge n \geq 0\} \\
\left\{ \begin{array}{c} \{\exists n. \mathbf{x} \mapsto n \wedge n \geq 0\} \\ \text{incr}(\mathbf{x}); \\ \{\exists n. \mathbf{x} \mapsto n \wedge n \geq 1\} \end{array} \right\} \parallel \left\{ \begin{array}{c} \{\exists n. \mathbf{x} \mapsto n \wedge n \geq 0\} \\ \text{incr}(\mathbf{x}); \\ \{\exists n. \mathbf{x} \mapsto n \wedge n \geq 1\} \end{array} \right\} \\
\{\exists n. \mathbf{x} \mapsto n \wedge n \geq 1\}
\end{array}$$

Figure 2.4: Reasoning about concurrent increments using interference abstraction.

2.5 Interference Abstraction

Jones [34] introduced *interference abstraction*, providing the rely/guarantee method as a way to improve the compositionality of the Owicki-Gries approach. To avoid the global **non-interference** condition, specifications explicitly constrain the interference from the concurrent context, and describe the interference that a thread may cause. To this end, each specification incorporates two relations—the *rely* and *guarantee* relations—that abstract the interference between threads. The rely relation abstracts the actions of other threads; each assertion in the derivation must be *stable* under all of these actions. The guarantee relation abstracts the actions in the derivation; each atomic update by the thread must be described by the guarantee.

Rely/guarantee specifications have the form $R, G \vdash_{\text{RG}} \{P\} \mathbb{C} \{Q\}$, where R and G are the rely and guarantee relations respectively. We denote the elements of the rely and guarantee relations in terms of actions $P \rightsquigarrow Q$ that describe the changes performed. When composing concurrent threads, the guarantee of each thread must be included in the rely of the other. The parallel composition rule is therefore adapted to:

$$\frac{\text{RG-PARALLEL} \quad R \cup G_2, G_1 \vdash_{\text{RG}} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad R \cup G_1, G_2 \vdash_{\text{RG}} \{P_2\} \mathbb{C}_2 \{Q_2\}}{R, G_1 \cup G_2 \vdash_{\text{RG}} \{P_1 \wedge P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 \wedge Q_2\}}$$

The rely/guarantee specifications for the **read** and **incr** operations are:

$$\begin{array}{l}
A, \emptyset \vdash_{\text{RG}} \{\exists n. \mathbf{x} \mapsto n\} \text{read}(\mathbf{x}) \{\exists n. \mathbf{x} \mapsto n \wedge \text{ret} \leq n\} \\
A, A \vdash_{\text{RG}} \{\exists n. \mathbf{x} \mapsto n\} \text{incr}(\mathbf{x}) \{\exists n. \mathbf{x} \mapsto n + 1 \wedge \text{ret} \leq n\}
\end{array}$$

where $A = \{\mathbf{x} \mapsto n \rightsquigarrow \mathbf{x} \mapsto n + 1 \mid n \in \mathbb{N}\}$. The **read** specification has an empty guarantee relation indicating that nothing is changed by the read. It has the rely relation A , stating that other threads can only increment the counter, and that they can do so as many times as they like. The **incr** relation has the same rely relation. Its guarantee relation is also A , stating that the increment can increase the value of the counter. The guarantee must be defined for all n , as the context can change the counter value. This means that we cannot express that the **incr** operation only does a single increment.

The rely/guarantee specification for the `wkIncr` operation is subtle. Recall that, intuitively, the `wkIncr` operation is intended to be used when no other threads are concurrently updating the counter. As a first try, we can give a simple specification with a rely condition that enforces this constraint:

$$\emptyset, G \vdash_{\text{RG}} \{x \mapsto n\} \text{wkIncr}(x) \{x \mapsto n + 1 \wedge \text{ret} = n\}$$

where $G \triangleq \{x \mapsto n \rightsquigarrow x \mapsto n + 1\}$. The rely relation is empty, so this specification cannot be used in a context where concurrent updates may occur. This means that the guarantee relation can be very precise, consisting of a single action. Effectively, the increment will appear as a single atomic operation.

Although this specification captures some of the intended behaviour of `wkIncr`, it is insufficient to reason about the ticket lock. With the ticket lock, it is possible for two invocations of the `wkIncr` operation to be executing concurrently. Only one thread can call `release` at any one time, because only one thread can have the lock. However, suppose one thread calling `release` has executed the body of `wkIncr`. Then, a second thread may correctly conclude that it now has the lock and release it, before the call of the first thread has returned. This results in a concurrent invocation of `wkIncr`. By ruling out all concurrent updates to the counter with an empty rely relation, the above specification does not allow this concurrent behaviour.

By changing the rely, we can allow such concurrent updates, but that would weaken the specification:

$$R, G \vdash_{\text{RG}} \{x \mapsto n\} \text{wkIncr}(x) \{\exists n' \geq n + 1. x \mapsto n' \wedge \text{ret} = n\}$$

where $R = \{x \mapsto m \rightsquigarrow x \mapsto m + 1 \mid m > n\}$ and G is as before. Notice that the rely states that concurrent increments can only happen when the value of the counter is above n . Also notice that, in weakening the rely, we must weaken the postcondition to make it stable.

In summary, this specification is again too weak to reason about the ticket lock. It is possible to instrument the code with auxiliary variables, as with the Owicki-Gries method, but this would again lead to a loss of modularity.

Thesis

The concept of *interference abstraction*, introduced in the rely/guarantee method, is important in specifying concurrent modules. By abstracting the interactions between different threads, specifications can express constraints on their concurrent contexts. This abstraction leads to more compositional reasoning: since the interference is part of the specification, we do not need to examine proofs in order to justify parallel composition. While they may specify it differently, some form of interference abstraction is generally present in subsequent concurrency verification approaches.

2.6 Resource Ownership

In the Owicki-Gries and rely/guarantee approaches, auxiliary variables provide a mechanism for reasoning about which threads can do what and when. For instance, auxiliary variables can be used to reason about the contribution of individual threads to the counter, as we demonstrated in §2.4, or that one thread can increment a counter after another. O'Hearn introduced a style of reasoning based on

resource ownership, developing concurrent separation logic [47] which provides an alternative, more modular approach to such reasoning.

Concurrent separation logic is a Hoare logic, with assertions describing data (such as heap cells or counter objects) treated as resources. Each operation acts on specific resources, with the precondition conferring ownership of the resources it represents. When threads operate on disjoint resources, they do not interfere and so their effects can be simply combined. This principle is embodied in the disjoint parallel composition rule:

$$\frac{\text{PARALLEL} \quad \frac{\{P_1\} \mathbb{C}_1 \{Q_1\} \quad \{P_2\} \mathbb{C}_2 \{Q_2\}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}}$$

where $P_1 * P_2$ describes the disjoint combination of the resources of P_1 and P_2 .

Remark 2.2 (On disjoint resources). The separating conjunction $*$ is based on separation logic [31, 51], which concurrent separation logic builds upon. Its key idea was to treat heap memory as a resource, which can be subdivided into disjoint resources. Heap operations, such as updating the value of a heap cell, require parts of the heap for their execution. These operations are *local* with respect to the specific resource they operate, as they do not affect other parts of the heap. This idea is expressed in the following rule:

$$\frac{\text{FRAME} \quad \frac{\{P\} \mathbb{C} \{Q\}}{\{P * R\} \mathbb{C} \{Q * R\}}}{\{P * R\} \mathbb{C} \{Q * R\}}$$

The frame rule of separation logic allows us to reason about programs in a local way, i.e. we can focus our reasoning on the resource that the program manipulates, and any additional resource, which would not be affected by the program, can be added on. This is possible because $*$ enforces disjointness. $\square$

We can think of ownership as embodying specialised notions of auxiliary state and interference abstraction. Ownership is a form of auxiliary state and an abstraction that we use for reasoning: the program does not explicitly record which threads own what resources. Ownership implements a simple interference abstraction: threads may update the resources that they own, and disjointness of ownership enforces that they cannot interfere with the resources of other threads.

In the original concurrent separation logic, it was only possible to reason about shared resource that had been transferred between threads through synchronisation. In [47], conditional critical regions provide the synchronisation mechanism. Subsequent approaches [63, 18, 15] added support for reasoning about fine-grained concurrency by incorporating various styles of rely/guarantee reasoning over shared resources. Building on this work, the concurrent abstract predicates (CAP) [12] approach introduces abstractions over these shared resources that may be split, effectively allowing concurrent manipulation at the abstract level. Let us illustrate this on the example of a counter.

We use the abstract predicate $\text{Counter}(x, n)$ to denote the existence of a counter at memory location x with the value n . In the counter implementation we are reasoning with, the abstract predicate is instantiated by the implementation as $x \mapsto n$.

Treating the abstract predicate $\text{Counter}(x, n)$ as a resource, we could use the original sequential specification as a concurrent one. However, for multiple threads to use the counter, they would have to transfer the resource between each other using some form of synchronisation. Such a specification

effectively enforces sequential access to the counter. This is because the client has no mechanism for dividing the resource: in particular,

$$\text{Counter}(\mathbf{x}, n) \implies \text{Counter}(\mathbf{x}, n) * \text{Counter}(\mathbf{x}, n)$$

does not hold.

Following Boyland [4], Bornat *et al.* [2] introduced *permission accounting* to separation logic. This allows shared resources to be divided by associating with them a fraction in the interval $(0, 1]$. Shared resources may be subdivided by splitting this fraction. For instance, we may associate fractions with our counter resource and declare the logical axiom:

$$\text{Counter}(\mathbf{x}, n, \pi_1 + \pi_2) \iff \text{Counter}(\mathbf{x}, n, \pi_1) * \text{Counter}(\mathbf{x}, n, \pi_2)$$

for $\pi_1 + \pi_2 \leq 1$. We can now modify our counter specification to give concurrent read access:

$$\begin{aligned} & \left\{ \text{Counter}(\mathbf{x}, n, \pi) \right\} \text{read}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n, \pi) * \text{ret} = n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, 1) \right\} \text{incr}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n + 1, 1) * \text{ret} = n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, 1) \right\} \text{wkIncr}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n + 1, 1) * \text{ret} = n \right\} \end{aligned}$$

Notice that we require full permission (the 1) in order to perform either increment operation. This means that only concurrent reads are permitted; concurrent updates must be synchronised with all other concurrent accesses (both increments and reads). If only partial permission were necessary, then the specification for `read` would be incorrect, since it could no longer guarantee that the value being read matched the resource it had.

It is possible to specify concurrent increments by changing how we interpret the counter predicate $\text{Counter}(\mathbf{x}, n, \pi)$. Now, the resource $\text{Counter}(\mathbf{x}, n, \pi)$ no longer asserts that the value of the counter is n , except if $\pi = 1$. Instead, it asserts that the thread is contributing n to the value of the counter; other threads may also have contributions. We can split this counter resource by declaring the logical axiom:

$$\text{Counter}(\mathbf{x}, n_1 + n_2, \pi_1 + \pi_2) \iff \text{Counter}(\mathbf{x}, n_1, \pi_1) * \text{Counter}(\mathbf{x}, n_2, \pi_2)$$

for $n_1, n_2 \in \mathbb{N}$ and $\pi_1, \pi_2 \in (0, 1]$. We then specify our counter operations as:

$$\begin{aligned} & \left\{ \text{Counter}(\mathbf{x}, n, \pi) \right\} \text{read}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n, \pi) * \text{ret} \geq n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, 1) \right\} \text{read}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n, 1) * \text{ret} = n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, \pi) \right\} \text{incr}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n + 1, \pi) * \text{ret} \geq n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, 1) \right\} \text{incr}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n + 1, 1) * \text{ret} = n \right\} \\ & \left\{ \text{Counter}(\mathbf{x}, n, 1) \right\} \text{wkIncr}(\mathbf{x}) \left\{ \text{Counter}(\mathbf{x}, n + 1, 1) * \text{ret} = n \right\} \end{aligned}$$

At last, we have a specification that allows concurrent reads and increments.

Figure 2.5 shows how it can be used to verify the example of two concurrent increments. Whereas in Figure 2.3 each thread was instrumented with different auxiliary code, here the code has not been

changed. Rather than each thread having an auxiliary variable to record its contribution to the counter, the contribution is recorded in auxiliary resources that are owned by the thread and encapsulated in the $\text{Counter}(\mathbf{x}, n, \pi)$ predicate. This idea of subjective auxiliary state is at the core of subjective concurrent separation logic (SCSL) [37] (and the subsequent FCSL [44, 52]).

This specification still has some weaknesses. The wkIncr operation must still be synchronised with the other operations. Also, sequenced reads will never see decreasing values of the counter (since the contribution is not changed and only provides the lower bound). It is possible to describe a more elaborate permission system that allows wkIncr in the presence of reads, and to extend the predicate to record the last known value as a lower bound for reads. This would give us a more useful, if somewhat cumbersome, specification. However, it would still not handle the ticket lock.

While a ticket lock has been verified using CAP [12], the proof depends on the atomicity of the underlying counter operations in order to synchronise access to shared resources. The proof does not work with any of our abstract specifications, since they simply do not embody the necessary atomicity.

Thesis

The concept of *resource ownership*, developed by concurrent separation logic and its successors, is important in specifying concurrent modules. The idiom of ownership can be seen as a form of auxiliary state, which critically embodies a notion of disjointness and interference abstraction. Various approaches have explored the power of ownership for reasoning about concurrency [12, 37, 59, 44, 55, 54, 52, 35]. While it is an effective concept, and can be used to give elegant specifications, something more is required to provide the strong specifications we are seeking.

2.7 Atomicity

Atomicity is the abstraction that an operation takes effect at a single, discrete instant in time. The concurrent behaviour of such operations is equivalent to a sequential interleaving of the operations, and a client can use such operations as if they were simple atomic operations. A well-known correctness condition for atomicity, which identifies when the operations of a concurrent module *appear* to behave atomically, is *linearisability* [26].

Using the linearisability approach, each operation is given a sequential specification. The operations are then proved to behave atomically *with respect to each other*. One way of seeing this is that there is an instant during the invocation of each operation at which that operation appears to take effect. This

$$\begin{array}{c}
 \{ \text{Counter}(\mathbf{x}, 0, 1) \} \\
 \{ \text{Counter}(\mathbf{x}, 0, 0.5) * \text{Counter}(\mathbf{x}, 0, 0.5) \} \\
 \{ \text{Counter}(\mathbf{x}, 0, 0.5) \} \quad \parallel \quad \{ \text{Counter}(\mathbf{x}, 0, 0.5) \} \\
 \text{incr}(\mathbf{x}) \quad \parallel \quad \text{incr}(\mathbf{x}) \\
 \{ \text{Counter}(\mathbf{x}, 1, 0.5) \} \quad \parallel \quad \{ \text{Counter}(\mathbf{x}, 1, 0.5) \} \\
 \{ \text{Counter}(\mathbf{x}, 1, 0.5) * \text{Counter}(\mathbf{x}, 1, 0.5) \} \\
 \{ \text{Counter}(\mathbf{x}, 2, 1) \}
 \end{array}$$

Figure 2.5: Ownership-based reasoning for concurrent increments.

instant is referred to as the *linearisation point*. With linearisability, the interference of every operation is tolerated at all times by any of the other operations. Consequently, the interference abstraction is deemed to be the module boundary.

Given our sequential specification for the counter in §2.3, is our implementation linearisable? If we only consider the `read` and `incr` operations, then yes, it is. However, the addition of the `wkIncr` operation breaks linearisability. The problem with `wkIncr` is that, for instance, two concurrent calls can result in the counter only being incremented once. This is not consistent with atomic behaviour.

The essence of the problem is that we only envisage calling `wkIncr` in a concurrent context where there are no other increments. In such a case, it would appear to behave atomically. By itself, the sequential specification cannot express this constraint. We need an interference abstraction that constrains the concurrent context.

Alternatively, we could adapt the sequential specification to abstract the changes to the counter value performed by `wkIncr`, allowing it to be used concurrently with other increments, as follows:

$$\begin{aligned}
& \vdash \{ \text{True} \} \text{makeCounter}() \{ \text{ret} \mapsto 0 \} \\
& \vdash \{ x \mapsto n \} \text{read}(x) \{ x \mapsto n \wedge \text{ret} = n \} \\
& \vdash \{ x \mapsto n \} \text{incr}(x) \{ x \mapsto n + 1 \wedge \text{ret} = n \} \\
& \vdash \{ x \mapsto n \} \text{wkIncr}(x) \{ \exists n. x \mapsto n \wedge \text{ret} \leq n \}
\end{aligned}$$

The new specification allows us to see the operations as if they were behaving atomically with respect to each other. However, the specification for `wkIncr` loses information about the state of the counter.

Linearisability is related to the notion of contextual refinement. With contextual refinement, the behaviour of program code is described by (more abstract) specification code.¹ Contextual refinement asserts that the specification code can be replaced by the program code in any context, without introducing new observable behaviours; we say that the program code contextually refines the specification code. Filipović *et al.* [19] have shown that, under certain assumptions about a programming language, linearisability implies contextual refinement for that language. For a linearisable module, each operation contextually refines the operation itself executed atomically. For instance, `incr(x)` contextually refines $\langle \text{incr}(x) \rangle$.

CaReSL [59] is a logic for proving contextual refinement of concurrent programs. CaReSL makes use of auxiliary state, interference abstraction and ownership in its proof technique. However, these concepts are not exposed in their specifications. This means that it is not obvious what a suitable specification of `wkIncr` in CaReSL should be.

Thesis

The concept of *atomicity*, put forward by linearisability, is important in specifying concurrent modules. Atomicity can be seen as a form of interference abstraction: it effectively guarantees that the only observable interference from an operation will occur at a single instant in its execution. This is a powerful abstraction, since a client need not consider intermediate states of an atomic operation (which, for non-atomic operations, might violate invariants) but only the overall transformation it performs.

¹In general, the specification code need not be directly executable, although it does have a semantics.

3 Informal Development

In chapter 2, we have considered a number of proof methods for verifying concurrent programs: Owicki-Gries, rely/guarantee, concurrent separation logics and linearisability. For each method, we have identified a particularly valuable contribution towards specifying concurrent modules.

In this chapter we combine these contributions to produce specifications that are both expressive and modular. We propose a solution that combines the virtues of each of the approaches. Specifically, we introduce a new *atomic triple* judgement for specifying abstract atomicity in a program logic. The simplest form of atomic triple judgement is

$$\vdash \langle P \rangle \mathbb{C} \langle Q \rangle$$

where P and Q are assertions in the style of separation logic, and $\mathbb{C}$ is a program. This judgement is read as “ $\mathbb{C}$ atomically updates P to Q ”. The program may actually take multiple steps, but each step before the atomic update from P to Q must preserve the assertion P . Before the atomic update occurs, the concurrent environment may also update the state, provided that the assertion P is preserved. As soon as the atomic update has happened, the environment can do whatever it likes; it is not constrained to preserve Q . At the same time, the program $\mathbb{C}$ may no longer have access to the resources in Q .

The atomicity of $\mathbb{C}$ is *only* expressed with respect to the abstraction defined by P . If the environment makes an observation at a lower level of abstraction, it may perceive multiple updates rather than this single atomic update. For example, suppose that a set module, which provides an atomic remove operation, is implemented using a linked list. The implementation might first mark a node as deleted before removing it from the list, and the environment could observe the change from “marked” to “removed”. This low-level step, however, does not change the abstract set; the change already occurred when the node was marked.

Atomic triples are our key contribution, as they allow us to overcome the limitations of the linearisability and CAP approaches. Atomic triples can be used to access shared resources concurrently, rather than relying on primitive atomic operations to do so. This makes it easier to build modules on top of each other. Atomic triples specify operations with respect to an abstraction, so they can be proved independently. This makes it possible to extend modules at a later date, and mix together both atomic and non-atomic operations, as well as operations working at different levels of abstraction. Atomic triples can specify clear constraints on how a client can use them. For instance, they can enforce that the unlock operation on a lock should not be called by two threads at the same time (§5.1). Furthermore, atomic triples can specify the transfer of resources between a client and a module. For instance, they can specify an operation that non-atomically stores the result of an atomic read into a buffer provided by a client (§5.2.3).

In order to illustrate our notion of atomicity for a program, we start by showing two simple programs

which manipulate the heap. First, let us consider a program $\mathbb{C}_1$ with the following specification:

$$\vdash \langle x \mapsto 0 \vee x \mapsto 1 \rangle \mathbb{C}_1 \langle x \mapsto 2 \rangle$$

By the precondition, we know that x can be only 0 or 1. The program is allowed to change x to 0 or 1 without this change being considered as part of the atomic update. This is because the abstraction boundary defined by the precondition allows for such interference to happen. The environment is required to maintain that condition until $\mathbb{C}_1$ assigns 2 to x . When $\mathbb{C}_1$ assigns 2 to x , that update is considered to be atomic and at that point, $\mathbb{C}_1$ no longer has access to the memory position x , as the environment is allowed to do anything with it.

Consider now a variation of the previous specification, featuring the TaDA-specific quantifier $\mathbb{W}$:

$$\vdash \mathbb{W}v \in \{0, 1\}. \langle x \mapsto v \rangle \mathbb{C}_2 \langle x \mapsto 2 \rangle$$

As in the previous example, the precondition here can either be $x \mapsto 0$ or $x \mapsto 1$, but the abstraction boundary introduced by the $\mathbb{W}$ quantifier is more precise. The environment is still allowed to change (multiple times) the value of x to any of those two values until the atomic update of $\mathbb{C}_2$ occurs. On the other hand, $\mathbb{C}_2$ is no longer allowed to change x , unless to perform the atomic update. The environment is the only one allowed to change the state for any v . At some point, $\mathbb{C}_2$ will perform its atomic update, writing 2 to x and at that point, as in the previous example, it will no longer have access to the memory position and the environment will be allowed to do anything with it.

We introduce TaDA by showing how to specify and verify the counter and the ticket lock modules shown in chapter 2.

3.1 Spin Counter

We consider the counter implementation from §2.1.1 and give for it a strong specification that allows us to prove the ticket lock. Using TaDA, we prove that the implementation satisfies this specification.

3.1.1 Atomic Specification

The counter operations are specified in terms of an abstract predicate [50] that represents the state of a counter: $\text{Counter}(s, x, n)$ asserts the existence of a counter at address x , with value n . The first parameter s ranges over an *abstract type* (in this case, $\mathbb{T}_1$), which captures implementation-specific information about the counter¹. This parameter serves a technical purpose that we shall address shortly. The predicate confers ownership of the counter: it is not possible to have more than one $\text{Counter}(s, x, n)$ for the same value of x .

The specification for the `makeCounter` operation is a simple Hoare triple:

$$\vdash \{\text{True}\} \text{makeCounter}() \{\exists s \in \mathbb{T}_1. \text{Counter}(s, \text{ret}, 0)\}$$

The operation creates a new counter, which is initially set to value 0, and returns its address. The specification says nothing about the granularity of the operation. In fact, the granularity is hardly

¹To the client, the type is opaque; the implementation realises the type appropriately.

relevant, since no concurrent environment can meaningfully observe the effects of `makeCounter` until its return value is known—that is, once the operation has been completed.

Remark 3.1 (On the abstractly-typed parameters). Many of the proofs of abstract atomicity conclude with a step that quantifies over some fixed parameters in the representation of the data-structure. This generally presumes a rule of the form:

$$\frac{\vdash \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\vdash \langle \exists x. P(x) \rangle \mathbb{C} \langle \exists x. Q(x) \rangle}$$

However, such a rule is unsound in general: in the conclusion, we assume that the environment is able to change the value of x , while in the premiss the value cannot be changed. To avoid this unsoundness, we instead expose the parameter (x in this case). The client does not need to know any particular information about x , only that it should not be changed; hence, the type of x can be abstracted. $\square$

The specification for the `incr(x)` operation uses an *atomic* triple:

$$\vdash \forall n \in \mathbb{N}. \langle \text{Counter}(s, \mathbf{x}, n) \rangle \text{read}(\mathbf{x}) \langle \text{Counter}(s, \mathbf{x}, n) * \text{ret} = n \rangle$$

Intuitively, this specification states that the `read` operation will read the state of the counter *atomically*, even in the presence of concurrent updates by the environment that may change the value of the counter, which are possible, as n is bound by $\forall$. However, the environment must preserve the counter and cannot, for instance, deallocate it.

This atomicity means that the resources in the specification may be *shared*—that is, concurrently accessible by multiple threads. Sharing in this way is not possible with ordinary Hoare triples, since they make no guarantee that intermediate steps preserve invariants on the resources. The atomic triple, by contrast, makes a strong guarantee: as long as the concurrent environment guarantees that the (possibly) shared resource `Counter( $s, \mathbf{x}, n$ )` is available for some n , the `read` operation will preserve `Counter( $s, \mathbf{x}, n$ )` until it reads it; after reading, the operation no longer requires `Counter( $s, \mathbf{x}, n$ )`, and is consequently oblivious to subsequent transformations by the environment (such as another thread incrementing the counter).

It is significant that the notion of atomicity is tied to the abstraction in the specification. The predicate `Counter( $s, \mathbf{x}, n$ )` could abstract multiple underlying states in the implementation. If we were to observe the underlying state, the operation might no longer appear to be atomic.

Specifying `incr` is similar:

$$\vdash \forall n \in \mathbb{N}. \langle \text{Counter}(s, \mathbf{x}, n) \rangle \text{incr}(\mathbf{x}) \langle \text{Counter}(s, \mathbf{x}, n + 1) * \text{ret} = n \rangle$$

The specification states that `incr` operation will increment the counter atomically and return its previous value. Note that when the operation returns, the counter might have been incremented by another thread if the counter is shared.

The last operation can be specified as follows:

$$\forall n \in \mathbb{N}. \vdash \langle \text{Counter}(s, \mathbf{x}, n) \rangle \text{wincr}(\mathbf{x}) \langle \text{Counter}(s, \mathbf{x}, n + 1) \rangle$$

The **wkincr** will atomically update the counter from n to $n + 1$, as long the environment guarantees that the shared counter will not change the value before the atomic update. The specification assumes that the environment will not change the value of the counter, since n is not bound by $\mathbb{W}$. This means that if the counter is shared, other threads can concurrently only perform **read** operations until the counter has been incremented. Technically, it is possible for other **incr** or **wkincr** operations to occur between the update and the return of the operation.

3.1.2 Implementation

To verify this implementation against the atomic specification, we must give a concrete interpretation of the abstract predicate **Counter**. For this, we need to introduce the notion of a *shared region*. A shared region encapsulates some resource that is available to multiple threads². In our example, this resource will be a heap cell that stores the value of the counter described by $x \mapsto n$. A shared region is associated with a protocol, which determines how its contents change over time. Following iCAP §4], the state of a shared region is abstracted, and protocols are expressed as transition systems over these abstract states. A thread may only change the abstract state of a region when it has the *guard* resource associated with the transition to be performed. An interpretation function associates each abstract state of a region with a concrete assertion. In summary, to specify a region we must supply the guards for the region, an abstract state transition system that is labelled by these guards, and a function interpreting abstract states as assertions.

In TaDA, guards consist of abstract resources taken from any separation algebra. This gives us more flexibility in specifying complex usage patterns for regions. For the counter, we need only a very simple guard separation algebra: there is a single, indivisible guard named **INC**, as well as the empty guard **0**. As a separation algebra, guard resources must have a partial composition operator that is associative and commutative. In this case, $\mathbf{0} \bullet x = x = x \bullet \mathbf{0}$, for $x \in \{\mathbf{0}, \mathbf{INC}\}$, and $\mathbf{INC} \bullet \mathbf{INC}$ is not allowed.

The transition system for the region will allow the counter to be incremented using the guard **INC**. This is specified by the labelled transition system:

$$\mathbf{INC} : \forall n. n \rightsquigarrow n + 1$$

It remains to give an interpretation for the abstract states of the transition system. To do so, we must have a name for the type of region we are defining; we shall use **Counter**. It is possible for multiple regions to be associated with the same region type name. To distinguish between them, each region has a unique region identifier, which is typically annotated as a subscript. A region specification may take some parameters that are used in the interpretation. With **Counter**, for instance, the address of the counter is such a parameter. We thus specify the type name, region identifier, parameters and state of a region in the form **Counter** _{a} (x, n).

The region interpretation for **CAPLock** is given by:

$$I(\mathbf{Counter}_a(x, n)) \triangleq x \mapsto n.$$

With this interpretation, the heap cell that contains the value of the counter is always in the region.

²Similar constructs exist in other logics such as [63, 11, 18, 12].

We can now give an interpretation to the predicate $\text{Counter}(s, x, n)$ and the abstract type $\mathbb{T}_1$:

$$\begin{aligned}\mathbb{T}_1 &\triangleq \text{Rld} \\ \text{Counter}(a, x, n) &\triangleq \mathbf{Counter}_a(x, n) * [\text{INC}]_a\end{aligned}$$

where Rld is the set of region identifiers. The abstract predicate asserts that there exists a region with identifier a and is in state n . It also states that there is a guard $[\text{INC}]_a$, which is used to update the region. Note that the first parameter of the Counter predicate fixes the region identifier.

To prove the implementations against our atomic specifications, we use TaDA's MAKEATOMIC rule. A slightly simplified version of this rule is as follows:

$$\frac{\begin{array}{c} \{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_t(G)^* \\ a : x \in X \rightsquigarrow Q(x) \vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * a \Rightarrow \blacklozenge \right\} \mathbb{C} \left\{ \exists x \in X, y \in Q(x). a \Rightarrow (x, y) \right\} \end{array}}{\vdash \mathbb{W}x \in X. \left\langle \mathbf{t}_a(\vec{z}, x) * [G]_a \right\rangle \mathbb{C} \left\langle \exists y \in Q(x). \mathbf{t}_a(\vec{z}, y) * [G]_a \right\rangle}$$

This rule establishes that $\mathbb{C}$ atomically updates the region a , from some state $x \in X$ to some state $y \in Q(x)$. To do so, it requires the guard G for the region, which must permit the update according to the appropriate transition system $\mathcal{T}_t(G)^*$, where $\mathbf{t}$ is the region type; this is established by the first premiss. In our case, the region type $\mathbf{t}$ will be Counter , the guard G will be INC , and the transition system will be $\mathcal{T}_t(\text{INC}) = \{(n, n+1) \mid x \in \mathbb{N}\}$. We use $*$ to denote reflexive-transitive closure.

We use $\mathbf{t}_a(\vec{z}, x)$ to represent a region with region type $\mathbf{t}$, identifier a , parameters $\vec{z}$ and abstract state x . In our example, the region is parametrised with the address of the counter and the abstract state is a natural number. The second premiss introduces two new notations. The first, $a : x \in X \rightsquigarrow Q(x)$, is called the *atomicity context*. The atomicity context records the abstract atomic action that is to be performed. The second, $a \Rightarrow -$, is the atomic tracking resource. The atomic tracking resource indicates whether the atomic update has occurred (the $a \Rightarrow \blacklozenge$ indicates it has not) and, if so, the state of the shared region immediately before and after (the $a \Rightarrow (x, y)$). The resource $a \Rightarrow \blacklozenge$ also plays two special roles that are normally filled by guards. Firstly, it limits the interference on region a : the environment may only update the state so long as it remains in the set X , as specified by the atomicity context. Secondly, it confers permission for the thread to update the region from state $x \in X$ to any state $y \in Q(x)$; in doing so, the thread also updates $a \Rightarrow \blacklozenge$ to $a \Rightarrow (x, y)$. This permission is expressed by the UPDATEREGION rule (see below), and ensures that the atomic update only happens once.

In essence, the second premiss is capturing the notion of atomicity (with respect to the abstraction in the conclusion) and expressing it as a proof obligation. Specifically, the region must be in the state x for some $x \in X$, which may be changed by the environment, until at some point the thread updates it to some $y \in Q(x)$. The atomic tracking resource bears witness to this.

The second key proof rule is the UPDATEREGION rule, which deals with using the atomicity tracking resource to update the region. A simplified version of this rule is as follows:

$$\frac{\vdash \mathbb{W}x \in X. \left\langle I(\mathbf{t}_a(\vec{z}, x)) * P(x) \right\rangle \mathbb{C} \left\langle \exists y \in Q(x). I(\mathbf{t}_a(\vec{z}, y)) * Q_1(x, y) \vee I(\mathbf{t}_a(\vec{z}, x)) * Q_2(x) \right\rangle}{a : x \in X \rightsquigarrow Q(x) \vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \right\} \mathbb{C} \left\{ \begin{array}{c} \exists x \in X. (\exists y \in Q(x). Q_1(x, y) * \\ a \Rightarrow (x, y) \vee \mathbf{t}_a(\vec{z}, x) * Q_2(x) * a \Rightarrow \blacklozenge) \end{array} \right\}}$$

Note that if $y = x$ in the postcondition, the abstract state of the region is not changed and we can either perform the atomic update or not.




Figure 3.1: Proof outline for the **read** operation.

The proof of the **read** implementation is given in Figure 3.1. The proof first massages the specification into a form where we can apply the **MAKEATOMIC** rule. The atomicity context allows the region a to be in any state. The **UpdateRegion** rule performs the atomic action, leaving the region in the same state, and recording the state in the atomic tracking resource. The proof of the **incr** implementation follows a similar style and is given in Figure 3.2. The main difference is that, when entering the loop, it first performs a read operation and stores the current value of the counter in v . The **OPENREGION** rule is used to open the region without changing its abstract state. We then use the **UPDATEREGION** rule to conditionally perform the atomic action. If the atomic compare-and-set operation succeeds, we transition the region from state n to $n + 1$ and update the atomic tracking component. Finally, the proof of the **wkincr** implementation fixes the value of the region and follows the same style as the previous proofs. Given the fact that the environment is not allowed to change the abstract state of the region until the update, the proof does not need to existentially quantify the state of the region in order to make the assertions stable.

3.2 Ticket Lock

We define a lock module with the operations **acquire** and **release** and a constructor **makeLock**. We consider the ticket lock implementation from §2.1.2 and show how to use the atomic specification from the counter to prove the correctness of the ticket lock. We recall the implementation in Figure 3.4.

3.2.1 CAP Specification

We start by specifying the lock module using a specification based on ownership transfer from CAP [12]. The specification provides two abstract predicates: $\text{IsLock}(x)$, which is a non-exclusive resource that




Figure 3.2: Proof outline for the `incr` operation.

allows a thread to compete for the lock; and `Locked(x)`, which is an exclusive resource that represents that the thread has acquired the lock, and allows it to release the lock. The lock is specified as follows:

$$\begin{aligned}
& \vdash \{ \text{True} \} \text{makeLock}() \{ \text{IsLock}(\text{ret}) \} \\
& \vdash \{ \text{Locked}(x) \} \text{release}(x) \{ \text{True} \} \\
& \vdash \{ \text{IsLock}(x) \} \text{acquire}(x) \{ \text{IsLock}(x) * \text{Locked}(x) \} \\
& \text{IsLock}(x) \iff \text{IsLock}(x) * \text{IsLock}(x) \\
& \text{Locked}(x) * \text{Locked}(x) \implies \text{False}
\end{aligned}$$

When a thread acquires the lock, it gets holds of the `Locked(x)` that can be used to subsequently release the lock. The last two axioms allow us to duplicate the non-exclusive resource describing the existence of a lock and guarantee that two threads cannot hold the `Locked(x)` resource at the same time.




Figure 3.3: Proof outline for the `wkincr` operation.

<pre> function makeLock() { next := makeCounter(); owner := makeCounter(); x := alloc(2); [x.next] := next; [x.owner] := owner; return x; } </pre>	<pre> function acquire(x) { next := [x.next]; owner := [x.owner]; t := incr(next); do { v := read(owner); } while (v ≠ t); } </pre>	<pre> function release(x) { owner := [x.owner]; wkIncr(owner); } where E.next $\stackrel{\text{def}}{=} E$ E.owner $\stackrel{\text{def}}{=} E + 1$. </pre>
--	---	--

Figure 3.4: A ticket lock implementation using the counter module.

3.2.2 Implementation

To verify this implementation against the atomic specification, we must first give a concrete interpretation of the abstract predicates. Let us recall that the ticket lock comprises two counters: the first counter records the next available ticket, whereas the second counter records the ticket which currently holds the lock. The lock is considered unlocked when the two counters are equal. In order for a thread to acquire the lock, it must obtain a ticket by incrementing the first counter and then needs to wait until the second counter reaches the value of the obtained ticket. To release the lock, a thread simply increments the second counter.

To verify the implementation, we introduce a new region type, **TLock**. The abstract state of the region will be a natural number n , and there are an infinite amount of them. The abstract state n represents the ticket that currently holds the lock. We associate three guards with the region type.

We have the guard $\text{PENDING}(n, m)$ that keeps track of how many tickets are being held by threads operating on the lock. We also have a unique guard $\text{KEY}(v)$, for each $n \leq v < m$ that represent each ticket, and the empty guard $\mathbf{0}$.

We define the partial composition operator in the following way: $\mathbf{0} \bullet x = x = x \bullet \mathbf{0}$ for $x \in \{\mathbf{0}\} \uplus \{\text{PENDING}(n) \mid n \in \mathbb{N}\} \uplus \{\text{KEY}(n) \mid n \in \mathbb{N}\}$. $\text{PENDING}(n_1, m_1) \bullet \text{PENDING}(n_2, m_2)$ is not allowed and $\text{KEY}(n) \bullet \text{KEY}(m)$ is only allowed for $n \neq m$. Moreover, $\text{PENDING}(n, m) \bullet \text{KEY}(v)$ is allowed if and only if $n \leq v < m$. This ensures that there can only be tickets for a particular v if they correspond to the ones tracked by the guard $\text{PENDING}(n, m)$. Finally, we define the composition to allow the creation and destruction of $\text{KEY}(v)$ guards as follows:

$$\begin{aligned} \text{PENDING}(n, m) &= \text{PENDING}(n, m+1) \bullet \text{KEY}(m), \text{ if } n \leq m \\ \text{PENDING}(n, m) \bullet \text{KEY}(n) &= \text{PENDING}(n+1, m), \text{ if } n < m \end{aligned}$$

The first equality allows us to create new guards and the second allows us to remove guards as long as they have the minimum value.

For the guards, we have constructed an instance of the authoritative monoid of Iris [35]. An alternative approach would be to have only KEY guards, one for each natural number, and define PENDING in terms of those KEY elements, such as the ones used in [7, 12].

The labelled transition system is as follows:

$$\text{KEY}(n) : n \rightsquigarrow n+1$$

It guarantees that a thread must hold the guard $\text{KEY}(n)$ in order to perform the transition.

We also give an interpretation to each abstract state as follows:

$$\begin{aligned} I(\mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n)) &\triangleq \exists m. x.\text{owner} \mapsto \text{owner} * x.\text{next} \mapsto \text{next} * \text{Counter}(s, \text{owner}, n) \\ &\quad * \text{Counter}(t, \text{next}, m) * [\text{PENDING}(n, m)]_a * n \leq m \end{aligned}$$

The region has five parameters, x is the address of the lock and allows us to retrieve both counter addresses, located at owner and next respectively. Moreover, we also have s and t , which correspond to the each counter abstract predicate.

We now define the interpretation of the predicates as follows:

$$\begin{aligned} \text{IsLock}(x) &\triangleq \exists a, s, t, \text{owner}, \text{next}, n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) \\ \text{Locked}(x) &\triangleq \exists a, s, t, \text{owner}, \text{next}, n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * [\text{KEY}(n)]_a \end{aligned}$$

The abstract predicate $\text{IsLock}(x)$ asserts there is a region with identifier a and the region is in some state n . $\text{Locked}(x)$ asserts the same as the previous and additionally states that there is a guard $[\text{KEY}(n)]_a$ which will be used to update the region. Note that by holding the $\text{Locked}(x)$ exclusively, we guarantee that the region abstract state cannot be changed by the environment, as they do not hold the necessary guards to perform such update.

It remains to prove the specifications for the operations and the axioms. The key proof is the




Figure 3.5: Proof outline for the **release** operation.

USEATOMIC rule. A simplified version of the rule is as follows:

$$\frac{\forall x \in X. (x, f(x)) \in \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \quad \vdash \mathbb{W}x \in X. \left\langle I(\mathbf{t}_a(\vec{z}, x)) * P(x) * [\mathbf{G}]_a \right\rangle \mathbb{C} \left\langle I(\mathbf{t}_a(\vec{z}, f(x))) * Q(x) \right\rangle}{\vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * P(x) * [\mathbf{G}]_a \right\} \mathbb{C} \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, f(x)) * Q(x) \right\}}$$

This rule allows a region a , with region type $\mathbf{t}$, to be opened so that it may be updated by $\mathbb{C}$, from some state $x \in X$ to state $f(x)$. In order to do so, the precondition must include a guard $\mathbf{G}$ that is sufficient to perform the update to the region, in accordance with the labelled transition system — this is established by the first premiss.

The proofs of the **release** and **acquire** operations are given in Figure 3.5 and Figure 3.6. In the **release** proof, the interesting step is the use of the USEATOMIC, where we first fix the region for some n by owning the guard $\text{KEY}(n)$ which guarantees that the assertion is stable. We then open the region and combine the $\text{KEY}(n)$ with the $\text{PENDING}(n, m)$ when the counter is atomically incremented. After this happens, the guard $\text{KEY}(n)$ is lost and as such we need to weaken to postcondition to allow an arbitrary n . In the end, we weaken the postcondition to **True** to satisfy the specification. It would also be possible to give a specification that has $\text{IsLock}(x)$ as the postcondition instead.

The **acquire** proof uses the $\mathbb{W}$ quantifier in the premiss of the USEATOMIC rule to account for the fact that, in the precondition, the lock could be in any state n . The first use of the USEATOMIC rule increments the counter and retrieves a $\text{KEY}(\mathbf{t})$ for the value read. After the read, because we own $\text{KEY}(\mathbf{t})$, we can guarantee that the state of the region cannot be larger than $\mathbf{t}$, i.e. that the environment does not have the necessary guards to perform such a transition. The loop then simply waits until the state of the region matches the ticket. When that happens, we know it cannot change as long as we own the guard $\text{KEY}(\mathbf{t})$ and as such we can satisfy the $\text{Locked}(x)$ predicate.

The axiom $\text{IsLock}(x) \iff \text{IsLock}(x) * \text{IsLock}(x)$ follows from the duplicability of region assertions: i.e. $\mathbf{TLock}_a(x, \text{owner}, \text{next}, n) \equiv \mathbf{TLock}_a(x, \text{owner}, \text{next}, n) * \mathbf{TLock}_a(x, \text{owner}, \text{next}, n)$. Finally, the axiom $\text{Locked}(x) * \text{Locked}(x) \implies \text{False}$ follows from the fact that $\text{KEY}(n) \bullet \text{KEY}(n)$ is undefined.

ABSTRACT; EXISTS: $a, s, t, \text{owner}, \text{next}$	$\{\text{IsLock}(x)\}$	
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n)\}$	
	$\text{next} := \lfloor x.\text{next} \rfloor;$	
	$\text{owner} := \lfloor x.\text{owner} \rfloor;$	
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * \text{next} = \text{next} * \text{owner} = \text{owner}\}$	
	USEATOMIC	$\mathbb{W}n.$
		$\langle \text{Counter}(s, \text{owner}, n) * \text{Counter}(t, \text{next}, m) * [\text{PENDING}(n, m)]_a \rangle$
		$\langle * n \leq m * \text{next} = \text{next} * \text{owner} = \text{owner} \rangle$
		$t := \text{incr}(x.\text{next});$
		$\langle \text{Counter}(s, \text{owner}, n) * \text{Counter}(t, \text{next}, m+1) * [\text{PENDING}(n, m+1)]_a \rangle$
		$\langle * n \leq m * [\text{KEY}(t)]_a * t = m * \text{next} = \text{next} * \text{owner} = \text{owner} \rangle$
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * [\text{KEY}(t)]_a * n \leq t * \text{next} = \text{next} * \text{owner} = \text{owner}\}$	
	do {	
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * [\text{KEY}(t)]_a * n \leq t * \text{next} = \text{next} * \text{owner} = \text{owner}\}$	
	USEATOMIC	$\mathbb{W}n.$
		$\langle \text{Counter}(s, \text{owner}, n) * \text{Counter}(t, \text{next}, m) * [\text{PENDING}(n, m)]_a \rangle$
		$\langle * n \leq m * [\text{KEY}(t)]_a * n \leq t * \text{next} = \text{next} * \text{owner} = \text{owner} \rangle$
		$v := \text{read}(x.\text{owner});$
		$\langle \text{Counter}(s, \text{owner}, n) * \text{Counter}(t, \text{next}, m) * [\text{PENDING}(n, m)]_a \rangle$
		$\langle * n \leq m * [\text{KEY}(t)]_a * n \leq t * n = v * \text{next} = \text{next} * \text{owner} = \text{owner} \rangle$
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * [\text{KEY}(t)]_a * n \leq t * n \geq v\}$	
	$\{ * \text{next} = \text{next} * \text{owner} = \text{owner} \}$	
	} while ($v \neq t$);	
	$\{\exists n. \mathbf{TLock}_a(x, s, t, \text{owner}, \text{next}, n) * [\text{KEY}(t)]_a * n = t\}$	
	$\{\text{Locked}(x)\}$	

Figure 3.6: Proof outline for the `acquire` operation.

The counter specifications shown in this section are strong: a client can derive the abstract disjoint specifications from them. Moreover, they are strong enough to support synchronisation: the correctness of the ticket lock can be justified from the counter specifications. These approaches to specification are expressive enough to enforce obligations on both the client and the implementation. By contrast, CAP specifications tend to unduly restrict the client (e.g. a counter specification cannot be used for synchronisation), while linearisability specifications tend to unduly restrict the implementation (e.g. a counter cannot provide a `wkincr` operation).

4 TaDA Logic

We formalise the TaDA logic and prove the soundness of the proof system. In order to do so, we first present the programming language (§4.1) and give its operational semantics (§4.2). We then formalise the program logic in §4.4 and present a system of proof rules for reasoning about programs. Finally, we provide the semantics for the logic (§4.5) and show the soundness of the proof system in §4.6.

4.1 Programming Language

We consider an imperative programming language with dynamically allocated heap and mutable local variables. The operations that manipulate the heap are atomic. Local variables have local scope, in that they can only be accessed within a function. Moreover, the language supports the spawning of new threads using a fork command.

Definition 4.1 (Variable Names). Assume a set of *variable names* Var , ranged over by $x, y, z, \dots$

Definition 4.2 (Expressions). The set of integer-valued *expressions* Expr , ranged over by $\mathbb{E}, \mathbb{E}_1, \dots$, is defined by the grammar:

$$\begin{array}{lcl} \mathbb{E} & ::= & x \\ & | & v \\ & | & \mathbb{E}_1 + \mathbb{E}_2 \\ & | & \mathbb{E}_1 - \mathbb{E}_2 \\ & | & \mathbb{E}_1 \times \mathbb{E}_2 \\ & | & \mathbb{E}_1 \div \mathbb{E}_2 \end{array}$$

where $v \in \mathbb{Z}$.

Expressions can be either variables or integer constants or the compositions of these two using standard binary arithmetic operations.

Definition 4.3 (Boolean Expressions). The set of *boolean expressions* BExpr , ranged over by $\mathbb{B}, \mathbb{B}_1, \dots$, is defined by the grammar:

$$\begin{array}{lcl} \mathbb{B} & ::= & \text{true} \\ & | & \text{false} \\ & | & \mathbb{E}_1 = \mathbb{E}_2 \\ & | & \mathbb{E}_1 < \mathbb{E}_2 \\ & | & \mathbb{B}_1 \text{ and } \mathbb{B}_2 \\ & | & \mathbb{B}_1 \text{ or } \mathbb{B}_2 \\ & | & \text{not } \mathbb{B} \end{array}$$

Boolean expressions can be either constants (**true** and **false**), or comparisons on expressions, or boolean combinations of boolean expressions.

Definition 4.4 (Function Names). Assume a set of *function names* Fun , ranged over by $f, g, \dots$

Definition 4.5 (Commands). The set of *commands* Cmd , ranged over by $\mathbb{C}, \mathbb{C}_1, \dots$, is defined by the following grammar:

$\mathbb{C} ::= \text{skip};$	Empty command
$\mathbb{C}_1 \mathbb{C}_2$	Sequencing
$\text{while } (\mathbb{B}) \{ \mathbb{C} \}$	Loop
$\text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}$	Conditional
$x := f(\vec{\mathbb{E}});$	Function call
$\text{fork } f(\vec{\mathbb{E}});$	Spawn thread
$x := \text{alloc}(\mathbb{E});$	Allocation
$x := \mathbb{E};$	Assignment
$x := [\mathbb{E}];$	Lookup
$[\mathbb{E}_1] := \mathbb{E}_2;$	Mutation
$x := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3);$	Compare-and-set

We identify commands up to associativity of sequencing: that is, $(\mathbb{C}_1 \mathbb{C}_2) \mathbb{C}_3 = \mathbb{C}_1 (\mathbb{C}_2 \mathbb{C}_3)$.

Note that we can encode the command $\text{do } \{ \mathbb{C} \} \text{ while } (\mathbb{B})$ as $\mathbb{C} \text{ while } (\mathbb{B}) \{ \mathbb{C} \}$.

Definition 4.6 (Function Bodies). The set of *function bodies* FunBody , ranged over by $\mathbb{F}, \mathbb{F}_1, \dots$, is defined as:

$$\mathbb{F} ::= \mathbb{C} \text{ return } \mathbb{E};$$

For simplicity, we restrict the return of a function to occur only at the end. When we are not interested on the return value we might omit it, in such cases consider it to be sugar syntax for a function with a return at the end with some expression.

Definition 4.7 (Programs). The set of *programs* Prog , ranged over by $\mathbb{P}, \mathbb{P}_1, \dots$, is defined as a finite sequence of function definitions:

$$\mathbb{P} ::= \epsilon \mid \mathbb{P}, \text{function } f(\vec{x}) \{ \mathbb{F} \}$$

where function names are assumed to be pairwise distinct.

4.2 Operational Semantics

We present the operational semantics of the programming language as a Views-style [13] labelled transition system. Transitions are labelled by atomic actions. The labelled transition system splits the control-flow aspect of the execution from the heap-transforming aspect of the execution which is represented by the labels of each transition.

Definition 4.8 (Program Values). Fix the set of *program values* $\text{Val} \stackrel{\text{def}}{=} \mathbb{Z}$ to be the integers, ranged over by $v, v_1, \dots$

Definition 4.9 (Variable Store). A *variable store* $\text{Store} \stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Val}$ is a total function from program variables to program values, ranged over by $\sigma, \sigma_1, \dots$

Definition 4.10 (Semantics of Expressions). The *semantics of expressions* $\mathcal{E}[-]_- : \text{Expr} \times \text{Store} \rightarrow \text{Val}$, is defined with respect to the variable store by:

$$\begin{aligned} \mathcal{E}[\mathbf{x}]_\sigma &\stackrel{\text{def}}{=} \sigma(\mathbf{x}) \\ \mathcal{E}[v]_\sigma &\stackrel{\text{def}}{=} v \\ \mathcal{E}[\mathbb{E}_1 + \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{E}[\mathbb{E}_1]_\sigma + \mathcal{E}[\mathbb{E}_2]_\sigma \\ \mathcal{E}[\mathbb{E}_1 - \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{E}[\mathbb{E}_1]_\sigma - \mathcal{E}[\mathbb{E}_2]_\sigma \\ \mathcal{E}[\mathbb{E}_1 \times \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{E}[\mathbb{E}_1]_\sigma \times \mathcal{E}[\mathbb{E}_2]_\sigma \\ \mathcal{E}[\mathbb{E}_1 \div \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \begin{cases} \mathcal{E}[\mathbb{E}_1]_\sigma \div \mathcal{E}[\mathbb{E}_2]_\sigma & \text{if } \mathcal{E}[\mathbb{E}_2]_\sigma \neq 0 \\ 42 & \text{otherwise.} \end{cases} \end{aligned}$$

Note that the store is a total function and, as such, if a variable is not initialised, then its value will be arbitrary but consistent in the sense that multiple accesses of the same variable returns the same value.

Definition 4.11 (Semantics of Boolean Expressions). The semantics of *boolean expressions* $\mathcal{B}[-]_- : \text{BExpr} \times \text{Store} \rightarrow \text{Bool}$, is defined with respect to the variable store by:

$$\begin{aligned} \mathcal{B}[\text{true}]_\sigma &\stackrel{\text{def}}{=} \text{True} \\ \mathcal{B}[\text{false}]_\sigma &\stackrel{\text{def}}{=} \text{False} \\ \mathcal{B}[\mathbb{E}_1 = \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{E}[\mathbb{E}_1]_\sigma = \mathcal{E}[\mathbb{E}_2]_\sigma \\ \mathcal{B}[\mathbb{E}_1 < \mathbb{E}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{E}[\mathbb{E}_1]_\sigma < \mathcal{E}[\mathbb{E}_2]_\sigma \\ \mathcal{B}[\mathbb{B}_1 \text{ and } \mathbb{B}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{B}[\mathbb{B}_1]_\sigma \wedge \mathcal{B}[\mathbb{B}_2]_\sigma \\ \mathcal{B}[\mathbb{B}_1 \text{ or } \mathbb{B}_2]_\sigma &\stackrel{\text{def}}{=} \mathcal{B}[\mathbb{B}_1]_\sigma \vee \mathcal{B}[\mathbb{B}_2]_\sigma \\ \mathcal{B}[\text{not } \mathbb{B}]_\sigma &\stackrel{\text{def}}{=} \neg \mathcal{B}[\mathbb{B}]_\sigma \end{aligned}$$

The programming language allows function calls inside a function body. This requires the creation of a new store that is used to evaluate the body of the function being called. We model the stack of stores and the current execution point using continuations.

In order to define the operational semantics over commands, we extend the commands to include execution stacks that occur as a result of a function called.

Definition 4.12 (Extended Commands). The set of extended commands ExtCmd , ranged over by $c, c_1, \dots$, is defined by the following grammar:

$$\begin{aligned} c &::= \mathbf{x} := (\sigma, c \text{ return } \mathbb{E}); \\ &\quad | \quad \mathbb{C} \\ &\quad | \quad c \mathbb{C} \end{aligned}$$

The first construct represents the execution stack. It is used to keep track of the point to which the function call occurred along with each store for each of the successive function calls. The second construct represents the current command being executed. Finally, the third construct allows us to sequence extended commands with commands.

Definition 4.13 (Continuations). The set of continuations Cont , ranged over by $k, k_1, \dots$, is defined

by the following grammar:

$$k ::= c \text{ return } \mathbb{E};$$

Definition 4.14 (Threads). A *thread* $\text{Thread} \stackrel{\text{def}}{=} \text{Store} \times \text{Cont}$ is a pair of variable store and a continuation, ranged over by $t, t_1, \dots$.

Definition 4.15 (Labels of Atomic Commands). The set of transition labels AAction , ranged over by α , is described by the following grammar:

$\alpha ::=$	id	Unchanged
	$ \text{ alloc}(v_1, v_2)$	Allocation
	$ \text{ read}(v_1, v_2)$	Lookup
	$ \text{ write}(v_1, v_2)$	Mutation
	$ \text{ cas}(v_1, v_2, v_3, v_4)$	Compare-and-set
	$ \text{ spawn}(\mathbf{f}, \vec{v})$	Fork thread
	$ \text{ } \not\downarrow$	Fault

The labels of atomic commands describe all possible operations that can manipulate the heap. Note that the labels $\text{alloc}(v_1, v_2)$, $\text{read}(v_1, v_2)$ and $\text{cas}(v_1, v_2, v_3, v_4)$ use the last argument to represent the returned values. In the label alloc , v_1 represents the size of heap allocated and v_2 its address. In the label read , v_1 is the heap address being read and v_2 the value stored at that address. Finally, in the label cas , v_1 is the heap address that the compare-and-set is performed on, v_2 is the expected value, v_3 the new value and v_4 represents if the operation was successful or not.

Definition 4.16 (Function Environment). A function environment $\text{FEnv} \stackrel{\text{def}}{=} \text{Fun} \rightarrow_{fn} (\text{Var}^* \times \text{FunBody})$ is a finite partial mapping from function names to parameters and bodies, ranged over by η . We denote a well-defined lookup as $\eta(\mathbf{f}) = (\vec{x}, \mathbb{F})$. Moreover, we refer to the parameters as $\text{vars}(\eta(\mathbf{f}))$ and to the body of the function as $\text{code}(\eta(\mathbf{f}))$. A function environment for a given program $\mathbb{P}$ must coincide with its function definitions.

Definition 4.17 (Thread Operational Semantics). The thread operational semantics is defined as a labelled transition relation of individual threads

$$- \xrightarrow{-} - : \text{Thread} \times \text{AAction} \times \text{FEnv} \times \text{Thread}$$

defined by the following rule:

$$\frac{\text{FUNCTIONSTEP} \quad (\sigma_1, c_1) \xrightarrow{\alpha}_{\eta} (\sigma_2, c_2)}{(\sigma_1, c_1 \text{ return } \mathbb{E};) \xrightarrow{\alpha}_{\eta} (\sigma_2, c_2 \text{ return } \mathbb{E};)}$$

The FUNCTIONSTEP reduces extended command in the function body until it is equal to $\text{skip};$.

Definition 4.18 (Extend Commands Operational Semantics). The extended commands operational semantics is defined as a labelled transition relation

$$(-, -) \xrightarrow{-} (-, -) : \text{Store} \times \text{ExtCmd} \times \text{AAction} \times \text{FEnv} \times \text{Store} \times \text{ExtCmd}$$

defined by the following rules:

$$\begin{array}{c}
\text{SKIP} \\
\hline
(\sigma, \text{skip}; c) \xrightarrow{\text{id}}_{\eta} (\sigma, c)
\end{array}
\quad
\begin{array}{c}
\text{SEQUENCING} \\
\hline
(\sigma_1, c_1) \xrightarrow{\alpha}_{\eta} (\sigma_2, c_2) \\
\hline
(\sigma_1, c_1 \text{ } \mathbb{C}) \xrightarrow{\alpha}_{\eta} (\sigma_2, c_2 \text{ } \mathbb{C})
\end{array}$$

$$\begin{array}{c}
\text{LOOPTRUE} \\
\hline
\mathcal{B}[\mathbb{B}]_{\sigma} = \text{True} \\
\hline
(\sigma, \text{while } (\mathbb{B}) \{ \mathbb{C} \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C} \text{ while } (\mathbb{B}) \{ \mathbb{C} \})
\end{array}
\quad
\begin{array}{c}
\text{LOOPFALSE} \\
\hline
\mathcal{B}[\mathbb{B}]_{\sigma} = \text{False} \\
\hline
(\sigma, \text{while } (\mathbb{B}) \{ \mathbb{C} \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \text{skip};)
\end{array}$$

$$\begin{array}{c}
\text{CONDITIONALTRUE} \\
\hline
\mathcal{B}[\mathbb{B}]_{\sigma} = \text{True} \\
\hline
(\sigma, \text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C}_1)
\end{array}
\quad
\begin{array}{c}
\text{CONDITIONALFALSE} \\
\hline
\mathcal{B}[\mathbb{B}]_{\sigma} = \text{False} \\
\hline
(\sigma, \text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C}_2)
\end{array}$$

$$\begin{array}{c}
\text{FUNCTIONCALL} \\
\hline
\mathcal{E}[\vec{\mathbb{E}}]_{\sigma_1} = \sigma_2(\text{vars}(\eta(\mathbf{f}))) \\
\hline
(\sigma_1, \mathbf{x} := \mathbf{f}(\vec{\mathbb{E}});) \xrightarrow{\text{id}}_{\eta} (\sigma_1, \mathbf{x} := (\sigma_2, \text{code}(\eta(\mathbf{f}))))
\end{array}$$

$$\begin{array}{c}
\text{FUNCTIONCALLSTEP} \\
\hline
(\sigma_1, c_1) \xrightarrow{\alpha}_{\eta} (\sigma_2, c_2) \\
\hline
(\sigma_3, \mathbf{x} := (\sigma_1, c_1 \text{ return } \mathbb{E};)) \xrightarrow{\alpha}_{\eta} (\sigma_3, \mathbf{x} := (\sigma_2, c_2 \text{ return } \mathbb{E};))
\end{array}$$

$$\begin{array}{c}
\text{RETURN} \\
\hline
(\sigma_1, \mathbf{x} := (\sigma_2, \text{skip}; \text{return } \mathbb{E};)) \xrightarrow{\text{id}}_{\eta} (\sigma_1[\mathbf{x} \mapsto \mathcal{E}[\mathbb{E}]_{\sigma_2}], \text{skip};)
\end{array}$$

$$\begin{array}{c}
\text{FORK} \\
\hline
(\sigma, \text{fork } \mathbf{f}(\vec{\mathbb{E}});) \xrightarrow{\text{spawn}(\mathbf{f}, \mathcal{E}[\vec{\mathbb{E}}]_{\sigma})}_{\eta} (\sigma, \text{skip};)
\end{array}
\quad
\begin{array}{c}
\text{ALLOCATION} \\
\hline
(\sigma, \mathbf{x} := \text{alloc}(\mathbb{E});) \xrightarrow{\text{alloc}(\mathcal{E}[\mathbb{E}]_{\sigma}, v)}_{\eta} (\sigma[\mathbf{x} \mapsto v], \text{skip};)
\end{array}$$

$$\begin{array}{c}
\text{ASSIGNMENT} \\
\hline
(\sigma, \mathbf{x} := \mathbb{E};) \xrightarrow{\text{id}}_{\eta} (\sigma[\mathbf{x} \mapsto \mathcal{E}[\mathbb{E}]_{\sigma}], \text{skip};)
\end{array}
\quad
\begin{array}{c}
\text{LOOKUP} \\
\hline
(\sigma, \mathbf{x} := [\mathbb{E}];) \xrightarrow{\text{read}(\mathcal{E}[\mathbb{E}]_{\sigma}, v)}_{\eta} (\sigma[\mathbf{x} \mapsto v], \text{skip};)
\end{array}$$

$$\begin{array}{c}
\text{MUTATION} \\
\hline
(\sigma, [\mathbb{E}_1] := \mathbb{E}_2;) \xrightarrow{\text{write}(\mathcal{E}[\mathbb{E}_1]_{\sigma}, \mathcal{E}[\mathbb{E}_2]_{\sigma})}_{\eta} (\sigma, \text{skip};)
\end{array}$$

$$(\sigma, \mathbf{x} := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3);) \xrightarrow{\text{cas}(\mathcal{E}[\mathbb{E}_1]_\sigma, \mathcal{E}[\mathbb{E}_2]_\sigma, \mathcal{E}[\mathbb{E}_3]_\sigma, v)}_\eta (\sigma[\mathbf{x} \mapsto v], \text{skip};)$$

The SKIP leaves the state unchanged and continues on the next command. The SEQUENCING performs a step of c_1 , reducing it to c_2 , when it is not equal to **skip**;. The LOOPTRUE reduces to a sequence containing the body of the loop, followed by the loop itself. This corresponds to executing once the loop body if the boolean expression holds. The LOOPFALSE reduces to **skip**; if the boolean expression does not hold, stopping the loop. The CONDITIONALTRUE and CONDITIONALFALSE simply reduce the conditional to one of its branches, without affecting the state, depending on whether the boolean expression holds. The FUNCTIONCALL creates a new store to execute the function body. The FUNCTIONCALLSTEP executes steps in a function previously called. Note that the function previously called may contain further function calls. The RETURN terminates a function previously called and destroys its store. The SPAWN spawns a new thread, which has no effect for the current thread. Finally, the remaining cases perform basic operations on the state of the program.

Definition 4.19 (Thread Identifiers). Assume a set of *thread identifiers* TId .

Definition 4.20 (Thread Pools). A thread pool $T \in \text{ThreadPool} \stackrel{\text{def}}{=} \text{TId} \rightarrow_{\text{fin}} \text{Thread}$ is a finite partial function from thread identifiers to threads.

Definition 4.21 (Heaps). A *heap* $\text{Heap} \stackrel{\text{def}}{=} \text{Addr} \rightarrow_{\text{fin}} \text{Val}$ is a finite partial function from addresses to values, ranged over by $h, h_1, \dots$. Heaps form a partial commutative monoid $(\text{Heap}, \uplus, \emptyset)$, where $\uplus$ is the disjoint union of partial functions, and $\emptyset$ is the partial function with the empty domain. The structure of the partial commutative monoid induces a resource order $\leq$, defined as follows:

$$h_1 \leq h_2 \stackrel{\text{def}}{\iff} \exists h_3. h_1 \uplus h_3 = h_2.$$

Note that $\leq$ is a partial order.

Definition 4.22 (Program States). A *program state* is a pair of a heap and a thread pool. Moreover, program states are enriched with an exceptional faulting state, denoted $\not\downarrow$, to represent the result of an invalid memory access:

$$\text{State} \stackrel{\text{def}}{=} (\text{Heap} \uplus \{\not\downarrow\}) \times \text{ThreadPool}$$

Definition 4.23 (Thread Pool Operational Semantics). The labelled transition relation of a thread pool

$$- \xrightarrow{-} - : \text{ThreadPool} \times \text{AAction} \times \text{FEnv} \times \text{ThreadPool}$$

is defined by:

$$\frac{\text{THREADRETURN}}{T \parallel (\sigma, \text{skip}; \text{return } \mathbb{E};) \xrightarrow{\text{id}}_\eta T} \quad \frac{\text{THREADSPAWN} \quad t \xrightarrow{\text{spawn}(\mathbf{f}, \vec{v})}_\eta t' \quad \sigma(\text{vars}(\eta(\mathbf{f}))) = \vec{v}}{T \parallel t \xrightarrow{\text{id}}_\eta T \parallel t' \parallel (\sigma, \text{code}(\eta(\mathbf{f})))}$$

$$\frac{\text{THREADSTEP} \quad t \xrightarrow{\alpha}_{\eta} t' \quad \alpha \notin \{\text{spawn}(\mathbf{f}, \vec{v}) \mid \mathbf{f} \in \text{Fun}, \vec{v} \in \text{Val}^*\}}{T \parallel t \xrightarrow{\alpha}_{\eta} T \parallel t'}$$

The **THREADRETURN** reduces the number of threads in the thread pool when the thread has finished its execution. A thread being executed can perform a fork operation spawning a new thread, while the other operations have no effect on the number of threads in the thread pool. The **THREADSPAWN** creates a new thread when an existing thread performs a fork operation. The newly created thread has a store σ with the parameters of the function initialised by the values passed by the fork operation. The **THREADSTEP** updates a thread in the thread pool, by performing a non-spawn action as defined by the thread operational semantics.

Definition 4.24 (Addresses). Fix the set of *addresses* $\text{Addr} \stackrel{\text{def}}{=} \mathbb{Z}^+$ to be the set of positive integers, where $\text{Addr} \subseteq \text{Val}$.

Definition 4.25 (Interpretation of Atomic Commands). The Interpretation of atomic action labels is given as a function

$$\llbracket - \rrbracket : \text{AAction} \rightarrow \text{Heap} \rightarrow \mathcal{P}(\text{Heap} \uplus \{\zeta\})$$

that associates each atomic command with a non-deterministic state transformer and is defined as follows:

$$\begin{aligned} \llbracket \text{id} \rrbracket &\stackrel{\text{def}}{=} \lambda h. \{h\} \\ \llbracket \text{alloc}(v_1, v_2) \rrbracket &\stackrel{\text{def}}{=} \lambda h. \begin{cases} \{h[\vec{v}_3 \mapsto \vec{v}_4] \mid \vec{v}_4 \in \text{Val}^{v_1}\} & \text{if } v_1 > 0 \text{ and } v_2 \in \text{Addr} \text{ and} \\ & \{v_2, \dots, v_2 + v_1 - 1\} \cap \text{dom}(h) = \emptyset \\ & \text{and } \vec{v}_3 = (v_2, \dots, v_2 + v_1 - 1) \\ \emptyset & \text{if } v_1 > 0 \text{ and } v_2 \in \text{Addr} \text{ and} \\ & \{v_2, \dots, v_2 + v_1 - 1\} \cap \text{dom}(h) \neq \emptyset \\ & \text{or } v_1 > 0 \text{ and } v_2 \notin \text{Addr} \\ \{\zeta\} & \text{otherwise.} \end{cases} \\ \llbracket \text{read}(v_1, v_2) \rrbracket &\stackrel{\text{def}}{=} \lambda h. \begin{cases} \{h\} & \text{if } v_1 \in \text{dom}(h) \text{ and } h(v_1) = v_2 \\ \emptyset & \text{if } v_1 \in \text{dom}(h) \text{ and } h(v_1) \neq v_2 \\ \{\zeta\} & \text{otherwise.} \end{cases} \\ \llbracket \text{write}(v_1, v_2) \rrbracket &\stackrel{\text{def}}{=} \lambda h. \begin{cases} \{h[v_1 \mapsto v_2]\} & \text{if } v_1 \in \text{dom}(h) \\ \{\zeta\} & \text{otherwise.} \end{cases} \\ \llbracket \text{cas}(v_1, v_2, v_3, v_4) \rrbracket &\stackrel{\text{def}}{=} \lambda h. \begin{cases} \{h[v_1 \mapsto v_3]\} & \text{if } v_1 \in \text{dom}(h) \text{ and } h(v_1) = v_2 \text{ and } v_4 \neq 0 \\ \{h\} & \text{if } v_1 \in \text{dom}(h) \text{ and } h(v_1) \neq v_2 \text{ and } v_4 = 0 \\ \{\zeta\} & v_1 \notin \text{dom}(h) \\ \emptyset & \text{otherwise.} \end{cases} \\ \llbracket \text{spawn}(\mathbf{f}, \vec{v}) \rrbracket &\stackrel{\text{def}}{=} \lambda h. \{h\} \\ \llbracket \zeta \rrbracket &\stackrel{\text{def}}{=} \lambda h. \{\zeta\} \end{aligned}$$

For a heap h , the set of states $\llbracket \alpha \rrbracket(h)$ is the set of possible outcomes of running the atomic command α . If the set is empty, then the command diverges. We lift non-deterministic state transformers to heaps or the faulting state by letting $\llbracket \alpha \rrbracket(\downarrow) = \{\downarrow\}$.

Definition 4.26 (Program Operational Semantics). The single step transition relation of a program

$$- \xrightarrow{-} - : \text{State} \times \text{AAction} \times \text{FEnv} \times \text{State}$$

is defined by the following rule:

$$\frac{\text{PROGRAMSTEP} \quad T \xrightarrow{\alpha}_{\eta} T' \quad h' \in \llbracket \alpha \rrbracket(h)}{(h, T) \rightarrow_{\eta} (h', T')}$$

The program operational semantics combines the execution of the thread pool and the updates performed by each thread on the heap, as described by the atomic commands α .

4.3 Assertion Language

We now present an assertion language for defining predicates. As well as standard separation logic connectives, the logic includes regions, guards, abstract predicates, atomicity tracking resources and view shifts.

Definition 4.27 (Logical Expressions). Assume a set of *logic expressions* LExp , ranged over by $e, e_1, \dots$.

Definition 4.28 (Assertions). The set of *assertions* Assn , ranged over by $P, Q, R, P_1, \dots$ are defined

inductively as follows:

$P, Q, R ::=$	False	Falsehood
	True	Truthfulness
	$P * Q$	Separating conjunction
	$x \mapsto y$	Heap cell
	$P \wedge Q$	Conjunction
	$P \vee Q$	Disjunction
	$\exists x. P$	Existential quantification
	$\forall x. P$	Universal quantification
	$P \implies Q$	Implication (material)
	$\mathbf{t}_a^\lambda(\vec{z}, x)$	Shared region a , of type $\mathbf{t}$, with level λ , parametrised by $\vec{z}$ and with abstract state x
	$I(\mathbf{t}_a^\lambda(\vec{z}, x))$	Interpretation of a shared region a , of type $\mathbf{t}$, with level λ , parametrised by $\vec{z}$ and with abstract state x
	$[G(\vec{z})]_a$	Guard G for region a , parametrised by $\vec{z}$
	$a \Rightarrow \blacklozenge$	Atomicity tracking resource for region a , allowing atomic update
	$a \Rightarrow \blacklozenge$	Atomicity tracking resource for region a , not allowing atomic update
	$a \Rightarrow (x, y)$	Atomicity tracking resource for region a , witnessing an atomic update from x to y
	$\mathbf{a}(\vec{z})$	Abstract predicate $\mathbf{a}$ parametrised by $\vec{z}$

where

- $x, y \in \text{LE Expr}$ ranges over logical expressions and $\vec{z}$ denotes vectors of logical expressions.
- $a \in \text{Rld}$ ranges over region identifiers,
- $\lambda \in \text{Level}$ ranges over levels,
- $\mathbf{a} \in \text{APName}$ ranges over abstract predicates names.

4.4 Program Logic

The TaDA proof system is presented here for the programming language in §4.1.

Definition 4.29 (Function Specification Context). Function specification contexts assign specifications to function symbols. They are generated by the following grammar, where $\vec{z}$ is a sequence of logical variables:

$$\Gamma \in \text{FunctionCtxt} ::= \epsilon \quad \mid \quad \Gamma, \left(\begin{array}{c} \mathbb{W}x \in X. \langle P_p(\vec{z}) \mid P(x, \vec{z}) \rangle \\ \lambda; \mathcal{A} \vdash \quad \mathbf{f}(\vec{z}) \\ \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle \end{array} \right)$$

Definition 4.30 (Atomic Judgements). The generalised form of atomic judgements in TaDA is:

$$\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

Here, P_p is the private precondition, $P(x)$ is the public precondition, $Q_p(x, y)$ is the private postcondition, and $Q(x, y)$ is the public postcondition. The private precondition is independent of x , since the environment can change x . The two parts of the postcondition are linked by y , which is chosen arbitrarily by the implementation when the atomic operation appears to take effect.

Formally, we only consider a single variable bound by $\mathbb{W}$ or $\exists$. However, it is often useful to consider multiple such variables. We treat this as syntactic sugar which can be interpreted using Cartesian product and projection functions. We also drop the binder altogether when the variable ranges over the set $\mathbf{1}$.

We define $\Gamma; \lambda; \mathcal{A} \vdash \{P\} \mathbb{C} \{Q\}$ as sugar for $\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in \mathbf{1}. \langle P \mid \text{True} \rangle \mathbb{C} \quad \exists y \in \mathbf{1}. \langle Q \mid \text{True} \rangle$.

We implicitly require the pre- and postcondition assertions in our judgements to be *stable*: that is, they must account for any updates other threads could have sufficient resources to perform. We will later define stability formally in the next section.

Definition 4.31 (Proof System). The derivation rules for atomic judgements of the form

$$\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

are given in Figure 4.1, Figure 4.2 and Figure 4.3.

The proof rules in Figure 4.1 are related to the programming language.

The SKIP rule asserts that the `skip`; does not change the state of the program. An alternative would be specify the pre- and postcondition as `True`.

The SEQUENCING rule allows us to compose programs that are executed sequentially.

The LOOP rule states that there is a loop invariant, P_p , which is preserved by the loop body $\mathbb{C}$. After the loop is finished, the invariant P_p still holds, and moreover $\mathbb{B}$ must be caused the loop to end.

The CONDITIONAL rule states that if both branches of the conditional establish Q_p , then the postcondition of the whole satisfies Q_p .

The FUNCTIONCALL rule and FORK rule establish the conditions for the functions to be called. Essentially, in order to call a function, we have to have enough resources. In the case of the fork operation, the ownership of P_p is transferred to the newly created thread.

The ALLOCATION rule allocates memory and returns the address on $\mathbf{x}$.

The ASSIGNMENT rule states that, after the assignment, the variable $\mathbf{x}$ has value $\mathbb{E}$ where each occurrence of $\mathbf{x}$ has been replaced by v . The MUTATION rule and the COMPAREANDSET rule are similar, but manipulate the heap instead.

The main proof rules to handle abstract atomicity are shown in Figure 4.2. The first three rules allow us to access the content of a shared region by using an atomic command. With all of the rules, the update to the shared region must be atomic, so its interpretation is in the public part in the premiss. The region is in the public part in the conclusion also, but may be moved by applying atomicity weakening. The last rule is used to prove that a command is atomic.

$$\begin{array}{c}
\text{SKIP} \\
\hline
\Gamma; \lambda; \mathcal{A} \vdash \{P_p\} \text{ skip}; \{P_p\}
\end{array}
\quad
\begin{array}{c}
\text{SEQUENCING} \\
\hline
\frac{\Gamma; \lambda; \mathcal{A} \vdash \{P_p\} \mathbb{C}_1 \{R_p\} \quad \Gamma; \lambda; \mathcal{A} \vdash \{R_p\} \mathbb{C}_2 \{Q_p\}}{\Gamma; \lambda; \mathcal{A} \vdash \{P_p\} \mathbb{C}_1 \mathbb{C}_2 \{Q_p\}}
\end{array}$$

$$\begin{array}{c}
\text{LOOP} \\
\hline
\frac{\Gamma; \lambda; \mathcal{A} \vdash \{P_p * \mathbb{B}\} \mathbb{C} \{P_p\}}{\Gamma; \lambda; \mathcal{A} \vdash \{P_p\} \text{ while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\}}
\end{array}$$

$$\begin{array}{c}
\text{CONDITIONAL} \\
\hline
\frac{\Gamma; \lambda; \mathcal{A} \vdash \{P_p * \mathbb{B}\} \mathbb{C}_1 \{Q_p\} \quad \Gamma; \lambda; \mathcal{A} \vdash \{P_p * \neg \mathbb{B}\} \mathbb{C}_2 \{Q_p\}}{\Gamma; \lambda; \mathcal{A} \vdash \{P_p\} \text{ if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \} \{Q_p\}}
\end{array}$$

$$\begin{array}{c}
\text{FUNCTIONCALL} \\
\hline
\frac{(\lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p(\vec{z}) \mid P(x, \vec{z}) \rangle \text{ f}(\vec{z}) \quad \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle) \in \Gamma}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p(\vec{z}) * \vec{z} = \vec{\mathbb{E}} \mid P(x, \vec{z}) \rangle}$$

$$\begin{array}{c}
\Gamma; \lambda; \mathcal{A} \vdash \\
\mathbf{x} := \text{f}(\vec{\mathbb{E}}); \\
\exists (y, r) \in Y \times \text{Val}. \langle Q_p(x, y, \vec{z}, r) * \mathbf{x} = r \mid Q(x, y, \vec{z}, r) \rangle
\end{array}$$

$$\begin{array}{c}
\text{FORK} \\
\hline
\frac{(\lambda; \emptyset \vdash \mathbb{W}x \in 1. \langle P_p(\vec{z}) \mid \text{True} \rangle \text{ f}(\vec{z}) \quad \exists y \in 1. \langle Q_p(\vec{z}, \text{ret}) \mid \text{True} \rangle) \in \Gamma}{\Gamma; \lambda; \mathcal{A} \vdash \{P_p(\vec{z}) * \vec{z} = \vec{\mathbb{E}}\} \text{ fork } \text{f}(\vec{\mathbb{E}}); \{\text{True}\}}
\end{array}$$

$$\begin{array}{c}
\text{ALLOCATION} \\
\hline
\Gamma; \lambda; \mathcal{A} \vdash \{\mathbb{E} = v * v > 0\} \mathbf{x} := \text{alloc}(\mathbb{E}); \{\mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + v - 1) \mapsto \_ \}
\end{array}$$

$$\begin{array}{c}
\text{ASSIGNMENT} \\
\hline
\Gamma; \lambda; \mathcal{A} \vdash \{\mathbf{x} = v\} \mathbf{x} := \mathbb{E}; \{\mathbf{x} = \mathbb{E}[v/\mathbf{x}]\}
\end{array}$$

$$\begin{array}{c}
\text{LOOKUP} \\
\hline
\Gamma; 0; \mathcal{A} \vdash \mathbb{W}v \in \text{Val}. \langle z = \mathbb{E} \mid z \mapsto v \rangle \mathbf{x} := [\mathbb{E}]; \quad \exists y \in 1. \langle \mathbf{x} = v \mid z \mapsto v \rangle
\end{array}$$

$$\begin{array}{c}
\text{MUTATION} \\
\hline
\Gamma; 0; \mathcal{A} \vdash \mathbb{W}v \in \text{Val}. \langle z_1 = \mathbb{E}_1 * z_2 = \mathbb{E}_2 \mid z_1 \mapsto v \rangle [\mathbb{E}_1] := \mathbb{E}_2; \quad \exists y \in 1. \langle \text{True} \mid z_1 \mapsto z_2 \rangle
\end{array}$$

$$\begin{array}{c}
\text{COMPAREANDSET} \\
\hline
\Gamma; 0; \mathcal{A} \vdash \mathbb{W}v_1 \in \text{Val}. \langle z = \mathbb{E}_1 * v_2 = \mathbb{E}_2 * v_3 = \mathbb{E}_3 \mid z \mapsto v_1 \rangle \\
\mathbf{x} := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3); \\
\exists y \in 1. \langle (v_1 = v_2 * \mathbf{x} \neq 0) \vee (v_1 \neq v_2 * \mathbf{x} = 0) \mid (v_1 = v_2 * z \mapsto v_3) \vee (v_1 \neq v_2 * z \mapsto v_1) \rangle
\end{array}$$

$$\begin{array}{c}
\text{FUNCTION} \\
\hline
\frac{\text{code}(\eta(\text{f})) = \mathbb{C} \text{ return } \mathbb{E}; \quad \text{vars}(\eta(\text{f})) = \vec{\mathbf{x}}}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p(\vec{z}) * \vec{\mathbf{x}} = \vec{z} \mid P(x, \vec{z}) \rangle}$$

$$\begin{array}{c}
\Gamma; \lambda; \mathcal{A} \vdash \\
\mathbb{C} \\
\exists (y, \text{ret}) \in Y \times \text{Val}. \langle Q_p(x, y, \vec{z}, \text{ret}) * \text{ret} = \mathbb{E} \mid Q(x, y, \vec{z}, \text{ret}) \rangle
\end{array}$$

$$\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p(\vec{z}) \mid P(x, \vec{z}) \rangle \text{ f}(\vec{z}) \quad \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle$$

Figure 4.1: Programming language proof rules of the TaDA logic.

$$\begin{array}{c}
\text{OPENREGION} \\
\frac{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) \rangle}{\Gamma; \lambda + 1; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) \rangle} \\
\\
\text{USEATOMIC} \\
\frac{a \notin \mathcal{A} \quad \forall x \in X. (x, f(x)) \in \mathcal{T}_t(\mathbf{G})^* \quad \Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\mathbf{G}]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) \rangle}{\Gamma; \lambda + 1; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\mathbf{G}]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) \rangle} \\
\\
\text{UPDATEREGION} \\
\frac{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \left\langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \right\rangle \mathbb{C} \quad \exists (y, w) \in Y \times W. \left\langle Q_p(x, y, w) \mid \begin{array}{l} I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) \\ \vee I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y) \end{array} \right\rangle}{\begin{array}{c} \mathbb{W}x \in X. \langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \rangle \\ \mathbb{C} \\ \Gamma; \lambda + 1; a : x \in X \rightsquigarrow W, \mathcal{A} \vdash \\ \exists (y, w) \in Y \times W. \left\langle Q_p(x, y, w) \mid \begin{array}{l} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * a \Rightarrow (x, w) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right\rangle \end{array}} \\
\\
\text{MAKEATOMIC} \\
\frac{a \notin \mathcal{A} \quad \{(x, y) \mid x \in X, y \in Y\} \subseteq \mathcal{T}_t(\mathbf{G})^* \quad \Gamma; \lambda'; a : x \in X \rightsquigarrow Y, \mathcal{A} \vdash \{P_p * \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * a \Rightarrow \blacklozenge\} \mathbb{C} \quad \{\exists x \in X, y \in Y. Q_p(x, y) * a \Rightarrow (x, y)\}}{\Gamma; \lambda'; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, y) * [\mathbf{G}]_a \rangle}
\end{array}$$

Figure 4.2: Atomicity proof rules of the TaDA logic.

The **OPENREGION** rule allows us to access the contents of a shared region without updating its abstract state. The command may change the concrete state of the region, so long as the abstract state is preserved.

The **USEATOMIC** rule allows us to update the abstract state of a shared region. To do so, it is necessary to have a guard for the region being updated, such that the change in state is permitted by this guard according to the transition system associated with the region. This rule takes a $\mathbb{C}$ which (abstractly) atomically updates the region a from some state $x \in X$ to the state $f(x)$. It requires the guard $\mathbf{G}$ for the region, which allows the update according to the transition system, as established by one of the premisses. Another premiss states that the command $\mathbb{C}$ performs the update described by the transition system of region a in an atomic way. This allows us to conclude that the region a is updated atomically by the command $\mathbb{C}$. Note that the command is not operating at the same level of abstraction as the region a . Instead it is working at a lower level of abstraction, which means that if it is atomic at that level it will also be atomic at the region a level.

The **UPDATEREGION** rule similarly allows us to update the abstract state of a shared region, but this time the authority comes from the atomicity context instead of a guard. In order to perform such an update, the atomic update to the region must not already have happened, indicated by $a \Rightarrow \blacklozenge$ in the precondition of the conclusion. In the postcondition, there are two cases: either the appropriate update happened, or no update happened. If it did happen, the new state of the region is some $w \in W$, and both x and w are recorded in the atomicity tracking resource. If it did not, then both the region's abstract state and the atomicity tracking resource are unchanged. The premiss requires the command to make a corresponding update to the concrete state of the region. The atomicity context and tracking resource are not present in the premiss; their purpose is rather to record information about the atomic

update that is performed for use further down the proof tree.

It is necessary for the update region rule to account for both the case where the update occurs and where it does not. One might expect that the case with no update could be dealt with by the open region rule, and the results combined using a disjunction rule. However, a general disjunction rule is not sound for atomic triples. (If we have $\langle P_1 \rangle \mathbb{C} \langle Q \rangle$ and $\langle P_2 \rangle \mathbb{C} \langle Q \rangle$, we may not have $\langle P_1 \vee P_2 \rangle \mathbb{C} \langle Q \rangle$ since $\mathbb{C}$ might rely on the environment not changing between P_1 and P_2 .) The proof of the atomic specification for the spin lock uses the conditional nature of the update region rule.

We revisit the **MAKEATOMIC** rule. As before, a guard in the conclusion must permit the update in accordance with the transition system for the region. This is replaced in the premiss by the atomicity context and atomicity tracking resource, which tracks the occurrence of the update. One difference is the inclusion of the private state, which is effectively preserved between the premiss and the conclusion. A second difference is the $\exists$ -binding of the resulting state of the atomic update. This allows the private state to reflect the result of the update.

Until now, we have elided a detail of the proof system: region levels. Each judgement of TaDA includes a region level λ in the context. This level is simply a number that indicates that only regions below level λ may be opened in the derivation of the judgement. For this to be meaningful, each region is associated with a level (indicated as a superscript) and rules that open regions require that the level of the judgement is higher than the level of the region being opened. The purpose of the levels is to ensure that a region can never be opened twice in a single branch of the proof tree, which could unsoundly duplicate resources. The rules that open regions enforce this by requiring the level of the conclusion ($\lambda + 1$) to be above the level of the region (λ), which is also the level of the premiss. For our examples, the level of each module's regions just needs to be greater than the levels of modules that it uses.

In all of our examples, the atomicity context describes an update to a single region. In the logic, there is no need to restrict in this way, and an atomicity context $\mathcal{A}$ may describe updates to multiple regions (although only one update to each). Both atomic and non-atomic judgements may have atomicity contexts.

The last set of rules is given in Figure 4.3. In order to define the **FRAME** rule, we introduce two auxiliary functions.

Definition 4.32 (Program Variables Sets). The set of program variables in an assertion is denoted by $\text{pvars}(P)$.

Definition 4.33 (Modified Sets). The set of modified variables, $\text{mods}(-) : \text{Cmd} \rightarrow \mathcal{P}(\text{Var})$, is defined

FRAME

$$\frac{\text{mods}(\mathbb{C}) \cap \text{pvars}(R') = \emptyset \quad \text{mods}(\mathbb{C}) \cap \text{pvars}(R(x)) = \emptyset \quad \Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle R' * P_p \mid R(x) * P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle R' * Q_p(x, y) \mid R(x) * Q(x, y) \rangle}$$

SUBSTITUTION

$$\frac{f : X' \rightarrow X \quad g : Y' \rightarrow Y \quad \Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y' \in Y'. \langle Q_p(x, g(y')) \mid Q(x, g(y')) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x' \in X'. \langle P_p \mid P(f(x')) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(f(x'), y) \mid Q(f(x'), y) \rangle}$$

CONSEQUENCE

$$\frac{\lambda; \mathcal{A} \vdash P_p \preceq P'_p \quad \forall x \in X, y \in Y. \lambda; \mathcal{A} \vdash Q'_p(x, y) \preceq Q_p(x, y) \quad \forall x \in X, y \in Y. \lambda; \mathcal{A} \vdash Q'(x, y) \preceq Q(x, y) \quad \Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q'_p(x, y) \mid Q'(x, y) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}$$

AEXISTS

$$\frac{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}z \in \mathbf{1}. \langle P_p \mid \exists x \in X. P(x) \rangle \mathbb{C} \quad \exists x \in X, y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}$$

AWEAKENING1

$$\frac{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P' * P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q'(x, y) * Q(x, y) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p * P' \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) * Q'(x, y) \mid Q(x, y) \rangle}$$

AWEAKENING2

$$\frac{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}{\forall x \in X. \Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}z \in \mathbf{1}. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}$$

AWEAKENING3

$$\frac{\lambda' \leq \lambda \quad \Gamma; \lambda'; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}{\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle}$$

Figure 4.3: Auxiliary proof rules of the TaDA logic.

for a given command as follows:

$$\begin{aligned} \text{mods}(\text{skip};) &= \emptyset \\ \text{mods}(\mathbb{C}_1 \mathbb{C}_2) &= \text{mods}(\mathbb{C}_1) \cup \text{mods}(\mathbb{C}_2) \\ \text{mods}(\text{while } (\mathbb{B}) \{ \mathbb{C} \}) &= \text{mods}(\mathbb{C}) \\ \text{mods}(\text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}) &= \text{mods}(\mathbb{C}_1) \cup \text{mods}(\mathbb{C}_2) \\ \text{mods}(\mathbf{x} := \mathbf{f}(\vec{\mathbb{E}});) &= \{ \mathbf{x} \} \\ \text{mods}(\text{fork } \mathbf{f}(\vec{\mathbb{E}});) &= \emptyset \\ \text{mods}(\mathbf{x} := \text{alloc}(\mathbb{E});) &= \{ \mathbf{x} \} \\ \text{mods}(\mathbf{x} := \mathbb{E};) &= \{ \mathbf{x} \} \\ \text{mods}(\mathbf{x} := [\mathbb{E}];) &= \{ \mathbf{x} \} \\ \text{mods}([\mathbb{E}_1] := \mathbb{E}_2;) &= \emptyset \\ \text{mods}(\mathbf{x} := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3);) &= \{ \mathbf{x} \} \end{aligned}$$

The functions are going to be used to guarantee that a command that modifies program variables does not invalidate the frame.

The FRAME rule, as in separation logic, allows us to add the same resources to the pre- and postcondition, which are untouched by the command. Our frame rule separately adds to both the private and public parts. Note that the frame for the public part may be parametrised by the $\mathbb{W}$ -bound variable x , additionally we require the frame to be stable as the pre- and postcondition assertions.

The SUBSTITUTION rule allows us to change the domain of $\mathbb{W}$ -bound variables. A consequence of this rule is that we can instantiate $\mathbb{W}$ -variables much like universally quantified variables, simply by choosing X' to be a single-element set.

The CONSEQUENCE rule allows to strengthen the precondition and weaken the postcondition, using view shifts which we define later in Definition 4.60.

The AEXISTS allows us to eliminate existential quantifiers in the public state in the premiss of the conclusion.

The AWEAKENING1 rule allows us to convert private state from the conclusion into public state in the premiss. The first effectively implements a non-atomic specification with an atomic specification, dividing the assertions between the public and private parts. This allows us to forget about the (abstract) atomicity of an operation.

The AWEAKENING2 rule allow us to widen the interference from the conclusion to the premiss. This rule combined with AWEAKENING1 allows us make a implement a fully non-atomic specification with an atomic specification that potentially supports more interference.

The AWEAKENING3 rule allow us to implement a specification with a command that requires a lower level.

Previously in chapter 3 we have introduced simplified versions of the MAKEATOMIC, UPDATEREGION and USEATOMIC rules. We will show that they can be derived from the existing rules.

We start with the MAKEATOMIC rule, recall the simplified rule as follows:

$$\frac{\{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_t(G)^* \quad a : x \in X \rightsquigarrow Q(x) \vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * a \Rightarrow \blacklozenge \right\} \mathbb{C} \left\{ \exists x \in X, y \in Q(x). a \Rightarrow (x, y) \right\}}{\vdash \mathbb{W}x \in X. \left\langle \mathbf{t}_a(\vec{z}, x) * [G]_a \right\rangle \mathbb{C} \left\langle \exists y \in Q(x). \mathbf{t}_a(\vec{z}, y) * [G]_a \right\rangle}$$

Take $\mathcal{A} = \emptyset$, then $a \notin \mathcal{A}$ holds trivially. Take $P_p = \text{True}$, $Q_p = \text{True}$, and $Y = Q(x)$ and we have established that the simplified version follows from the MAKEATOMIC rule.

The simplified version of the UPDATEREGION rule is as follows (renaming y as w):

$$\frac{\vdash \mathbb{W}x \in X. \left\langle I(\mathbf{t}_a(\vec{z}, x)) * P(x) \right\rangle \mathbb{C} \left\langle \exists w \in Q(x). I(\mathbf{t}_a(\vec{z}, w)) * Q_1(x, w) \vee I(\mathbf{t}_a(\vec{z}, x)) * Q_2(x) \right\rangle}{a : x \in X \rightsquigarrow Q(x) \vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \right\} \mathbb{C} \left\{ \begin{array}{l} \exists x \in X. (\exists w \in Q(x). Q_1(x, w) * \\ a \Rightarrow (x, w) \vee \mathbf{t}_a(\vec{z}, x) * Q_2(x) * a \Rightarrow \blacklozenge) \end{array} \right\}}$$

Starting from the UPDATEREGION rule, take $\mathcal{A} = \emptyset$, $P_p = \text{True}$, $Q_p = \text{True}$, $Y = \mathbf{1}$, $W = Q(x)$, we have:

$$\begin{array}{c} \mathbb{W}x \in X. \left\langle \text{True} \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \right\rangle \\ \mathbb{C} \\ \Gamma; \lambda + 1; a : x \in X \rightsquigarrow Q(x) \vdash \\ \exists w \in Q(x). \left\langle \text{True} \mid \begin{array}{l} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, w) * a \Rightarrow (x, w) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x) * a \Rightarrow \blacklozenge \end{array} \right\rangle \end{array}$$

We apply the CONSEQUENCE rule to weaken the postcondition, forgetting about the region in the case there was an atomic update.

$$\begin{array}{c} \mathbb{W}x \in X. \langle \text{True} \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \rangle \\ \Gamma; \lambda + 1; a : x \in X \rightsquigarrow Q(x) \vdash \\ \mathbb{C} \\ \exists w \in Q(x). \langle \text{True} \mid \begin{array}{l} Q_1(x, w) * a \Rightarrow (x, w) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x) * a \Rightarrow \blacklozenge \end{array} \rangle \end{array}$$

We apply the AEXISTS rule to weaken the abstract state of the region x and similarly the CONSEQUENCE rule for the postcondition we have:

$$\begin{array}{c} \mathbb{W}z \in \mathbf{1}. \langle \text{True} \mid \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \rangle \\ \Gamma; \lambda + 1; a : x \in X \rightsquigarrow Q(x) \vdash \\ \mathbb{C} \\ \exists y \in Q(x). \langle \text{True} \mid \begin{array}{l} \exists x \in X. Q_1(x, y) * a \Rightarrow (x, y) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x) * a \Rightarrow \blacklozenge \end{array} \rangle \end{array}$$

Finally, we apply the CONSEQUENCE rule to the postcondition to move the existential quantification of $w \in Q(x)$ and AWEAKENING1 to move the whole public state to the private state, essentially converting the atomic triple into a normal Hoare triple.

Finally, the simplified USEATOMIC rule is as follows:

$$\frac{\forall x \in X. (x, f(x)) \in \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \quad \vdash \mathbb{W}x \in X. \langle I(\mathbf{t}_a(\vec{z}, x)) * P(x) * [\mathbf{G}]_a \rangle \mathbb{C} \langle I(\mathbf{t}_a(\vec{z}, f(x))) * Q(x) \rangle}{\vdash \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, x) * P(x) * [\mathbf{G}]_a \right\} \mathbb{C} \left\{ \exists x \in X. \mathbf{t}_a(\vec{z}, f(x)) * Q(x) \right\}}$$

Starting from the USEATOMIC rule, take $\mathcal{A} = \emptyset$, $P_p = \text{True}$, $Q_p = \text{True}$ and $Y = \mathbf{1}$ we have:

$$\Gamma; \lambda + 1; \emptyset \vdash \mathbb{W}x \in X. \langle \text{True} \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\mathbf{G}]_a \rangle \mathbb{C} \quad \exists y \in \mathbf{1}. \langle \text{True} \mid \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x) \rangle$$

We apply the AEXISTS rule to weaken the abstract state of the region x and similarly the CONSEQUENCE rule for the postcondition we have:

$$\begin{array}{c} \mathbb{W}z \in \mathbf{1}. \langle \text{True} \mid \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\mathbf{G}]_a \rangle \\ \Gamma; \lambda + 1; \emptyset \vdash \\ \mathbb{C} \\ \exists y \in \mathbf{1}. \langle \text{True} \mid \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x) \rangle \end{array}$$

Finally, we apply the AWEAKENING1 to move the whole public state to the private state, essentially converting the atomic triple into a normal Hoare triple.

4.5 Model

We now present the semantic model of the TaDA program logic, which we will use to define the semantics for the assertions and judgements in the logic.

Definition 4.34 (Variable Names). Assume a set of *logic variable names* LVar , ranged over by $x, y, \dots$.

Definition 4.35 (Guards and Guard Algebras). Assume a set **Guard** that will contain all guards that we might wish to use, ranged over by $g, g_1, \dots$. A *guard algebra* $\zeta = (\mathbf{G}, \bullet, \mathbf{0}, \mathbf{1})$ consists of:

- a carrier set $\mathbf{G} \subseteq \mathbf{Guard}$;
- an associative, commutative partial binary operator $\bullet : \mathbf{G} \times \mathbf{G} \rightharpoonup \mathbf{G}$;
- an identity element $\mathbf{0} \in \mathbf{G}$, with $\mathbf{0} \bullet g = g$ for all $g \in \mathbf{G}$; and
- a maximal element $\mathbf{1} \in \mathbf{G}$, with $g \leq \mathbf{1}$ for all $g \in \mathbf{G}$,

where

$$g_1 \leq g_2 \stackrel{\text{def}}{\iff} \exists g_3. g_1 \bullet g_3 = g_2.$$

Let **GAlg** denote the set of all guard algebras, ranged over by ζ and $\mathbf{G}_\zeta$ denote the carrier set for guard algebra ζ .

Note that a guard algebra is a separation algebra (in the sense of [13]) with a single unit, $\mathbf{0}$.

Definition 4.36 (Abstract States and Transition Systems). Assume a set **AState** that will contain all abstract region states that we might wish to use. For a given guard algebra ζ , a *guard-labelled transition system* $\mathcal{T} : \mathbf{G}_\zeta \rightarrow_{\text{mon}} \mathcal{P}(\mathbf{AState} \times \mathbf{AState})$ is a mapping from guards to relations between abstract states. Let $\mathcal{T}(g)^*$ denote the reflexive-transitive closure of the relation $\mathcal{T}(g)$ for $g \in \mathbf{G}_\zeta$. Let **ASTS** $_\zeta$ denote the set of all ζ -labelled transition systems.

The mapping is monotone with respect to the resource ordering ($\leq_\zeta$) and subset ordering ($\subseteq$), meaning that having more guard resource permits more transitions.

Definition 4.37 (Abstract Region Types). Assume a set **RTName** of region type names. An abstract region typing

$$t \in \mathbf{ARType} \stackrel{\text{def}}{=} \mathbf{RTName} \rightarrow \coprod_{\zeta \in \mathbf{GAlg}} \mathbf{ASTS}_\zeta$$

maps region type names to pairs of guard algebras and guard-labelled transition systems.

Definition 4.38 (Abstract Predicates). Assume a set **APName** of abstract predicate names. An abstract predicate $a \in \mathbf{APName} \times \mathbf{Val}^*$ consists of an abstract predicate name and a list of parameters. An abstract predicate bag $b \in \mathbf{APBag} \stackrel{\text{def}}{=} \mathcal{M}_{\text{fin}}(\mathbf{APName} \times \mathbf{Val}^*)$ is a finite multiset of abstract predicates. Abstract predicate bags form a separation algebra $(\mathbf{APBag}, \cup, \emptyset)$, where $\cup$ is multiset union, and $\emptyset$ is the empty multiset. Abstract predicate bags are ordered by the usual subset order $\subseteq$, which corresponds to the resource order.

Definition 4.39 (Levels). Assume a level $\lambda \in \mathbf{Level} \stackrel{\text{def}}{=} \mathbb{N}$. Levels are ordered by the usual well-founded ordering on natural numbers.

Remark 4.1 (On levels). It would be possible to take the levels from a more general well-founded order. This might be useful if we need some kind of unbounded nesting of regions.

Definition 4.40 (Region Identifiers). Assume a countably infinite set of region identifiers, **RId**.

Definition 4.41 (Region Assignments). A region assignment $r \in \mathbf{RAss} \stackrel{\text{def}}{=} \mathbf{RId} \rightarrow_{\text{fin}} \mathbf{Level} \times \mathbf{RTName} \times \mathbf{Val}^*$ is a finite partial function from region identifiers to levels and parametrised region type names. Region assignments are ordered by extension ordering: $r_1 \leq r_2 \stackrel{\text{def}}{\iff} \forall a \in \text{dom}(r_1). r_2(a) = r_1(a)$.

For the following semantic definitions, we assume a fixed abstract region typing $t \in \mathbf{ARType}$.

Definition 4.42 (Guard Assignments). Given a region assignment r , a guard assignment

$$\gamma \in \mathbf{GAss}_r \stackrel{\text{def}}{=} \prod_{a \in \text{dom}(r)} \mathbf{G}_{\zeta(t(r(a)))}$$

is a mapping from the regions declared in r to guards of the appropriate type for each region. Guard assignments form a separation algebra $(\mathbf{GAss}_r, \bullet, \lambda a. \mathbf{0}_{\zeta(t(r(a)))})$, where $\bullet$ is the pointwise lift of the guard combination operators:

$$\gamma_1 \bullet \gamma_2 \stackrel{\text{def}}{=} \lambda a. \gamma_1(a) \bullet \gamma_2(a)$$

For $\gamma_1 \in \mathbf{GAss}_{r_1}$, $\gamma_2 \in \mathbf{GAss}_{r_2}$ with $r_1 \leq r_2$, guards assignments are ordered pointwise-extensionally:

$$\gamma_1 \leq \gamma_2 \stackrel{\text{def}}{\iff} \forall a \in \text{dom}(\gamma_1). \gamma_1(a) \leq \gamma_2(a).$$

Definition 4.43 (Region States). Given a region assignment r , a region state

$$\rho \in \mathbf{RState}_r \stackrel{\text{def}}{=} \text{dom}(r) \rightarrow \mathbf{AState}$$

is a mapping from the regions declared in r to abstract states. For $\rho_1 \in \mathbf{RState}_{r_1}$, $\rho_2 \in \mathbf{RState}_{r_2}$ with $r_1 \leq r_2$, region states are ordered extensionally: $\rho_1 \leq \rho_2 \stackrel{\text{def}}{\iff} \forall a \in \text{dom}(\rho_1). \rho_1(a) = \rho_2(a)$.

Definition 4.44 (Worlds). A *world*

$$w \in \mathbf{World} \stackrel{\text{def}}{=} \prod_{r \in \mathbf{RAss}} (\mathbf{Heap} \times \mathbf{APBag} \times \mathbf{GAss}_r \times \mathbf{RState}_r)$$

consists of a region assignment, a heap, an abstract predicate bag, a guard assignment, and a region state. Worlds can be combined, provided they agree on the region assignment and region state, by combining the remaining components in the appropriate separation algebras. Thus, worlds form a (multi-unit) separation algebra $(\mathbf{World}, \bullet, \text{emp})$ where

$$(r, h_1, b_1, \gamma_1, \rho) \bullet (r, h_2, b_2, \gamma_2, \rho) \stackrel{\text{def}}{=} (r, h_1 \uplus h_2, b_1 \cup b_2, \gamma_1 \bullet \gamma_2, \rho)$$

$$\text{emp} \stackrel{\text{def}}{=} \{(r, \emptyset, \emptyset, \lambda a. \mathbf{0}_{\zeta(t(r(a)))}, \rho) \mid r \in \mathbf{RAss}, \rho \in \mathbf{RState}_r\}$$

Worlds are also ordered by the product order. If $w_1 \leq w_2$, then w_2 may be obtained from w_1 by introducing new regions (with arbitrary associated type name and state) and adding heap, abstract-predicate and guard resources.

Definition 4.45 (World Predicates). A world predicate $P \in \mathbf{WPred} \stackrel{\text{def}}{=} \mathcal{P}^\uparrow(\mathbf{World})$ is a set of worlds that is upwards closed with respect to the world ordering. That is, if $w \in P$ and $w \leq w'$ then $w' \in P$. This is also sometimes called intuitionistic interpretation.

The composition operator on worlds is lifted to world predicates:

$$P_1 * P_2 \stackrel{\text{def}}{=} \{w \mid \exists w_1 \in P_1, w_2 \in P_2. w = w_1 \bullet w_2\}$$

(That the results is upwards closed is not difficult to check: any extension to the composition of two worlds can be tracked back and applied to one of the components.) The $*$ operator is associative and commutative with identity World . To denote $*$ iterated over a finite set X , we write $\bigotimes_{x \in X} P(x)$.

Remark 4.2 (Intuitionistic Interpretation). An alternative approach would be to define world predicates using a classical interpretation. This would require changes in the proof rules, such as the MAKEATOMIC , to make sure that they do not forget regions.

Definition 4.46 (Worlds with Atomic Tracking). The atomic tracking separation algebra is defined to be $((\text{AState} \times \text{AState}) \uplus \{\blacklozenge, \blacklozenge\}, \bullet, (\text{AState} \times \text{AState}) \cup \{\blacklozenge\})$, where $\bullet$ is defined by

$$\begin{aligned}\blacklozenge \bullet \blacklozenge &= \blacklozenge = \blacklozenge \bullet \blacklozenge \\ \blacklozenge \bullet \blacklozenge &= \blacklozenge \\ (x, y) \bullet (x, y) &= (x, y)\end{aligned}$$

and undefined in all other cases. The resource ordering on this separation algebra is characterised by the two rules: $k \leq k$ (for all $k \in (\text{AState} \times \text{AState}) \uplus \{\blacklozenge, \blacklozenge\}$) and $\blacklozenge \leq \blacklozenge$.

Given a finite set of region identifiers $\mathcal{R} \subseteq_{\text{fin}} \text{Rld}$, a world with atomic tracking $\varphi \in \text{AWorld}_{\mathcal{R}} \stackrel{\text{def}}{=} \text{World} \times (\mathcal{R} \rightarrow (\text{AState} \times \text{AState}) \uplus \{\blacklozenge, \blacklozenge\})$ consists of a world together with a mapping that associates atomic tracking resources with each region in $\mathcal{R}$. The mapping records if an atomic update has taken place on a region, and, if so, what state change the region underwent in the update. Specifically, $\blacklozenge$ and $\blacklozenge$ record that the atomic update has not yet happened, while (x, y) records that the update has happened, and it entailed updating the abstract state from x to y . The difference between $\blacklozenge$ and $\blacklozenge$ is that $\blacklozenge$ embodies a right to perform the update, while $\blacklozenge$ does not.

By lifting $\bullet$ to maps, the maps form a separation algebra. Consequently, by combining the operators of its components, $\text{AWorld}_{\mathcal{R}}$ is also an ordered separation algebra. We consider that $\text{World} = \text{AWorld}_{\emptyset}$. As with worlds, we consider predicates over worlds with atomic tracking $P \in \text{AWPred}_{\mathcal{R}} \stackrel{\text{def}}{=} \mathcal{P}^{\uparrow}(\text{AWorld}_{\mathcal{R}})$ to be upwards-closed sets. These predicates similarly have a $*$ operator.

Definition 4.47 (Atomicity Context). An atomicity context $\mathcal{A} \in \text{AContext} \stackrel{\text{def}}{=} \text{Rld} \rightarrow_{\text{fin}} \text{AState} \rightarrow \mathcal{P}(\text{AState})$ is a (finite) partial mapping from region identifiers to partial, non-deterministic abstract state transformers. In the context of proving that an operation is abstractly atomic, the atomicity context records the abstract operation to be performed. This has implications in terms of both how the thread performing the operation and the environment can update the region mentioned in the context.

Definition 4.48 (Rely Relation). For a given atomicity context $\mathcal{A} \in \text{AContext}$, with $\mathcal{R} = \text{dom}(\mathcal{A})$, the rely relation $R_{\mathcal{A}} \subseteq \text{AWorld}_{\mathcal{R}} \times \text{AWorld}_{\mathcal{R}}$ is the smallest reflexive-transitive relation that satisfies the following rules:

$$\begin{aligned}& \frac{g \# g' \quad (s, s') \in \mathcal{T}_{t(n)}(g')^* \quad (d(a) \in \{\blacklozenge, \blacklozenge\} \Rightarrow s' \in \text{dom}(\mathcal{A}(a)))}{(r[a \mapsto n], h, b, \gamma[a \mapsto g], \rho[a \mapsto s], d) R_{\mathcal{A}} (r[a \mapsto n], h, b, \gamma[a \mapsto g], \rho[a \mapsto s'], d)} \\& \frac{(s, s') \in \mathcal{A}(a)}{(r[a \mapsto n], h, b, \gamma, \rho[a \mapsto s], d[a \mapsto \blacklozenge]) R_{\mathcal{A}} (r[a \mapsto n], h, b, \gamma, \rho[a \mapsto s'], d[a \mapsto (s, s')])}\end{aligned}$$

Interference by the environment is abstracted by the rely relation.

The first rule expresses that the environment may make any update to a region for which it can have a guard that permits it in the corresponding transition system. (It can only have such a guard if it is compatible with the guard held by the thread, expressed as $g \# g'$.) The exception to this is that, if an atomic update is pending then the environment must not take the state outside of those on which the atomic operation is set to perform.

The second rule expresses that having the $\blacklozenge$ entitles one to perform an update corresponding to that expressed in the atomicity context.

Note that interference is explicitly confined to the shared regions and atomic tracking resources. Furthermore, extending the atomicity context decreases the possible interference of the environment.

Definition 4.49 (Stable Predicates). Given an atomicity context $\mathcal{A} \in \text{AContext}$, the stable predicates are those which are closed under the associated rely relation:

$$\mathcal{A} \models P \text{ stable} \stackrel{\text{def}}{\iff} R_{\mathcal{A}}(P) \subseteq P.$$

We call the stable predicates *views* (as in [13]) and denote the set of views (in atomicity context $\mathcal{A}$) by $\text{View}_{\mathcal{A}}$. Note that if $\mathcal{A}'$ is an extension of $\mathcal{A}$ then $\text{View}_{\mathcal{A}} \subseteq \text{View}_{\mathcal{A}'}$. We drop the subscript when the empty atomicity context is intended.

If $\mathcal{A}'$ is an extension of $\mathcal{A}$, we have a coercion from $\text{View}_{\mathcal{A}}$ to $\text{View}_{\mathcal{A}'}$ by extending the atomicity tracking component for the additional regions in every possible way.

Stable predicates are closed under $*$. That is

$$\mathcal{A} \models P \text{ stable} \wedge \mathcal{A} \models Q \text{ stable} \implies \mathcal{A} \models P * Q \text{ stable}$$

Definition 4.50 (Region Interpretation). A region interpretation $I \in \text{RInterp} \stackrel{\text{def}}{=} \text{Level} \times \text{RTName} \times \text{Val}^* \times \text{Rld} \times \text{AState} \rightarrow \text{View}$ associates a view with each abstract state of each parametrised region type.

The parameters are used to specify, for example, the address of a data structure contained in the region. The region identifier is often a necessary parameter as it is common for a region interpretation to refer to guards for the region. Here, we have purposely avoided having region interpretations directly referring to region interpretations. Impredicative CAP [54] does support this by constructing the relevant domains in the topos of trees. We opt for a simpler, if less powerful, alternative: breaking self-reference by indirection through region type names.

Definition 4.51 (Abstract Predicate Interpretation). An abstract predicate interpretation $\iota \in \text{APInterp} \stackrel{\text{def}}{=} \text{APName} \times \text{Val}^* \rightarrow \text{View}$ associates a view with each abstract predicate.

For the following, assume a fixed region interpretation I and abstract predicate interpretation ι .

Definition 4.52 (Region Collapse). Given a level $\lambda \in \text{Level}$, the region collapse of a world $\varphi \in \text{AWorld}_{\mathcal{R}'}$ is a set of worlds given by:

$$\varphi \downarrow_{\lambda} \stackrel{\text{def}}{=} \left\{ \varphi \cdot (w', \emptyset) \mid w' \in \bigotimes_{\{a \mid \exists \lambda' < \lambda. r_{\varphi}(a) = (\lambda', -, -)\}} I(r_{\varphi}(a), a, \rho_{\varphi}(a)) \right\}$$

where r_{φ} and ρ_{φ} are projections of region assignments and region states respectively, from φ . This operation is lifted to predicates in a straightforward manner: $P \downarrow_{\lambda} \stackrel{\text{def}}{=} \bigcup_{\varphi \in P} \varphi \downarrow_{\lambda}$.

Definition 4.53 (Abstract Predicate Collapse). The one-step abstract predicate collapse of a world is a set of worlds given by:

$$(r, h, b, \gamma, \rho, d)|_1 \stackrel{\text{def}}{=} \left\{ (r, h, \emptyset, \gamma, \rho, d) \cdot (w, \emptyset) \mid w \in \bigotimes_{a \in b} \iota(a) \right\}$$

This is lifted to predicates: $P|_1 \stackrel{\text{def}}{=} \bigcup_{\varphi \in P} \varphi|_1$. The one-step collapse is iterated to give the multi-step collapse: $P|_{n+1} \stackrel{\text{def}}{=} (P|_n)|_1$.

The abstract predicate collapse of a predicate applies the multi-step collapse to collapse all abstract predicates:

$$P| \stackrel{\text{def}}{=} \{\varphi \mid \exists n. \varphi \in P|_n \wedge b_\varphi = \emptyset\}$$

where b_φ is a projection of abstract predicate bags from φ .

Remark 4.3 (On the interpretation of abstract predicates). This approach to interpreting abstract predicates is different from the usual one. It effectively gives a step-indexed interpretation to the predicates: the concrete interpretation is given by the finite unfolding. If a predicate cannot be made fully concrete by finite unfolding, then its semantics will be **False**.

Definition 4.54 (Reification). The reification operation on worlds collapses the regions and the abstract predicates, and then considers only the heap portion:

$$[\varphi]_\lambda \stackrel{\text{def}}{=} \{h_{\varphi'} \mid \varphi' \in \varphi \downarrow_\lambda\}$$

This operation is lifted to predicates in the usual manner.

Definition 4.55 (Guarantee Relation). Given a level $\lambda \in \text{Level}$, and atomicity context $\mathcal{A} \in \text{AContext}$, the guarantee relation $G_{\lambda; \mathcal{A}} \subseteq \text{AWorld}_{\mathcal{R}'} \times \text{AWorld}_{\mathcal{R}'}$ is defined as:

$$\begin{aligned} \varphi G_{\lambda; \mathcal{A}} \varphi' &\stackrel{\text{def}}{\iff} \forall a. (\exists \lambda' \geq \lambda. r_\varphi(a) = (\lambda', -, -)) \implies \rho_\varphi(a) = \rho_{\varphi'}(a) \wedge \\ &\forall a \in \text{dom } \mathcal{A}. \left(\begin{aligned} &(d_\varphi(a) = d_{\varphi'}(a) \wedge \rho_\varphi(a) = \rho_{\varphi'}(a)) \vee \\ &\left(d_\varphi(a) = \blacklozenge \wedge d_{\varphi'}(a) = (\rho_\varphi(a), \rho_{\varphi'}(a)) \right) \right. \\ &\quad \left. \wedge (\rho_\varphi(a), \rho_{\varphi'}(a)) \in \mathcal{A}(a) \right) \end{aligned} \right) \end{aligned}$$

The guarantee relation enforces that regions with level λ or higher cannot be modified. It also enforces that regions mentioned in the atomicity context can only be updated using the atomicity context.

Remark 4.4 (On reification). It will be necessary to enforce that each execution step preserves regions above a certain level, because these regions will simply be dropped by the reification. If we didn't constrain them in this way, a thread could change them as it liked (resources permitting) without even making a concrete update!

4.5.1 Semantic Assertions

Logical variables in assertions are interpreted over the set of values Val and we assume that $\text{RId} \subseteq \text{Val}$, so that variables may range over region identifiers. In the semantics of assertions, free variables are evaluated using a *variable interpretation*, which maps variables to values.

Definition 4.56 (Variable Interpretations). The set of variable Interpretations is $\text{Interp} \stackrel{\text{def}}{=} \text{LVar} \rightarrow \text{Val}$, and is ranged over by $i, i_1, \dots$.

Definition 4.57 (Logical Expressions Semantics). We assume an appropriate semantics for logical expressions

$$\llbracket - \rrbracket_- : \text{LExp} \times \text{Interp} \rightarrow \text{Val}$$

Definition 4.58 (Assertion Semantics). The semantics of assertions

$$\llbracket - \rrbracket_{(-, -, -)} : \text{Assn} \times \text{Interp} \times \text{RInterp} \times \text{APInterp} \rightarrow \text{View}_{\mathcal{A}}$$

is inductively defined as follows:

$$\begin{aligned} \llbracket \text{False} \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \emptyset \\ \llbracket \text{True} \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \text{AWorld} \\ \llbracket P \wedge Q \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \mid \varphi \in \llbracket P \rrbracket_{(i, I, \iota)} \text{ and } \varphi \in \llbracket Q \rrbracket_{(i, I, \iota)} \} \\ \llbracket P \vee Q \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \mid \varphi \in \llbracket P \rrbracket_{(i, I, \iota)} \text{ or } \varphi \in \llbracket Q \rrbracket_{(i, I, \iota)} \} \\ \llbracket \exists x. P \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \bigcup_{v \in \text{Val}} \llbracket P \rrbracket_{(i[x \mapsto v], I, \iota)} \\ \llbracket \forall x. P \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \bigcap_{v \in \text{Val}} \llbracket P \rrbracket_{(i[x \mapsto v], I, \iota)} \\ \llbracket P * Q \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \mid \text{there exists } \varphi_1 \in \llbracket P \rrbracket_{(i, I, \iota)}, \varphi_2 \in \llbracket Q \rrbracket_{(i, I, \iota)} \text{ s.t. } \varphi = \varphi_1 \bullet \varphi_2 \} \\ \llbracket x \mapsto y \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid h_\varphi(\llbracket x \rrbracket_i) = \llbracket y \rrbracket_i \} \\ \llbracket P \implies Q \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \mid \varphi \in \llbracket P \rrbracket_{(i, I, \iota)} \implies \varphi \in \llbracket Q \rrbracket_{(i, I, \iota)} \} \\ \llbracket \mathbf{t}_a^\lambda(\vec{z}, x) \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid r_\varphi(\llbracket a \rrbracket_i) = (\llbracket \lambda \rrbracket_i, \mathbf{t}, \llbracket \vec{z} \rrbracket_i) \text{ and } \rho_\varphi(\llbracket a \rrbracket_i) = \llbracket x \rrbracket_i \} \\ \llbracket I(\mathbf{t}_a^\lambda(\vec{z}, x)) \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} I(\mathbf{t}_a^{\llbracket \lambda \rrbracket_i}(\llbracket \vec{z} \rrbracket_i, \llbracket x \rrbracket_i)) \\ \llbracket \mathbf{G}(\vec{z}) \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid \mathbf{G}(\llbracket \vec{z} \rrbracket_i) \leq \gamma_\varphi(\llbracket a \rrbracket_i) \} \\ \llbracket a \Rightarrow \blacklozenge \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid d_\varphi(\llbracket a \rrbracket_i) = \blacklozenge \} \\ \llbracket a \Rightarrow \blacklozenge \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid d_\varphi(\llbracket a \rrbracket_i) = \blacklozenge \text{ or } d_\varphi(\llbracket a \rrbracket_i) = \blacklozenge \} \\ \llbracket a \Rightarrow (x, y) \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \{ \varphi \in \text{AWorld} \mid d_\varphi(\llbracket a \rrbracket_i) = (\llbracket x \rrbracket_i, \llbracket y \rrbracket_i) \} \\ \llbracket \mathbf{a}(\vec{z}) \rrbracket_{(i, I, \iota)} &\stackrel{\text{def}}{=} \iota(\mathbf{a}(\llbracket \vec{z} \rrbracket_i)) \end{aligned}$$

The **False** assertion is not satisfied by any world. The **True** assertion is satisfied by any world. Conjunction and disjunction have the usual semantics. The existential quantifier is satisfied by all worlds in which the predicate P holds for some assignment of x . The universal quantifier is satisfied by all worlds in which the predicate P holds for all assignments of x . The $P * Q$ assertion is interpreted for worlds with atomicity tracking in a similar style as explained in Definition 4.45. The $x \mapsto y$ assertion is satisfied by any world that contains a heap cell with address x and value y . The $P \implies Q$ is satisfied by the usual first order logic semantics. The $\mathbf{t}_a^\lambda(\vec{z}, x)$ assertion is satisfied by any world that contains

a region with identifier a such that its level, type, parameters and abstract state match the region assertion. The semantics of region interpretation is described in Definition 4.52. The $[G(\vec{z})]_a$ assertion is satisfied by any world that contains guards for the region with identifier a that are at least G . The assertions relating to the atomicity tracking component are described in Definition 4.46. Finally, the abstract predicate semantics is described in Definition 4.53.

4.5.2 Semantic Judgements

In the Views Framework [13], primitive atomic actions are abstracted to relations on views by means of an atomic satisfaction judgement. Here, we have an analogous judgement, but which is more complex as it expresses the role of an action in performing an abstractly-atomic operation.

Definition 4.59 (Primitive Atomic Satisfaction Judgement). The primitive atomic satisfaction judgement $\lambda; \mathcal{A} \models \langle P \rangle \alpha \langle Q \rangle$, where $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, $\alpha \in \text{AAction}$ and $P, Q \in \text{View}_{\mathcal{A}}$, is defined as:

$$\lambda; \mathcal{A} \models \langle P \rangle \alpha \langle Q \rangle \stackrel{\text{def}}{\iff} \forall R \in \text{View}_{\mathcal{A}}. \forall \varphi \in P * R. \forall h \in \llbracket \varphi \rrbracket_{\lambda}. \forall h' \in \llbracket \alpha \rrbracket(h). \\ \exists \varphi'. \varphi \text{ G}_{\lambda; \mathcal{A}} \varphi' \wedge h' \in \llbracket \varphi' \rrbracket_{\lambda} \wedge \varphi' \in Q * R.$$

The primitive atomic satisfaction judgement incorporates two assertions: P , the precondition; and Q , the post condition which the atomic update will satisfy. Note that the action α changes the worlds according to the guarantee and level.

Definition 4.60 (View Shift). We define view shifts as:

$$\lambda; \mathcal{A} \models P \preceq Q \stackrel{\text{def}}{\iff} \lambda; \mathcal{A} \models \langle P \rangle \text{id} \langle Q \rangle$$

View shifts allows us to change views, such as allocating regions or weakening the type of a variable to a supertype, as long as the underlying program state remains the same.

For constructing a new region we can use the following view shift:

$$\frac{\text{REGIONCREATION} \quad G \in \mathcal{G}_{\mathbf{t}} \quad \forall a. \lambda'; \mathcal{A} \models P * [G]_a \preceq I(\mathbf{t}_a^{\lambda}(\vec{z}, x)) * Q(a)}{\lambda'; \mathcal{A} \models P \preceq \exists a. \mathbf{t}_a^{\lambda}(\vec{z}, x) * Q(a)}$$

The premiss states that if we pick a guard G , from the guard algebra for the region type we want to create, and combined with predicate P it implies the interpretation of the region for abstract state x . The predicate $Q(a)$ may contain guards from the newly created region. Then we can create a new region, which includes creating G as well.

View shifts can also be used to repartition where no concrete state changes. We use view shifts in the rule of consequence to move resources between regions. Finally, the view shift includes material implication and as such it subsumes the traditional one.

Definition 4.61 (Semantic Judgement). The semantic judgement

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \text{ c } \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

where

- $\eta \in \text{FEnv}$ is a function environment;
- $\lambda \in \text{Level}$ is a level strictly greater than that of any region that will be affected by the program;
- $\mathcal{A} \in \text{AContext}$ is the atomicity context, which constrains updates to regions on which an abstractly atomic update is to be performed;
- $P_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ is the private part of the precondition, which does not correspond to resources in some opened shared region, and is parametrised by the valuation of program variables;
- $P \in X \rightarrow \text{View}_{\mathcal{A}}$ is the public part of the precondition, which may correspond to resources from some opened shared regions, and is parametrised by $x \in X$ that tracks the precondition at the atomic update and by the valuation of program variables;
- $c \in \text{ExtCmd}$ is the program under consideration;
- $Q_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ is the private part of the postcondition, which is parametrised by $x \in X$ that tracks the precondition at the atomic update, by $y \in Y$ that tracks the postcondition at the atomic update, and by the valuation of program variables;
- $Q \in X \times Y \rightarrow \text{View}_{\mathcal{A}}$ is the public part of the postcondition, which is similarly parametrised by $x \in X$ and $y \in Y$, and by the valuation of program variables;

is defined to be the most-general judgement that holds when the following conditions hold:

- For all $\sigma, \sigma_1 \in \text{Store}$, $c_1 \in \text{ExtCmd}$, $\alpha \in \text{AAction}$ with $(\sigma, c) \xrightarrow{\alpha}_{\eta} (\sigma_1, c_1)$, for all $x \in X$, there exists $P'_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$, $P''_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ such that

$$\eta; \lambda; \mathcal{A} \models \langle P_p(\sigma) * P(x) \rangle \alpha \langle P'_p(\sigma_1) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * Q(x, y) \rangle \quad (4.1)$$

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle_{c_1} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle, \quad (4.2)$$

$$\text{and for all } y \in Y, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\}_{c_1} \{Q_p(x, y)\}. \quad (4.3)$$

- For all $\sigma, \sigma_1 \in \text{Store}$, $c_1 \in \text{ExtCmd}$, $\mathbf{f}$, $\vec{v}$ with $(\sigma, c) \xrightarrow{\text{spawn}(\mathbf{f}, \vec{v})}_{\eta} (\sigma_1, c_1)$, for all $x \in X$, there exist $P'_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$, $P''_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ and $P_{\mathbf{f}} \in \text{Store} \rightarrow \text{View}$ such that for all $\sigma_{\mathbf{f}} \in \text{Store}$ with $\text{vars}(\eta(\mathbf{f})) = \vec{x}$, $\sigma_{\mathbf{f}}(\vec{x}) = \vec{v}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}$,

$$\lambda; \mathcal{A} \models P_p(\sigma) * P(x) \preceq P'_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) * P(x) \vee \exists y \in Y. P''_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) * Q(x, y), \quad (4.4)$$

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle_{c_1} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle, \quad (4.5)$$

$$\text{for all } y \in Y, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\}_{c_1} \{Q_p(x, y)\}, \quad (4.6)$$

$$\begin{aligned} & \mathbb{W}x \in 1. \langle \lambda \sigma_{\mathbf{f}}. P_{\mathbf{f}}(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \text{True} \rangle \\ \text{and } \eta; \lambda; \emptyset \models & \quad \mathbb{C} \quad . \end{aligned} \quad (4.7)$$

$$\exists(y, \text{ret}) \in 1 \times \text{Val}. \langle \lambda \sigma_{\mathbf{f}}. Q_{\mathbf{f}}(\vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \text{True} \rangle$$

- If $c = \text{skip}$; then, for all $\sigma \in \text{Store}$, $x \in X$, there exists $y \in Y$ such that

$$\lambda; \mathcal{A} \models P_p(\sigma) * P(x) \preceq Q_p(x, y, \sigma) * Q(x, y). \quad (4.8)$$

Here, we adopt the syntax $\eta; \lambda; \mathcal{A} \models \{P_p\} \text{ c } \{Q_p\}$ as shorthand for

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in 1. \langle P_p \mid \text{True} \rangle \text{ c } \exists y \in 1. \langle Q_p \mid \text{True} \rangle.$$

The semantic judgement breaks down into three mutually-exclusive cases: two progressing and one terminating. The first case covers normal progress, where the thread performs some atomic action (possibly id) described by (4.1). The action may or may not perform the atomic update: the two new private views express the outcome of each case. In the case where the atomic update is not performed, the continuation (4.2) takes up this obligation. In the case where the atomic update is performed, the continuation (4.3) loses responsibility for the public part.

The second case covers forking a new thread. This is just like the first case, taking the action id (4.4), but with an additional obligation on the semantics of the new thread: we must split the private part to give a precondition for both the continuation (4.6) and the newly-forked thread (4.7). Note that the forked thread does not participate in the atomic action of the original thread.

The third case covers ordinary termination. In this case, the atomic action must be performed by the id action (since the thread is not going to perform any further actions) enforced by (4.8).

Definition 4.62 (Function Environment and Function Context Agreement). The function environment and function context agreement, given by

$$- \models - : \text{FEnv} \times \text{FunctionCtx} \rightarrow \text{Bool}$$

is defined inductively as follows:

$$\begin{aligned} \eta \models \epsilon &\stackrel{\text{def}}{\iff} \text{True} \\ \eta \models \Gamma, &\left(\begin{array}{c} \mathbb{W}x \in X. \langle P_p(\vec{z}) \mid P(x, \vec{z}) \rangle \\ \lambda; \mathcal{A} \vdash \quad \mathbf{f}(\vec{z}) \\ \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle \end{array} \right) \\ &\stackrel{\text{def}}{\iff} \eta \models \Gamma \text{ and } \eta; \lambda; \mathcal{A} \models \begin{array}{c} \mathbb{W}x \in X. \langle \lambda\sigma. P_p(\vec{z}) * \sigma(\vec{x}) = \vec{z} \mid P(x, \vec{z}) \rangle \\ \mathbb{C} \\ \exists(y, \text{ret}) \in Y \times \text{Val}. \langle \lambda\sigma. Q_p(x, y, \vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_\sigma \mid Q(x, y, \vec{z}, \text{ret}) \rangle \end{array} \\ &\text{where } \text{vars}(\eta(\mathbf{f})) = \vec{x} \text{ and } \text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}; \end{aligned}$$

Definition 4.63 (Semantic Judgement with Function Context).

$$\Gamma; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \quad c \quad \mathbb{W}y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

$$\stackrel{\text{def}}{\iff} \left(\forall \eta \in \text{FEnv}. \eta \models \Gamma \implies \eta; \lambda; \mathcal{A} \models \begin{array}{c} \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \\ c \\ \mathbb{W}y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle \end{array} \right)$$

4.6 Soundness

We prove that the TaDA program logic is sound with respect to the operational semantics.

Theorem 4.1 (Soundness of Commands). If

$$\Gamma; \lambda; \mathcal{A} \vdash \mathbb{W}\vec{x} \in X. \langle P_p \mid P(\vec{x}) \rangle \quad \mathbb{C} \quad \mathbb{W}\vec{y} \in Y. \langle Q_p(\vec{x}, \vec{y}) \mid Q(\vec{x}, \vec{y}) \rangle$$

is provable in the logic, then

$$\Gamma; \lambda; \mathcal{A} \models \mathbb{W}\vec{x} \in X. \langle P_p \mid P(\vec{x}) \rangle \quad \mathbb{C} \quad \mathbb{W}\vec{y} \in Y. \langle Q_p(\vec{x}, \vec{y}) \mid Q(\vec{x}, \vec{y}) \rangle$$

holds semantically.

To prove the Soundness of Commands theorem, we show that the proof rules hold semantically.

Lemma 4.2 (SKIP Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \text{ skip}; \{P_p\} \tag{4.9}$$

Proof. To establish (4.9), it is sufficient to establish the sub-condition (4.8):

$$\lambda; \mathcal{A} \models P_p(\sigma) \preceq P_p(\sigma) \tag{4.10}$$

The (4.10) sub-condition follows from the reflexivity of view shift. $\square$

Lemma 4.3 (SEQUENCING Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, suppose that

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \quad \mathbb{C}_1 \quad \{R_p\} \tag{4.11}$$

and

$$\eta; \lambda; \mathcal{A} \models \{R_p\} \quad \mathbb{C}_2 \quad \{Q_p\} \tag{4.12}$$

then

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \quad \mathbb{C}_1 \quad \mathbb{C}_2 \quad \{Q_p\} \tag{4.13}$$

Proof. The proof is by coinduction. We have two cases to consider, when $\mathbb{C}_1 = \text{skip}$; or otherwise.

Consider the first case where $\mathbb{C}_1 = \text{skip}$. The only possible reduction is $(\sigma, \text{skip}; \mathbb{C}_2) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C}_2)$, and so the first case of the semantic judgement applies. We have to establish the following sub-conditions,

by finding a suitable P'_p and P''_p :

$$\lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \text{ id } \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma) \rangle \quad (4.14)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} \mathbb{C}_2 \{Q_p\} \quad (4.15)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} \mathbb{C}_2 \{Q_p\} \quad (4.16)$$

Take $P'_p = \lambda\sigma. R_p$ and $P''_p = \lambda\sigma. \text{False}$, the (4.16) sub-condition becomes trivial, while the (4.15) follows from the assumption (4.12). The (4.14) sub-condition follows the assumption (4.11) and the definition of semantic judgment case (4.8) when $c = \text{skip}$.

Consider the second case where $\mathbb{C}_1$ is not a skip. It must be the case that $(\sigma, \mathbb{C}_1 \mathbb{C}_2) \xrightarrow{\alpha}_\eta (\sigma_1, c_3 \mathbb{C}_2)$, where $(\sigma, \mathbb{C}_1) \xrightarrow{\alpha}_\eta (\sigma_1, c_3)$ for some α and c_3 . We consider cases on whether α is a **spawn** action.

If $\alpha \notin \{\text{spawn}(\mathbf{f}, \vec{v}) \mid \mathbf{f} \in \text{Fun}, \vec{v} \in \text{Val}^*\}$, the first condition of the semantic judgement is the only one that may apply. We have to establish the following sub-conditions to show (4.13), by finding a suitable P'_p and P''_p :

$$\lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \alpha \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma_1) \rangle \quad (4.17)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} c_3 \mathbb{C}_2 \{Q_p\} \quad (4.18)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} c_3 \mathbb{C}_2 \{Q_p\} \quad (4.19)$$

From the assumption (4.11) and the definition of semantic judgement we have that:

$$\lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \alpha \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma_1) \rangle \quad (4.20)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} c_3 \{R_p\} \quad (4.21)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} c_3 \{R_p\} \quad (4.22)$$

The (4.17) sub-condition follows from (4.20). To establish (4.18), we have from (4.21)

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} c_3 \{R_p\}$$

and from (4.12)

$$\eta; \lambda; \mathcal{A} \models \{R_p\} \mathbb{C}_2 \{Q_p\}$$

By the coinductive hypothesis, we have

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} c_3 \mathbb{C}_2 \{Q_p\}$$

To establish (4.19), we have from (4.22) for all $y \in \mathbf{1}$

$$\eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} c_3 \{R_p\}$$

and from (4.12)

$$\eta; \lambda; \mathcal{A} \models \{R_p\} \mathbb{C}_2 \{Q_p\}$$

By the coinductive hypothesis, we have

$$\eta; \lambda; \mathcal{A} \models \left\{ P_p''(x, y) \right\} c_3 \mathbb{C}_2 \left\{ Q_p \right\}$$

If $\alpha = \text{spawn}(\mathbf{f}, \vec{v})$, where $\mathbf{f} \in \text{Fun}$ and $\vec{v} \in \text{Val}^*$, the second condition of the semantic judgement is the only one that may apply. We have to establish the following sub-conditions to show (4.13), where $\sigma_{\mathbf{f}} \in \text{Store}$ with $\text{vars}(\eta(\mathbf{f})) = \vec{x}$, $\sigma_{\mathbf{f}}(\vec{x}) = \vec{v}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}$; by finding a suitable P'_p , $P_{\mathbf{f}}$ and P_p'' :

$$\lambda; \mathcal{A} \models P_p(\sigma) \preceq P'_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \vee \exists y \in \mathbf{1}. P_p''(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \quad (4.23)$$

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} c_3 \mathbb{C}_2 \left\{ Q_p \right\} \quad (4.24)$$

$$\text{for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \left\{ P_p''(x, y) \right\} c_3 \mathbb{C}_2 \left\{ Q_p \right\} \quad (4.25)$$

$$\begin{aligned} & \mathbb{W}x \in \mathbf{1}. \left\langle \lambda \sigma_{\mathbf{f}}. P_{\mathbf{f}}(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \lambda \sigma_{\mathbf{f}}. \text{True} \right\rangle \\ \text{and } \eta; \lambda; \emptyset \models & \mathbb{C} \\ & \exists!(y, \text{ret}) \in \mathbf{1} \times \text{Val}. \left\langle \lambda \sigma_{\mathbf{f}}. Q_{\mathbf{f}}(\vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \lambda \sigma_{\mathbf{f}}. \text{True} \right\rangle \end{aligned} \quad (4.26)$$

From the assumption (4.11) and the definition of semantic judgement we have that:

$$\lambda; \mathcal{A} \models P_p(\sigma) \preceq P'_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \vee \exists y \in \mathbf{1}. P_p''(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \quad (4.27)$$

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} c_3 \left\{ R_p \right\} \quad (4.28)$$

$$\text{for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \left\{ P_p''(x, y) \right\} c_3 \left\{ R_p \right\} \quad (4.29)$$

$$\begin{aligned} & \mathbb{W}x \in \mathbf{1}. \left\langle \lambda \sigma_{\mathbf{f}}. P_{\mathbf{f}}(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \lambda \sigma_{\mathbf{f}}. \text{True} \right\rangle \\ \text{and } \eta; \lambda; \emptyset \models & \mathbb{C} \\ & \exists!(y, \text{ret}) \in \mathbf{1} \times \text{Val}. \left\langle \lambda \sigma_{\mathbf{f}}. Q_{\mathbf{f}}(\vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \lambda \sigma_{\mathbf{f}}. \text{True} \right\rangle \end{aligned} \quad (4.30)$$

The (4.23) sub-condition follows from (4.27) and the (4.26) sub-condition follows from (4.30). To establish (4.24), we have from (4.28)

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} c_3 \left\{ R_p \right\}$$

and from (4.12)

$$\eta; \lambda; \mathcal{A} \models \left\{ R_p \right\} \mathbb{C}_2 \left\{ Q_p \right\}$$

By the coinductive hypothesis, we have

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} c_3 \mathbb{C}_2 \left\{ Q_p \right\}$$

Finally, to establish (4.25), we have from (4.29) for all $y \in \mathbf{1}$

$$\eta; \lambda; \mathcal{A} \models \left\{ P_p''(x, y) \right\} c_3 \left\{ R_p \right\}$$

and from (4.12)

$$\eta; \lambda; \mathcal{A} \models \{R_p\} \mathbb{C}_2 \{Q_p\}$$

By the coinductive hypothesis, we have

$$\eta; \lambda; \mathcal{A} \models \{P_p''(x, y)\} c_3 \mathbb{C}_2 \{Q_p\}$$

□

Lemma 4.4 (LOOP Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, suppose that

$$\eta; \lambda; \mathcal{A} \models \{P_p * \mathbb{B}\} \mathbb{C} \{P_p\} \quad (4.31)$$

then

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \text{while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\} \quad (4.32)$$

Proof. Since $\text{while } (\mathbb{B}) \{ \mathbb{C} \}$ has two possible reductions, both with transition id , to show (4.32), it is sufficient to establish:

$$\eta; \lambda; \mathcal{A} \models \{P_p * \mathbb{B}\} \mathbb{C} \text{while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\} \quad (4.33)$$

$$\eta; \lambda; \mathcal{A} \models \{P_p * \neg \mathbb{B}\} \text{skip}; \{P_p * \neg \mathbb{B}\} \quad (4.34)$$

This is sufficient since the first condition of the semantic judgement is the only one that may apply. For reduction $(\sigma, \text{while } (\mathbb{B}) \{ \mathbb{C} \}) \xrightarrow{\text{id}}_\eta (\sigma, \mathbb{C} \text{while } (\mathbb{B}) \{ \mathbb{C} \})$ (which requires $\mathcal{B}[\mathbb{B}]_\sigma = \text{True}$), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \text{id} \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma) \rangle \quad (4.35)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} \mathbb{C} \text{while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\} \quad (4.36)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} \mathbb{C} \text{while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\} \quad (4.37)$$

Take $P'_p = P_p * \mathbb{B}$ and $P''_p = \lambda\sigma. \text{False}$. The (4.35) and (4.37) sub-conditions become trivial, while the (4.36) reduces to (4.33). For the reduction $(\sigma, \text{while } (\mathbb{B}) \{ \mathbb{C} \}) \xrightarrow{\text{id}}_\eta (\sigma, \text{skip};)$ (which requires $\mathcal{B}[\mathbb{B}]_\sigma = \text{False}$), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \text{id} \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma) \rangle \quad (4.38)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} \text{skip}; \{P_p * \neg \mathbb{B}\} \quad (4.39)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} \text{skip}; \{P_p * \neg \mathbb{B}\} \quad (4.40)$$

Take $P'_p = P_p * \neg \mathbb{B}$ and $P''_p = \lambda\sigma. \text{False}$. Similarly, the (4.38) and (4.40) sub-conditions are trivial and the (4.39) reduces to (4.34).

To establish (4.33), we have from (4.31)

$$\eta; \lambda; \mathcal{A} \models \{P_p * \mathbb{B}\} \mathbb{C} \{P_p\}$$

By the coinductive hypothesis, we have

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \text{ while } (\mathbb{B}) \{ \mathbb{C} \} \{P_p * \neg \mathbb{B}\}$$

Now (4.33) follows from the above by the SEQUENCING Rule Lemma.

It is trivial to establish (4.34) by applying the SKIP Rule Lemma. $\square$

Lemma 4.5 (CONDITIONAL Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \{P_p * \mathbb{B}\} \mathbb{C}_1 \{Q_p\} \quad (4.41)$$

and

$$\eta; \lambda; \mathcal{A} \models \{P_p * \neg \mathbb{B}\} \mathbb{C}_2 \{Q_p\} \quad (4.42)$$

then

$$\eta; \lambda; \mathcal{A} \models \{P_p\} \text{ if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \} \{Q_p\} \quad (4.43)$$

Proof. Since $\text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}$ has two possible reductions, both with transition id , the first condition of the semantic judgement is the only one that may apply. For reduction

$$(\sigma, \text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C}_1)$$

(which requires $\mathcal{B}[\mathbb{B}]_{\sigma} = \text{True}$), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\eta; \lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \text{ id } \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma) \rangle \quad (4.44)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} \mathbb{C}_1 \{Q_p\} \quad (4.45)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} \mathbb{C}_1 \{Q_p\} \quad (4.46)$$

Take $P'_p = P_p * \mathbb{B}$ and $P''_p = \lambda\sigma.\text{False}$, the (4.44) and (4.46) sub-conditions become trivial, while the (4.45) follows from (4.41). For the reduction

$$(\sigma, \text{if } (\mathbb{B}) \{ \mathbb{C}_1 \} \text{ else } \{ \mathbb{C}_2 \}) \xrightarrow{\text{id}}_{\eta} (\sigma, \mathbb{C}_2)$$

(which requires $\mathcal{B}[\mathbb{B}]_{\sigma} = \text{False}$), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\eta; \lambda; \mathcal{A} \models \langle P_p(\sigma) \rangle \text{ id } \langle P'_p(\sigma) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma) \rangle \quad (4.47)$$

$$\eta; \lambda; \mathcal{A} \models \{P'_p\} \mathbb{C}_2 \{Q_p\} \quad (4.48)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \{P''_p(x, y)\} \mathbb{C}_2 \{Q_p\} \quad (4.49)$$

Take $P'_p = P_p * \neg \mathbb{B}$ and $P''_p = \lambda\sigma.\text{False}$. Similarly, the (4.47) and (4.49) sub-conditions are trivial and the (4.48) follows from (4.42). $\square$

Lemma 4.6 (Stack Frame). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, suppose that

$$\begin{aligned} & \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad c \\ & \exists!(y, r) \in Y \times \text{Val}. \langle \lambda\sigma. Q_p(x, y, r) * r = \mathcal{E}[\![\mathbb{E}]\!]_\sigma \mid Q(x, y, r) \rangle \end{aligned} \quad (4.50)$$

then

$$\begin{aligned} & \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad \mathbf{x} := (\sigma, c \text{ return } \mathbb{E}); \\ & \exists!(y, r) \in Y \times \text{Val}. \langle Q_p(x, y, r) * \mathbf{x} = r \mid Q(x, y, r) \rangle \end{aligned} \quad (4.51)$$

Proof. The proof is by coinduction. We have three cases to consider, one for each condition of the semantic judgement.

Consider the case where $c = \text{skip}$. To establish (4.51) we have to show that:

$$\lambda; \mathcal{A} \models P_p * P(x) \preceq Q_p(x, y, r) * \mathbf{x} = R * Q(x, y, r) \quad (4.52)$$

From (4.50) and the third condition of the semantic judgement we have that:

$$\lambda; \mathcal{A} \models P_p * P(x) \preceq Q_p(x, y, r) * r = \mathcal{E}[\![\mathbb{E}]\!]_\sigma * Q(x, y, r) \quad (4.53)$$

We are perform a view shift, corresponding to the id action

$$(\sigma_1, \mathbf{x} := (\sigma, \text{skip}; \text{return } \mathbb{E});) \xrightarrow{\text{id}}_\eta (\sigma_1[\mathbf{x} \mapsto \mathcal{E}[\![\mathbb{E}]\!]_\sigma], \text{skip};) \quad (4.54)$$

To establish (4.52), we have $r = \mathcal{E}[\![\mathbb{E}]\!]_\sigma$ from (4.53) and $\mathbf{x} = \mathcal{E}[\![\mathbb{E}]\!]_\sigma$ from (4.54).

Consider the case where c is not skip and performs an action $\alpha \notin \{\text{spawn}(\mathbf{f}, \vec{v}) \mid \mathbf{f} \in \text{Fun}, \vec{v} \in \text{Val}^*\}$, the first condition of the semantic judgement is the only one that may apply. For reduction $(\sigma_1, \mathbf{x} := (\sigma, c \text{ return } \mathbb{E});) \xrightarrow{\alpha}_\eta (\sigma_1, \mathbf{x} := (\sigma_2, c_2 \text{ return } \mathbb{E});)$, where $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_2, c_2)$, we have to establish the following sub-conditions to show (4.51), by finding a suitable P'_p and P''_p :

$$\lambda; \mathcal{A} \models \langle P_p * P(x) \rangle \alpha \langle P'_p * P(x) \vee \exists(y, r) \in Y \times \text{Val}. P''_p(x, y, r) * \mathbf{x} = R * Q(x, y, r) \rangle \quad (4.55)$$

$$\begin{aligned} & \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad \mathbf{x} := (\sigma_2, c_2 \text{ return } \mathbb{E}); \\ & \exists!(y, r) \in Y \times \text{Val}. \langle P'_p(x, y, r) \mid Q(x, y, r) \rangle \end{aligned} \quad (4.56)$$

$$\text{and for all } (y, r) \in Y \times \text{Val}, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y, r) \right\} \mathbf{x} := (\sigma', c_2 \text{ return } \mathbb{E}); \left\{ Q_p(x, y, r) * \mathbf{x} = r \right\} \quad (4.57)$$

We have from (4.50) and the first condition of the semantic judgement, for $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_2, c_2)$, we

have that:

$$\lambda; \mathcal{A} \models \langle P_p * P(x) \rangle \alpha \langle P'_p * P(x) \vee \exists(y, r) \in Y \times \text{Val}. P''_p(x, y, r) * r = \mathcal{E}[\mathbb{E}]_\sigma * Q(x, y, r) \rangle \quad (4.58)$$

$$\begin{aligned} & \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad c_2 \\ & \exists(y, r) \in Y \times \text{Val}. \langle P'_p(x, y, r) \mid Q(x, y, r) \rangle \end{aligned} \quad (4.59)$$

$$\text{and for all } (y, r) \in Y \times \text{Val}, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y, r) \right\} c_2 \left\{ \lambda\sigma. Q_p(x, y, r) * r = \mathcal{E}[\mathbb{E}]_\sigma \right\} \quad (4.60)$$

Take P'_p and P''_p from (4.58), then we have established (4.55) sub-condition. We have (4.58), by the coinductive hypothesis, we have (4.56) sub-condition. Similarly, we have (4.60), by the coinductive hypothesis, we have (4.57) sub-condition.

Finally, the case where c performs a fork action follows the same reasoning. $\square$

Lemma 4.7 (FUNCTIONCALL Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, $\Gamma \in \text{FunctionCtx}$, suppose that $\eta \models \Gamma$ and

$$\left(\lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \langle P_p(\vec{z}) \mid P(x, \vec{z}) \rangle \mathbf{f}(\vec{z}) \quad \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle \right) \in \Gamma \quad (4.61)$$

then

$$\begin{aligned} & \mathbb{W}x \in X. \langle P_p(\vec{z}) * \vec{z} = \vec{\mathbb{E}} \mid P(x, \vec{z}) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad \mathbf{x} := \mathbf{f}(\vec{\mathbb{E}}); \\ & \exists(y, r) \in Y \times \text{Val}. \langle Q_p(x, y, \vec{z}, r) * \mathbf{x} = r \mid Q(x, y, \vec{z}, r) \rangle \end{aligned} \quad (4.62)$$

Proof. Since $\mathbf{x} := \mathbf{f}(\vec{\mathbb{E}});$ has only one possible reduction with transition id , the first condition of the semantic judgement is the only one that may apply. For reduction

$$\left(\sigma, \mathbf{x} := \mathbf{f}(\vec{\mathbb{E}}); \right) \xrightarrow{\text{id}}_\eta \left(\sigma, \mathbf{x} := (\sigma_{\mathbf{f}}, \text{code}(\eta(\mathbf{f}))) ; \right),$$

where $\mathcal{E}[\vec{\mathbb{E}}]_\sigma = \sigma_{\mathbf{f}}(\text{vars}(\eta(\mathbf{f})))$, to show (4.62), we have to establish the following sub-conditions, where $\sigma_{\mathbf{f}} \in \text{Store}$ with $\text{vars}(\eta(\mathbf{f})) = \vec{\mathbf{x}}_{\mathbf{f}}$, $\sigma_{\mathbf{f}}(\vec{\mathbf{x}}) = \vec{z}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}_{\mathbf{f}};$, by finding a suitable P'_p and P''_p :

$$\eta; \lambda; \mathcal{A} \models \langle P_p(\vec{z}, \sigma) * P(x, \vec{z}) \rangle \text{id} \langle P'_p(\vec{z}, \sigma) * P(x, \vec{z}) \vee \exists(y, r) \in Y \times \text{Val}. P''_p(x, y, \vec{z}, \sigma) * Q(x, y, \vec{z}) \rangle \quad (4.63)$$

$$\begin{aligned} & \mathbb{W}x \in X. \langle P'_p(\vec{z}) \mid P(x, \vec{z}) \rangle \\ \eta; \lambda; \mathcal{A} \models & \quad \mathbf{x} := (\sigma_{\mathbf{f}}, \mathbb{C} \text{ return } \mathbb{E}_{\mathbf{f}}); \\ & \exists(y, r) \in Y \times \text{Val}. \langle Q_p(x, y, \vec{z}) * \mathbf{x} = r \mid Q(x, y, \vec{z}) \rangle \end{aligned} \quad (4.64)$$

$$\text{and for all } (y, r) \in Y \times \text{Val}, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y, \vec{z}) \right\} \mathbf{x} := (\sigma_{\mathbf{f}}, \mathbb{C} \text{ return } \mathbb{E}_{\mathbf{f}}); \left\{ Q_p(x, y, \vec{z}) * \mathbf{x} = r \right\} \quad (4.65)$$

Take $P'_p = \lambda\sigma. P_p$ and $P''_p = \lambda\sigma. \text{False}$. Recall that P_p does not depend on the store σ . The (4.63) sub-condition follows from the reflexivity of view shift and (4.65) sub-condition is trivial.

To establish (4.64), we have from (4.61)

$$\left(\lambda; \mathcal{A} \vdash \mathbb{W}x \in X. \left\langle P_p(\vec{z}) \mid P(x, \vec{z}) \right\rangle \mathbf{f}(\vec{z}) \quad \exists y \in Y. \left\langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \right\rangle \right) \in \Gamma$$

By definition (4.62), we have that

$$\begin{aligned} & \mathbb{W}x \in X. \left\langle \lambda\sigma_{\mathbf{f}}. P_p(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}_{\mathbf{f}}) = \vec{z} \mid P(x, \vec{z}) \right\rangle \\ \eta; \lambda; \mathcal{A} \models & \quad \mathbb{C} \\ & \exists(y, r) \in Y \times \text{Val}. \left\langle \lambda\sigma_{\mathbf{f}}. Q_p(x, y, \vec{z}, r) * r = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid Q(x, y, \vec{z}, r) \right\rangle \end{aligned} \quad (4.66)$$

where $\text{vars}(\eta(\mathbf{f})) = \vec{x}_{\mathbf{f}}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E};$.

The (4.64) sub-condition follows from Lemma (4.6), given we have (4.66). $\square$

Lemma 4.8 (FORK Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, $\Gamma \in \text{FunctionCtx}$, suppose that $\eta \models \Gamma$ and

$$\left(\lambda; \emptyset \vdash \mathbb{W}x \in \mathbf{1}. \left\langle P_p(\vec{z}) \mid \text{True} \right\rangle \mathbf{f}(\vec{z}) \quad \exists y \in \mathbf{1}. \left\langle Q_p(\vec{z}, \text{ret}) \mid \text{True} \right\rangle \right) \in \Gamma \quad (4.67)$$

then

$$\Gamma; \lambda; \mathcal{A} \models \left\{ P_p(\vec{z}) * \vec{z} = \vec{\mathbb{E}} \right\} \text{fork } \mathbf{f}(\vec{\mathbb{E}}); \left\{ \text{True} \right\} \quad (4.68)$$

Proof. Since $\text{fork } \mathbf{f}(\vec{\mathbb{E}});$ has only one possible reduction with transition $\alpha = \text{spawn}(\mathbf{f}, \vec{v})$, where $\mathbf{f} \in \text{Fun}$ and $\vec{v} \in \text{Val}^*$, the second condition of the semantic judgement is the only one that may apply. For reduction $(\sigma, \text{fork } \mathbf{f}(\vec{\mathbb{E}});) \xrightarrow{\text{spawn}(\mathbf{f}, \vec{v})}_{\eta} (\sigma, \text{skip};)$, where $\vec{v} = \mathcal{E}[\vec{\mathbb{E}}]_{\sigma}$, to show (4.68), we have to establish the following sub-conditions, where $\sigma_{\mathbf{f}} \in \text{Store}$ with $\text{vars}(\eta(\mathbf{f})) = \vec{x}$, $\sigma_{\mathbf{f}}(\vec{x}) = \vec{v}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E};$, by finding a suitable P'_p , $P_{\mathbf{f}}$ and P''_p and $Q_{\mathbf{f}}$:

$$\lambda; \mathcal{A} \models P_p(\vec{z}, \sigma) \preceq P'_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \vee \exists y \in \mathbf{1}. P''_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) \quad (4.69)$$

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} \text{skip}; \left\{ \text{True} \right\} \quad (4.70)$$

$$\text{for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y) \right\} \text{skip}; \left\{ \text{True} \right\} \quad (4.71)$$

$$\begin{aligned} & \mathbb{W}x \in \mathbf{1}. \left\langle \lambda\sigma_{\mathbf{f}}. P_{\mathbf{f}}(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \text{True} \right\rangle \\ \text{and } \eta; \lambda; \emptyset \models & \quad \mathbb{C} \\ & \exists(y, r) \in \mathbf{1} \times \text{Val}. \left\langle \lambda\sigma_{\mathbf{f}}. Q_{\mathbf{f}}(\vec{z}, r) * r = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \text{True} \right\rangle \end{aligned} \quad (4.72)$$

Take $P'_p = \lambda\sigma. \text{True}$, $P_{\mathbf{f}} = \lambda\sigma_{\mathbf{f}}. P_p(\vec{z})$, $P''_p = \lambda\sigma. \text{True}$ and $Q_{\mathbf{f}} = \lambda\sigma_{\mathbf{f}}. Q_p(\vec{z}, r)$, the (4.70) sub-condition follows from the reflexivity of view shift. The (4.70) and (4.71) sub-conditions follow directly from the SKIP Rule Lemma.

To establish (4.72), we have from (4.67)

$$\left(\lambda; \emptyset \vdash \mathbb{W}x \in \mathbf{1}. \left\langle P_p(\vec{z}) \mid \text{True} \right\rangle \mathbf{f}(\vec{z}) \quad \exists y \in \mathbf{1}. \left\langle Q_p(\vec{z}, \text{ret}) \mid \text{True} \right\rangle \right) \in \Gamma$$

By definition (4.62), we have that

$$\begin{aligned} \eta; \lambda; \mathcal{A} \models & \quad \mathbb{W}x \in X. \left\langle \lambda\sigma_{\mathbf{f}}. P_p(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \text{True} \right\rangle \\ & \quad \mathbb{C} \\ & \quad \exists(y, \text{ret}) \in \mathbf{1} \times \text{Val}. \left\langle \lambda\sigma_{\mathbf{f}}. Q_p(\vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \text{True} \right\rangle \end{aligned} \quad (4.73)$$

where $\text{vars}(\eta(\mathbf{f})) = \vec{x}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}$; The (4.72) sub-condition follows directly from (4.73). $\square$

Lemma 4.9 (ALLOCATION Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \left\{ \mathbb{E} = w * w > 0 \right\} \mathbf{x} := \text{alloc}(\mathbb{E}); \left\{ \mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + w - 1) \mapsto \_ \right\} \quad (4.74)$$

Proof. Since $\mathbf{x} := \text{alloc}(\mathbb{E});$ has only one possible reduction

$$(\sigma, \mathbf{x} := \text{alloc}(\mathbb{E});) \xrightarrow{\text{alloc}(\mathcal{E}[\mathbb{E}]_{\sigma}, v)}_{\eta} (\sigma[\mathbf{x} \mapsto v], \text{skip};)$$

To show (4.74), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\eta; \lambda; \mathcal{A} \models \left\langle \mathbb{E} = w * w > 0 \right\rangle \text{alloc}(\mathcal{E}[\mathbb{E}]_{\sigma}, v) \left\langle P'_p(\sigma[\mathbf{x} \mapsto v]) \vee \exists y \in \mathbf{1}. P''_p(x, y, \sigma[\mathbf{x} \mapsto v]) \right\rangle \quad (4.75)$$

$$\eta; \lambda; \mathcal{A} \models \left\{ P'_p \right\} \text{skip}; \left\{ \mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + w - 1) \mapsto \_ \right\} \quad (4.76)$$

$$\text{and for all } y \in \mathbf{1}, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y) \right\} \text{skip}; \left\{ \mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + w - 1) \mapsto \_ \right\} \quad (4.77)$$

Take $P'_p = \lambda\sigma. \mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + w - 1) \mapsto \_$ and $P''_p = \lambda\sigma. \text{False}$, To establish (4.75), fix $R \in \text{View}_{\mathcal{A}}$, $\varphi \in \mathbb{E} = w * w > 0 * R$, $h \in [\varphi]_{\lambda}$, $h' \in [\text{alloc}(\mathcal{E}[\mathbb{E}]_{\sigma}, v)](h)$. Let

$$\varphi' = (r_{\varphi}, h', b_{\varphi}, \gamma_{\varphi}, \rho_{\varphi}, d_{\varphi}).$$

By the interpretation of atomic commands, given that $w > 0$, it must be that

$$h' = \emptyset[v \mapsto k_0, \dots, v + w - 1 \mapsto k_{w-1}] \uplus h$$

and

$$\varphi' \in \mathbf{x} \mapsto \_ * \dots * (\mathbf{x} + w - 1) \mapsto \_ * R$$

Finally, the (4.77) becomes trivial and (4.76) follows from the SKIP Rule Lemma. $\square$

Lemma 4.10 (ASSIGNMENT Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \left\{ \mathbf{x} = v \right\} \mathbf{x} := \mathbb{E}; \left\{ \mathbf{x} = \mathbb{E}[v/\mathbf{x}] \right\} \quad (4.78)$$

Proof. The proof follows directly from the operational semantics and is similar to that for the ALLOCATION rule. $\square$

Lemma 4.11 (LOOKUP Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; 0; \mathcal{A} \models \mathbb{W}v \in \text{Val}. \langle z = \mathbb{E} \mid z \mapsto v \rangle \quad \mathbf{x} := [\mathbb{E}]; \quad \exists y \in 1. \langle \mathbf{x} = v \mid z \mapsto v \rangle \quad (4.79)$$

Proof. The proof follows directly from the operational semantics and is similar to that for the ALLOCATION rule. $\square$

Lemma 4.12 (MUTATION Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; 0; \mathcal{A} \models \mathbb{W}v \in \text{Val}. \langle z_1 = \mathbb{E}_1 * z_2 = \mathbb{E}_2 \mid z_1 \mapsto v \rangle \quad [\mathbb{E}_1] := \mathbb{E}_2; \quad \exists y \in 1. \langle \text{True} \mid z_1 \mapsto z_2 \rangle$$

Proof. The proof follows directly from the operational semantics and is similar to that for the ALLOCATION rule. $\square$

Lemma 4.13 (COMPAREANDSET Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\begin{aligned} \eta; 0; \mathcal{A} \models & \quad \mathbb{W}v_1 \in \text{Val}. \langle z = \mathbb{E}_1 * v_2 = \mathbb{E}_2 * v_3 = \mathbb{E}_3 \mid z \mapsto v_1 \rangle \\ & \quad \mathbf{x} := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3); \\ & \quad \exists y \in 1. \langle (v_1 = v_2 * \mathbf{x} \neq 0) \vee (v_1 \neq v_2 * \mathbf{x} = 0) \mid (v_1 = v_2 * z \mapsto v_3) \vee (v_1 \neq v_2 * z \mapsto v_1) \rangle \end{aligned}$$

Proof. The proof follows directly from the operational semantics and is similar to that for the ALLOCATION rule. $\square$

Lemma 4.14 (OPENREGION Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \rangle \quad c \quad \exists y \in Y. \langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) \rangle \quad (4.80)$$

then

$$\eta; \lambda + 1; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \rangle \quad c \quad \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) \rangle \quad (4.81)$$

Proof. Consider the case where c is not skip and performs an action $\alpha \notin \{\text{spawn}(\mathbf{f}, \vec{v}) \mid \mathbf{f} \in \text{Fun}, \vec{v} \in \text{Val}^*\}$, the first condition of the semantic judgement is the only one that may apply. For reduction $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_1, c_1)$, to show (4.81), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\eta; \lambda + 1; \mathcal{A} \models \left\langle P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \right\rangle \alpha \left\langle \begin{array}{c} P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \vee \\ \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) \end{array} \right\rangle \quad (4.82)$$

$$\eta; \lambda + 1; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \rangle \quad c_1 \quad \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) \rangle \quad (4.83)$$

$$\text{and for all } y \in Y, \eta; \lambda + 1; \mathcal{A} \models \{P''_p(x, y)\} \quad c_1 \quad \{Q_p(x, y)\} \quad (4.84)$$

By (4.80), there are P'_p and P''_p with:

$$\eta; \lambda; \mathcal{A} \models \left\langle P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \right\rangle \alpha \left\langle \begin{array}{c} P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \vee \\ \exists y \in Y. P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) \end{array} \right\rangle \quad (4.85)$$

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \left\langle P'_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \right\rangle c_1 \quad \exists y \in Y. \left\langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) \right\rangle \quad (4.86)$$

$$\text{and for all } y \in Y, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y) \right\} c_1 \left\{ Q_p(x, y) \right\} \quad (4.87)$$

To establish (4.82), fix $x \in X$, $R \in \mathbf{View}_{\mathcal{A}}$, $\varphi \in P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * R$, $h \in \lfloor \varphi \rfloor_{\lambda+1}$, $h' \in \llbracket \alpha \rrbracket(h)$. There will be some $\bar{\varphi} \in P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * R$, with $r_\varphi = r_{\bar{\varphi}}$ and $\rho_\varphi = \rho_{\bar{\varphi}}$, and $\lfloor \bar{\varphi} \rfloor_\lambda = \lfloor \varphi \rfloor_{\lambda+1}$, and so $h \in \lfloor \bar{\varphi} \rfloor_\lambda$. The $\bar{\varphi}$ corresponds to φ with all regions at level λ collapsed. (Recall that r_φ and ρ_φ are projections of region states and region assignments from φ .)

By (4.85), there is some $\bar{\varphi}'$ with $\bar{\varphi} \text{ G}_{\lambda; \mathcal{A}} \bar{\varphi}'$, $h' \in \lfloor \bar{\varphi}' \rfloor_\lambda$ and

$$\bar{\varphi}' \in \left(P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) \right) * R$$

We have the following cases for $\bar{\varphi}'$:

- $\bar{\varphi}' \in P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * R$. In this case, $\bar{\varphi}' = \varphi'' \bullet \bar{\bar{\varphi}}'$ where

$$\begin{array}{c} \bar{\bar{\varphi}}' \in I(\mathbf{t}_a^\lambda(\vec{z}, x)) * \quad \quad \quad \bigotimes \quad \quad \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad a' \in \text{Rld} \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad a' \neq a \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad r_\varphi(a') = (\lambda, -, -) \end{array}$$

and $\varphi'' \in P'_p(\sigma_1) * P(x) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}).$$

Hence, by the guarantee, $\varphi' \in P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * R$, and by construction $\lfloor \varphi' \rfloor_{\lambda+1} = \lfloor \bar{\varphi}' \rfloor_\lambda$ and $\varphi \text{ G}_{\lambda+1; \mathcal{A}} \varphi'$.

- $\bar{\varphi}' \in P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q(x, y) * R$ for some $y \in Y$. In this case, $\bar{\varphi}' = \varphi'' \bullet \bar{\bar{\varphi}}'$ where

$$\begin{array}{c} \bar{\bar{\varphi}}' \in I(\mathbf{t}_a^\lambda(\vec{z}, x)) * \quad \quad \quad \bigotimes \quad \quad \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad a' \in \text{Rld} \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad a' \neq a \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad r_\varphi(a') = (\lambda, -, -) \end{array}$$

and $\varphi'' \in P''_p(x, y, \sigma_1) * Q(x, y) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}).$$

Hence, by the guarantee, $\varphi' \in P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) * R$, and by construction $\lfloor \varphi' \rfloor_{\lambda+1} = \lfloor \bar{\varphi}' \rfloor_\lambda$ and $\varphi \text{ G}_{\lambda+1; \mathcal{A}} \varphi'$.

In each case we have φ' which satisfies $\varphi \text{ G}_{\lambda+1;\mathcal{A}} \varphi'$, $h' \in \llbracket \varphi' \rrbracket_{\lambda+1}$ and

$$\varphi' \in \left(P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * Q(x, y) \right) * R.$$

So we have established (4.82).

The (4.86) sub-condition follows from (4.86) and the coinductive hypothesis. The (4.84) sub-condition follows from (4.87) and the AWEAKENING3 Rule Lemma.

The remaining cases are simpler, or follow a similar reasoning. $\square$

Lemma 4.15 (USEATOMIC Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, $a \notin \mathcal{A}$ and $\forall x \in X. (x, f(x)) \in \mathcal{T}_t(\text{G})^*$:

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \left\langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\text{G}]_a \right\rangle c \quad \exists y \in Y. \left\langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) \right\rangle \quad (4.88)$$

then

$$\eta; \lambda + 1; \mathcal{A} \models \mathbb{W}x \in X. \left\langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\text{G}]_a \right\rangle c \quad \exists y \in Y. \left\langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) \right\rangle \quad (4.89)$$

Proof. Consider the case where c is not skip and performs an action $\alpha \notin \{\text{spawn}(\mathbf{f}, \vec{v}) \mid \mathbf{f} \in \text{Fun}, \vec{v} \in \text{Val}^*\}$, the first condition of the semantic judgement is the only one that may apply. For reduction $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_1, c_1)$, to show (4.89), we have to establish the following sub-conditions, by finding a suitable P'_p and P''_p :

$$\eta; \lambda + 1; \mathcal{A} \models \left\langle P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\text{G}]_a \right\rangle \alpha \left\langle \begin{array}{c} P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\text{G}]_a \vee \\ \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) \end{array} \right\rangle \quad (4.90)$$

$$\eta; \lambda + 1; \mathcal{A} \models \mathbb{W}x \in X. \left\langle P'_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\text{G}]_a \right\rangle c_1 \quad \exists y \in Y. \left\langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) \right\rangle \quad (4.91)$$

$$\text{and for all } y \in Y, \eta; \lambda + 1; \mathcal{A} \models \left\{ P''_p(x, y) \right\} c_1 \left\{ Q_p(x, y) \right\} \quad (4.92)$$

By (4.88), there are P'_p and P''_p with:

$$\eta; \lambda; \mathcal{A} \models \left\langle P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\text{G}]_a \right\rangle \alpha \left\langle \begin{array}{c} P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\text{G}]_a \vee \\ \exists y \in Y. P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) \end{array} \right\rangle \quad (4.93)$$

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \left\langle P'_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\text{G}]_a \right\rangle c_1 \quad \exists y \in Y. \left\langle Q_p(x, y) \mid I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) \right\rangle \quad (4.94)$$

$$\text{and for all } y \in Y, \eta; \lambda; \mathcal{A} \models \left\{ P''_p(x, y) \right\} c_1 \left\{ Q_p(x, y) \right\} \quad (4.95)$$

To establish (4.90), fix $x \in X$, $R \in \text{View}_{\mathcal{A}}$, $\varphi \in P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\text{G}]_a * R$, $h \in \llbracket \varphi \rrbracket_{\lambda+1}$, $h' \in \llbracket \alpha \rrbracket(h)$. There will be some $\bar{\varphi} \in P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\text{G}]_a * R$, with $r_\varphi = r_{\bar{\varphi}}$ and $\rho_\varphi = \rho_{\bar{\varphi}}$,

and $[\overline{\varphi}]_\lambda = [\varphi]_{\lambda+1}$, and so $h \in [\overline{\varphi}]_\lambda$. By (4.93), there is some $\overline{\varphi}'$ with $\overline{\varphi} \text{ G}_{\lambda;\mathcal{A}} \overline{\varphi}'$, $h' \in [\overline{\varphi}']_\lambda$ and

$$\overline{\varphi}' \in \left(P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\mathbf{G}]_a \vee \exists y \in Y. P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) \right) * R$$

We have the following cases for $\overline{\varphi}'$:

- $\overline{\varphi}' \in P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * [\mathbf{G}]_a * R$. In this case, $\overline{\varphi}' = \varphi'' \bullet \overline{\overline{\varphi}'}$ where

$$\begin{aligned} \overline{\overline{\varphi}'} \in I(\mathbf{t}_a^\lambda(\vec{z}, x)) * & \quad \bigotimes \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ & a' \in \text{Rld} \\ & a' \neq a \\ & r_\varphi(a') = (\lambda, -, -) \end{aligned}$$

and $\varphi'' \in P'_p(\sigma_1) * P(x) * [\mathbf{G}]_a * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}).$$

Hence, by the guarantee, $\varphi' \in P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * [\mathbf{G}]_a * R$, and by construction $[\varphi']_{\lambda+1} = [\overline{\varphi}']_\lambda$ and $\varphi \text{ G}_{\lambda+1;\mathcal{A}} \varphi'$.

- $\overline{\varphi}' \in P''_p(x, y, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * Q(x, y) * R$ for some $y \in Y$. In this case, $\overline{\varphi}' = \varphi'' \bullet \overline{\overline{\varphi}'}$ where

$$\begin{aligned} \overline{\overline{\varphi}'} \in I(\mathbf{t}_a^\lambda(\vec{z}, f(x))) * & \quad \bigotimes \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ & a' \in \text{Rld} \\ & a' \neq a \\ & r_\varphi(a') = (\lambda, -, -) \end{aligned}$$

and $\varphi'' \in P''_p(x, y, \sigma_1) * Q(x, y) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}[a \mapsto f(x)], d_{\varphi''}).$$

Hence, by the guarantee, $\varphi' \in P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) * R$, and by construction $[\varphi']_{\lambda+1} = [\overline{\varphi}']_\lambda$ and $\varphi \text{ G}_{\lambda+1;\mathcal{A}} \varphi'$.

In each case we have φ' which satisfies $\varphi \text{ G}_{\lambda+1;\mathcal{A}} \varphi'$, $h' \in [\varphi']_{\lambda+1}$ and

$$\varphi' \in \left(P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, f(x)) * Q(x, y) \right) * R.$$

So we have established (4.90).

The (4.94) sub-condition follows from (4.94) and the coinductive hypothesis. The (4.92) sub-condition follows from (4.95) and the AWEAKENING3 Rule Lemma.

The remaining cases are simpler, or follow a similar reasoning. □

Lemma 4.16 (UPDATEREGION Rule). Suppose that $a \notin \mathcal{A}$ and

$$\begin{array}{c} \mathbb{W}x \in X. \left\langle P_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \right\rangle \\ \mathbb{C} \\ \eta; \lambda; \mathcal{A} \models \exists(y, w) \in Y \times W. \left\langle Q_p(x, y, w) \mid \begin{array}{l} I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) \\ \vee I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y) \end{array} \right\rangle \end{array} \quad (4.96)$$

Then

$$\begin{array}{c} \mathbb{W}x \in X. \left\langle P_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \right\rangle \\ \mathbb{C} \\ \eta; \lambda + 1; \mathcal{A}' \models \exists(y, w) \in Y \times W. \left\langle Q_p(x, y, w) \mid \begin{array}{l} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * a \Rightarrow (x, w) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right\rangle \end{array} \quad (4.97)$$

where $\mathcal{A}' = (a : x \in X \rightsquigarrow W, \mathcal{A})$.

Proof. Suppose that $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_1, c_1)$ with $\alpha \in \mathbf{AAction}$.

Fix $x \in X$. From our assumption (4.96), there are P'_p and $\overline{P''_p}$ with

$$\begin{array}{c} \langle P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \rangle \\ \alpha \\ \lambda; \mathcal{A} \models \left\langle \begin{array}{l} P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \vee \\ \exists(y, w) \in Y \times W. \overline{P''_p}(x, y, \sigma_1) * (I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) \vee I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y)) \end{array} \right\rangle \end{array} \quad (4.98)$$

$$\begin{array}{c} \mathbb{W}x \in X. \left\langle P'_p \mid I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \right\rangle \\ \mathbb{C} \\ \lambda; \mathcal{A} \models \exists(y, w) \in Y \times W. \left\langle Q_p(x, y, w) \mid I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) \vee I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y) \right\rangle \end{array} \quad (4.99)$$

$$\forall(y, w) \in Y \times W. \quad \lambda; \mathcal{A} \models \left\{ \overline{P''_p}(x, y) \right\} \quad c_1 \quad \left\{ Q_p(x, y) \right\} \quad (4.100)$$

We will show that these P'_p and $P''_p(x, y, w) = \overline{P''_p}(x, y, w) * a \Rightarrow (x, w)$ work to establish our goal.

Fix $R \in \mathbf{View}_{\mathcal{A}'}$, $\varphi \in P_p(\sigma) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * a \Rightarrow \blacklozenge * R$, $h \in \lfloor \varphi \rfloor_{\lambda+1}$, $h' \in \llbracket \alpha \rrbracket(h)$.

Let $\overline{R} \in \mathbf{View}_{\mathcal{A}}$ be such that

$$\overline{R} = \text{removedone}_a \left(R * \begin{array}{c} \text{\textcircled{*}} \\ a' \in \mathbf{RId} \\ a' \neq a \\ r_\varphi(a') = (\lambda, -, -) \end{array} I(r_\varphi(a'), a', \rho_\varphi(a')) \right).$$

That is, we open all regions at level λ (except a) with their states as given by φ and remove the atomicity tracking for a .

There will be some $\overline{\varphi} \in P_p * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * \overline{R}$ with $r_\varphi = r_{\overline{\varphi}}$ and $\rho_\varphi = \rho_{\overline{\varphi}}$, and $\lfloor \overline{\varphi} \rfloor_\lambda = \lfloor \varphi \rfloor_{\lambda+1}$,

and so $h \in [\bar{\varphi}]_\lambda$. By (4.98), there is some $\bar{\varphi}'$ with $\bar{\varphi} \text{ G}_{\lambda;\mathcal{A}} \bar{\varphi}'$, $h' \in [\bar{\varphi}']_\lambda$ and

$$\bar{\varphi}' \in \left(\begin{array}{c} P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) \vee \\ \exists(y, w) \in Y \times W. \bar{P}_p''(x, y, w, \sigma_1) * \left(I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) \vee I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y) \right) \end{array} \right) * \bar{R}$$

We have the following cases for $\bar{\varphi}'$:

- $\bar{\varphi}' \in P'_p(\sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * P(x) * \bar{R}$. In this case, $\bar{\varphi}' = \varphi'' \bullet \bar{\bar{\varphi}'}$ where

$$\begin{array}{c} \bar{\bar{\varphi}'} \in I(\mathbf{t}_a^\lambda(\vec{z}, x)) * \quad \quad \quad \bigotimes \quad \quad \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ a' \in \text{Rld} \\ a' \neq a \\ r_\varphi(a') = (\lambda, -, -) \end{array}$$

and $\varphi'' \in P'_p(\sigma_1) * P(x) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}[a \mapsto \blacklozenge]).$$

Hence, by the guarantee, $\varphi' \in P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * R$, and by construction $[\varphi']_{\lambda+1} = [\bar{\varphi}']_\lambda$. Also $\varphi \text{ G}_{\lambda+1;\mathcal{A}'} \varphi'$.

- $\bar{\varphi}' \in \bar{P}_p''(x, y, w, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, w)) * Q_1(x, y, w) * \bar{R}$ for some $(y, w) \in Y \times W$. In this case, $\bar{\varphi}' = \varphi'' \bullet \bar{\bar{\varphi}'}$ where

$$\begin{array}{c} \bar{\bar{\varphi}'} \in I(\mathbf{t}_a^\lambda(\vec{z}, w)) * \quad \quad \quad \bigotimes \quad \quad \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ a' \in \text{Rld} \\ a' \neq a \\ r_\varphi(a') = (\lambda, -, -) \end{array}$$

and $\varphi'' \in \bar{P}_p''(x, y, w, \sigma_1) * Q_1(x, y, w) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}[a \mapsto w], d_{\varphi''}[a \mapsto (x, w)]).$$

Hence, by the guarantee, $\varphi' \in P''_p(x, y, w, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * R$, and by construction $[\varphi']_{\lambda+1} = [\bar{\varphi}']_\lambda$. Also $\varphi \text{ G}_{\lambda+1;\mathcal{A}'} \varphi'$.

- $\bar{\varphi}' \in \bar{P}_p''(x, y, w, \sigma_1) * I(\mathbf{t}_a^\lambda(\vec{z}, x)) * Q_2(x, y) * \bar{R}$ for some $(y, w) \in Y \times W$. In this case, $\bar{\varphi}' = \varphi'' \bullet \bar{\bar{\varphi}'}$ where

$$\begin{array}{c} \bar{\bar{\varphi}'} \in I(\mathbf{t}_a^\lambda(\vec{z}, x)) * \quad \quad \quad \bigotimes \quad \quad \quad I(r_\varphi(a'), a', \rho_\varphi(a')) \\ a' \in \text{Rld} \\ a' \neq a \\ r_\varphi(a') = (\lambda, -, -) \end{array}$$

and $\varphi'' \in \bar{P}_p''(x, y, w, \sigma_1) * Q_2(x, y) * R$. Let

$$\varphi' = (r_{\varphi''}, h_{\varphi''}, b_{\varphi''}, \gamma_{\varphi''}, \rho_{\varphi''}, d_{\varphi''}[a \mapsto \blacklozenge]).$$

Hence, by the guarantee, $\varphi' \in P''_p(x, y, w, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * R$, and by construction $[\varphi']_{\lambda+1} = [\bar{\varphi}']_\lambda$. Also $\varphi \text{ G}_{\lambda+1;\mathcal{A}'} \varphi'$.

In each case we have φ' which satisfies $\varphi \text{ G}_{\lambda+1;\mathcal{A}} \varphi'$, $h' \in \llbracket \varphi' \rrbracket_{\lambda+1}$ and

$$\varphi' \in \left(\begin{array}{c} P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \vee \\ \exists(y, w) \in Y \times W. P''_p(x, y, w, \sigma_1) * \left(\begin{array}{c} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * a \Rightarrow (x, w) \\ \vee \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right) \end{array} \right) * R.$$

So we have established that

$$\lambda + 1; \mathcal{A}' \models \left\langle \begin{array}{c} P'_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \\ \alpha \\ P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \vee \\ \exists(y, r) \in Y \times W. P''_p(x, y, w, \sigma_1) * \left(\begin{array}{c} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * a \Rightarrow (x, w) \vee \\ \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right) \end{array} \right\rangle.$$

By (4.99), it follows from the coinductive hypothesis that

$$\lambda + 1; \mathcal{A}' \models \mathbb{W}x \in X. \left\langle P'_p \left| \mathbf{t}_a^\lambda(\vec{z}, x) * P(x) * a \Rightarrow \blacklozenge \right. \right\rangle \\ \exists(y, w) \in Y \times W. \left\langle Q_p(x, y, w) \left| \begin{array}{c} \mathbf{t}_a^\lambda(\vec{z}, w) * Q_1(x, y, w) * a \Rightarrow (x, w) \vee \\ \mathbf{t}_a^\lambda(\vec{z}, x) * Q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right. \right\rangle$$

Fix $y \in Y$ and $w \in W$. Because $P''_p(x, y, w) = \overline{P''_p}(x, y, w) * a \Rightarrow (x, w)$, we can extend the atomicity context and have

$$\lambda + 1; \mathcal{A}' \models \left\{ P''_p(x, y, w) \right\} c_1 \left\{ Q_p(x, y, w) \right\}$$

□

Lemma 4.17 (Drop Context). If, for $P \in \text{View}_{\mathcal{A}}$, Q , $x \in X$, $y \in Y$

$$\eta; \lambda; a : x \in X \rightsquigarrow Y, \mathcal{A} \models \left\{ P * a \Rightarrow (x, y) \right\} c \left\{ \exists x, y. Q(x, y) * a \Rightarrow (x, y) \right\}$$

then

$$\eta; \lambda; \mathcal{A} \models \left\{ P \right\} c \left\{ Q(x, y) \right\}$$

Lemma 4.18 (MAKEATOMIC Rule). Suppose that

$$\{(x, y) \mid x \in X, y \in Y\} \subseteq \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \\ \eta; \lambda'; \mathcal{A} \models \left\{ P_p * \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * a \Rightarrow \blacklozenge \right\} c \left\{ \exists x \in X, y \in Y. Q_p(x, y) * a \Rightarrow (x, y) \right\}$$

where $\mathcal{A} = a : x \in X \rightsquigarrow Y, \mathcal{A}'$ and $a \notin \text{dom}(\mathcal{A}')$. Then

$$\eta; \lambda'; \mathcal{A}' \models \mathbb{W}x \in X. \left\langle P_p \left| \mathbf{t}_a^\lambda(\vec{z}, x) * [G]_a \right. \right\rangle c \exists y \in Y. \left\langle Q_p(x, y) \left| \mathbf{t}_a^\lambda(\vec{z}, y) * [G]_a \right. \right\rangle$$

Proof. The proof is by coinduction. Consider the case where c performs an action. Suppose that $(\sigma, c) \xrightarrow{\alpha}_\eta (\sigma_1, c_1)$ with $\alpha \in \mathbf{AAction}$. It is sufficient to show that, for all $x \in X$ there exist P'_p and P''_p such that

$$\eta; \lambda'; \mathcal{A} \models \langle P_p(\sigma) * P(x) \rangle \alpha \langle P'_p(\sigma_1) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * Q(x, y) \rangle \quad (4.101)$$

$$\eta; \lambda'; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle c_1 \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle, \quad (4.102)$$

$$\text{and for all } y \in Y, \eta; \lambda'; \mathcal{A} \models \{P''_p(x, y)\} c_1 \{Q_p(x, y)\}. \quad (4.103)$$

By the premiss, there must be some $\overline{P'_p}$ with

$$\lambda'; \mathcal{A} \models \langle P_p(\sigma) * \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * a \Rightarrow \blacklozenge \rangle \alpha \langle \overline{P'_p}(\sigma_1) \rangle \quad (4.104)$$

$$\lambda'; \mathcal{A} \models \{\overline{P'_p}\} c_1 \{\exists x \in X, y \in Y. Q_p(x, y) * a \Rightarrow (x, y)\}. \quad (4.105)$$

Fix $x \in X$. Fix $R \in \text{View}_{\mathcal{A}'}$. Fix $\varphi \in P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a * R$.

Let $P'_p = \lambda\sigma. \{\varphi \in \text{AWorld}_{\text{dom } \mathcal{A}'} \mid \varphi \bullet a \Rightarrow \blacklozenge \in \overline{P'_p}(\sigma)\}$.

Let $P''_p(x, y) = \lambda\sigma. \{\varphi \in \text{AWorld}_{\text{dom } \mathcal{A}'} \mid \varphi \bullet a \Rightarrow (x, y) \in \overline{P'_p}(\sigma)\}$.

Let $\overline{R} = R * [\mathbf{G}]_a * a \Rightarrow -$. ($\overline{R}$ is stable with respect to $\mathcal{A}$ since the additional interference will be $a : x \in X \rightsquigarrow Y$, and the subset of $\overline{R}$ that is compatible with $[\mathbf{G}]_a$ must be closed under this.) Let $\overline{\varphi} = \varphi \bullet a \Rightarrow \blacklozenge$. By construction, $[\varphi]_{\lambda'} = [\overline{\varphi}]_{\lambda'}$. We have that $\overline{\varphi} \in (P_p(\sigma) * \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * a \Rightarrow \blacklozenge) * \overline{R}$.

By (4.104) there exists $\overline{\varphi'}$ with

$$\overline{\varphi} \text{ G}_{\lambda'; \mathcal{A}} \overline{\varphi'}, \quad (4.106)$$

$$h' \in [\overline{w'}]_{\lambda'}, \quad (4.107)$$

$$\text{and } \overline{\varphi'} \in \overline{P'_p}(\sigma_1) * \overline{R}. \quad (4.108)$$

From (4.106) we can be sure that $d_{\overline{\varphi'}} \neq \blacklozenge$. Indeed, since $d_{\overline{\varphi}} = \blacklozenge$ and $\rho_{\overline{\varphi}} = x$, it must be that either $d_{\overline{\varphi'}} = \blacklozenge$ or $d_{\overline{\varphi}} = (x, y)$ for some $y \in Y$.

Let φ' be such that $\varphi \in \varphi' * a \Rightarrow -$. Now

$$\varphi' \in P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a \vee \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, y)$$

since $\overline{\varphi'} \in \overline{P'_p}(\sigma_1) * \overline{R}$, by (4.108). By (4.106) and definitions, we get $\varphi \text{ G}_{\lambda'; \mathcal{A}} \varphi'$. By construction $[\varphi']_{\lambda'} = [\overline{\varphi'}]_{\lambda'}$ so $h' \in [\varphi']_{\lambda'}$ by (4.107). Hence, we have established

$$\lambda'; \mathcal{A} \models \langle P_p(\sigma) * \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a \rangle \alpha \langle P'_p(\sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a \vee \exists y \in Y. P''_p(x, y, \sigma_1) * \mathbf{t}_a^\lambda(\vec{z}, y) \rangle.$$

We have that $P'_p * \exists x \in X. \mathbf{t}_a^\lambda(\vec{z}, x) * a \Rightarrow \blacklozenge \models \overline{P'_p}$ and is stable with respect to $\mathcal{A}$. From (4.105), by left consequence and the coinductive hypothesis, we have

$$\lambda'; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid \mathbf{t}_a^\lambda(\vec{z}, x) * [\mathbf{G}]_a \rangle c_1 \exists y \in Y. \langle Q_p(x, y) \mid \mathbf{t}_a^\lambda(\vec{z}, y) * [\mathbf{G}]_a \rangle$$

Finally, from (4.105) and Lemma 4.17, we have, for all $y \in Y$

$$\lambda'; \mathcal{A}' \models \{P''_p(x, y)\} c_1 \{Q_p(x, y)\}.$$

The remaining cases are simpler, or follow similar reasoning. $\square$

Lemma 4.19 (FRAME Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$, suppose $\text{mods}(c) \cap \text{pvars}(R') = \emptyset$, $\text{mods}(c) \cap \text{pvars}(R(x)) = \emptyset$ and

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle R' * P_p \mid R(x) * P(x) \rangle c \quad \exists y \in Y. \langle R' * Q_p(x, y) \mid R(x) * Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

Lemma 4.20 (SUBSTITUTION Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$, suppose $f : X' \rightarrow X$, $g : Y' \rightarrow Y$ and

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle c \quad \exists y' \in Y'. \langle Q_p(x, g(y')) \mid Q(x, g(y')) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x' \in X'. \langle P_p \mid P(f(x')) \rangle c \quad \exists y \in Y. \langle Q_p(f(x'), y) \mid Q(f(x'), y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

Lemma 4.21 (CONSEQUENCE Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$, suppose $\lambda; \mathcal{A} \models P_p \preceq P'_p$, $\forall x \in X, y \in Y. \lambda; \mathcal{A} \models Q'_p(x, y) \preceq Q_p(x, y)$, $\forall x \in X, y \in Y. \lambda; \mathcal{A} \models Q'(x, y) \preceq Q(x, y)$ and

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle c \quad \exists y \in Y. \langle Q'_p(x, y) \mid Q'(x, y) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule using the definition of view shift. $\square$

Lemma 4.22 (AEXISTS Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}z \in 1. \langle P_p \mid \exists x \in X. P(x) \rangle c \quad \exists x \in X, y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

Lemma 4.23 (AWEAKENING1 Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P' * P(x) \rangle c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q'(x, y) * Q(x, y) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p * P' \mid P(x) \rangle \ c \quad \exists y \in Y. \langle Q_p(x, y) * Q'(x, y) \mid Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

Lemma 4.24 (AWEAKENING2 Rule). For all $\eta \in \text{FEnv}$, $\lambda \in \text{Level}$ and $\mathcal{A} \in \text{AContext}$:

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \ c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

then

$$\forall x \in X. \eta; \lambda; \mathcal{A} \vdash \mathbb{W}z \in 1. \langle P_p \mid P(x) \rangle \ c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

Lemma 4.25 (AWEAKENING3 Rule). For all $\eta \in \text{FEnv}$, $\lambda, \lambda' \in \text{Level}$, $\mathcal{A} \in \text{AContext}$, suppose $\lambda' \leq \lambda$ and:

$$\eta; \lambda'; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \ c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

then

$$\eta; \lambda; \mathcal{A} \models \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \ c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

Proof. The proof follows directly from the definition of semantic judgement and follows a similar reasoning to that for the OPENREGION rule. $\square$

5 Using TaDA

In this chapter, we use the TaDA logic on a series of examples. We first specify an atomic lock module (§5.1). From this specification, we then derive the CAP-style lock specification given in §3.2, which illustrates the weakening of the atomic specification to a specific use case. We also prove that a spin lock implementation satisfies the atomic lock specification. We show how the logic supports vertical reasoning about modules, by verifying an implementation of multiple-compare-and-set (MCAS) using the lock specification (§5.2), and an implementation of a concurrent double-ended queue (deque) using the MCAS specification (§5.3). Moreover, we demonstrate how to apply the logic to fit the different protocols required by each concurrent module.

The examples demonstrate that TaDA combines the benefits of abstract atomicity and abstract disjointness within a single program logic. Further examples can be found in [64] where we use TaDA to prove a concurrent skiplist map implementation and various clients.

5.1 Spin Lock

We define a lock module with the operations `acquire` and `release` and a constructor `makeLock`, similarly to the one shown in §3.2.

5.1.1 Atomic Specification

The lock operations are specified using abstract predicates that represent the state of a lock: $\text{Lock}(s, x, v)$ asserts the existence of a lock, addressed by x , that is in the unlocked state when $v = 0$ and is in the locked state when $v = 1$. These predicates confer ownership of the lock: it is not possible to have more than one $\text{Lock}(s, x, v)$ for the same value of x . This contrasts with the style of specification given with CAP [12] shown in §3.2.1, but we shall see how the CAP specification can be derived using the atomic specification given here.

The specification for the `makeLock` operation is a simple Hoare triple:

$$\lambda; \mathcal{A} \vdash \{ \text{True} \} \text{makeLock}() \{ \exists s \in \mathbb{T}_1. \text{Lock}(s, \text{ret}, 0) \}$$

The operation allocates a new lock, which is initially unlocked, and returns its address. The specification says nothing about the granularity of the operation. In fact, the granularity is hardly relevant, since no concurrent environment can meaningfully observe the effects of `makeLock` until its return value is known—that is, once the operation has completed.

In the context of the specification, λ is a constant number that indicates the *level* of the specification. To a client, λ is abstract (much like $\text{Lock}(s, x, v)$ is); different implementations may fix different values for λ . The purpose of the level in our specifications is to prevent certain unsound circularities. This is achieved by requiring the level to be reduced at each step where such a circularity could be introduced.

The specification for the **release** operation uses an *atomic* triple:

$$\lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, 1) \rangle \text{ release}(x) \langle \text{Lock}(s, x, 0) \rangle$$

Intuitively, this specification means that **release** will atomically take the lock x from the locked to unlocked state. The atomic triple makes a strong guarantee: as long as the concurrent environment guarantees that the (possibly) shared resource $\text{Lock}(s, x, 1)$ is available, the **release** operation will preserve $\text{Lock}(s, x, 1)$ until it transforms it into $\text{Lock}(s, x, 0)$; after the transformation, the operation no longer requires $\text{Lock}(s, x, 0)$, and is consequently oblivious to subsequent transformations by the environment (such as another thread acquiring the lock).

It is significant that the notion of atomicity is tied to the abstraction in the specification. The predicate $\text{Lock}(s, x, v)$ could abstract multiple underlying states in the implementation. If we were to observe the underlying state, the operation might no longer appear to be atomic.

Specifying the **acquire** operation is more subtle. It can be called regardless of whether the lock is in the locked or unlocked state, and always results in setting it to the locked state (if it ever terminates). A first attempt at a specification might, therefore, be:

$$\lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, 0) \vee \text{Lock}(s, x, 1) \rangle \text{ acquire}(x) \langle \text{Lock}(s, x, 1) \rangle$$

This specification has two significant flaws. Firstly, it allows **acquire** to do nothing at all when the lock is already locked. This is contrary to what it should do, which is wait for it to become unlocked and then (atomically) lock it. Secondly, as the level of abstraction given by the precondition is $\text{Lock}(s, x, 0) \vee \text{Lock}(s, x, 1)$, an implementation could change the state of the lock arbitrarily *without appearing to have done anything*. In particular, an implementation could transition between the two states any number of times, so long as it is in the $\text{Lock}(s, x, 1)$ state when it finishes.

A second attempt to overcome these issues might be:

$$\begin{aligned} \lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, 1) \rangle \text{ acquire}(x) \langle \text{False} \rangle \\ \lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, 0) \rangle \text{ acquire}(x) \langle \text{Lock}(s, x, 1) \rangle \end{aligned}$$

In the left-hand triple, the lock is initially locked; the implementation may not terminate, nor change the state of the lock. In the right-hand triple, the lock is initially unlocked; the implementation may only make one atomic transformation from unlocked to locked. These specifications have a subtle flaw, in that they assume that the environment will not change the state of the lock. They prevent us from having multiple threads competing to acquire the lock, which is the essential purpose of a lock.

An equivalent specification makes use of a boolean logical variable:

$$\forall v \in \{0, 1\}. \lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, v) \rangle \text{ acquire}(x) \langle \text{Lock}(s, x, 1) * v = 0 \rangle$$

The variable v records the state of the lock when the atomic operation takes effect. In particular, it cannot take effect unless the lock is already unlocked.

These specifications do not express the subtlety that the interference permitted before the atomic update is different for the environment and the operation. The environment should be allowed to change the value of v (i.e. acquire and release the lock) but the **lock** operation should not. The correct

specification expresses this by binding the variable v in a new way:

$$\lambda; \mathcal{A} \vdash \mathbb{W}v \in \{0, 1\}. \langle \text{Lock}(s, x, v) \rangle \text{acquire}(x) \langle \text{Lock}(s, x, 1) * v = 0 \rangle$$

The special role of v (indicated by the pseudo-quantifier $\mathbb{W}$) is in distinguishing the constraints on the environment and on the thread before the atomic operation takes effect. Specifically, the environment is at liberty to change the value of v for which the precondition holds (that is, lock and unlock the lock), but the thread executing the operation must preserve the value of v (that is, it cannot lock or unlock the lock except by performing the atomic operation).

5.1.2 CAP Specification

The atomic specification of the lock captures its essence as a synchronisation primitive. In practice, a lock is often used to protect some resource. We demonstrate how a CAP-style lock specification [12], which views the lock as a mechanism for protecting a resource invariant, can be derived from the atomic specification. This illustrates a typical use of a TaDA specification: first prove a strong abstract-atomic specification; then specialise to whatever is required by the client.

The specification provides two abstract predicates: $\text{IsLock}(x)$, which is a non-exclusive resource that allows a thread to compete for the lock; and $\text{Locked}(x)$, which is an exclusive resource that represents that the thread has acquired the lock, and allows it to release the lock. The lock is specified as follows:

$$\begin{aligned} \lambda'; \mathcal{A} &\vdash \{ \text{True} \} \text{makeLock}() \{ \text{IsLock}(\text{ret}) \} \\ \lambda'; \mathcal{A} &\vdash \{ \text{Locked}(x) \} \text{release}(x) \{ \text{True} \} \\ \lambda'; \mathcal{A} &\vdash \{ \text{IsLock}(x) \} \text{acquire}(x) \{ \text{IsLock}(x) * \text{Locked}(x) \} \\ \text{IsLock}(x) &\iff \text{IsLock}(x) * \text{IsLock}(x) \\ \text{Locked}(x) * \text{Locked}(x) &\implies \text{False} \end{aligned}$$

To implement this specification, we must provide an interpretation for the abstract predicates IsLock and Locked . For this, we need to introduce a shared region. As in CAP, a shared region encapsulates some resource that is available to multiple threads. In our example, this resource will be the predicates $\text{Lock}(s, x, v)$, plus some additional guard resource (described below). A shared region is associated with a protocol, which determines how its contents change over time. Following iCAP [54], the state of a shared region is abstracted, and protocols are expressed as transition systems over these abstract states. A thread may only change the abstract state of a region when it has the *guard* resource associated with the transition to be performed. An interpretation function associates each abstract state of a region with a concrete assertion. In summary, to specify a region we must supply the guards for the region, an abstract state transition system that is labelled by these guards, and a function interpreting abstract states as assertions.

For the **CAPLock**, we need only a very simple guard separation algebra: there is a single, indivisible guard named UNLOCK , as well as the empty guard $\mathbf{0}$. As a separation algebra, guard resources must have a partial composition operator that is associative and commutative. In this case, $\mathbf{0} \bullet x = x = x \bullet \mathbf{0}$ for all $x \in \{\mathbf{0}, \text{UNLOCK}\}$, and $\text{UNLOCK} \bullet \text{UNLOCK}$ is not defined.

The transition system for the region will have two states: 0 and 1, corresponding to unlocked and locked states respectively. Intuitively, any thread should be allowed to lock the lock, if it is unlocked, but only the thread holding the ‘key’ should be able to unlock it. This is specified by the labelled transition system:

$$\begin{aligned} \mathbf{0} : 0 &\rightsquigarrow 1 \\ \text{UNLOCK} : 1 &\rightsquigarrow 0 \end{aligned}$$

It remains to give an interpretation for the abstract states of the transition system. To do so, we must have a name for the type of region we are defining; we shall use **CAPLock**. It is possible for there to be multiple regions associated with the same region type name. To distinguish them, each region has a unique region identifier, which is typically annotated as a subscript. Each region is also associated with a level, annotated as a superscript, to avoid circularity. A region specification may take some parameters that are used in the interpretation. With **CAPLock**, for instance, the address of the lock is such a parameter. We thus specify the type name, region identifier, the region level, parameters and state of a region in the form $\mathbf{CAPLock}_a^{\lambda'}(s, x, v)$.

The region interpretation for **CAPLock** is given by:

$$\begin{aligned} I(\mathbf{CAPLock}_a^{\lambda'}(s, x, 0)) &\triangleq \text{Lock}(s, x, 0) * [\text{UNLOCK}]_a \\ I(\mathbf{CAPLock}_a^{\lambda'}(s, x, 1)) &\triangleq \text{Lock}(s, x, 1) \end{aligned}$$

With this interpretation, the guard UNLOCK is in the region when it is in the unlocked state. This means that, when a thread acquires the lock, it takes ownership of the guard and the lock invariant by removing them from the region. Having the guard UNLOCK allows the thread to subsequently release the lock, returning the guard and invariant to the region.

We can now give an interpretation to the predicates $\text{IsLock}(x)$ and $\text{Locked}(x)$:

$$\begin{aligned} \text{IsLock}(x) &\triangleq \exists a, s \in \mathbb{T}_1, v \in \{0, 1\}. \mathbf{CAPLock}_a^{\lambda'}(s, x, v) \\ \text{Locked}(x) &\triangleq \exists a, s \in \mathbb{T}_1. \mathbf{CAPLock}_a^{\lambda'}(s, x, 1) * [\text{UNLOCK}]_a \end{aligned}$$

We also interpret the level λ' as $\lambda' \triangleq \lambda + 1$, where λ is the level of the atomic lock specification.

It remains to prove the specifications for the operations and the axioms.

The proofs of the **release** and **acquire** operations are given in Figure 5.1 and Figure 5.2. In the **release** proof, note that the immediate postcondition of the USEATOMIC rule is not stable, since it is possible for the environment to acquire the lock. For illustrative purposes, we weaken it minimally to a stable assertion, although it could be weakened to **True** directly.

While a region is opened, it must not be re-opened. This is enforced by requiring the level in the context of the conclusion of USEATOMIC to be higher than the level of the region being opened, which is also the level in the context of the premiss. We see this in the proof for **release** as $\lambda' = \lambda + 1$ is higher than the level of the **CAPLock** region, namely λ . The atomic specification for **release** indicates that it only opens regions below level λ . In particular, this assures us that it will not reopen the **CAPLock** region. Since the level constraints are typically straightforward, we shall omit them from subsequent proof outlines.

$$\begin{array}{c}
\{ \text{Locked}(x) \} \\
\hline
\text{ABSTRACT} \quad \left\{ \exists a, s. \mathbf{CAPLock}_a^{\lambda'}(s, x, 1) * [\text{UNLOCK}]_a \right\} \\
\hline
\text{EXISTS: } a, s \quad \left\langle \mathbf{CAPLock}_a^{\lambda'}(s, x, 1) * [\text{UNLOCK}]_a \right\rangle \\
\hline
\text{USEATOMIC} \quad \left\langle \text{Lock}(s, x, 1) * [\text{UNLOCK}]_a \right\rangle \\
\hline
\text{FRAME: } [\text{UNLOCK}]_a \quad \left\langle \text{Lock}(s, x, 1) \right\rangle \\
\text{release}(x) \\
\left\langle \text{Lock}(s, x, 0) \right\rangle \\
\hline
\left\langle \text{Lock}(s, x, 0) * [\text{UNLOCK}]_a \right\rangle \\
\left\langle \mathbf{CAPLock}_a^{\lambda'}(s, x, 0) \right\rangle \\
// \text{ weaken to stabilise} \\
\left\{ \exists a, s, v \in \{0, 1\}. \mathbf{CAPLock}_a^{\lambda'}(s, x, v) \right\} \\
\hline
\{ \text{True} \}
\end{array}$$

Figure 5.1: Derivation of the **release** specification.

$$\begin{array}{c}
\{ \text{IsLock}(x) \} \\
\hline
\text{ABSTRACT; EXISTS } a, s \quad \left\{ \exists v \in \{0, 1\}. \mathbf{CAPLock}_a^{\lambda'}(s, x, v) \right\} \\
\hline
\text{USEATOMIC} \quad \mathbb{W}v \in \{0, 1\}. \\
\left\langle \text{Lock}(s, x, v) * v = 1 \vee (\text{Lock}(s, x, v) * [\text{UNLOCK}]_a) * v = 0 \right\rangle \\
\hline
\text{FRAME: } v = 0 \uparrow [\text{UNLOCK}]_a \quad \mathbb{W}v \in \{0, 1\}. \\
\left\langle \text{Lock}(s, x, v) \right\rangle \\
\text{acquire}(x) \\
\left\langle \text{Lock}(s, x, 1) * v = 0 \right\rangle \\
\hline
\left\langle \text{Lock}(s, x, 1) * [\text{UNLOCK}]_a \right\rangle \\
\left\{ \mathbf{CAPLock}_a^{\lambda'}(s, x, 1) * [\text{UNLOCK}]_a \right\} \\
\hline
\{ \text{IsLock}(x) * \text{Locked}(x) \}
\end{array}$$

Figure 5.2: Derivation of the **acquire** specification.

<pre> function makeLock() { v := alloc(1); [v] := 0; return v; } </pre>	<pre> function release(x) { [x] := 0; } </pre>	<pre> function acquire(x) { do { b := CAS(x, 0, 1); } while (b = 0); } </pre>
---	--	---

Figure 5.3: Spin lock operations.

The **lock** proof uses the $\forall$ quantifier in the premiss of the **USEATOMIC** rule to account for the fact that, in the precondition, the lock could be in either state. The proof uses the frame rule, with a frame that is conditional on the state of the lock. To derive the final postcondition, we use the fact that region assertions, since they refer to shared resource, are freely duplicable: *i.e.* $\mathbf{CAPLock}_a^{\lambda'}(s, x, 1) \equiv \mathbf{CAPLock}_a^{\lambda'}(s, x, 1) * \mathbf{CAPLock}_a^{\lambda'}(s, x, 1)$.

The axiom $\mathbf{IsLock}(x) \iff \mathbf{IsLock}(x) * \mathbf{IsLock}(x)$ follows from the duplicability of region assertions. Finally, the axiom $\mathbf{Locked}(x) * \mathbf{Locked}(x) \implies \mathbf{False}$ follows from **UNLOCK** • **UNLOCK** being undefined.

Note that neither of the bad specifications for **acquire**(**x**) could be used in this derivation: the first because there would be no way to express that the frame $[\mathbf{UNLOCK}]_a$ is conditional on the state of the lock; and the second because we could not combine both cases in a single derivation.

5.1.3 Implementation

We consider a spin lock implementation of the atomic lock specification. The code is given in Figure 5.3. To verify this implementation against the atomic specification, we must give a concrete interpretation of the abstract predicates. To do this, we introduce a new region type, **SLock**, with only one non-empty guard for a **SLock** region, named **G**, much as for **CAPLock**. There are two states for an **SLock** region: 0 and 1, representing unlocked and locked, respectively. The key difference from **CAPLock** is that transitions in both directions are guarded by **G**. The labelled transition system is as follows:

$$\begin{aligned} \mathbf{G} &: 0 \rightsquigarrow 1 \\ \mathbf{G} &: 1 \rightsquigarrow 0 \end{aligned}$$

We also give an interpretation to each abstract state:

$$\begin{aligned} I(\mathbf{SLock}_a^0(x, 0)) &\triangleq x \mapsto 0 \\ I(\mathbf{SLock}_a^0(x, 1)) &\triangleq x \mapsto 1 \end{aligned}$$

We now define the interpretation of the abstract predicate:

$$\mathbf{Lock}(a, x, v) \triangleq \mathbf{SLock}_a^0(x, v) * [\mathbf{G}]_a$$

The abstract predicate $\mathbf{Lock}(a, x, v)$ asserts there is a region with identifier a and that the region is in state v , where v is 0 when the region is unlocked and 1 when the region is locked. It also states that there is a guard $[\mathbf{G}]_a$, which will be used to update the region. Note that the **SLock** region here is at level 0, since it is not necessary to open any further regions when the **SLock** region is opened.

To prove the implementation against our atomic specification, we use the **MAKEATOMIC** rule. The proof of the **acquire**(**x**) implementation is given in Figure 5.4. The proof first massages the specification into a form where we can apply the **MAKEATOMIC** rule. The atomicity context allows the region a to be in either state, but insists that it must have been in the unlocked state when the atomic operation takes effect. The **UPDATEREGION** rule conditionally performs the atomic action—transitioning the region from state 0 to 1, and recording this in the atomic tracking resource—if the atomic compare-and-set operation succeeds. The **MAKEATOMIC** rule achieves this, by requiring that the implementation can only update the abstract region state once and also that the environment cannot invalidate the




Figure 5.4: Proof outline for the `acquire` operation.

precondition for any $v \in \{0, 1\}$, until the implementation performs the atomic update. This means that the environment is allowed to change the abstract region state to any state, as long as it is 0 or 1.

The proofs for `makeLock` and `release` are similar, and as such, we omit them.

Remark 5.1 (On alternative implementations). It is possible to prove the following alternative implementation of `release` with the same atomic specification:

$$\lambda; \mathcal{A} \vdash \langle \text{Lock}(s, x, 1) \rangle [x] := 1; [x] := 0; \langle \text{Lock}(s, x, 0) \rangle$$

The first write to x has no effect, since the specification asserts that the lock must be locked initially. This code would clearly not be atomic in a different context; it would not satisfy the specification $\vdash \langle \text{Lock}(s, x, 0) \vee \text{Lock}(s, x, 1) \rangle \text{release}(x) \langle \text{Lock}(s, x, 0) \rangle$, for example. Since the specification constrains the client, it allows flexibility in the implementation. $\square$

5.2 Multiple Compare-and-Set (MCAS)

We look at an interface over the heap which provides atomic double- (`dcas`) and triple-compare-and-set (`3cas`) operations, in addition to the basic read, write and compare-and-set operations.

5.2.1 Atomic Specification

Our atomic specification makes use of two abstract predicates: $\text{MCL}(s, l)$ to represent an instance of the MCAS module with address l ; and $\text{MCP}(s, l, x, v)$ to represent the “MCAS heap cell” at address x with value v , protected by instance l . There is an abstract disjointness, as we can view each heap cell as disjoint from the others at the abstract level, even if that is not the case with the implementation

itself. The specification for creating the interface, transferring memory cells to and from it as well as manipulating it is given in Figure 5.5.

$$\begin{aligned}
& \lambda; \mathcal{A} \vdash \{\text{True}\} \text{makeMCL}() \{ \text{MCL}(s, \text{ret}) \} \\
& \lambda; \mathcal{A} \vdash \{ \text{MCL}(s, 1) * x \mapsto v \} \text{makeMCP}(1, x) \{ \text{MCP}(s, 1, x, v) \} \\
& \lambda; \mathcal{A} \vdash \{ \text{MCP}(s, 1, x, v) \} \text{unmakeMCP}(1, x) \{ x \mapsto v \} \\
& \lambda; \mathcal{A} \vdash \mathbb{W}v. \langle \text{MCP}(s, 1, x, v) \rangle \text{read}(1, x) \langle \text{ret} = v * \text{MCP}(s, 1, x, v) \rangle \\
& \lambda; \mathcal{A} \vdash \mathbb{W}v. \langle \text{MCP}(s, 1, x, v) \rangle \text{write}(1, x, w) \langle \text{MCP}(s, 1, x, w) \rangle \\
& \lambda; \mathcal{A} \vdash \mathbb{W}v. \langle \text{MCP}(s, 1, x, v) \rangle \text{cas}(1, x, v1, v2) \left\langle \begin{array}{l} \text{if } v = v1 \text{ then } \text{ret} = 1 * \text{MCP}(s, 1, x, v2) \\ \text{else } \text{ret} = 0 * \text{MCP}(s, 1, x, v) \end{array} \right\rangle \\
& \lambda; \mathcal{A} \vdash \mathbb{W}v, w. \quad \left\langle \begin{array}{l} \text{MCP}(s, 1, x, v) * \text{MCP}(s, 1, y, w) \\ \text{b} := \text{dcas}(1, x, y, v1, w1, v2, w2); \\ \text{if } v = v1 * w = w1 \\ \text{then } \text{b} = 1 * \text{MCP}(s, 1, x, v2) * \text{MCP}(s, 1, y, w2) \\ \text{else } \text{b} = 0 * \text{MCP}(s, 1, x, v) * \text{MCP}(s, 1, y, w) \end{array} \right\rangle \\
& \lambda; \mathcal{A} \vdash \mathbb{W}v, w, u. \quad \left\langle \begin{array}{l} \text{MCP}(s, 1, x, v) * \text{MCP}(s, 1, y, w) * \text{MCP}(s, 1, z, u) \\ \text{b} := \text{3cas}(1, x, y, z, v1, w1, u1, v2, w2, u2); \\ \text{if } v = v1 * w = w1 * u = u1 \\ \text{then } \text{b} = 1 * \text{MCP}(s, 1, x, v2) * \text{MCP}(s, 1, y, w2) * \text{MCP}(s, 1, z, u2) \\ \text{else } \text{b} = 0 * \text{MCP}(s, 1, x, v) * \text{MCP}(s, 1, y, w) * \text{MCP}(s, 1, z, u) \end{array} \right\rangle \\
& \text{MCL}(s, l) \iff \text{MCL}(s, l) * \text{MCL}(s, l) \\
& \text{MCP}(s, l, x, v) * \text{MCP}(s, l, x, w) \implies \text{False}
\end{aligned}$$

Figure 5.5: The abstract specification for the MCAS module.

5.2.2 Implementation

We give a straightforward coarse-grained implementation of the MCAS specification. The full code is given in Figure 5.6. The operation `makeMCL` creates a lock which protects updates to pointers under the control of the module. The other operations simply acquire the lock, perform the appropriate reads and writes, and then release the lock. The module supports different kinds of compare-and-set operations, that allow more than one pointer to be updated atomically. The `makeMCP` and `unmakeMCP` operations allow pointers to be moved to the module and from the module. The `read` operation has to acquire the lock as well, this prevents reading intermediate states of the compare-and-set operations that update more than one pointer. Moreover, the `unmakeMCP` also has to acquire the lock to guarantee that no other operation is attempting to read the pointer. By acquiring the lock it ensures full ownership over the pointer which can then be moved away from the module.

We interpret the abstract predicates using a single shared region, with type name **MCAS**. The abstract states of the region are *partial heaps*, which represent the part of the heap that is protected by the module. For instance, the abstract state $x \mapsto v \bullet y \mapsto w$ indicates that heap cells x and y are under the protection of the module, with logical values v and w respectively. Note that the physical

```

function makeMCL() {
  l := makeLock();
  return l;
}

function makeMCP(l, x) {
  skip;
}

function read(l, x) {
  acquire(l);
  v := [x];
  release(l);
  return v;
}

function write(l, x, v) {
  acquire(l);
  [x] := v;
  release(l);
}

function dcas(l, x, y, v1, w1, v2, w2) {
  acquire(l);
  v := [x];
  w := [y];
  if (v = v1 and w = w1) {
    [x] := v2;
    [y] := w2;
    r := 1;
  } else {
    r := 0;
  }
  release(l);
  return r;
}

function unmakeMCP(l, x) {
  acquire(l);
  release(l);
}

function cas(l, x, v1, v2) {
  acquire(l);
  v := [x];
  if (v = v1) {
    [x] := v2;
    r := 1;
  } else {
    r := 0;
  }
  release(l);
  return r;
}

function 3cas(l, x, y, z, v1, w1, u1, v2, w2, u2) {
  acquire(l);
  v := [x];
  w := [y];
  u := [z];
  if (v = v1 and w = w1 and u = u1) {
    [x] := v2;
    [y] := w2;
    [z] := u2;
    r := 1;
  } else {
    r := 0;
  }
  release(l);
  return r;
}

```

Figure 5.6: Multiple compare-and-set module operations.

values at x and y need not be the same as their logical values, specifically when the lock has been acquired and they are being modified.

For the **MCAS** region, there are five kinds of guards. The $\text{OWN}(x)$ guard confers ownership of the heap cell at address x under the control of the region. This guard is used by all operations of the module that access the heap cell x . In order to ensure that there can only be one instance of $\text{OWN}(x)$ we require $\text{OWN}(x) \bullet \text{OWN}(x)$ to be undefined. We amalgamate the OWN guards for heap cells that are not currently under the protection of the module into $\text{OWNED}(X)$, where X is the set of all cells that *are* protected. We define the following equality on guards:

$$\text{OWNED}(X) = \text{OWNED}(X \uplus \{x\}) \bullet \text{OWN}(x)$$

Initially, the set X will be empty. When we add an element $x \mapsto v$ to the region, we get a guard $\text{OWN}(x)$ that allows us to manipulate the abstract state for that particular x . There can be only one OWNED guard, we enforce this by requiring $\text{OWNED}(X) \bullet \text{OWNED}(Y)$ is not defined.

The remaining guards are effectively used as auxiliary state. When a thread acquires the lock, it removes some heap cells from the shared region in order to access them. The $\text{LOCKED}(h)$ guard will

be used to record that the heap cells in h have been removed in this way. The thread that acquired the lock will have a corresponding $\text{KEY}(h)$ guard. When it releases the lock, the two guards will be reunited inside the region to form the UNLOCKED guard. This is expressed by the following equivalence:

$$\text{UNLOCKED} = \text{LOCKED}(h) \bullet \text{KEY}(h)$$

The transition system for the region is parametric in each heap cell. It allows anyone to add the resource $x \mapsto v$ to the region. There is no need to guard this action, as the resource is unique and as such only one thread can do it for a particular value of x . It allows the value of x to be updated using the guard $\text{OWN}(x)$. Finally, given the guard $\text{OWN}(x)$, $x \mapsto v$ can be removed from the region. We formally define the transition system as follows:

$$\begin{aligned} \mathbf{0} & : \quad \forall h, x, v. h \rightsquigarrow x \mapsto v \bullet h \\ \text{OWN}(x) & : \quad \forall h, v, w. x \mapsto v \bullet h \rightsquigarrow x \mapsto w \bullet h \\ \text{OWN}(x) & : \quad \forall h, x, v. x \mapsto v \bullet h \rightsquigarrow h \end{aligned}$$

We define the interpretation of abstract states for the **MCAS** region:

$$\begin{aligned} I(\mathbf{MCAS}_a(s', l, h)) & \triangleq [\text{OWNED}(\text{dom}(h))]_a * (\text{Lock}(s', l, 0) * h * [\text{UNLOCKED}]_a \vee \\ & \quad \exists h_1, h_2. \text{Lock}(s', l, 1) * h_1 * [\text{LOCKED}(h_2)]_a * h = h_1 \bullet h_2) \end{aligned}$$

Internally, the region may be in one of two states, indicated by the disjunction. Either the lock l is unlocked, and the heap cells corresponding to the abstract state of the region are actually in the region, as well as the UNLOCKED guard, or the lock l is locked, and some portion h_1 of the abstract heap is in the region, while the remainder h_2 has been removed, together with the $\text{KEY}(h_2)$ guard, leaving behind the $\text{LOCKED}(h_2)$ guard. In both cases, the $\text{OWNED}(\text{dom}(h))$ guard belongs to the region, encapsulating the OWN guards for heap addresses that are not protected.

We now give an interpretation to the predicates as follows:

$$\begin{aligned} \text{MCL}((a, s'), l) & \triangleq \exists h. \mathbf{MCAS}_a(s', l, h) \\ \text{MCP}((a, s'), l, x, v) & \triangleq \exists h. \mathbf{MCAS}_a(s', l, x \mapsto v \bullet h) * [\text{OWN}(x)]_a \end{aligned}$$

The predicate $\text{MCL}(s, l)$ states the existence of the shared region, but makes no assumptions about its state. The predicate $\text{MCP}(s, l, x, v)$ states that there is x with value v , which it owns, and possibly other heap cells in the region. The axiom $\text{MCP}(s, l, x, v) * \text{MCP}(s, l, x, w) \implies \text{False}$ follows from the fact that $\text{OWN}(x) \bullet \text{OWN}(x)$ is not defined.

We can now prove that the specification is satisfied by the implementation. We show the **dcas** command in Fig. 5.7. The other commands have very similar proofs.

5.2.3 Resource Transfer

Consider an addition to the **MCAS** module: the **readTo** operation takes an **MCAS** heap cell and an ordinary heap cell and copies the value of the former into the latter. Such an operation could be



Figure 5.7: Proof of the dcas implementation.

implemented as follows:

```

function readTo(l, x, y) {
  v := read(l, x);
  [y] := v;
}

```

This implementation atomically reads the MCAS cell at x , then writes the value to the cell at y . The overall effect is non-atomic, in the sense that a concurrent environment could update x and then witness y being updated to the old value of x . However, if the environment's interaction is confined to the MCAS cell, the effect *is* atomic.

TaDA allows us to specify this kind of partial atomicity by splitting the pre- and postcondition of an atomic judgement into a *private* and a *public* part. The private part will contain resources that are particular to the thread—in this example, the heap cell at y . When the atomic triple is used to update a region (*e.g.* with the `USEATOMIC` rule), these private resources cannot form part of the region’s invariant. The public part will contain resources that can form part of a region’s invariant — in this example, the MCAS cell at x .

The `readTo` operation can be specified as follows:

$$\vdash \mathbb{W}v, w. \langle y \mapsto w \mid \text{MCP}(s, 1, x, v) \rangle \text{ readTo}(1, x, y) \langle y \mapsto v \mid \text{MCP}(s, 1, x, v) \rangle$$

One way of understanding such specifications is in terms of ownership transfer between a client and a module, as in [22]: ownership of the private precondition is transferred from the client; ownership of the private postcondition is transferred to the client. In this example, the same resources (albeit modified) are transferred in and out, but this need not be the case in general. For instance, an operation could allocate a fresh location in which to store the retrieved value, which is then transferred to the client.

5.3 Deque

We show how to use TaDA to specify a double-ended queue (deque) and verify its fine-grained implementation. A deque allows elements to be inserted and removed from both ends of a list.

This example shows that TaDA can scale to multiple levels of abstraction: the deque uses MCAS, which uses the lock, which is based on primitive atomic heap operations. This proof development would not be possible with CAP, since atomicity is central to the abstractions at each level. It would also not be possible using traditional approaches to linearisability, since separation of resources between and within abstraction layers is also crucial.

We define a constructor `makeDeque()`, that creates a new empty deque. The deque supports the insertion and removal of elements from the left side and the right side. On the left side we can insert using `pushLeft(d, v)` and remove using `popLeft(d)`. The right side has similar operations, `pushRight(d, v)` and `popRight(d)`.

5.3.1 Atomic specification

We represent the deque state by the abstract predicate `Deque(s, d, vs)`. It asserts that there is a deque at address d with list of elements vs .

The `makeDeque()` operation creates an empty deque and returns its address. It has the following specification:

$$\lambda; \mathcal{A} \vdash \{ \text{True} \} \text{ makeDeque}() \{ \exists s \in \mathbb{T}_8. \text{ Deque}(s, \text{ret}, []) \}$$

The operations `pushLeft(d, v)` and `popLeft(d)` are specified to update the state of the deque atomically:

$$\begin{aligned} & \lambda; \mathcal{A} \vdash \mathbb{W}vs. \langle \text{Deque}(s, d, vs) \rangle \text{ pushLeft}(d, v) \langle \text{Deque}(s, d, v : vs) \rangle \\ & \lambda; \mathcal{A} \vdash \mathbb{W}vs. \langle \text{Deque}(s, d, vs) \rangle \text{ popLeft}(d) \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } v = 0 * \text{Deque}(s, d, vs) \\ \text{else } vs = v : vs' * \text{ret} = v * \text{Deque}(s, d, vs') \end{array} \right\rangle \end{aligned}$$

The `pushLeft(d, v)` operation adds the value v to the left of the deque. The `popLeft(d)` operation tries to remove an element from the left end of the deque. However, if the deque is empty, then it returns 0 and does not change its state. Otherwise, it removes the element at the left, updating the state of the deque, and returns the removed value. The `pushRight` and `popRight` operations have analogous specifications, operating on the right end of the deque.

5.3.2 Implementation

We consider an implementation that represents the deque as a doubly-linked list of nodes, based on the *Snark* linked-list deque [16]. The code is shown in Figure 5.8.

An example of the shape of the data structure is shown in Figure 5.9. Each node consists of a left-link pointer, a right-link pointer, and a value. There are two anchor variables, *left hat* and *right hat* ($\hat{l}$ and $\hat{r}$ in the figure), that generally point to the leftmost and rightmost node in the list, except when the deque is empty. When the deque is not empty, its leftmost node's left-link and the rightmost node's right-link point to a so-called *dead* node—a node whose left- and right-links point to itself (e.g. node a in the figure). When the deque is empty, then the left hat and the right hat point to dead nodes.

We focus on the `popLeft` implementation. This implementation first reads the left hat value to a local variable. It then reads the left-link of the node referenced by that variable. If both values are the same, it means that the node is dead and the list might be empty. It is necessary to recheck the left hat to confirm, since the node might have died since the left hat was first read. If the deque is indeed empty, the operation returns 0; otherwise it is restarted. If the left node is not dead, it tries to atomically update the left hat to point to the node to its right, and, at the same time, update the left node to be dead. This could fail, in which case the operation restarts. An example of such an update is shown in Figure 5.9. In order to update three pointers atomically, the implementation makes use of the `3cas` command described in §5.2.

To verify the `popLeft` operation, we introduce a new region type, **Deque**. The region has two parameters, d standing for the deque address and L for the MCAS address. There is only one non-empty guard for the region, named G . We represent the abstract state by a tuple (ns, ds) where: ns is a list of pairs of node addresses and values, the values representing the elements stored in the deque; and ds is a set of pairs of nodes addresses and values that were part of the deque, but are now dead. We maintain the set of dead nodes to guarantee that after a node is removed from the deque, its value can still be read. In order to change the abstract state of the deque, we require the guard G . The labelled transition system is as follows:

$$\begin{aligned} G &: \forall n, v, ns, ds. (ns, ds) \rightsquigarrow ((n, v) : ns, ds) \\ G &: \forall n, v, ns, ds. (ns, ds) \rightsquigarrow (ns : (n, v), ds) \\ G &: \forall n, v, ns, ds. ((n, v) : ns, ds) \rightsquigarrow (ns, ds \uplus \{(n, v)\}) \\ G &: \forall n, v, ns, ds. (ns : (n, v), ds) \rightsquigarrow (ns, ds \uplus \{(n, v)\}) \end{aligned}$$

In order to provide an interpretation for the abstract state, we first define a number of auxiliary predicates. A node predicate at address n in the deque is defined using the MCAS cells predicates:

$$\text{node}(s, L, n, l, r, v) \equiv \text{MCP}(s, L, n.\text{left}, l) * \text{MCP}(s, L, n.\text{right}, r) * n.\text{value} \mapsto v$$

```

function makeDeque() {
  L := makeMCL();
  n := makeNode(0, 0, 0);
  write(L, n.left, n);
  write(L, n.right, n);
  p := alloc(3);
  [p.left] := n; // left hat
  [p.right] := n; // right hat
  [p.mcl] := L; // pointer to MCAS interface
  [p.dummy] := n; // pointer to dummy node
  makeMCP(L, p.left);
  makeMCP(L, p.right);
  return p;
}

function makeNode(L, l, r, v) {
  n := alloc(3);
  [n.left] := l;
  [n.right] := r;
  [n.value] := v;
  makeMCP(L, l);
  makeMCP(L, r);
  return n;
}

function pushLeft(d, v) {
  L := [d.mcl];
  m := [d.dummy];
  n := makeNode(L, m, 0, v);
  while (true) {
    lh := read(L, d.left);
    lhL := read(L, lh.left);
    if (lhL = lh) {
      write(L, n.right, m);
      rh := read(L, d.right);
      b := dcas(L, d.left, lh.right, lh, rh, n, n);
      if (b = 1) {
        return ;
      }
    } else {
      write(L, n.right, lh);
      b := dcas(L, d.left, lh.left, lh, lhL, n, n);
      if (b = 1) {
        return ;
      }
    }
  }
}

function popLeft(d) {
  L := [d.mcl];
  while (true) {
    lh := read(L, d.left);
    lhR := read(L, lh.right);
    lhL := read(L, lh.left);
    if (lhL = lh) { // left hat seems dead
      lh2 := read(L, d.left);
      if (lh2 = lhL) { // left hat confirmed dead
        return 0;
      } // left hat not dead — try again
    } else {
      b := 3cas(L, d.left, lh.right, lh.left,
                lh, lhR, lhL, lhR, lh, lh);
      if (b = 1) {
        v := [lh.value];
        return v;
      }
    }
  }
}

```

where

$$\mathbb{E}.left \stackrel{\text{def}}{=} \mathbb{E} \quad \mathbb{E}.right \stackrel{\text{def}}{=} \mathbb{E} + 1 \quad \mathbb{E}.value \stackrel{\text{def}}{=} \mathbb{E} + 2 \quad \mathbb{E}.mcl \stackrel{\text{def}}{=} \mathbb{E} + 2 \quad \mathbb{E}.dummy \stackrel{\text{def}}{=} \mathbb{E} + 3.$$

Figure 5.8: Snark deque operations.




Figure 5.9: Examples of a deque before and after performing `popLeft`, which uses `3cas` to updated pointers `c`, `d` and `e`.

Here, l and r are the left- and right-link addresses. The L parameter is the address of the MCAS lock. A dead node is defined as:

$$\text{dead}(s, L, n, v) \equiv \text{node}(s, L, n, n, n, v)$$

We also define a predicate to stand for the doubly-linked list that contains all the elements in the list, (i.e. the shaded nodes in the figure).

$$\begin{aligned} \text{dlseg}(s, L, l, r, n, m, ns) &\equiv ns = [] * l = m * r = n \vee \\ &\exists v, ns', p. ns = (l, v) : ns' * \text{node}(s, L, l, n, p, v) * \text{dlseg}(s, L, p, r, l, m, ns') \end{aligned}$$

We define a predicate to include the dead nodes (ds) as well as the doubly-linked list:

$$\text{dls}(s, L, l, r, ns, ds) \equiv \exists a, b. (a, -), (b, -) \in ds * \text{dlseg}(s, L, l, r, a, b, ns) * \bigotimes_{(n,v) \in ds} \text{dead}(s, L, n, v)$$

Note that there must be at least one dead node in ds .

Our last auxiliary predicate is used to represent the whole deque: the double linked list; the anchors left hat and right hat; and the reference to the MCAS interface.

$$\begin{aligned} \text{deque}(s, d, L, ns, ds) &\equiv \exists l, r. \text{dls}(s, L, l, r, ns, ds) * \\ &\text{MCP}(s, L, d.\text{left}, l) * \text{MCP}(s, L, d.\text{right}, r) * d.\text{mcl} \mapsto L * \text{MCL}(s, L) \end{aligned}$$

We now define the interpretation of abstract states as follows:

$$I(\mathbf{Deque}_a(s, d, L, (ns, ds))) \triangleq \mathbf{deque}(s, d, L, ns, ds)$$

We define the interpretation of the **Deque** predicate as follows:

$$\mathbf{Deque}((a, s', L), d, vs, (ns, ds)) \triangleq \mathbf{Deque}_a(s', d, L, ns, ds) * [G]_a * vs = \mathit{proj}_2(ns)$$

where proj_2 is the second projection over the list of pairs ns .

We show the proof of the **popLeft** operation in Figure 5.10. To prove the implementation against our atomic specifications, we use the **MAKEATOMIC** rule again. The **UPDATEREGION** rule is applied at two different points. The first one is when performing a read to confirm that the left hat is dead. We require the **UPDATEREGION** rule as there is a possibility that the deque is empty and as such this corresponds to the atomic update, which requires updating the atomicity tracking component. The second point is when performing the **3cas** operation to remove an element of the deque, in the case of success it corresponds to the atomic update. The remaining proofs are similar.

We have shown that by using TaDA, we can prove atomicity for fine-grained implementations that use other concurrent implementations. We achieve scalability by using atomicity to abstract the updates at each abstraction level.

We have shown how we can use atomicity in our specifications to build up abstraction levels. The deque's abstract atomic specification is implemented using the abstract atomicity of the MCAS, which is, in turn, implemented using the abstract atomicity of the lock. With CAP, we could not do this because we cannot specify atomicity, only weaker specifications.

$\mathbb{W}vs.$

$\langle \text{Deque}(d, vs) \rangle$

ABSTRACT: $s = (a, s', L), t = (ns, ds)$	MAKEATOMIC	$ \begin{array}{l} \mathbb{W}(ns, vs). \\ \langle \text{Deque}_a(s', d, L, ns, ds) * [G]_a * vs = \text{proj}_2(ns) \rangle \\ a : (ns, ds) \rightsquigarrow \text{if } ns = [] \text{ then } (ns, ds) \text{ else } (ns', (n, v) : ds) * ns = (n, v) : ns' \vdash \\ \quad \{ \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge \} \\ L := [d.mcl]; \\ \text{while (true) } \{ \\ \quad \{ \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge * L = L \} \\ \quad \text{lh} := \text{read}(L, d.\text{left}); \text{lhR} := \text{read}(L, \text{lh}.\text{right}); \text{lhL} := \text{read}(L, \text{lh}.\text{left}); \\ \quad \left\{ \begin{array}{l} \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge * L = L * \\ \text{if } \text{lh} = \text{lhL} \text{ then } (\text{lh}, -) \in ds \\ \quad \text{else } \{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \end{array} \right\} \\ \quad \text{if } (\text{lhL} = \text{lh}) \{ // \text{left hat seems dead} \\ \quad \quad \{ \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge * L = L * (\text{lhL}, -) \in ds \} \\ \quad \quad \text{UPDATEREGION} \left\{ \begin{array}{l} \mathbb{W}ns, ds. \\ \langle \text{deque}(s', d, L, ns, ds) * L = L * (\text{lhL}, -) \in ds \rangle \\ \text{lh2} := \text{read}(L, d.\text{left}); \\ \langle \text{deque}(s', d, L, ns, ds) * L = L * \rangle \\ \langle (\text{lh2} = \text{lhL} \rightarrow ns = []) \rangle \end{array} \right\} \\ \quad \quad \left\{ \begin{array}{l} \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * L = L * \\ \text{if } \text{lh2} = \text{lhL} \text{ then } a \Rightarrow ([], ds), ([], ds) \text{ else } a \Rightarrow \blacklozenge \end{array} \right\} \\ \quad \quad \text{if } (\text{lh2} = \text{lhL}) \{ // \text{left hat confirmed dead} \\ \quad \quad \quad \text{return } 0; \\ \quad \quad \quad \{ \exists ds. \text{ret} = 0 * a \Rightarrow ([], ds), ([], ds) \} \\ \quad \quad \} // \text{left hat not dead — try again} \\ \quad \} \text{ else } \{ \\ \quad \quad \left\{ \begin{array}{l} \exists ns, ds. \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge * L = L * \\ \{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \end{array} \right\} \\ \quad \quad \text{UPDATEREGION} \left\{ \begin{array}{l} \mathbb{W}ns, ds. \\ \langle \text{deque}(s', d, L, ns, ds) * L = L * \\ \{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \rangle \\ \text{b} := 3\text{cas}(L, d.\text{left}, \text{lh}.\text{right}, \text{lh}.\text{left}, \text{lh}, \text{lhR}, \text{lhL}, \text{lhR}, \text{lh}, \text{lh}); \\ \left\langle \begin{array}{l} \exists ns', v. \text{if } b = 1 \text{ then } \left(\begin{array}{l} \text{deque}(s', d, L, ns', (\text{lh}, v) : ds) * \\ L = L * (\text{lh}, v) \in ds * ns = (\text{lh}, v) : ns' \end{array} \right) \\ \text{else } \text{deque}(s', d, L, ns, ds) * L = L \end{array} \right\rangle \\ \left\{ \begin{array}{l} \exists ns, ds, v. \text{if } b = 1 \text{ then } \left(\begin{array}{l} a \Rightarrow ((\text{lh}, v) : ns, ds), (ns, (\text{lh}, v) : ds) \\ * L = L * (\text{lh}, v) \in ds \end{array} \right) \\ \text{else } \text{Deque}_a(s', d, L, ns, ds) * a \Rightarrow \blacklozenge * L = L \end{array} \right\} \\ \text{if } (b = 1) \{ \\ \quad v := [\text{lh}.\text{value}]; \text{return } v; \\ \quad \{ \exists ns, ds. \text{ret} = v * a \Rightarrow ((\text{lh}, v) : ns, ds), (ns, (\text{lh}, v) : ds) \} \\ \} \} \} \\ \quad \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{ret} = 0 * \text{Deque}_a(s', d, L, ns, ds) * [G]_a \\ \text{else } \left(\begin{array}{l} \exists ns', v. ns = (n, v) : ns' * \text{ret} = v * \\ \text{Deque}_a(s', d, L, ns', (n, v) : ds) * [G]_a * vs' = \text{proj}_2(ns') \end{array} \right) \end{array} \right\rangle \\ \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{ret} = 0 * \text{Deque}(d, vs) \\ \text{else } \exists vs', v. vs = v : vs' * \text{ret} = v * \text{Deque}(d, vs') \end{array} \right\rangle \end{array} $
--	------------	--

Figure 5.10: Proof of the popLeft implementation.

6 Related Work

TaDA inherits from a family of logics deriving from concurrent separation logic [47]: RGSep [63], Deny-Guarantee [15], CAP [12], Higher-Order CAP (HOCAP) [55] and Impredicative CAP (iCAP) [54]. In particular, it makes use of dynamic *shared regions* with capability resources (called *guards* in TaDA) that determine how the regions may be updated. Following iCAP, TaDA eschews the use of boxed assertions to describe the state of shared regions and instead represents regions by *abstract states*. The protocol for updating the region is specified as a transition system on these abstract states, labelled by guards. This use of transition systems to describe protocols derives from previous work by Dreyer *et al.* [17], and also appears in Turon *et al.* [61] as “local life stories”.

By treating the abstract state-space of a region as a separation algebra, it is possible to localise updates on it, as in the MCAS example (§5.2). Such locality is in the spirit of local life stories [61], and can be seen as an instance of Ley-Wild and Nanevski’s “subjective auxiliary state” [37].

While HOCAP and iCAP do not support abstract atomic specifications, they support an approach to atomicity introduced by Jacobs and Piessens [32] that achieves similar effects. In their work, operations may be parametrised by an update to auxiliary state that is performed when the abstract atomic operation appears to take effect. This update is performed atomically by the implementation, and can therefore involve shared regions. This approach is inherently higher-order, which has the disadvantage of leading to complex specifications, which are parametrised by the specification of the auxiliary code that performs the atomic update.

Recently, Iris [35] combined TaDA with iCAP [54] into a new higher order program logic, which encodes TaDA’s proof rules in logic. Atomic triples have been encoded using view shifts in Iris [35] by interpreting them as specifications in the Jacobs-Piessens style. This captures the intensional meaning behind atomic triples—that is, what they can be used for—which in TaDA is expressed through the proof rules for using atomic triples. Because it is higher-order, it has expressive power to handle higher-order programs and reentrancy. TaDA takes a first-order approach, leading to simpler specifications. An alternative treatment of TaDA, in a refinement calculus setting, can be found in [45].

There has been extensive work understanding and generalising linearisability, especially in light of work on separation logic. Vafeiadis [63] has combined the ownership given by his RGSep reasoning with linearisability. Gotsman and Yang [22] have generalised linearisability to include ownership transfer of memory between a client and a module, which is also supported by our approach. We can do ownership transfer with the private and the public components. Filipovic *et al.* [19] have demonstrated that linearisability can be viewed as a particular proof technique for contextual refinement. Turon *et al.* [59] have introduced CaReSL, a logic that combines contextual refinement and Hoare-style reasoning to prove higher-order concurrent programs. Like linearisability, contextual refinement requires a whole-module approach.

7 Reasoning about Termination

Throughout this thesis, we have proved properties of programs using the partial correctness interpretation: if a program is run in a state satisfying the precondition and the program terminates, then the resulting state will be described by the postcondition. In some programs, it is also important to know *that* they terminate. To prove such a property is especially challenging for concurrent programs. When multiple threads are changing some shared resource, knowing if each thread terminates can often depend on the behaviour of the other threads and even on the scheduler that decides which thread should run at a particular moment.

We have applied TaDA to reason about fine-grained concurrency, which is characterised by the use of low-level synchronisation operations (such as compare-and-set). A well-known class of fine-grained concurrent programs is that of *non-blocking* algorithms. With non-blocking algorithms, suspension of a thread cannot halt the progress of other threads: the progress of a single thread cannot require another thread to be scheduled. Thus, if the interference from the environment is suitably restricted, the operations are guaranteed to terminate.

If we prove that a program produces the correct results and also always completes in a finite time, we establish *total correctness*. Turing [58] and Floyd [20] introduced the use of well-founded relations, combined with partial-correctness arguments, to prove the termination of sequential programs. The same technique is general enough to prove concurrent programs, too. However, previous applications of this technique in the concurrent setting, which we discuss in §7.5.4, do not support straightforward reasoning about clients.

In this chapter, we extend TaDA with well-founded termination reasoning. With the resulting logic, Total-TaDA, we can prove total correctness of fine-grained concurrent programs. The novelty of our approach is in using TaDA’s abstraction mechanisms to specify constraints on the environment necessary to ensure termination. It retains the modularity of TaDA and abstracts the internal termination arguments. We demonstrate our approach on a counter module and a stack module.

We observe that Total-TaDA can be used to verify standard non-blocking properties of algorithms. However, our specifications capture more: we propose the concept of *non-impedance* that our specifications suggest. We say that one operation *impedes* another if the second can be prevented from terminating by repeated concurrent invocations of the first. This concept seems important to the design and use of non-blocking algorithms, where we have some expectation about how clients use the algorithm, and what progress guarantees they expect.

Using an atomic triple, an increment operation of a counter is specified as:

$$\vdash \forall n \in \mathbb{N}. \langle \text{Counter}(s, x, n) \rangle \text{incr}(x) \langle \text{Counter}(s, x, n+1) * \text{ret} = n \rangle$$

The internal structure of the counter is abstracted using the abstract predicate [50] $\text{Counter}(s, x, n)$, which states that there is a counter at address x with value n , whereas s abstracts implementation-

specific information about the counter. The specification says that the `incr` atomically increments the counter by 1. The environment is allowed to update the counter to any value of n as long as it is a natural number. The specification enforces obligations on both the client and the implementation: the client must guarantee that the counter is not destroyed and that its value is a natural number until the atomic update occurs; and the implementation must guarantee that it does not change the value of the counter until it performs the specified atomic action. Working at the abstraction of the counter means that each operation can be verified without knowing the rest of the operations of the module. Consequently, modules can be extended with new operations without having to re-verify the existing operations. Additionally, the implementation of `incr` can be replaced by another implementation that satisfies the same specification, without needing to re-verify the clients that make use of the counter. While atomic triples are expressive, they do not guarantee termination. In particular, an implementation could block, deadlock or live-lock and still be considered correct.

Non-blocking Algorithms

In general, guaranteeing the termination of concurrent programs is a difficult problem. In particular, termination could depend on the behaviour of the scheduler (whether or not it is *fair*) and on other threads that might be competing for resources. We focus on non-blocking programs. Non-blocking programs have the benefit that their termination is not dependent on the behaviour of the scheduler.

There are two common non-blocking properties: *wait-freedom* [25] and *lock-freedom* [41]. Wait-freedom requires that operations complete irrespective of the interference caused by other threads: termination cannot depend on the amount of interference caused by the environment. Lock-freedom is less restrictive. It requires that, when multiple threads are performing operations, then at least one of them must make progress. This means that a thread might never terminate if the amount of interference caused by the environment is unlimited.

TaDA is well suited to reasoning about interference between threads. In particular, we can write specifications that limit the amount of interference caused by the client, and so guarantee termination of lock-free algorithms in Total-TaDA. We will see how both wait-freedom and lock-freedom can be expressed in Total-TaDA.

Moreover, there is another non-blocking property in the literature called *obstruction-freedom* [27]. It is the weakest requirement, guaranteeing progress only if a thread is run in isolation for a sufficiently long amount of time. In the presence of other operations, it can cause a livelock with each thread undoing the work done by other threads.

We cannot express in general obstruction-freedom in Total-TaDA, as these programs do not may not terminate without further constraints. It is possible to reason about this class of algorithms in Total-TaDA if one limits the amount of interference to guarantee termination.

Termination

Well-founded relations provide a general way to prove termination. In particular, Floyd [20] used well-founded relations to prove the termination of sequential programs. In fact, it is sufficient to use ordinal numbers [6] without losing expressivity. A LOOP rule, using ordinals and adapted from Floyd's

work, has the form:

$$\text{LOOP} \quad \frac{\forall \gamma \leq \alpha. \vdash_{\tau} \{P(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{\exists \beta. P(\beta) \wedge \beta < \gamma\}}{\vdash_{\tau} \{P(\alpha)\} \text{ while } (\mathbb{B}) \{\mathbb{C}\} \{\exists \beta. P(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\}}$$

The loop invariant $P(\gamma)$ is parametrised by an ordinal γ (the *variant*) which is decreased by every execution of the loop body $\mathbb{C}$. Because ordinals cannot have infinite descending chains, the loop must terminate in a finite number of steps. This proof rule allows termination reasoning to be localised to the individual loops in the program. In this chapter, we extend TaDA with termination based on ordinal numbers, using the LOOP rule given above.

Total-TaDA

We obtain the program logic Total-TaDA by modifying TaDA to have a total-correctness semantics. The details are given in §7.2. With Total-TaDA, we can specify and verify non-blocking algorithms. Wait-free operations always terminate, independently of the operations performed by the environment. For lock-free operations, however, we need to restrict the amount of interference the environment can cause in order to guarantee termination. Our key insight is that, as well as bounding the number of iterations of loops, ordinals can bound the interference on a module. This allows us to give total-correctness specifications for lock-free algorithms. In §7.1, we specify and verify lock-free implementations of a counter. The specification introduces ordinals to bound the number of times a client may update the counter. This makes it possible to guarantee that the lock-free increment operation will terminate, since either it will succeed or some other concurrent increment will succeed. As the number of increments is bounded, the operation must eventually succeed.

Total-TaDA retains the modularity of TaDA. In particular, we can verify the termination of clients of modules using the total-correctness specifications, without reference to the implementation. We show an example of this in §7.1.2. Since the client only depends on the specification, we can replace the implementation. In §7.1.3, we show that two different implementations of a counter satisfy the same total-correctness specification. With Total-TaDA, we can verify the operations of a module independently, exploiting locality.

As a case study for Total-TaDA, we show how to specify and verify both functional correctness and termination of Treiber's stack in §7.3. In §7.4, we present the total-correctness semantics and the soundness proof of Total-TaDA. In §7.5, we show how lock-freedom and wait-freedom can be expressed with Total-TaDA specifications. We also introduce the concept of non-impedance in §7.5.3 and argue for its value in specifying non-blocking algorithms.

7.1 Motivating Examples

We introduce Total-TaDA by providing specifications of the operations of a counter module. We justify the specifications by using them to reason about two clients, one sequential and one concurrent. We show how two different implementations can be proved to satisfy the specification.

Consider a counter module with a constructor `makeCounter` and two operations: `incr`, which increments the value of the counter by 1 and returns its previous value; and `read`, which returns the

value of the counter. We give an implementation of the spin counter operations in Figure 7.1, and an alternative implementation of `incr` in Figure 7.5.

7.1.1 Atomic Specification

The Total-TaDA specification for the `makeCounter()` operation is the following Hoare triple with a total-correctness interpretation:

$$\forall \alpha. \vdash_{\tau} \{ \text{True} \} \text{makeCounter}() \{ \exists s. \text{Counter}(s, \text{ret}, 0, \alpha) \}$$

The counter predicate is extended with an ordinal parameter α that bounds the amount of interference the counter can sustain. When the value of the counter is updated, the ordinal α must decrease.

The `makeCounter()` operation allocates a new counter with value 0, and allows the client to pick an initial ordinal α . If a finite bound on the number of updates can be determined beforehand, then that bound is an appropriate choice for the ordinal. However, it could be the case that that bound is determined by subsequent (non-deterministic) operations, in which case an infinite ordinal should be used. For example, consider the following client program:

```
x := makeCounter();
m := random();
while (m > 0) { incr(x); m := m - 1; }
```

Here, the number of increments is bounded by the (finite) value returned by `random`, but it is not determined when the counter is constructed. Choosing $\alpha = \omega$ (the first infinite ordinal) is appropriate in this case: the first increment can decrease the ordinal from ω to $m - 1$, while subsequent increments simply decrement the ordinal by 1. We will later see examples where higher ordinals must be used.

The increment operation is specified as follows:

$$\forall \beta. \vdash_{\tau} \forall n \in \mathbb{N}, \alpha. \langle \text{Counter}(s, x, n, \alpha) * \alpha > \beta(n, \alpha) \rangle \text{incr}(x) \langle \text{Counter}(s, x, n + 1, \beta(n, \alpha)) * \text{ret} = n \rangle$$

The specification resembles the partial-correctness specification given in the introduction, but with the addition of the ordinal α and the function β . The client chooses how to decrease the ordinal by providing a function β that determines the new ordinal in terms of the old ordinal and previous value of the counter. The condition $\alpha > \beta(n, \alpha)$ requires the client to guarantee that such a decrease is possible. (So, for example, the client could not use the specification in a situation where the concurrent environment might reduce the ordinal to zero.) The implementation may rely on the fact that a

<pre>function makeCounter() { x := alloc(1); [x] := 0; return x; }</pre>	<pre>function incr(x) { b := 0; while (b = 0) { v := [x]; b := CAS(x, v, v + 1); } return v; }</pre>	<pre>function read(x) { v := [x]; return v; }</pre>
--	--	---

Figure 7.1: Spin counter operations.

counter's ordinal cannot be increased to guarantee termination.

The read operation is specified as follows:

$$\vdash_{\tau} \forall n \in \mathbb{N}, \alpha. \langle \text{Counter}(s, \mathbf{x}, n, \alpha) \rangle \text{ read}(\mathbf{x}) \langle \text{Counter}(s, \mathbf{x}, n, \alpha) * \text{ret} = n \rangle$$

Unlike the increment, the read operation does not affect the ordinal, meaning that the client is not bounded with respect to the number of reads it performs. Such a specification is possible for operations that do not impede the progress of other operations. In this case, **read** does not impede **incr** or **read**.

Finally, we have an axiom that allows the client to decrease the ordinal without requiring any physical operation to be performed:

$$\forall s, n, \alpha, \beta < \alpha. \text{Counter}(s, \mathbf{x}, n, \alpha) \implies \text{Counter}(s, \mathbf{x}, n, \beta)$$

This is possible because the ordinals do not have any concrete representation in memory. They are just a logical mechanism to limit the amount of interference over a resource.

The ordinal parameter is exposed in the specification of the counter to allow the implementation to guarantee that its loops terminate. In a wait-free implementation—where the termination of each operation is independent of the concurrent operations—it would not be necessary to expose the ordinal parameter. For this counter, the read operation is wait-free, while the increment operation is lock-free, since termination depends on bounding the number of interfering increments.

7.1.2 Clients

A Sequential Client

Consider a program that creates a counter and contains two nested loops. As in the previous example, the outer loop runs a finite but randomly determined number of times. The inner loop also runs a randomly determined number of times, and increments the counter on each iteration. Figure 7.2 shows this client, together with its total-correctness proof.

The LOOP rule is used for each of the loops: for the outer loop, the variant is **n**; for the inner loop, the variant is **m**. Since the number of iterations of each loop is determined before it is run, the variants need only be considered up to finite ordinals (*i.e.* natural numbers). We could modify the code to use a single loop that conditionally decrements **n** (and randomises **m**) or decrements **m**. This variation would require a transfinite ordinal for the variant.

As well as enforcing loop termination, ordinals play a role as a parameter to the **Counter** predicate, which must be decreased on each increment. When we create the counter, we choose ω^2 as the initial ordinal. We have seen that ω allows us to decrement the counter a non-deterministic (but finite) number of times. We want to repeat this a non-deterministic (but finite) number of times, so $\omega \cdot \omega = \omega^2$ is the appropriate ordinal. Once the number **n** of iterations of the outer loop is determined, we decrease this to $\omega \cdot \mathbf{n}$ by using the axiom provided by the counter module. Similarly, when **m** is chosen, we decrease the ordinal from $\omega \cdot \mathbf{n} = \omega \cdot (\mathbf{n} - 1) + \omega$ to $\omega \cdot (\mathbf{n} - 1) + \mathbf{m}$.

A Concurrent Client

Consider a program that creates two threads, each of which increments the counter a finite, but unbounded, number of times. We again prove this client using the abstract specification of the counter, as shown in Figure 7.3. For presentation purposes, we present the example using parallel composition rather than using fork. One could rewrite the example to use fork and encode the join of threads using auxiliary heap state.

We reuse the PARALLEL from §2.6:

$$\frac{\text{PARALLEL} \quad \frac{\{P_1\} \mathbb{C}_1 \{Q_1\} \quad \{P_2\} \mathbb{C}_2 \{Q_2\}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}}$$

In this example, the counter is shared between the two threads, which may concurrently update it. To reason about sharing, we use a *shared region*.

As in TaDA, a shared region encapsulates some resource that is available to multiple threads. Threads can access the resource when performing (abstractly) atomic operations, such as `incr`. The region presents an abstract state, and defines a protocol that determines how the region may be updated. Ghost resources, called *guards*, are associated with transitions in the protocol. The guards for a region form a partial commutative monoid with the operation $\bullet$, which is lifted to $*$ in assertions. In order for a thread to make a particular update, it must have ownership of a guard associated with the

```

{True}
x := makeCounter();
{∃s. Counter(s, x, 0, ω2)}
n := random();
{∃s. Counter(s, x, 0, ω · n) * (α = n)}
while (n > 0) {
  ∀γ.
  {∃s, v. Counter(s, x, v, ω · n) * (γ = n) * n > 0}
  m := random();
  {∃s, v. Counter(s, x, v, ω · (n - 1) + m) * (γ = n) * n > 0}
  0
  ∧
  n
  *
  n
  =
  n
  (γ
  =
  n)
  FRAME:
  {
    while (m > 0) {
      ∀δ.
      {∃s, v. Counter(s, x, v, ω · (n - 1) + m) * (δ = m) * m > 0}
      incr(x);
      {∃s, v. Counter(s, x, v, ω · (n - 1) + m - 1) * (δ = m) > 0}
      m := m - 1;
      {∃s, ζ, v. Counter(s, x, v, ω · (n - 1) + m) * (ζ = m) * ζ < δ}
    }
  }
  {∃s, v. Counter(s, x, v, ω · (n - 1)) * (γ = n) * n > 0}
  n := n - 1;
  {∃s, β, v. Counter(s, x, v, ω · n) * (β = n) * β < γ}
}
{∃s, v. Counter(s, x, v, 0)}

```

Figure 7.2: Proof of a sequential client of the counter.

corresponding transition. All guards are allocated along with the region they are associated with.

For the concurrent client, we introduce a region with type name **CClient**. This region encapsulates the shared counter. Accordingly, the region type is parametrised by the address of the counter. The abstract state of the region records the current value of the counter.

There are two types of guard resources associated with **CClient** regions. The guard $\text{INC}(m, \beta, \pi)$ provides capability to increment the counter. Conceptually, multiple threads may have INC guards, and a fractional permission $\pi \in (0, 1]$ (in the style of [4]) is used to keep track of these capabilities. The parameter m expresses the *local contribution* to the value of the counter — the actual value is the sum of the local contributions. The ordinal parameter β represents a local bound on the number of increments. Again, the actual bound is a sum of the local bounds. Standard ordinal addition is inconvenient since it is not commutative; we use the natural (or Hessenberg) sum [28], denoted $\oplus$, which is associative, commutative, and monotone in its arguments.

To allow the INC guard to be shared among threads, we impose the following equivalence on guards:

$$\text{INC}(n + m, \alpha \oplus \beta, \pi_1 + \pi_2) = \text{INC}(n, \alpha, \pi_1) \bullet \text{INC}(m, \beta, \pi_2)$$

where $n \geq 0$, $m \geq 0$ and $1 \geq \pi_1 + \pi_2 > 0$. This equivalence expresses that INC guards can be split (or joined), preserving the total contribution to the value of the counter, ordinal bound and permission.

The second type of guard resource is $\text{TOTAL}(n, \alpha)$, which tracks the actual value of the counter n and ordinal α . These values should match the totals for the INC guards, which we enforce by requiring the following implication to hold:

$$\text{TOTAL}(n, \alpha) \bullet \text{INC}(m, \beta, 1) \text{ defined} \implies n = m \wedge \alpha = \beta$$

We wish to allow the contributions recorded in INC guards to change, but to do so we must simultaneously update the TOTAL guard, as expressed by the following equivalence:

$$\text{TOTAL}(n + m, \alpha \oplus \beta) \bullet \text{INC}(m, \beta, \pi) = \text{TOTAL}(n + m', \alpha \oplus \beta') \bullet \text{INC}(m', \beta', \pi)$$

The possible states of **CClient** regions are the natural numbers $\mathbb{N}$, representing the value of the shared counter, together with the distinguished state $\circ$, representing that the region is no longer required. The protocol for a region is specified by a guarded transition system, which describes how the abstract state may be updated in atomic steps, and which guard resources are required to do so. The transitions for **CClient** regions are as follows:

$$\text{INC}(m, \gamma, \pi) : n \rightsquigarrow n + 1 \quad \text{INC}(m, \gamma, 1) : m \rightsquigarrow \circ$$

This specifies that any thread with an INC guard may increment the value of the counter, and a thread owning the full INC guard may dispose of the region.

It remains to define the interpretation of the region states:

$$\begin{aligned} I(\mathbf{CClient}_a(s, x, n)) &\triangleq \exists \alpha. \text{Counter}(s, x, n, \alpha) * [\text{TOTAL}(n, \alpha)]_a \\ I(\mathbf{CClient}_a(s, x, \circ)) &\triangleq \text{True} \end{aligned}$$

By interpreting the state $\circ$ as **True**, we allow a thread transitioning into that state to acquire the counter that previously belonged to the region. This justifies the last step of the proof in Figure 7.3.

The proof rule that allows us to use the atomic specification of the **incr** operation to update the shared region is the **USEATOMIC** rule, inherited from TaDA. The $\{\}$ -assertions in Total-TaDA are also required to be stable. That is, the region states must account for the changes that the concurrent environment could make, under the assumption that it has guards that are compatible with those of the thread. This is why the state of the **CClient** region is always existentially quantified in Figure 7.3.

$$\begin{array}{c}
\{\text{True}\} \\
\mathbf{x} := \text{makeCounter}(); \\
\{\exists s. \text{Counter}(s, \mathbf{x}, 0, \omega \oplus \omega)\} \\
\{\exists s, a. \text{CClient}_a(s, \mathbf{x}, 0) * [\text{INC}(0, \omega \oplus \omega, 1)]_a\} \\
\{\exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(0, \omega, \frac{1}{2})]_a * 0 \leq v\} \\
\mathbf{n} := \text{random}(); \\
\mathbf{i} := 0; \\
\{\exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(\mathbf{i}, \mathbf{n}, \frac{1}{2})]_a * 0 \leq v * \mathbf{i} = 0\} \\
\text{while } (\mathbf{i} < \mathbf{n}) \{ \\
\quad \forall \beta. \left\{ \begin{array}{l} \exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(\mathbf{i}, \beta, \frac{1}{2})]_a * \mathbf{i} \leq v \\ * \mathbf{i} < \mathbf{n} * \beta = \mathbf{n} - \mathbf{i} \end{array} \right\} \\
\quad \text{incr}(\mathbf{x}); \mathbf{i} := \mathbf{i} + 1; \\
\quad \left\{ \begin{array}{l} \exists s, \delta, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(\mathbf{i}, \delta, \frac{1}{2})]_a * \mathbf{i} \leq v \\ * \mathbf{i} \leq \mathbf{n} * \delta = \mathbf{n} - \mathbf{i} * \delta < \beta \end{array} \right\} \\
\} \\
\{\exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(\mathbf{n}, 0, \frac{1}{2})]_a\} \\
\{\exists s, a. \text{CClient}_a(s, \mathbf{x}, \mathbf{n} + \mathbf{m}) * [\text{INC}(\mathbf{n} + \mathbf{m}, 0, 1)]_a\} \\
\{\exists s. \text{Counter}(s, \mathbf{x}, \mathbf{n} + \mathbf{m}, 0)\}
\end{array}
\parallel
\begin{array}{c}
\{\exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(0, \omega, \frac{1}{2})]_a\} \\
\mathbf{m} := \text{random}(); \\
\mathbf{j} := 0; \\
\text{while } (\mathbf{j} < \mathbf{m}) \{ \\
\quad \text{incr}(\mathbf{x}); \\
\quad \mathbf{j} := \mathbf{j} + 1; \\
\} \\
\{\exists s, v. \text{CClient}_a(s, \mathbf{x}, v) * [\text{INC}(\mathbf{m}, 0, \frac{1}{2})]_a\}
\end{array}$$

Figure 7.3: Proof of a concurrent client of the counter.

7.1.3 Implementations

We prove the total correctness of the two distinct increment implementations against the abstract specification given in §7.1.1.

Spin Counter Increment

Consider **incr** as shown in Figure 7.1. Note that the read, write and compare-and-set operations are atomic. We want to prove the total correctness of **incr** against the atomic specification. The first step is to give a concrete interpretation of the abstract predicate $\text{Counter}(s, x, n, \alpha)$. We introduce a new region type, **Counter**, with only one non-empty guard, **G**. The abstract states of the region are pairs of the form (n, α) , where n is the value of the counter and α is a bound on the number of increments. All transitions are guarded by **G** with the transition:

$$G : \forall n \in \mathbb{N}, m \in \mathbb{N}, \alpha > \beta. (n, \alpha) \rightsquigarrow (n + m, \beta)$$

The transition requires that updates to the state of the region must decrease the ordinal. This allows us to effectively bound interference, necessary to guarantee the termination of the loop in **incr**.

The interpretation of the **Counter** region states is defined as follows:

$$I(\mathbf{Counter}_a(x, n, \alpha)) \triangleq x \mapsto n$$

The expression $x \mapsto n$ asserts that there exists a heap cell with address x and value n . Note that α is not represented in the concrete heap, as it is not part of the program. We use it solely to ensure that the number of operations is finite.

We define the interpretation of the abstract predicate as follows:

$$\mathbf{Counter}(a, x, n, \alpha) \triangleq \mathbf{Counter}_a(x, n, \alpha) * [G]_a$$

The abstract predicate $\mathbf{Counter}(a, x, n, \alpha)$ asserts that there is a **Counter** region with identifier a , address x , and with abstract state (n, α) . Furthermore, it encapsulates exclusive ownership of the guard G , and so embodies exclusive permission to update the counter. (Note that the type of the first parameter of **Counter**, which is abstract to the client, is instantiated as **RId**.)

The specification for the increment is atomic and as such, we use the **MAKEATOMIC** rule which is just the same as that of TaDA. The only difference is that termination is enforced. Whereas in TaDA it would be possible for an abstract atomic operation to loop forever without performing its atomic update, in Total-TaDA it is guaranteed to eventually perform the update.

A proof of the increment implementation is shown in Figure 7.4. The atomicity context allows the environment to modify the abstract state of the counter, without restriction on the number of times. The **Counter** transition system enforces that the ordinal α must decrease every time the value of the counter is increased. This means that the number of times the region's abstract state is updated is finite. Our loop invariant is parametrised with a variant γ that takes the value of α at the beginning of each loop iteration. When we first read the value of the counter n , we can assert: $n > v \implies \gamma > \alpha$.

If the compare-and-set operation fails, the value of the counter has changed. This can only happen in accordance with the region's transition system, and so the ordinal parameter α must have decreased. As such, the invariant still holds but for a lower ordinal, $\alpha < \gamma$. We are localising the termination argument for the loop, by relating the local variant with the ordinal parametrising the region.

If the compare-and-set succeeds, then we record our update from (v, α) to $(v + 1, \beta(v, \alpha))$, where β is the function chosen by the client that determines how the ordinal is reduced. The **MAKEATOMIC** rule allows us to export this update in the postcondition of the whole operation.

Backoff Increment

Consider a different implementation of the increment operation, given in Figure 7.5, which loops attempting to perform the operation, like the previous increment. However, if the compare-and-set fails due to contention, it waits for a random number of iterations before retrying.

Despite the differences to the previous increment, the specification is the same. In fact, we can give the same interpretation for the abstract predicate $\mathbf{Counter}(s, x, n, \alpha)$, and the same guards and regions that were used for the previous implementation. (Since this is the case, a counter module could provide *both* of these operations: the proof system guarantees that they work correctly together.)

The main difference in the proof is that each iteration of the loop depends on not only the amount of interference on the counter, but also on the variable n that is randomised when the compare-and-set




Figure 7.4: Proof of total correctness of increment.

fails. Any random number will be smaller than ω , and the maximum amount of times that the compare-and-set can fail is α , the parameter of the **Counter** predicate. This is because α is a bound on the number of times the counter can be incremented. We therefore use $\omega \cdot \alpha + \mathbf{n}$ as the upper bound on the number of loop iterations.

Let γ be equal to $\omega \cdot \alpha + \mathbf{n}$ at the start of the loop iteration. At each loop iteration, we have two cases, when $\mathbf{n} = 0$ or otherwise. In the first case we try to perform the increment by doing a compare-and-set. If the compare-and-set succeeds, then the increment occurs and the loop will exit. If it fails, then the environment must have decreased α . This means that $\gamma \geq \omega \cdot \alpha + \omega$ for the new value of α . We then set $\mathbf{n}$ to be a new random number, which is less than ω , and end up with $\gamma > \omega \cdot \alpha + \mathbf{n}$. In the second case of the loop iteration, we simply decrement $\mathbf{n}$ by 1 and we know that $\gamma > \omega \cdot \alpha + \mathbf{n}$ for the new value of $\mathbf{n}$. The proof of the backoff increment is shown in Figure 7.6.

```

function incr(x) {
  n := 0;
  b := 0;
  while (b = 0) {
    if (n = 0) {
      v := [x];
      b := CAS(x, v, v + 1);
      n := random();
    } else {
      n := n - 1;
    }
  }
  return v;
}

```

Figure 7.5: Backoff increment.

7.2 Logic

Total-TaDA is a Hoare logic which can be used to prove total correctness for fine-grained non-blocking concurrent programs. The logic is essentially the same as for TaDA, simply adapted to incorporate termination analysis using ordinals in a standard way.

The Total-TaDA proof judgement has the form:

$$\lambda; \mathcal{A} \vdash_{\tau} \mathbb{W}x \in X. \langle P_p \mid P(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle.$$

The meaning of the judgement is similar to TaDA's, however it has a total correctness interpretation, which additionally enforces that $\mathbb{C}$ will eventually terminate. The pre- and postconditions are split into a private part (the P_p and $Q_p(x, y)$) and a public part (the $P(x)$ and $Q(x, y)$). The idea is that the command may make multiple, non-atomic updates to the private part, but must only make a single atomic update to the public part. Before the atomic update, the environment is allowed to change the public part of the state, but only by changing the parameter x of P which must remain within X . After the atomic update, the specification makes no constraint on how the environment modifies the public state. All that is known is that, immediately after the atomic update, the public and private parts satisfy the postcondition for a common value of y . The private assertions in our judgements must be *stable*: that is, they must account for any updates other threads could have sufficient resources to perform. The main difference is that the judgement has a total-correctness interpretation, i.e. $\mathbb{C}$ must eventually terminate. We use the τ subscript to emphasise the total-correctness interpretation.

The proof rules are mostly the same as in TaDA, except for the LOOP rule and, FUNCTION rule, which use variants to prohibit non-termination as follows:

$\forall \beta. \mathbb{W}n \in \mathbb{N}, \alpha.$



Figure 7.6: Proof of total correctness of backoff increment.

$$\begin{array}{c}
\text{LOOP} \\
\frac{\forall \gamma \leq \alpha. \Gamma; \lambda; \mathcal{A} \vdash_{\tau} \{P(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{ \exists \beta. P(\beta) \wedge \beta < \gamma \}}{\Gamma; \lambda; \mathcal{A} \vdash_{\tau} \{P(\alpha)\} \text{ while } (\mathbb{B}) \{ \mathbb{C} \} \{ \exists \beta. P(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha \}} \\
\\
\text{FUNCTION} \\
\frac{\begin{array}{c} \Gamma = \Gamma', \forall \gamma < \alpha. \lambda; \mathcal{A} \vdash_{\tau} \mathbb{W}x \in X. \langle P_p(\vec{z}, \gamma) \mid P(x, \vec{z}) \rangle \mathbf{f}(\vec{z}) \quad \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle \\ \text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E}; \quad \text{vars}(\eta(\mathbf{f})) = \vec{x} \\ \mathbb{W}x \in X. \langle P_p(\vec{z}, \alpha) * \vec{x} = \vec{z} \mid P(x, \vec{z}) \rangle \\ \Gamma'; \lambda; \mathcal{A} \vdash_{\tau} \quad \mathbb{C} \\ \exists (y, \text{ret}) \in Y \times \text{Val}. \langle Q_p(x, y, \vec{z}, \text{ret}) * \text{ret} = \mathbb{E} \mid Q(x, y, \vec{z}, \text{ret}) \rangle \end{array}}{\Gamma; \lambda; \mathcal{A} \vdash_{\tau} \mathbb{W}x \in X. \langle P_p(\vec{z}, \alpha) \mid P(x, \vec{z}) \rangle \mathbf{f}(\vec{z}) \quad \exists y \in Y. \langle Q_p(x, y, \vec{z}, \text{ret}) \mid Q(x, y, \vec{z}, \text{ret}) \rangle}
\end{array}$$

The LOOP rule enforces that the number of times that the loop body can run is finite. The rule allows us to perform a while loop if we can guarantee that each loop iteration decreases the ordinal parametrising the invariant P . By the finite-chain property of ordinals, there cannot be an infinite number of iterations.

The FUNCTION rule enforces that the number of times that function body can call itself or other functions is finite. This is enforced similarly to the loop, by decreasing the ordinal parametrising the precondition of the function. Note that this has a direct impact on the FUNCTIONCALL rule and FORK rule, as they make use of the function specifications. We are guaranteeing that both function calls and forks occur a finite amount of times within a program.

Moreover, we add a new proof rule to the non-deterministic function used throughout this chapter:

$$\begin{array}{c}
\text{RANDOM} \\
\hline
\Gamma; \lambda; \mathcal{A} \vdash_{\tau} \{ \text{True} \} \mathbf{x} := \text{random}(); \{ \exists n \in \mathbb{N}. \mathbf{x} = n \}
\end{array}$$

7.3 Case Study: Treiber's Stack

We now consider a version of Treiber's stack [57] to demonstrate how Total-TaDA can be applied to verify the total correctness of larger modules.

$$\begin{array}{c}
\forall \alpha. \vdash_{\tau} \{ \text{True} \} \text{makeStack}() \{ \exists s \in \mathbb{T}_1, t \in \mathbb{T}_2. \text{Stack}(s, \text{ret}, [], t, \alpha) \} \\
\forall \beta. \vdash_{\tau} \mathbb{W}vs, t, \alpha. \langle \text{Stack}(s, \mathbf{x}, vs, t, \alpha) * \alpha > \beta(vs, \alpha) \rangle \text{push}(\mathbf{x}, \mathbf{v}) \langle \exists t'. \text{Stack}(s, \mathbf{x}, \mathbf{v} : vs, t', \beta(vs, \alpha)) \rangle \\
\vdash_{\tau} \mathbb{W}vs, t, \alpha. \langle \text{Stack}(s, \mathbf{x}, vs, t, \alpha) \rangle \text{pop}(\mathbf{x}) \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{Stack}(s, \mathbf{x}, vs, t, \alpha) * \text{ret} = 0 \\ \text{else } \exists vs', t'. \text{Stack}(s, \mathbf{x}, vs', t', \alpha) * vs = \text{ret} : vs' \end{array} \right\rangle
\end{array}$$

Figure 7.7: Stack operation specifications.

7.3.1 Atomic Specification

In Figure 7.7, we give the specification of the lock-free stack operations. This is a Total-TaDA specification satisfiable by a reasonable non-blocking implementation. As with the counter, the predicate representing the stack is parametrised by an ordinal that bounds the number of operations on the stack, in order to guarantee termination. The $\text{Stack}(s, x, vs, t, \alpha)$ predicate has five parameters: the address of the stack x ; its contents vs ; an ordinal α that decreases every time a **push** operation is performed; and two parameters, s and t that range over abstract types $\mathbb{T}_1$ and $\mathbb{T}_2$ respectively. These last two parameters encapsulate implementation-specific information about the configuration of the stack (s is invariant, while t may vary) and hence their types are abstract to the client. Note that it is possible to abstract the parameter t using the AEXISTS rule in this case.

The constructor returns an empty stack, parametrised by an arbitrary ordinal chosen by the client. The **push** operation atomically adds an element to the head of the stack. The **pop** operation atomically removes one element from the head of the stack, if one is available (*i.e.* the stack is non-empty); otherwise it will simply return 0. (As this stack is non-blocking, it would not be possible for the **pop** operation to wait for the stack to become non-empty.)

Note that the ordinal parametrising the stack is not required to decrease when popping the stack. This means that the stack operations cannot be starved by an unbounded number of **pop** invocations. This need not be the case in general for a lock-free stack, but it is true for Treiber's stack. We discuss the ramifications of this kind of specification further in §7.5.3.

7.3.2 Implementation

Figure 7.8 gives an implementation of the stack operations based on Treiber's stack [7]. The stack is represented as a heap cell containing a pointer (the head pointer) to a singly-linked list of the values on the stack. Values are pushed onto the stack by allocating a new node holding the value to be pushed and a pointer to the old head of the stack. A compare-and-set operation updates the old head of the stack to point to the new node. If the operation fails, it will be because the head of the stack has changed, and so the operation is retried. Values are popped from the stack by moving the head pointer one step along the list. Again, a compare-and-set operation is used for this update, so if the head of the stack changes the operation can be retried. If the stack is empty (*i.e.* the head points to 0), then **pop** simply returns 0, without affecting the stack.

To prove correctness of the implementation, we introduce predicates to represent the linked list:

$$\begin{aligned} \text{list}(x, ns) &\triangleq (x = 0 * ns = []) \vee (\exists v, l. \text{node}(x, v, l) * \text{list}(l, ns') * ns = (x, v) : ns') \\ \text{node}(n, v, l) &\triangleq n.\text{value} \mapsto v * n.\text{next} \mapsto l \end{aligned}$$

It is important for the correctness of the algorithm that nodes that have been popped can never reappear as the head of the stack. Otherwise, a **pop** operation could have seen the node when it was previously the head, and update the head to point to what was then the next node, but may no longer be. To account for this, in our representation of the stack we track the set of previously popped nodes, and ensure that they are disjoint from the nodes in the stack. The $\text{stack}(x, ns, ds)$ predicate, therefore,

consists of a list starting at address x , with contents ns , and a disjoint set of discarded nodes ds :

$$\text{stack}(x, ns, ds) \triangleq \text{list}(x, ns) * \bigotimes_{(n,v) \in ds} \text{node}(n, v, -)$$

We define a region type **TStack** to hold the shared data-structure. The type is parametrised by the address of the stack, and its abstract state consists of a list of nodes in the stack ns , a set of popped nodes ds , and an ordinal α . The **TStack** region type has the following interpretation:

$$I(\mathbf{TStack}_a(x, ns, ds, \alpha)) \triangleq \exists y. x \mapsto y * \text{stack}(y, ns, ds)$$

We use a single guard G to give threads permissions to push and pop the stack. The transition system is given as follows:

$$\begin{aligned} G : \forall n, v, ns, ds, \alpha, \beta < \alpha. (ns, ds, \alpha) &\rightsquigarrow ((n, v) : ns, ds, \beta) \\ G : \forall n, v, ns, ds, \alpha. ((n, v) : ns, ds, \alpha) &\rightsquigarrow (ns, (n, v) \uplus ds, \alpha) \end{aligned}$$

The first action allows us to add an element to the head of the stack. The second action allows us to remove the top element of the stack, adding it to the set of discarded nodes. There is no explicit transition for the pop on the empty stack, since this operation does not change the abstract state.

Note that for every transition $(ns, ds, \alpha) \rightsquigarrow (ns', ds', \alpha')$, we have $2 \cdot \alpha + |ns| > 2 \cdot \alpha' + |ns'|$. Pushing decreases the ordinal, but extends the length of the stack by 1; popping maintains the ordinal, but decreases the length of the stack. This property allows us to use $2 \cdot \alpha + |ns|$ as a variant in the compare-and-set loops, since it is guaranteed to decrease under any interference.

The abstract predicate $\text{Stack}(s, x, vs, t, \alpha)$ combines the region and the guard:

$$\text{Stack}(a, x, vs, (ns, ds), \alpha) \triangleq \mathbf{TStack}_a(x, ns, ds, \alpha) * [G]_a * vs = \text{snds}(ns)$$

The function snds returns the list of elements of the second elements of the list of pairs ns . Consequently, vs is the list of values on the stack, rather than address-value pairs.

The proof for the **pop** operation is given in Figure 7.9. When the stack is non-empty, if the compare-

<pre>function makeStack() { x := alloc(1); [x] := 0; return x; }</pre>	<pre>function push(x, v) { y := alloc(2); [y.value] := v; do { z := [x]; [y.next] := z; b := CAS(x, z, y); } while (b = 0); }</pre>	<pre>function pop(x) { do { y := [x]; if (y = 0) { return 0; } z := [y.next]; b := CAS(x, y, z); } while (b = 0); v := [y.value]; return v; }</pre>
--	---	---

where

$$\mathbb{E}.\text{value} \stackrel{\text{def}}{=} \mathbb{E} \qquad \mathbb{E}.\text{next} \stackrel{\text{def}}{=} \mathbb{E} + 1.$$

Figure 7.8: Treiber's stack operations.

and-set fails then another thread must have succeeded in updating the stack, and so reduced the ordinal or the length of the stack; by basing the loop variant on the ordinal and stack length, we can guarantee that the operation will eventually succeed.

The proof of the **push** operation is in Figure 7.10. If the compare-and-set fails, then another thread must have updated the stack by removing or adding an element. As in the **pop** operation proof, updating the stack reduces the ordinal or the length of the stack. We related the loop variant on the ordinal and the stack length to guarantee that the operation will eventually succeed.

7.4 Semantics and Soundness

The semantics of Total-TaDA are mostly similar to that for TaDA, we will focus on the differences.

We first extend the operational semantics of the language to include the new function as follows:

$$\frac{\text{RANDOM} \quad n \in \mathbb{N}}{(\sigma, \mathbf{x} := \text{random}();) \xrightarrow{\text{id}}_{\eta} (\sigma[\mathbf{x} \mapsto n], \text{skip};)}$$

We use the same model for assertions as that for TaDA. We also use a similar semantic judgement, $\models_{\tau}$, which ensures that the concrete behaviours of programs simulate the abstract behaviours represented by the specifications.

Definition 7.1 (Semantic Judgement). The semantic judgement

$$\eta; \lambda; \mathcal{A} \models_{\tau} \forall x \in X. \langle P_p \mid P(x) \rangle c \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle$$

where

- $\eta \in \text{FEnv}$ is a function environment;
- $\lambda \in \text{Level}$ is a level strictly greater than that of any region that will be affected by the program;
- $\mathcal{A} \in \text{AContext}$ is the atomicity context, which constrains updates to regions on which an abstractly atomic update is to be performed;
- $P_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ is the private part of the precondition, which does not correspond to resources in some opened shared region, and is parametrised by the valuation of program variables;
- $P \in X \rightarrow \text{View}_{\mathcal{A}}$ is the public part of the precondition, which may correspond to resources from some opened shared regions, and is parametrised by $x \in X$ that tracks the precondition at the atomic update and by the valuation of program variables;
- $c \in \text{ExtCmd}$ is the program under consideration;
- $Q_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ is the private part of the postcondition, which is parametrised by $x \in X$ that tracks the precondition at the atomic update, by $y \in Y$ that tracks the postcondition at the atomic update, and by the valuation of program variables;
- $Q \in X \times Y \rightarrow \text{View}_{\mathcal{A}}$ is the public part of the postcondition, which is similarly parametrised by $x \in X$ and $y \in Y$, and by the valuation of program variables;

$$\begin{array}{c}
\mathbb{W}vs, t, \alpha. \\
\langle \text{Stack}(s, x, vs, t, \alpha) \rangle \\
\langle \text{TStack}_a(x, ns, ds, \alpha) * [G]_a * vs = \text{snds}(ns) \rangle \\
a : (ns, ds, \alpha) \rightsquigarrow \text{if } ns = [] \text{ then } (ns, ds, \alpha) \text{ else } (ns', (n, v) \uplus ds, \alpha) \wedge ns = (n, v) : ns' \vdash_\tau \\
\{ \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge \} \\
\text{do } \{ \\
\quad \forall \gamma. \\
\quad \{ \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \gamma \geq 2 \cdot \alpha + |ns| \} \\
\quad \text{UPDATEREGION} \left\langle \begin{array}{l} \mathbb{W}ns, ds, \alpha. \\ \langle \exists w. x \mapsto w * \text{stack}(w, ns, ds) * \gamma \geq 2 \cdot \alpha + |ns| \rangle \\ y := [x]; \\ \langle x \mapsto y * \text{stack}(y, ns, ds) * \gamma \geq 2 \cdot \alpha + |ns| * \\ \text{if } y = 0 \text{ then } ns = [] \text{ else } \exists v. (y, v) = \text{head}(ns) \rangle \end{array} \right\rangle \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{if } y = 0 \text{ then } a \Rightarrow (([], ds, \alpha), ([], ds, \alpha)) \\ \text{else } \exists v. \left(\text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * (y, v) \in ns \uplus ds * \right. \\ \left. \gamma \geq 2 \cdot \alpha + |ns| * \text{head}(ns) \neq (y, v) \implies \gamma > 2 \cdot \alpha + |ns| \right) \end{array} \right\} \\
\quad \text{if } (y = 0) \{ \\
\quad \quad \text{return } 0; \\
\quad \quad \{ \exists ds, \alpha. a \Rightarrow (([], ds, \alpha), ([], ds, \alpha)) * \text{ret} = 0 \} \\
\quad \} \\
\quad \left\{ \begin{array}{l} \exists ns, ds, v, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * (y, v) \in ns \uplus ds * \\ \gamma \geq 2 \cdot \alpha + |ns| * \text{head}(ns) \neq (y, v) \implies \gamma > 2 \cdot \alpha + |ns| \end{array} \right\} \\
\quad z := [y.\text{next}]; \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \gamma \geq 2 \cdot \alpha + |ns| * \\ \left((\exists v, v', ns'. ns = [(y, v), (z, v')] \uplus ns') \vee (\exists v. ns = [(y, v)] * z = 0) \right) \\ \vee (\exists v. (y, v) \in ns \uplus ds * \text{head}(ns) \neq (y, v) * \gamma > 2 \cdot \alpha + |ns|) \end{array} \right\} \\
\quad \text{UPDATEREGION} \left\langle \begin{array}{l} \mathbb{W}ns, ds, \alpha. \\ \langle \exists w. x \mapsto w * \text{stack}(w, ns, ds) * \gamma \geq 2 \cdot \alpha + |ns| * \\ \left((\exists v, v', ns'. ns = [(y, v), (z, v')] \uplus ns') \vee (\exists v. ns = [(y, v)] * z = 0) \right) \\ \vee (\exists v. (y, v) \in ns \uplus ds * \text{head}(ns) \neq (y, v) * \gamma > 2 \cdot \alpha + |ns|) \rangle \end{array} \right\rangle \\
\quad b := \text{CAS}(x, y, z); \\
\quad \left\langle \begin{array}{l} \text{if } b = 0 \text{ then } \exists w. x \mapsto w * \text{stack}(w, ns, ds) * \gamma > 2 \cdot \alpha + |ns| \\ \text{else } \exists v, ns'. x \mapsto z * \text{stack}(z, ns', (y, v) \uplus ds) * ns = (y, v) : ns' \end{array} \right\rangle \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \gamma \geq 2 \cdot \alpha + |ns| * \\ \text{if } b = 0 \text{ then } \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \gamma > 2 \cdot \alpha + |ns| \\ \text{else } \left(\exists v, ns', ds', \alpha'. (y, v) \in ds' * \text{TStack}_a(x, ns', ds', \alpha') \right) \\ * a \Rightarrow ((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds), \alpha \end{array} \right\} \\
\quad \} \text{ while } (b = 0); \\
\quad \left\{ \begin{array}{l} \exists v, ns, ds, \alpha, ns', ds', \alpha'. (y, v) \in ds' * \text{TStack}_a(x, ns', ds', \alpha') \\ * a \Rightarrow ((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds, \alpha) \end{array} \right\} \\
\quad v := [y.\text{value}]; \\
\quad \{ \exists ns, ds, \alpha. a \Rightarrow ((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds, \alpha) \} \\
\quad \text{return } v; \\
\quad \{ \exists y, ns, ds, \alpha. a \Rightarrow ((y, \text{ret}) : ns, ds, \alpha), (ns, (y, \text{ret}) \uplus ds), \alpha \} \\
\quad \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{TStack}_a(x, ns, ds, \alpha) * [G]_a * vs = \text{snds}(ns) * \text{ret} = 0 \\ \text{else } \exists ns', vs', y. \text{TStack}_a(x, ns', (y, \text{ret}) \uplus ds, \alpha) * [G]_a * vs' = \text{snds}(ns') * ns = (y, \text{ret}) : ns' \end{array} \right\rangle \\
\langle \text{if } vs = [] \text{ then } \text{Stack}(s, x, vs, t, \alpha) * \text{ret} = 0 \text{ else } \exists vs', t'. \text{Stack}(s, x, vs', t', \alpha) * vs = \text{ret} : vs' \rangle
\end{array}$$

Figure 7.9: Proof of total correctness of Treiber's stack pop operation.

ABSTRACT: $s = a, t = (ns, ds)$	MAKEATOMIC $\forall \beta. \forall vs, t, \alpha.$ $\langle \text{Stack}(s, x, vs, t, \alpha) * \alpha > \beta(\alpha, vs) \rangle$ $\langle \text{TStack}_a(x, ns, ds, \alpha) * [G]_a * vs = \text{snds}(ns) * \alpha > \beta(\alpha, \text{snds}(ns)) \rangle$ $a : (ns, ds, \alpha) \wedge \alpha > \beta(\alpha, \text{snds}(ns)) \rightsquigarrow ((n, v) : ns, ds, \beta(\alpha, \text{snds}(ns))) \vdash_\tau$ $\{ \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \alpha > \beta(\alpha, \text{snds}(ns)) \}$ $y := \text{alloc}(2);$ $\{ \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \text{node}(y, -, -) * \alpha > \beta(\alpha, \text{snds}(ns)) \}$ $[y] := v;$ $\{ \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \text{node}(y, v, -) * \alpha > \beta(\alpha, \text{snds}(ns)) \}$ $\text{do } \{$ $\quad \forall \gamma.$ $\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \text{node}(y, v, -) \\ * \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns \end{array} \right\}$ OPENREGION $\forall ns, ds, \alpha.$ $\langle \exists y. x \mapsto y * \text{stack}(y, ns, ds) * \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns \rangle$ $z := [x];$ $\langle \exists y. x \mapsto y * \text{stack}(y, ns, ds) * z = y * \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns \rangle$ $\left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \text{node}(y, v, -) \\ * \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns * \\ ((ns = [] * z = 0) \vee \text{head}(ns) = (z, -) \vee \gamma > 2 \cdot \alpha + ns) \end{array} \right\}$ $[y.\text{next}] := z;$ $\left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \text{node}(y, v, z) * \\ \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns * \\ ((ns = [] * z = 0) \vee \text{head}(ns) = (z, -) \vee \gamma > 2 \cdot \alpha + ns) \end{array} \right\}$ UPDATEREGION $\forall ns, ds, \alpha.$ $\left\langle \begin{array}{l} \exists w. x \mapsto w * \text{stack}(w, ns, ds) * \text{node}(y, v, z) * \\ \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma \geq 2 \cdot \alpha + ns * \\ ((ns = [] * z = 0) \vee \text{head}(ns) = (z, -) \vee \gamma > 2 \cdot \alpha + ns) \end{array} \right\rangle$ $b := \text{CAS}(x, z, y);$ $\left\langle \begin{array}{l} \text{if } b = 0 \text{ then } \exists w. x \mapsto w * \text{stack}(w, ns, ds) * \text{node}(y, v, -) * \gamma > 2 \cdot \alpha + ns \\ \text{else } x \mapsto y * \text{stack}(y, ((y, v) : ns, ds)) \end{array} \right\rangle$ $\left\{ \begin{array}{l} \text{if } b = 0 \text{ then } \left(\begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_a(x, ns, ds, \alpha) * a \Rightarrow \blacklozenge * \\ \text{node}(y, v, -) * \alpha > \beta(\alpha, \text{snds}(ns)) * \gamma > 2 \cdot \alpha + ns \end{array} \right) \\ \text{else } \exists ns, ds, \alpha. a \Rightarrow ((ns, ds, \alpha), ((y, v) : ns, ds, \beta(\alpha, \text{snds}(ns)))) \end{array} \right\}$ $\} \text{ while } (b = 0);$ $\{ \exists ns, ds, \alpha. a \Rightarrow ((ns, ds, \alpha), ((y, v) : ns, ds, \beta(\alpha, \text{snds}(ns)))) \}$ $\langle \exists n. \text{TStack}_a(x, (n, v) : ns, ds, \beta(\alpha, vs)) * [G]_a * v : vs = \text{snds}((n, v) : ns) \rangle$ $\langle \exists t'. \text{Stack}(s, x, v : vs, t', \beta(\alpha, vs)) \rangle$
---------------------------------	--

Figure 7.10: Proof of total correctness of Treiber's stack push operation.

is defined to be the least-general judgement that holds when the following conditions hold:

- For all $\sigma, \sigma_1 \in \text{Store}$, $c_1 \in \text{ExtCmd}$, $\alpha \in \text{AAction}$ with $(\sigma, c) \xrightarrow{\alpha}_{\eta} (\sigma_1, c_1)$, for all $x \in X$, there exists $P'_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$, $P''_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ such that

$$\eta; \lambda; \mathcal{A} \models \langle P_p(\sigma) * P(x) \rangle \alpha \langle P'_p(\sigma_1) * P(x) \vee \exists y \in Y. P''_p(x, y, \sigma_1) * Q(x, y) \rangle \quad (7.1)$$

$$\eta; \lambda; \mathcal{A} \models_{\tau} \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle c_1 \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle, \quad (7.2)$$

$$\text{and for all } y \in Y, \eta; \lambda; \mathcal{A} \models_{\tau} \{ P''_p(x, y) \} c_1 \{ Q_p(x, y) \}. \quad (7.3)$$

- For all $\sigma, \sigma_1 \in \text{Store}$, $c_1 \in \text{ExtCmd}$, $\mathbf{f}$, $\vec{v}$ with $(\sigma, c) \xrightarrow{\text{spawn}(\mathbf{f}, \vec{v})}_{\eta} (\sigma_1, c_1)$, for all $x \in X$, there exist $P'_p \in \text{Store} \rightarrow \text{View}_{\mathcal{A}}$, $P''_p \in X \times Y \rightarrow \text{Store} \rightarrow \text{View}_{\mathcal{A}}$ and $P_{\mathbf{f}} \in \text{Store} \rightarrow \text{View}$ such that for all $\sigma_{\mathbf{f}} \in \text{Store}$ with $\text{vars}(\eta(\mathbf{f})) = \vec{x}$, $\sigma_{\mathbf{f}}(\vec{x}) = \vec{v}$ and $\text{code}(\eta(\mathbf{f})) = \mathbb{C} \text{ return } \mathbb{E};$,

$$\lambda; \mathcal{A} \models P_p(\sigma) * P(x) \preceq P'_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) * P(x) \vee \exists y \in Y. P''_p(\sigma_1) * P_{\mathbf{f}}(\sigma_{\mathbf{f}}) * Q(x, y), \quad (7.4)$$

$$\eta; \lambda; \mathcal{A} \models_{\tau} \mathbb{W}x \in X. \langle P'_p \mid P(x) \rangle c_1 \quad \exists y \in Y. \langle Q_p(x, y) \mid Q(x, y) \rangle, \quad (7.5)$$

$$\text{for all } y \in Y, \eta; \lambda; \mathcal{A} \models_{\tau} \{ P''_p(x, y) \} c_1 \{ Q_p(x, y) \}, \quad (7.6)$$

$$\begin{aligned} & \mathbb{W}x \in 1. \langle \lambda \sigma_{\mathbf{f}}. P_{\mathbf{f}}(\vec{z}) * \sigma_{\mathbf{f}}(\vec{x}) = \vec{z} \mid \text{True} \rangle \\ \text{and } \eta; \lambda; \emptyset \models_{\tau} & \quad \mathbb{C} \quad . \quad (7.7) \\ & \exists (y, \text{ret}) \in 1 \times \text{Val}. \langle \lambda \sigma_{\mathbf{f}}. Q_{\mathbf{f}}(\vec{z}, \text{ret}) * \text{ret} = \mathcal{E}[\mathbb{E}]_{\sigma_{\mathbf{f}}} \mid \text{True} \rangle \end{aligned}$$

- If $c = \text{skip}$; then, for all $\sigma \in \text{Store}$, $x \in X$, there exists $y \in Y$ such that

$$\lambda; \mathcal{A} \models P_p(\sigma) * P(x) \preceq Q_p(x, y, \sigma) * Q(x, y). \quad (7.8)$$

The key distinction is that, whereas in TaDA the judgement (4.61) is defined coinductively (as a greatest fixed point), in Total-TaDA the judgement is defined inductively (as a least fixed point). By the Knaster-Tarski theorem [56], we know it has a least fixed-point. This means that TaDA admits executions that never terminate, while Total-TaDA requires executions to always terminate: that is, reach a base case of the inductive definition.

The proof of soundness of Total-TaDA is similar to that for TaDA. The soundness proof consists of lemmas that justify each of the proof rules for the semantic judgement. Most of the Total-TaDA rules have similar proofs to the corresponding TaDA rules, but proceed by induction instead of coinduction. Of course, the LOOP and FUNCTION rules are different, since termination does not follow trivially. We give the proof for LOOP.

Lemma 7.1 (LOOP Rule). Let α be an ordinal. If, for all $\gamma \leq \alpha$,

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{ P_p(\gamma) \wedge \mathbb{B} \} \mathbb{C} \{ \exists \beta. P_p(\beta) \wedge \beta < \gamma \} \quad (7.9)$$

then

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{ P_p(\alpha) \} \text{ while } (\mathbb{B}) \{ \mathbb{C} \} \{ \exists \beta. P_p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha \}. \quad (7.10)$$

Proof. The proof is by transfinite induction on the ordinal α . As the inductive hypothesis, we shall assume that the lemma holds for all $\delta < \alpha$. Since **while** ($\mathbb{B}$) $\{\mathbb{C}\}$ has two possible reductions, both with transition **id**, to show (7.10), it is sufficient to establish:

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{P_p(\alpha) \wedge \mathbb{B}\} \mathbb{C} \text{ while } (\mathbb{B}) \{\mathbb{C}\} \{ \exists \beta. P_p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha \} \quad (7.11)$$

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{P_p(\alpha) \wedge \neg \mathbb{B}\} \text{ skip}; \{ \exists \beta. P_p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha \} \quad (7.12)$$

This is sufficient since the first condition of the semantic judgement is the only one that may apply. For the reduction $(\sigma, \text{while } (\mathbb{B}) \{\mathbb{C}\}) \xrightarrow{\text{id}} (\sigma, \mathbb{C} \text{ while } (\mathbb{B}) \{\mathbb{C}\})$ (which requires $\mathbb{B}(\sigma)$), take $P'_p = P_p(\alpha) \wedge \mathbb{B}$ and $P''_p = \text{False}$. The first and third sub-conditions become trivial, while the second reduces to (7.11). For the reduction $(\sigma, \text{while } (\mathbb{B}) \{\mathbb{C}\}) \xrightarrow{\text{id}} (\sigma, \text{skip};)$ (which requires $\neg \mathbb{B}(\sigma)$), take $P'_p = P_p(\alpha) \wedge \neg \mathbb{B}$ and $P''_p = \text{False}$. Similarly, the first and third sub-conditions are trivial and the second reduces to (7.12).

To establish (7.11), we have from (7.9), for $\gamma = \alpha$ and renaming β to δ :

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{P_p(\alpha) \wedge \mathbb{B}\} \mathbb{C} \{ \exists \delta. P_p(\delta) \wedge \delta < \alpha \}.$$

Next, for all $\delta < \alpha$, we have, by the inductive hypothesis¹, that

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{P_p(\delta)\} \text{ while } (\mathbb{B}) \{\mathbb{C}\} \{ \exists \beta. P_p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \delta \},$$

and hence

$$\eta; \lambda; \mathcal{A} \models_{\tau} \{ \exists \delta. P_p(\delta) \wedge \delta < \alpha \} \text{ while } (\mathbb{B}) \{\mathbb{C}\} \{ \exists \beta. P_p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha \}.$$

Now (7.11) follows from the above by the sequencing lemma.

On the other hand, it is trivial to establish (7.12) by choosing α as the witness for β . $\square$

7.5 Non-blocking Properties

Non-blocking properties are used to characterise concurrent algorithms that guarantee progress. A *lock-free* algorithm guarantees global progress: an individual thread might fail to make progress, but only because some other thread does make progress. A *wait-free* algorithm guarantees local progress: every thread makes progress when it is scheduled. We consider how non-blocking properties can be formalised using Total-TaDA.

7.5.1 Lock-freedom

We have described lock-freedom in terms of an informal notion of “progress”. In order to properly characterise modules as lock-free, we need a more formal definition. We can characterise global progress for a module as follows: at any time, eventually either a pending operation will be completed or another operation will be begun. If we assume that the number of threads is bounded, then as long as there are pending module operations, some operation will eventually complete. If the number of threads

¹Since we have that $\forall \gamma \leq \alpha. \eta; \lambda; \mathcal{A} \models_{\tau} \{P_p(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{ \exists \beta. P_p(\beta) \wedge \beta < \gamma \}$, we certainly have that $\forall \gamma \leq \delta. \lambda; \mathcal{A} \models_{\tau} \{P_p(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{ \exists \beta. P_p(\beta) \wedge \beta < \gamma \}$, since $\delta < \alpha$. This allows us to obtain the consequent of the inductive hypothesis.

is unbounded, then there is no guarantee that any operation will complete, even if it is scheduled arbitrarily often, since additional operations can always begin.

Based on this observation, Gotsman *et al.* [23] reduced lock-freedom to the termination of a simple class of programs, the bounded most-general clients (BMGCs) of a module. Hoffmann *et al.* [30] generalised the result to apply to algorithms where the identity or the number of threads is significant. An (m, n) -bounded general client consists of m threads which each invoke n module operations in sequence. If all such bounded general clients (for all n and m)² terminate, then the module is lock-free.

Definition 7.2 ((m, n) -bounded General Client). Consider a module $\mathcal{M}$ with initialiser `init` and a set of operations O . We define the following sets of programs:

$$T_n = \{\text{op}_1; \dots; \text{op}_n \mid \text{op}_i \in O\}$$

$$C_{m,n} = \{\text{init}; (t_1 \parallel \dots \parallel t_m) \mid t_i \in T_n\}.$$

Theorem 7.2 (Hoffmann *et al.* [30]). Given a module $\mathcal{M}$, if, for all m and n , every program $c \in C_{m,n}$ terminates, then $\mathcal{M}$ is lock free.

Using this theorem, we define a specification pattern for Total-TaDA that guarantees lock-freedom and follows from the typical specifications we establish for lock-free modules.

Theorem 7.3 (Lock-freedom). Given a module $\mathcal{M}$ and some abstract predicate $\mathbf{M}$ (with two abstract parameters and an ordinal parameter), suppose that the following specifications are provable:

$$\forall \alpha. \vdash_\tau \{\text{True}\} \text{ init } \{\exists s, u. \mathbf{M}(s, u, \alpha)\}$$

$$\forall \text{op} \in O. \forall \beta. \vdash_\tau \mathbb{W}u, \alpha. \langle \mathbf{M}(s, u, \alpha) \wedge \alpha > \beta(\alpha) \rangle \text{ op } \langle \exists u'. \mathbf{M}(s, u', \beta(\alpha)) \rangle.$$

Then, $\mathcal{M}$ is lock-free.

Proof. By Theorem 7.2, it is sufficient to show that, for arbitrary m, n and $c \in C_{m,n}$, the program c terminates. Fix the number of threads m .

We define a region type $\mathbf{M}$ whose abstract states consist of vectors $\bar{x} \in \mathbb{N}^m$. We denote by x_i , for $1 \leq i \leq m$, the i -th component of vector $\bar{x}$. We denote by $\sum \bar{x}$ the sum $\sum_{i=1}^m x_i$. Region states are interpreted as follows:

$$I(\mathbf{M}_a(s, \bar{x})) \triangleq \exists u. \mathbf{M}(s, u, \sum \bar{x}).$$

The guard algebra for $\mathbf{M}$ consists of m distinct guards $G_1, \dots, G_m$. Formally, the guards are subsets of $\{G_i \mid 1 \leq i \leq m\}$ under disjoint union. The state transition system for $\mathbf{M}$ allows a thread holding guard G_i to decrease the i -th component of the abstract state:

$$G_i : (\forall j \neq i. x_j = y_j) \wedge x_i > y_i \wedge \bar{x} \rightsquigarrow \bar{y}.$$

²The bounded *most-general* client may be seen as the program which non-deterministically chooses among all bounded general clients.

For $1 \leq i \leq m$, arbitrary n , and $\text{op} \in O$, we have

$$\begin{array}{c}
\left\{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_i]_a * x_i = n + 1 \right\} \\
\left| \begin{array}{c} \text{EXISTS} \\ \left| \begin{array}{c} \text{USEATOMIC} \\ \left| \begin{array}{c} \text{AEXISTS} \\ \left| \begin{array}{c} \langle \exists u. \mathbf{M}(s, u, k + n + 1) \rangle \\ \forall u, k. \\ \langle \mathbf{M}(s, u, k + n + 1) \rangle \\ \text{op} \\ \langle \exists u'. \mathbf{M}(s, u', k + n) \rangle \\ \langle \exists u'. \mathbf{M}(s, u', k + n) * [\mathbf{G}_i]_a * x_i = n \rangle \end{array} \end{array} \end{array} \end{array} \right. \\
\left\{ \exists \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_i]_a * x_i = n \right\} \\
\left\{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_i]_a * x_i = n \right\}
\end{array}$$

Applying this specification repeatedly (by induction), we have for arbitrary $t \in T_n$

$$\vdash_{\tau} \left\{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_i]_a * x_i = n \right\} t \left\{ \text{True} \right\}$$

Let $c = \text{init}; (t_1 \parallel \dots \parallel t_m) \in C_{m,n}$ be arbitrary. We derive $\vdash_{\tau} \left\{ \text{True} \right\} c \left\{ \text{True} \right\}$ by choosing $n \cdot m$ as the initial ordinal and creating an $\mathbf{M}$ -region with initial state $(n, \dots, n)$ as follows:

$$\begin{array}{c}
\left\{ \text{True} \right\} \\
\text{init;} \\
\left\{ \exists s, u. \mathbf{M}(s, u, n \cdot m) \right\} \\
// \text{ create region} \\
\left\{ \exists a, s. \mathbf{M}_a(s, (n, \dots, n)) * [\mathbf{G}_1]_a * \dots * [\mathbf{G}_m]_a \right\} \\
\left\{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_1]_a * x_1 = n \right\} t_1 \left\{ \text{True} \right\} \parallel \dots \parallel \left\{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_m]_a * x_m = n \right\} t_m \left\{ \text{True} \right\} \\
\left\{ \text{True} \right\}
\end{array}$$

Consequently, c terminates, as required. $\square$

It is straightforward to apply Theorem 7.3 to the modules we have considered.

7.5.2 Wait-freedom

Whereas lock-freedom only requires that *some* thread makes progress, wait-freedom requires that *every* thread makes progress (provided that it is not permanently descheduled). In terms of operations, this requires that each operation of a module should complete within a finite number of steps. Since Total-TaDA specifications guarantee that operations terminate, it is simple to describe a specification that implies that a module is wait-free.

Theorem 7.4 (Wait-freedom). Given a module $\mathcal{M}$ and some abstract predicate $\mathbf{M}$ (with two abstract

parameters), suppose that the following specifications are provable:

$$\begin{aligned} & \vdash_{\tau} \{ \text{True} \} \text{ init } \{ \exists s, t. \mathbf{M}(s, u) \} \\ & \forall \text{op} \in O. \vdash_{\tau} \mathbb{W}u. \langle \mathbf{M}(s, u) \rangle \text{ op } \langle \exists u'. \mathbf{M}(s, u') \rangle. \end{aligned}$$

Then $\mathcal{M}$ is wait-free.

Proof. The specifications imply that $\mathbf{M}$ is an invariant which is established by the initialiser and preserved at all times by the module operations. Furthermore, all of the module operations terminate, assuming the environment maintains $\mathbf{M}$ invariant. Consequently, all of the module operations terminate in the context of an environment calling module operations: the module is wait-free. $\square$

Lock-freedom can only be applied to a module as a whole, since it relates to global progress. Wait-freedom, by contrast, relates to local progress — that the operations of *each* thread terminate — and so it is meaningful to consider an individual operation to be wait-free in a context where other operations may be lock-free or even blocking. By combining (partial-correctness) TaDA and Total-TaDA specifications (indicated by $\vdash$ and $\vdash_{\tau}$ respectively), we can give a specification pattern that guarantees wait-freedom for a specific module operation.

Theorem 7.5 (Wait-free Operation). Given a module $\mathcal{M}$ and some abstract predicate $\mathbf{M}$ (with two abstract parameters), suppose that the following specifications are provable:

$$\begin{aligned} & \vdash \{ \text{True} \} \text{ init } \{ \exists s, u. \mathbf{M}(s, u) \} \\ & \vdash_{\tau} \mathbb{W}u. \langle \mathbf{M}(s, u) \rangle \text{ op } \langle \exists u'. \mathbf{M}(s, u') \rangle \\ & \forall \text{op}' \in O. \vdash \mathbb{W}u. \langle \mathbf{M}(s, u) \rangle \text{ op}' \langle \exists u'. \mathbf{M}(s, u') \rangle \end{aligned}$$

Then op is wait-free.

Proof. As before, $\mathbf{M}$ is a module invariant; op is guaranteed to terminate with this invariant, therefore it is wait-free. $\square$


The specifications required by Theorem 7.5 do not follow from those given for our examples. However, where applicable, the proofs can easily be adapted. For instance, to show that the **read** operation of the counter is wait-free, we would remove the ordinals from the region definition, and abstract the value of the counter. This breaks the termination proof for the increment operations, but we can adapt it to a partial-correctness proof in TaDA. The termination proof for **read** does not depend on the ordinal parameter of the region, and so we can still establish total correctness, as required.

7.5.3 Non-impedance

Recall the counter specification from §7.1.1. If we abstract the value and address of the counter (which are irrelevant to termination), the specification becomes:

$$\begin{aligned} \forall \alpha. \vdash_{\tau} \{ \text{True} \} \text{ x} := \text{makeCounter}(); \{ \exists s \in \mathbb{T}_1, u \in \mathbb{T}_2. \text{Counter}(s, u, \alpha) \} \\ \vdash_{\tau} \mathbb{W}u, \alpha. \langle \text{Counter}(s, u, \alpha) \rangle \text{ read}(\text{x}) \langle \text{Counter}(s, u, \alpha) \rangle \\ \forall \beta. \vdash_{\tau} \mathbb{W}u, \alpha. \langle \text{Counter}(s, u, \alpha) * \alpha > \beta(\alpha) \rangle \text{ incr}(\text{x}) \langle \exists u'. \text{Counter}(s, u', \beta(\alpha)) \rangle \end{aligned}$$

As the **read** operation does not change the ordinal, it implies that both the **read** and **incr** operations will terminate in a concurrent environment that performs an unbounded number of **reads**. This suggests an alternative approach to characterising lock-free modules in terms of which operations *impede* each other—that is, which operations may prevent the termination of an operation if infinitely many of them are invoked during a (fair) execution of the operation. Our specification implies that **read** does not impede either **read** or **incr**. This is expressed by edges 1 and 2 in the following non-impedance graph:



Note that the above specifications for the counter do not by themselves imply that **incr** does not impede **read** (edge 3). This can be demonstrated by considering an alternative implementation of **read**, that satisfies the specification but is not wait-free:

```

do {
  v := [x];
  w := [x];
} while (v ≠ w);
return v;
  
```

Recall that we can prove that **read** is wait-free by giving a different specification as in Theorem 7.5. An operation is wait-free precisely when every operation does not impede it. For **read**, this is expressed by edges 1 and 3 in the above graph.

The stack specification in Figure 7.7, much like the counter specification, implies that **pop** does not impede either **push** or **pop**:



The **pop** operation, however, may be impeded by **push**.

The non-impedance relationships implied by the stack specification are important for clients. For instance, consider a producer-consumer scenario in which the stack is used to communicate data from producers to consumers. When no data is available, consumers may simply loop attempting to pop the stack. If the **pop** operation could impede **push**, then producers might be starved by consumers. In this situation, we could not guarantee that the system would make progress. This suggests that non-impedance, which is captured by Total-TaDA specifications, can be an important property of non-blocking algorithms.

7.5.4 Related Work

Hoffmann *et al.* [30] introduced a concurrent separation logic for verifying total correctness. By adapting the most-general-client approach of Gotsman *et al.* [23], they establish that modules are lock-free. They do not, however, establish functional correctness. This method involves a thread passing “tokens” to other threads whose lock-free operations are impeded by modifications to the shared state. Subsequent approaches [3, 40] also use some form of tokens that are used up in loops or function calls. These approaches require special proof rules for the tokens. When these approaches restrict to dealing with finite numbers of tokens, support for unbounded non-determinism (as in the backoff increment example of Figure 7.6) is limited. In our extension of TaDA to reason about termination, named Total-TaDA, such token passing is not necessary. Instead, we require the client to provide a general (ordinal) limit on the amount of impeding interference. Consequently, we can guarantee the termination of loops with standard proof rules.

Jia *et al.* have introduced loop depth counters to prove lock-freedom [33]. Their approach adds auxiliary code to the loops that counts the number of times progress towards an operation’s completion is made. Their approach was applied to stack and queue modules using a tool.

Liang *et al.* [40] have developed a proof theory for termination-preserving refinement, applying it to verify linearisability and lock-freedom. Their approach constrains impedance by requiring that impeding actions correspond to progress at the abstract level. In Total-TaDA, such constraints are made by requiring that impeding actions decrease an ordinal associated with a shared region. Their approach does not freely combine lock-free and wait-free specifications, whereas with Total-TaDA we can reason about lock- and wait-freedom in combination, as well as about more subtle conditions, such as non-impedance. For example, we can show when a read operation of a lock-free data-structure is wait-free. Their specifications establish termination-preserving refinement (given a context, if the abstract program is guaranteed to terminate, then so is the concrete), whereas Total-TaDA specifications establish termination (in a context, the program will terminate).

Boström and Müller [3] have introduced an approach that can verify termination and progress properties of concurrent programs. The approach supports blocking concurrency and non-terminating programs, which Total-TaDA does not. However, the approach does not aim at racy concurrent programs and cannot deal with any of the examples shown in this chapter. Furthermore, the relationship between termination and lock- and wait-freedom is not considered.

Of the above approaches, none covers total functional correctness for fine-grained concurrent programs. With Total-TaDA, we can reason about clients that use modules without knowing their implementation details. Moreover, with Total-TaDA it is easy to verify module operations independently, with respect to a common abstraction, rather than considering a whole module at once. Finally, our approach to specification is unique in supporting lock- and wait-freedom simultaneously, as well as expressing more subtle conditions, such as non-impedance.

8 Extending the logic

We have formally verified that implementations of concurrent data-structures with abstractly atomic operations. However, implementations can use techniques that make verification difficult. Among the most challenging of these are *helping* and *speculation* [1].

Typically, we would expect that an implementation of an abstractly atomic operation performs some concrete atomic update during its execution that effects the abstract operation. However, in general it is possible for this update to be performed by a different thread, which is in the process of performing another operation. We refer to this scenario as *helping*: one thread is performing an abstract operation on behalf of another.

Helping is particularly common in *wait-free* algorithms, where each operation is guaranteed to terminate. Suppose that one operation has partially completed when another begins. The second operation cannot be performed while the first is partially complete, so it must either wait for the first operation, undo the progress of the first operation, or help complete the first operation. The first two options do not guarantee termination, so helping becomes the only viable option.

We might also expect that, at the point the abstractly atomic operation of a thread is performed, we would know this to be the case. However, in general we may only know that it might be the case, and future behaviour will determine whether it was or not. We refer to this situation as *speculation*: we consider speculatively the cases where the abstract operation was and was not performed, and decide which applies later in the execution.

As an example of speculation, consider a data-structure that implements an abstractly atomic read operation as follows: it performs two concretely atomic reads and then randomly decides which result to return. The abstractly atomic read happens at one of the two concretely atomic reads, but exactly which is determined later. This example is illustrative but is, however, contrived. In practice, speculation does not arise as a result of random decisions of a single thread but from the non-deterministic behaviour of the concurrent environment.

We show how to verify abstractly atomic operations of concurrent data-structures that use helping and speculation by extending TaDA. To prove an atomic specification in TaDA, it is necessary to show that a single atomic step in the implementation performs the corresponding abstract atomic action. This is achieved by a special resource that allows the thread to perform the action, and witnesses when it is completed. However, in TaDA this resource cannot be transferred to other threads, and so it does not support helping. Our extension, TaDA 2, supports helping by allowing such resources (called *proxies*, which embody permission to perform the abstract action, and *witnesses*, which guarantee that the action has been performed) to be transferred between threads.

This may seem like a simple change, but it requires significant changes to the semantics and soundness proof of the logic. In particular, the TaDA semantics of atomic specifications ties the abstractly atomic update to a concrete atomic operation in the implementation, thus precluding helping.

We support speculation in TaDA 2 by introducing *speculative resources*. These speculative resources

consist of a set of possible (speculative) configurations of proxy and witness resources. When we speculatively perform the action of a proxy, we introduce cases in the set corresponding to whether or not the action is performed. We can resolve the speculation by choosing a particular configuration from the set.

8.1 Motivating Examples

We motivate our work with two examples that implement abstract atomicity using helping and speculation. Our examples build on an abstract counter module described in §3.1. We start by revising the ticket lock module (§8.1.1), but now we are going to specify it with an atomic specification. The ticket lock illustrates helping: when a thread releases the lock, it acquires the lock on behalf of the next thread waiting for the lock (if any). The second is a counter algorithm that internally increments in two steps (§8.1.2). The two-step counter illustrates speculation: when multiple threads have partially completed their increment operations, it will appear that half of them have abstractly incremented the counter; however, which ‘half’ is determined by the order in which they complete the operation.

8.1.1 Ticket Lock

We consider the ticket lock module from §3.2. The lock uses two counters, **next** and **owner**, which initially have value 0. The counter **next** is used to generate the tickets and the counter **owner** is used to indicate which ticket holds the lock. When their values match the lock is considered to be unlocked.

The **acquire** operation starts by incrementing the counter **next**; the previous value of the counter represents the ticket held by the thread. It then waits until the value in the counter **owner** matches its ticket. When they match, the thread has acquired the lock.

The **release** operation simply increments the value of the counter **owner**. Since only one thread can hold a lock at a time, the **owner** counter is not subject to concurrent increments, and so the **wkIncr** operation is appropriate.

Atomic Specification

We introduce an abstract predicate $\text{TLock}(s, x, l, t)$ to represent the lock as a resource. The second parameter, $x \in \text{Addr}$, is the address of the lock object. The third parameter, $l \in \{0, 1\}$, indicates the state of the lock: 0 for unlocked, 1 for locked. The two parameters, $s \in \mathbb{T}_2$ and $t \in \mathbb{T}_3$, range over *abstract types*. The types $\mathbb{T}_2$ and $\mathbb{T}_3$ will capture implementation-specific information about the configuration of the lock. The ticket lock has the following specification:

$$\begin{aligned} &\vdash \{ \text{True} \} \text{makeLock}() \{ \exists s \in \mathbb{T}_2, t \in \mathbb{T}_3. \text{TLock}(s, \text{ret}, 0, t) \} \\ &\vdash \mathbb{W}(l, t) \in \{0, 1\} \times \mathbb{T}_3. \langle \text{TLock}(s, x, l, t) \rangle \text{acquire}(x) \langle \text{TLock}(s, x, 1, t) * l = 0 \rangle \\ &\vdash \langle \text{TLock}(s, x, 1, t) \rangle \text{release}(x) \langle \exists t' \in \mathbb{T}_3. \text{TLock}(s, x, 0, t') \rangle \end{aligned}$$

Since the **release** operation makes use of **wkIncr**, it is important this operation does not race with any other increment of the **owner** counter. Without the t parameter, we could imagine specifying an operation **nop** that increments both the **owner** and **next** counters in a single atomic step. As this would

not conceptually change whether the lock is in the locked or unlocked state, we might specify it as:

$$\vdash \forall l \in \{0, 1\}. \langle \mathbf{TLock}(s, \mathbf{x}, l) \rangle \text{nop}(\mathbf{x}) \langle \mathbf{TLock}(s, \mathbf{x}, l) \rangle$$

However, such an operation causes havoc: a thread waiting to acquire the lock may proceed, even though the thread that actually holds the lock has not released it! The implementation ties the additional parameter t to the value of the **owner** counter. The specification for the **release** operation consequently prevents the environment from concurrently modifying this counter (as it is not bound by $\forall$).

Implementation

In order to verify the specification, we must provide an interpretation for the abstract predicate $\mathbf{TLock}(s, x, l, t)$. For this, we introduce shared regions. As in TaDA, a shared region encapsulates some resources that may be shared by multiple threads, with the proviso that they can only be accessed by atomic operations. A shared region has an abstract state (which has a concrete interpretation), and a protocol that determines how the state can change. A shared region is also equipped with resources called guards, which determine how threads can update the region's shared state.

For the ticket lock, we define a region type **TLock**. The abstract states of **TLock** regions consist of pairs $(l, n) \in \{0, 1\} \times \mathbb{N}$, where l indicates whether the lock is currently locked and n indicates the current value of the **owner** counter. The protocol for the **TLock** region is specified by the following transition system, labelled by guards:

$$\begin{aligned} G : \forall n \in \mathbb{N}. (1, n) &\rightsquigarrow (0, n + 1) \\ G : \forall n \in \mathbb{N}. (0, n) &\rightsquigarrow (1, n) \end{aligned}$$

The first of these transitions allows a thread holding the (unique) guard resource G for the region to release the lock, increasing the **owner** counter in doing so. The second transition allows a thread holding G to acquire the lock, without updating the **owner** counter.

In addition to the G guard, we define two other types of guards associated with **TLock** regions. These guards do not play a role in the protocol, but are used as ghost resources to manage bookkeeping. The first of these is the $\text{TICKET}(v, i)$ guard, where $v \in \mathbb{N}$ is a number and $i \in \text{TId}$ is a thread identifier. This guard represents thread i 's claim to hold the ticket number v , while waiting to acquire the lock. The final guard $\text{QUEUE}(S)$ accounts for the tickets that have been issued, where $S \subseteq \mathbb{N} \times \text{TId}$. These guards are defined to obey the following equation:

$$\text{QUEUE}(S) = \text{QUEUE}(S \uplus \{(v, i)\}) \bullet \text{TICKET}(v, i) \text{ if } v \notin \text{proj}_1(S), i \notin \text{proj}_2(S)$$

As in the previous chapters, $\bullet$ is the composition operator on guard resources, which lifts to $*$ in assertions. It is associative, with $\mathbf{0}$ as the neutral element. The following compositions are not allowed:

$$\begin{aligned} &G \bullet G \\ &\text{QUEUE}(S) \bullet \text{QUEUE}(S') \\ &\text{TICKET}(v, i) \bullet \text{TICKET}(v', i') \quad \text{if } v = v' \text{ or } i = i' \\ &\text{QUEUE}(S) \bullet \text{TICKET}(v, i) \quad \text{if } (v, i) \notin S \end{aligned}$$

It remains to define the state interpretation for **TLock** regions. However, since the ticket lock algorithm uses helping, we must first describe how atomic specifications are proved in TaDA 2. There are two proof rules that are key to atomicity proofs. The first is the **MAKEATOMIC** rule, which allows us to prove that an operation can be seen as abstractly atomic. A slightly simplified version of this rule is as follows:

$$\frac{\{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \quad a:i \vdash \left\{ \exists x \in X. \mathbf{t}_a(\bar{z}, x) * \langle \mathbb{W}x \in X \rightsquigarrow Q(x) \rangle_a^i \right\} \mathbb{C} \left\{ \exists x \in X, y \in Q(x). |x \rightsquigarrow y\rangle_a^i \right\}}{\vdash \mathbb{W}x \in X. \langle \mathbf{t}_a(\bar{z}, x) * [\mathbf{G}]_a \rangle \mathbb{C} \langle \exists y \in Q(x). \mathbf{t}_a(\bar{z}, y) * [\mathbf{G}]_a \rangle}$$

The conclusion of the rule specifies that $\mathbb{C}$ performs an atomic update, transforming the state of the region a (of type $\mathbf{t}$ and with parameters $\bar{z}$) from $x \in X$ to $y \in Q(x)$. The environment may change the state of the region before the atomic update occurs, so long as the state remains in the set X . Any update to a shared region has to be justified by a guard for that region; in the conclusion of the rule, this is the guard resource $[\mathbf{G}]_a$. The first premiss of the rule requires that the update from x to y is permitted by the transition system $\mathcal{T}_{\mathbf{t}}$ for region type $\mathbf{t}$ given guard $\mathbf{G}$; the superscript $*$ indicates, again, the reflexive-transitive closure.

While the first premiss ensures that the update is permitted, the second premiss ensures that $\mathbb{C}$ actually performs the atomic update. This is achieved by providing a *proxy* resource $\langle \mathbb{W}x \in X \rightsquigarrow Q(x) \rangle_a^i$ in the precondition, which $\mathbb{C}$ must update to a *witness* resource $|x \rightsquigarrow y\rangle_a^i$, featured in the postcondition. The proxy resource is conceptually a proxy for $[\mathbf{G}]_a$, in that it permits the update to the region by deriving its authority from the guard resource (which is not present in the premiss). The proxy records the atomic update to be performed, the region a in which it is to be performed, and has a unique identifier i : no two proxies or witnesses with the same identifier can coexist. Since the premiss is not an atomic triple, $\mathbb{C}$ may take multiple steps. However, the proxy resource can only be used once, and in doing so is updated to a witness resource that records the specific update that took place. This guarantees that $\mathbb{C}$ effectively performs an atomic update.

The second key proof rule is the **UPDATEREGION** rule, which deals with using a proxy to update a region. A simplified version of this rule is as follows:

$$\frac{a:i \vdash \mathbb{W}x \in X. \langle I(\mathbf{t}_a(\bar{z}, x)) * p(x) \rangle \mathbb{C} \left\langle \exists y \in Q(x). \overset{a}{\underset{x \rightsquigarrow y}{\Diamond}} (I(\mathbf{t}_a(\bar{z}, y)) * q(x, y)) \right\rangle}{a:i \vdash \mathbb{W}x \in X. \langle \mathbf{t}_a(\bar{z}, x) * p(x) \rangle \mathbb{C} \langle \exists y \in Q(x). \mathbf{t}_a(\bar{z}, y) * q(x, y) \rangle}$$

In this rule, the region $\mathbf{t}_a(\bar{z}, x)$ in the conclusion is replaced by its interpretation $I(\mathbf{t}_a(\bar{z}, x))$ in the premiss. This allows the resources from the region to be accessed for the duration of an atomic update. To account for the change to the abstract state of the region (from x to y), it is necessary to update a proxy to a witness corresponding to the update. This is achieved by the $\overset{a}{\underset{x \rightsquigarrow y}{\Diamond}}$ modality, which appears in the postcondition of the premiss. Intuitively, the assertion $\overset{a}{\underset{x \rightsquigarrow y}{\Diamond}} P$ should be read as “it is possible to obtain P by updating proxy resources to witnesses, corresponding to the update from x to y ”. The **UPDATEREGION** rule subsumes the rule with the same name from TaDA, but also the **OPENREGION** rule, as **OPENREGION** is now just a special case of the new rule.

$$P \Longrightarrow \underset{x \rightsquigarrow x}{\overset{a}{\Diamond}} P \quad (8.1)$$

$$\left(\underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} P \right) * R \Longrightarrow \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} (P * R) \quad (8.2)$$

$$\forall x \in X, y \in Q(x). \langle \forall z \in X \rightsquigarrow Q(z) \rangle_a^i \Longrightarrow \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} |x \rightsquigarrow y\rangle_a^i \quad (8.3)$$

$$\underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} \underset{y \rightsquigarrow z}{\overset{a}{\Diamond}} P \Longrightarrow \underset{x \rightsquigarrow z}{\overset{a}{\Diamond}} P \quad (8.4)$$

$$\underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} (P \vee Q) \Longleftrightarrow \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} P \vee \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} Q \quad (8.5)$$

$$\exists z. \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} P(z) \Longleftrightarrow \underset{x \rightsquigarrow y}{\overset{a}{\Diamond}} (\exists z. P(z)) \quad (8.6)$$

Figure 8.1: Axioms for proxy modalities.

Figure 8.1 gives axioms for working with the proxy modality. Axiom (8.1) means that we do not need to update a proxy to perform the identity transition. Axiom (8.2) means that we can move a frame across the proxy modality. Axiom (8.3) allows us to introduce a proxy modality by updating a suitable proxy to a witness corresponding to the required update; this is the basic rule for using proxies. Axiom (8.4) allows us to chain proxies together transitively: if we have a proxy that allows a transition from x to y and another that allows a transition from y to z , then we can use both together to transition from x to z . Finally, axioms (8.5) and (8.6) mean that the proxy modality distributes over disjunction and existential quantification.

The proxies and witnesses of TaDA 2 play a similar role to the atomic tracking resources of TaDA. The most important difference here, however, is that they can be transferred between threads. This allows one thread to perform the atomic action on behalf of another, by using its proxy.

In the ticket lock, such helping occurs in the **release** operation. Conceptually, **release** releases the lock by incrementing the **owner** counter. However, if another thread is waiting to acquire the lock, and has incremented the **next** counter, the lock does not in fact enter the unlocked state. We resolve this incongruity by viewing the atomic increment to the **owner** counter as performing both the unlock action on behalf of its own thread and the lock action on behalf of the waiting thread. That is, the unlocking thread *helps* the locking thread to acquire the lock.

The shared state of the **TLock** region, therefore, holds proxies and witnesses for threads that are waiting to acquire the lock. We define an auxiliary predicate T to represent the shared resources belonging to the **TLock** region in Figure 8.2.

Lines (8.7) and (8.8) state that the **next** and **owner** counters are at addresses y and z , and have values nxt and own . Line (8.9) gives the guard resource $[\text{QUEUE}(S)]_r$, tracking tickets currently held by threads waiting to acquire the lock. A thread that increments the **next** counter from k to $k+1$ will acquire a ticket resource $[\text{TICKET}(k, i)]_r$; by doing so, the pair (k, i) is added to S . In the process, a thread gives up its proxy resource (with identifier i) so that another thread can lock on its behalf. The condition on S asserts that there is a ticket for every thread that incremented the **next** counter but has yet to acquire the lock, and that all tickets are below nxt (the next available ticket). Line (8.10) states that there are witnesses for each ticket that successfully acquired the lock, but has not yet reclaimed its witness. Line (8.11) states that there are proxies for each ticket waiting to acquire the lock.

$$T(a, x, s, t, y, z, \text{next}, \text{own}, S) \triangleq x.\text{next} \mapsto y * \text{Counter}(s, y, \text{next}) \quad (8.7)$$

$$* x.\text{owner} \mapsto z * \text{Counter}(t, z, \text{own}) \quad (8.8)$$

$$* [\text{QUEUE}(S)]_a * \{j \mid \text{own} < j < \text{next}\} \subseteq \text{proj}_1(S) \subseteq \{j \mid j < \text{next}\} \quad (8.9)$$

$$* \bigotimes_{(v,i) \in S} \left(v \leq \text{own} \implies |(0, v) \rightsquigarrow (1, v)|_a^i \right) \quad (8.10)$$

$$* \bigotimes_{(v,i) \in S} \left(v > \text{own} \implies \langle \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N} \rightsquigarrow (1, n) \wedge b = 0 \rangle_a^i \right) \quad (8.11)$$

Figure 8.2: Auxiliary predicate describing the shared resources of the ticket lock.

We define the interpretation of abstract states for the ticket lock region:

$$I(\mathbf{TLock}_a(x, s, t, y, z, (l, n))) \triangleq \exists m, S. T(a, s, t, x, y, z, m, n, S) * m \geq n * (l = 1 \iff m > n)$$

We can now give the interpretation of the abstract types and abstract predicate as follows:

$$\mathbb{T}_2 \triangleq \text{Rld} \times \mathbb{T}_1 \times \mathbb{T}_1 \times \text{Addr} \times \text{Addr}$$

$$\mathbb{T}_3 \triangleq \mathbb{N}$$

$$\mathbf{TLock}((a, s', t', y, z), x, l, n) \triangleq \mathbf{TLock}_a(x, s', t', y, z, (l, n)) * [\mathbf{G}]_r$$

where **Rld** is the set of region identifiers and **Addr** is the set of addresses.

It remains to prove that the implementations of the operations satisfy the specifications, given this interpretation. The proof for **acquire** is given in Figure 8.3, and the proof for **release** in Figure 8.4.

8.1.2 Two-step Counter

Speculation can be observed in many fine-grained implementations, such as the Michael-Scott queue [43] or Heller *et al.* lazy set [24]. However, we do not need to look at complex implementations to find speculation. We propose a simple counter module implemented using the counter previously presented:

<pre>function readVal(x) { v := read(x); return v / 2; }</pre>	<pre>function incrVal(x) { incr(x); incr(x); }</pre>
--	--

Atomic Specification

The **readVal** operation consists of reading the current value of the counter and returns the floor of its value divided by 2. The **incrVal** operation performs two increments, where each increment is atomic. This means that in between each increment, other threads can perform reads or other increments. We

$$\begin{array}{c}
\mathbb{W}(l, t) \in \{0, 1\} \times \mathbb{T}_3. \\
\langle \text{TLock}(s, \mathbf{x}, l, t) \rangle \\
\begin{array}{|l}
\mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N}. \\
\langle \text{TLock}_a(\mathbf{x}, s', t', y, z, (l, n)) * [G]_a \rangle \\
a : i \vdash \\
\left\{ \exists l \in \{0, 1\}, n \in \mathbb{N}. \text{TLock}_a(\mathbf{x}, s', t', y, z, (l, n)) * \langle \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N} \rightsquigarrow (1, n) \wedge b = 0 \rangle_a^i \right\} \\
\text{next} := [\mathbf{x}.\text{next}]; \\
\text{owner} := [\mathbf{x}.\text{owner}]; \\
\left\{ \exists l \in \{0, 1\}, n \in \mathbb{N}. \text{TLock}_a(\mathbf{x}, s', t', \text{next}, \text{owner}, (l, n)) \right\} \\
* \langle \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N} \rightsquigarrow (1, n) \wedge b = 0 \rangle_a^i \\
\text{UPDATEREGION} \quad \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N}. \\
\left\{ \exists m, S. T(a, \mathbf{x}, s', t', \text{next}, \text{owner}, m, n, S) * \langle \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N} \rightsquigarrow (1, n) \wedge b = 0 \rangle_a^i \right\} \\
* m \geq n * (l = 1 \iff m > n) \\
\mathbf{t} := \text{incr}(\text{next}); \\
\left\langle \begin{array}{c} a \\ \diamond \\ (l, n) \rightsquigarrow (1, n) \end{array} \left(\exists m, S. T(a, \mathbf{x}, s', t', \text{next}, \text{owner}, m + 1, n, S \uplus \{(m, i)\}) \right) \right\rangle \\
* [\text{TICKET}(m, i)]_a * m + 1 > n * \mathbf{t} = m \\
\left\{ \exists l \in \{0, 1\}, n \in \mathbb{N}. \text{TLock}_a(\mathbf{x}, s', t', \text{next}, \text{owner}, (l, n)) * [\text{TICKET}(\mathbf{t}, i)]_a \right\} \\
\text{do} \{ \\
\left\{ \exists l \in \{0, 1\}, n \in \mathbb{N}. \text{TLock}_a(\mathbf{x}, s', t', \text{next}, \text{owner}, (l, n)) * [\text{TICKET}(\mathbf{t}, i)]_a \right\} \\
\text{UPDATEREGION} \quad \mathbb{W}(l, n) \in \{0, 1\} \times \mathbb{N}. \\
\langle \exists m, S. T(a, \mathbf{x}, s', t', \text{next}, \text{owner}, m, n, S \uplus \{(\mathbf{t}, i)\}) * [\text{TICKET}(\mathbf{t}, i)]_a * m \geq n * m > \mathbf{t} \rangle \\
\mathbf{v} := \text{read}(\text{owner}); \\
\left\langle \begin{array}{c} a \\ \diamond \\ (l, n) \rightsquigarrow (l, n) \end{array} \left(\mathbf{v} = n * \underline{\text{if}} \ n = \mathbf{t} \ \underline{\text{then}} \ \left(\exists m, S. T(a, \mathbf{x}, s', t', \text{next}, \text{owner}, m, n, S) \right) \right. \right. \\
\left. \left. * |(0, \mathbf{t}) \rightsquigarrow (1, \mathbf{t})|_a^i * m > n \right) \right. \\
\left. \underline{\text{else}} \ \left(\exists m, S. T(a, \mathbf{x}, s', t', \text{next}, \text{owner}, m, n, S \uplus \{(\mathbf{t}, i)\}) \right) \right. \\
\left. * [\text{TICKET}(\mathbf{t}, i)]_a * m \geq n \right) \right\rangle \\
\left\{ \underline{\text{if}} \ \mathbf{v} = \mathbf{t} \ \underline{\text{then}} \ |(0, \mathbf{t}) \rightsquigarrow (1, \mathbf{t})|_a^i \right. \\
\left. \underline{\text{else}} \ \exists l \in \{0, 1\}, n \in \mathbb{N}. \text{TLock}_a(\mathbf{x}, s', t', \text{next}, \text{owner}, (l, n)) * [\text{TICKET}(\mathbf{t}, i)]_a \right\} \\
\} \text{ while } (\mathbf{v} \neq \mathbf{t}); \\
\left\{ |(0, \mathbf{t}) \rightsquigarrow (1, \mathbf{t})|_a^i \right\} \\
\langle \text{TLock}_a(\mathbf{x}, s', t', y, z, (1, n)) * [G]_a * b = 0 \rangle \\
\langle \text{TLock}(s, \mathbf{x}, 1, t) * l = 0 \rangle
\end{array}
\end{array}$$

Figure 8.3: Proof of correctness of the `acquire` operation.

specify the counter module in a similar way as the counter we use to implement it, as follows:

$$\begin{array}{l}
\vdash \{ \text{True} \} \text{makeCounter}() \{ \exists s \in \mathbb{T}_2. \text{TSCounter}(s, \text{ret}, 0) \} \\
\vdash \mathbb{W}n \in \mathbb{N}. \langle \text{TSCounter}(s, \mathbf{x}, n) \rangle \text{readVal}(\mathbf{x}) \langle \text{TSCounter}(s, \mathbf{x}, n) * \text{ret} = n \rangle \\
\vdash \mathbb{W}n \in \mathbb{N}. \langle \text{TSCounter}(s, \mathbf{x}, n) \rangle \text{incrVal}(\mathbf{x}) \langle \text{TSCounter}(s, \mathbf{x}, n + 1) \rangle
\end{array}$$

Implementation

We use the notation $\mathbb{S}_a(X)$, where a is a region identifier and X is a set of configurations (lists containing proxies and witnesses) to denote that we can speculatively choose between any of the configurations. When we create a region, we have $\mathbb{S}_a(\{\emptyset\})$ —only the empty configuration. If we have a proxy or witness separate from $\mathbb{S}_a(X)$, then we can merge it in by adding it to all of the configurations. We can go from $\mathbb{S}_a(X)$ to $\mathbb{S}_a(Y)$ if $Y \subseteq X$. If we can pull out a specific proxy or witness from every configuration in the set, then we can pull it out from the overall speculation.




Figure 8.4: Proof of correctness of the **release** operation.

$$C_a(X, 0, v) \stackrel{\text{def}}{=} \left\{ \bigotimes_{i \in X} \langle \forall n \in \mathbb{N}. n \rightsquigarrow n+1 \rangle_a^i \right\}$$

$$C_a(X, m+1, v) \stackrel{\text{def}}{=} \left\{ \langle v - m - 1 \rightsquigarrow v - m \rangle_a^i * x \mid i \in X, x \in C_a(X \setminus \{i\}, m, v) \right\}$$

Figure 8.5: Auxiliary predicate describing the set of configurations for the two-step counter.

To describe the set of configurations for the two-step counter, we define $C_a(X, m, v)$ in Figure 8.5 where X is a set of proxy/witness configurations, $0 \leq m \leq |X|$ is the number of witnesses to have and v is the current value of the counter.

We extend the proxy modality to handle speculative sets of proxy/witnesses with the axiom:

$$\frac{\forall p' \in P'. \exists p \in P. p \text{ R}_{x,y} p'}{\mathbb{S}_a(P) \Rightarrow \overset{a}{\underset{x \rightsquigarrow y}{\Diamond}} \mathbb{S}_a(P')}$$

The premiss requires that we justify the update for each element in the set. Essentially, if we update the region, we have enough proxies to justify such an update in all configurations. The relation R enforces this and is defined as follows:

$$\frac{}{p \text{ R}_{x,x} p} \tag{8.12}$$

$$\frac{p \text{ R}_{x,y} p' \quad p' \text{ R}_{y,z} p''}{p \text{ R}_{x,z} p''} \tag{8.13}$$

$$\frac{p \text{ R}_{x,y} p'}{p * r \text{ R}_{x,z} p' * r} \tag{8.14}$$

$$\frac{\forall x \in X, y \in Q(x)}{\langle \forall z \in X \rightsquigarrow Q(z) \rangle_a^i \text{ R}_{x,y} \mid x \rightsquigarrow y \rangle_a^i} \tag{8.15}$$

The relation resembles the axioms for proxies, where (8.12) means that we do not need to update a

proxy to perform the identity transition. (8.13) allows us to chain proxies together transitively: if we have a proxy that allows a transition from x to y and another that allows a transition from y to z , then we can use both together to transition from x to z . (8.14) means that we can frame resources. Finally, (8.15) allows us to update a suitable proxy to a witness corresponding to the required update.

Note that the **MAKEATOMIC** rules enforces that the speculation for an operation must be resolved before the operation terminates.

For the two-step counter, we introduce a region with type name **TSCounter**. This region has three types of guard resources. The unique non-empty guard **INC** provides the capability to increment the counter. The second type of guard resource is **ACTIVE**(X), which tracks the set X of proxy or witness identifiers of the threads incrementing the counter; these are the threads which have performed the first increment, but not the second. The last type of guard resource is the **INC**(i), which expresses that the identifier i exists in the set X described by the guard **ACTIVE**(X). We wish to allow threads to add and remove identifiers from the set X . We enforce this using the following equivalence:

$$\mathbf{ACTIVE}(X) = \mathbf{ACTIVE}(X \uplus \{i\}) \bullet \mathbf{INC}(i)$$

The possible states of **TSCounter** are natural numbers, representing the value of the counter. All transitions over abstract states of the region are guarded by **INC**, which has the following labelled transition system:

$$\mathbf{INC} : \forall n \in \mathbb{N}. n \rightsquigarrow n + 1$$

The transition system requires that every update must increase the value of the counter. This is necessary to guarantee the correctness of the **incrVal** operation.

$$\begin{aligned} I(\mathbf{TSCounter}_a(s', x, n)) &\triangleq \exists m, X. \mathbf{Counter}(s', x, m) * [\mathbf{ACTIVE}(X)]_a \\ &\quad * \mathbb{S}_a \left(C_a \left(X, \left\lfloor \frac{|X|}{2} \right\rfloor, \left\lfloor \frac{m}{2} \right\rfloor \right) \right) * (m = 2n \vee m = 2n + 1) \end{aligned}$$

The interpretation enforces that the threads performing increments ($|X|$ of them) delay the choice of picking which increment corresponds to which thread until after they have finished.

We now define the interpretation for the abstract types and abstract predicates:

$$\begin{aligned} \mathbb{T}_2 &\triangleq \mathbf{Rld} \\ \mathbf{TSCounter}((r, s'), x, n) &\triangleq \mathbf{TSCounter}_a(s', x, n) * [\mathbf{INC}]_r \end{aligned}$$

Proofs for the **readVal** and **incrVal** operations are shown in Figure 8.6 and Figure 8.7.

8.2 Case Study: Michael-Scott Queue

We now consider the `java.util.concurrent` package [21] variant of the Michael-Scott queue [43] to illustrate how we can apply this technique to verify larger modules.

$$\begin{array}{c}
\mathbb{W}n \in \mathbb{N}. \\
\langle \text{TSCOUNTER}(s, \mathbf{x}, n) \rangle \\
\hline
\text{ABSTRACT: } s = (a, s') \\
\hline
\text{MAKEATOMIC} \\
\hline
\mathbb{W}n \in \mathbb{N}. \\
\langle \text{TSCOUNTER}_a(s', \mathbf{x}, n) * [\text{INC}]_a \rangle \\
\hline
a : i \vdash \\
\left\{ \exists n. \text{TSCOUNTER}_a(s', \mathbf{x}, n) * \langle \mathbb{W}n \in \mathbb{N}. n \rightsquigarrow n \rangle_a^i \right\} \\
\hline
\text{UPDATEREGION} \\
\mathbb{W}n \in \mathbb{N}. \\
\left\langle \exists m, X. \text{Counter}(s', \mathbf{x}, m) * [\text{ACTIVE}(X)]_a * \mathbb{S}_a \left(C_a \left(X, \left\lfloor \frac{|X|}{2} \right\rfloor, \left\lfloor \frac{m}{2} \right\rfloor \right) \right) \right\rangle \\
* (m = 2n \vee m = 2n + 1) * \langle \mathbb{W}n \in \mathbb{N}. n \rightsquigarrow n \rangle_a^i \\
\mathbf{v} := \text{read}(\mathbf{x}); \\
\left\langle \underset{n \rightsquigarrow n}{\overset{a}{\Diamond}} \left(\exists m, X. \text{Counter}(s', \mathbf{x}, m) * [\text{ACTIVE}(X)]_a * \mathbb{S}_a \left(C_a \left(X, \left\lfloor \frac{|X|}{2} \right\rfloor, \left\lfloor \frac{m}{2} \right\rfloor \right) \right) \right) \right\rangle \\
* (m = 2n \vee m = 2n + 1) * \langle n \rightsquigarrow n \rangle_a^i * \mathbf{v} = m \\
\left\{ \exists n. (\mathbf{v} = 2n \vee \mathbf{v} = 2n + 1) * \langle n \rightsquigarrow n \rangle_a^i \right\} \\
\text{return } \mathbf{v} / 2; \\
\langle \text{TSCOUNTER}_a(s', \mathbf{x}, n) * [\text{INC}]_a * \text{ret} = n \rangle \\
\langle \text{TSCOUNTER}(s, \mathbf{x}, n) * \text{ret} = n \rangle
\end{array}$$

Figure 8.6: Proof of correctness of the `readVal` operation.

$$\begin{array}{c}
\mathbb{W}n \in \mathbb{N}. \\
\langle \text{TSCOUNTER}(s, \mathbf{x}, n) \rangle \\
\hline
\text{ABSTRACT: } s = (a, s') \\
\hline
\text{MAKEATOMIC} \\
\hline
a : i \vdash \\
\left\{ \exists n. \text{TSCOUNTER}_a(s', \mathbf{x}, n) * \langle \mathbb{W}n \in \mathbb{N}. n \rightsquigarrow n + 1 \rangle_a^i \right\} \\
\hline
\text{UPDATEREGION} \\
\mathbb{W}n \in \mathbb{N}. \\
\left\langle \exists m, X. \text{Counter}(s', \mathbf{x}, m) * [\text{ACTIVE}(X)]_a * \mathbb{S}_a \left(C_a \left(X, \left\lfloor \frac{|X|}{2} \right\rfloor, \left\lfloor \frac{m}{2} \right\rfloor \right) \right) \right\rangle \\
* (m = 2n \vee m = 2n + 1) * \langle \mathbb{W}n \in \mathbb{N}. n \rightsquigarrow n + 1 \rangle_a^i \\
\text{incr}(\mathbf{x}); \\
\left\langle \underset{\left\lfloor \frac{m}{2} \right\rfloor \rightsquigarrow \left\lfloor \frac{m+1}{2} \right\rfloor}{\overset{a}{\Diamond}} \left(\exists m, X. \text{Counter}(s', \mathbf{x}, m + 1) * [\text{ACTIVE}(X \uplus \{i\})]_a * \mathbb{S}_a \left(C_a \left(X \uplus \{i\}, \left\lfloor \frac{|X \uplus \{i\}|}{2} \right\rfloor, \left\lfloor \frac{m+1}{2} \right\rfloor \right) \right) \right) \right\rangle \\
* (m + 1 = 2n + 1 \vee m + 1 = 2(n + 1)) * [\text{INC}(i)]_a \\
\left\{ \exists n. \text{TSCOUNTER}_a(s', \mathbf{x}, n) * [\text{INC}(i)]_a \right\} \\
\hline
\text{UPDATEREGION} \\
\mathbb{W}n \in \mathbb{N}. \\
\left\langle \exists m, X. \text{Counter}(s', \mathbf{x}, m) * [\text{ACTIVE}(X \uplus \{i\})]_a * \mathbb{S}_a \left(C_a \left(X \uplus \{i\}, \left\lfloor \frac{|X \uplus \{i\}|}{2} \right\rfloor, \left\lfloor \frac{m}{2} \right\rfloor \right) \right) \right\rangle \\
* (m = 2n \vee m = 2n + 1) * [\text{INC}(i)]_a \\
\text{incr}(\mathbf{x}); \\
\left\langle \underset{\left\lfloor \frac{m}{2} \right\rfloor \rightsquigarrow \left\lfloor \frac{m+1}{2} \right\rfloor}{\overset{a}{\Diamond}} \left(\exists m, X. \text{Counter}(s', \mathbf{x}, m + 1) * [\text{ACTIVE}(X)]_a * \mathbb{S}_a \left(C_a \left(X, \left\lfloor \frac{|X|}{2} \right\rfloor, \left\lfloor \frac{m+1}{2} \right\rfloor \right) \right) \right) \right\rangle \\
* (m + 1 = 2n + 1 \vee m + 1 = 2(n + 1)) * \exists v. \langle v \rightsquigarrow v + 1 \rangle_a^i \\
\left\{ \exists v. \langle v \rightsquigarrow v + 1 \rangle_a^i \right\} \\
\langle \text{TSCOUNTER}_a(s', \mathbf{x}, n + 1) * [\text{INC}]_a \rangle \\
\langle \text{TSCOUNTER}(s, \mathbf{x}, n + 1) \rangle
\end{array}$$

Figure 8.7: Proof of correctness of the `incrVal` operation.

8.2.1 Atomic Specification

We start by giving an abstract specification for a non-blocking queue as follows:

$$\begin{array}{l}
\vdash \{ \text{True} \} \text{makeQueue}() \left\{ \exists s \in \mathbb{T}_1, t \in \mathbb{T}_2. \text{Queue}(s, \text{ret}, [], t) \right\} \\
\vdash \mathbb{W}(vs, t) \in \mathbb{L}_1 \times \mathbb{T}_2. \quad \left\langle \text{Queue}(s, x, vs, t) * v \neq 0 \right\rangle \\
\quad \text{enqueue}(x, v) \\
\quad \left\langle \exists vs', t'. \text{Queue}(s, x, vs', t') * vs' = vs ++ [v] \right\rangle \\
\vdash \mathbb{W}(vs, t) \in \mathbb{L}_1 \times \mathbb{T}_2. \quad \left\langle \text{Queue}(s, x, vs, t) \right\rangle \\
\quad \text{dequeue}(x) \\
\quad \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{Queue}(s, x, vs, t) * \text{ret} = 0 \\ \text{else } \exists vs', t'. \text{Queue}(s, x, vs', t') * vs = \text{ret} : vs' \end{array} \right\rangle
\end{array}$$

We represent the queue by the abstract predicate $\text{Queue}(s, x, vs, t)$. It asserts that there is a queue at address x and with contents vs , where $vs \in \mathbb{L}_1$, and $\mathbb{L}_1$ is a list containing non-zero integers. The parameters $s \in \mathbb{T}_1$ and $t \in \mathbb{T}_2$ capture implementation-specific information of the queue.

The constructor returns an empty queue. The **enqueue** operation atomically adds an element to the end of the queue; this element cannot be 0. This restricts all the elements in the queue from being equal to 0. The **dequeue** operation atomically removes an element from the front of the queue, if one is available (*i.e.* the queue is not empty); otherwise it will return 0.

8.2.2 Implementation

The queue, given in Figure 8.8, is represented as a non-empty linked list. The first node in the list is considered to be outside of the queue. Its value is not part of the abstract state of the queue. Additionally, there are two anchor pointers: the head and the tail. The head pointer always points to the sentinel node. The sentinel node is a node that is not part of the queue, but points to the node that contains the first element of the queue. The tail pointer points to a node that is or was in the queue, as long as the true tail is accessible from that node.

An **enqueue** operation appends nodes to the end of the list, by updating the next pointer of the true tail from 0 to the address of the new node. A **dequeue** operation removes nodes from the front of the queue by updating the head pointer with a compare-and-set operation to the next node from the sentinel node. If the queue is empty, then the sentinel node in the list is the only node accessible from the head pointer and its next pointer has value 0. Note that when a **dequeue** operation removes a node from the queue, it is still possible to reach any node in the queue from that node. Moreover, we set the value of the node to 0 after removing it from the queue. This is an optimisation to improve garbage collection in the original algorithm and does not affect the correctness of the implementation.

In the **enqueue** operation, the atomic update always happens if $\text{CAS}(t.\text{next}, s, n)$ succeeds. Otherwise, the algorithm retries. There are additional compare-and-set operations used to update the tail pointer. The outcome of such operations does not affect the abstract state of the queue.

In the **dequeue** operation, if the queue is not empty and it succeeds in advancing the sentinel node, the atomic update occurs on the $\text{CAS}(x.\text{head}, h, \text{first})$. The reading and nullifying the value of the node does not affect the abstract state of the queue. If the queue is empty, the **dequeue** operation's atomic

```

function makeQueue() {
  n := alloc(2);
  [n.value] := 0;
  [n.next] := 0;
  x := alloc(2);
  [x.head] := n;
  [x.tail] := n;
  return x;
}

function enqueue(x, v) {
  n := alloc(2);
  [n.value] := v;
  [n.next] := 0;
  while (true) {
    t := [x.tail];
    s := [t.next];
    t' := [x.tail];
    if (t = t') {
      if (s = 0) {
        b := CAS(t.next, s, n);
        if (b = 1) {
          CAS(x.tail, t, n)
          return
        }
      } else {
        CAS(x.tail, t, s)
      }
    }
  }
}

function dequeue(x) {
  while (true) {
    h := [x.head];
    t := [x.tail];
    first := [h.next];
    h' := [x.head];
    if (h = h') {
      if (h = t) {
        if (first = 0) {
          return 0;
        } else {
          CAS(x.tail, t, first)
        }
      } else {
        b := CAS(x.head, h, first);
        if (b = 1) {
          v := [first.value];
          if (v ≠ 0) {
            [first.value] := 0;
            return v;
          }
        }
      }
    }
  }
}

```

where $\mathbb{E}.\text{head} \stackrel{\text{def}}{=} \mathbb{E}$ $\mathbb{E}.\text{tail} \stackrel{\text{def}}{=} \mathbb{E} + 1$ $\mathbb{E}.\text{value} \stackrel{\text{def}}{=} \mathbb{E}$ $\mathbb{E}.\text{next} \stackrel{\text{def}}{=} \mathbb{E} + 1$.

Figure 8.8: Michael-Scott Queue implementation.

update depends on a future event. In particular, when the operation performs `first := [h.next]`, there are two cases: either `h` value matches the current head of the queue, or it doesn't. We can only be sure that `h` matches the head of the queue when the queue is empty, and in that case, `first` has value 0. However, we cannot determine if the atomic update occurs at this particular point, as it is dependent on the next read `h' := [x.head]`; getting the same value as `h`. If it is not, then the entire operation needs to start over.

We want to prove the correctness of the implementation. In both operations, the atomic updates are always performed by the thread performing the operation. There is no helping involved for the Michael-Scott queue. However, the `dequeue` operation can experience speculation when dequeuing from an empty queue. This happens because we cannot determine if the operation will decide to retry or not, until a future operation is performed.

In order to prove the correctness of the queue, we define a region type **MSQueue**. The type is parametrised by the address of the queue and its abstract state. The abstract state of **MSQueue** regions consists of a pair of lists $(ds, ns) \in \mathbb{L}_2 \times \mathbb{L}_3$, where ds is a list of nodes that have been in the queue plus the initial sentinel node and ns is the list of nodes containing the values currently in the queue. We define their types as $\mathbb{L}_2 \triangleq (\text{Addr} \times \mathbb{Z} \times \text{Addr} \uplus \{0\})^+$ and $\mathbb{L}_3 \triangleq (\text{Addr} \times \{v \in \mathbb{Z} \mid v \neq 0\} \times \text{Addr} \uplus \{0\})^*$. Each element of each of the lists contains a node's address, its value and next pointer. Note that the

ds list cannot be empty, since its last element is the sentinel node. Furthermore, the list ns cannot contain empty values. It is important to the correctness of the algorithm that nodes that have been removed from the queue cannot reappear again. To account for this, we track all the nodes that have been previously removed from the head of the queue in the ds list. The protocol that governs how the **MSQueue** region can be updated is specified by the transition system as follows:

$$\begin{aligned}
G : \forall n, m, w, ds. \quad & (ds ++ [(n, v, 0)], []) \rightsquigarrow (ds ++ [(n, v, m)], [(m, w, 0)]) \\
G : \forall n, m, w, ds, ns. \quad & (ds, ns ++ [(n, v, 0)]) \rightsquigarrow (ds, ns ++ [(n, v, m)] ++ [(m, w, 0)]) \\
G : \forall n, m, ds, ns. \quad & (ds, (n, v, m) : ns) \rightsquigarrow (ds ++ [(n, v, m)], ns)
\end{aligned}$$

There is a single guard G that gives threads permissions to enqueue and dequeue. The first two transitions allow a thread to enqueue a new value, in the case the queue is empty and otherwise. The last transition allows a thread to dequeue a value, by moving the node from the list ns to the list ds .

In order to track operations that have observed the queue as empty, but have not yet decided what action to take, we introduce two guards: $\text{EMPTY}(i)$ and $\text{DEQ}(S)$. The guard $\text{EMPTY}(i)$, where $i \in \text{Tld}$, states the thread i has observed the queue as empty when performing a **dequeue** operation. The guard $\text{DEQ}(S)$ accounts for the $\text{EMPTY}(i)$ guards that have been issued, where $S \subseteq \text{Tld}$. These guards are defined to obey the following equation:

$$\text{DEQ}(S) = \text{DEQ}(S \uplus \{i\}) \bullet \text{EMPTY}(i) \quad \text{if } i \notin \text{proj}_1(S)$$

When the **dequeue** succeeds in removing a node from the queue, it nullifies the value of a node. We need to guarantee that the only thread doing it is the thread that actually removes the node from the queue. Otherwise, another operation could potentially set the value of the node to 0 before the thread dequeuing read its value. For this purpose, we introduce a guard $\text{VALUE}(n, v)$, where $n \in \text{Addr}$ and $v \in \{v \in \mathbb{Z} \mid v \neq 0\}$, that tracks the value v of the node n when it is part of the queue, and prevents other threads from manipulating its value. Moreover, we have an additional guard $\text{OUT}(S)$ to track how many $\text{VALUE}(n, v)$ have been issued. The guards must satisfy the following equation:

$$\text{OUT}(S) = \text{OUT}(S \uplus \{(n, v)\}) \bullet \text{VALUE}(n, v) \quad \text{if } n \notin \text{proj}_1(S)$$

We now introduce auxiliary predicates to represent both lists, as well as the data structure of the queue in Figure 8.9. The $\text{node}(n, v, m)$ represents a node at address n , with value v and the next pointer with value m . The $\text{outList}(a, x, ds, y)$ represents the ds linked list, where the sentinel node when the queue was created has address x , and the most recent node in the list has a next pointer with value y . We require that the nodes that have been nullified must have the guard $\text{VALUE}(x, v)$ to record the initial value. The $\text{inList}(x, ns)$ represents the ns linked list, where the last node of the ds list points to its first node with value x , and the last node of the list always has next pointer with value 0. The predicate $\text{queue}(a, x, ds, ns)$ represents the data structure of the queue: The **head** pointer points to the last node of the list ds , and the **tail** pointer points either to the last element of the list ds or one node in the list ns ¹. Moreover, we enforce the ds list to have at least one node and that its last node has a next pointer to the first element of the ns list. The final predicate is used in the proofs

¹This contrasts the original Michael-Scott queue algorithm, where the tail pointer points to one of the last two nodes.

$$\begin{aligned}
\text{node}(n, v, m) &\stackrel{\text{def}}{=} n.\text{value} \mapsto v * n.\text{next} \mapsto m \\
\text{outList}(a, x, ds, y) &\stackrel{\text{def}}{=} (x = y * ds = []) \vee \exists v, n, ds'. \left((\text{node}(x, v, n) \vee \text{node}(x, 0, n) * [\text{VALUE}(x, v)]_r) \right. \\
&\quad \left. * \text{outList}(a, n, ds', y) * ds = (x, v, n) : ds' \right) \\
\text{inList}(x, ns) &\stackrel{\text{def}}{=} (x = 0 * ns = []) \vee (\exists v, n, ns'. \text{node}(x, v, n) * \text{inList}(n, ns') * ns = (x, v, n) : ns') \\
\text{queue}(a, x, ds, ns) &\stackrel{\text{def}}{=} \exists h, t, y, z, v, ds', S. x.\text{head} \mapsto h * x.\text{tail} \mapsto t * \text{outList}(a, y, ds, z) * \text{inList}(z, ns) \\
&\quad * [\text{OUT}(S)]_a * S = \{(n, v) \mid (n, v, m) \in ds\} * ds = ds' ++ [(h, v, z)] \\
&\quad * t \in \text{proj}_1([(h, v, z)] ++ ns) \\
\text{after}(x, y, vs) &\stackrel{\text{def}}{=} \exists vs'. vs = v : vs' * (x = v * y \in vs \vee x \neq v * \text{after}(x, y, vs'))
\end{aligned}$$

Figure 8.9: Auxiliary predicates to describe the shared resources of the Michael-Scott queue.

$$\begin{aligned}
\text{PE}(i) &= \left\langle \mathbb{W}(ds, ns) \rightsquigarrow \begin{array}{l} \text{then } \exists ds''. ds = ds'' ++ [(n, v, 0)] \wedge ds' = ds'' ++ [(n, v, m)] \wedge ns' = [(m, w, 0)] \\ \text{else } \exists ns''. ds = ds' \wedge ns = ns'' ++ [(n, v, 0)] \wedge ns' = ns'' ++ [(n, v, m), (m, w, 0)] \end{array} \right\rangle_a^i \\
\text{WE}(i, m, w, ds, ns, ds', ns') &= \left(ds, ns \rightsquigarrow \begin{array}{l} \text{then } \left(\begin{array}{l} (ds', ns') \wedge \text{if } ns = [] \\ \exists ds''. ds = ds'' ++ [(n, v, 0)] \wedge ds' = ds'' ++ [(n, v, m)] \\ \wedge ns' = [(m, w, 0)] \end{array} \right) \\ \text{else } \left(\begin{array}{l} \exists ns''. ds = ds' \wedge ns = ns'' ++ [(n, v, 0)] \\ \wedge ns' = ns'' ++ [(n, v, m), (m, w, 0)] \end{array} \right) \end{array} \right) \right\rangle_a^i \\
\text{PD}(i) &= \langle \mathbb{W}(ds, ns) \rightsquigarrow \text{if } ns = [] \text{ then } (ds, ns) \text{ else } (ds ++ [(n, v, m)], ns') \wedge ns = (n, v, m) : ns' \rangle_a^i \\
\text{WD1}(i) &= \exists ds. |(ds, []) \rightsquigarrow (ds, [])|_a^i \\
\text{WD2}(i, n, v, ds, ns, ds', ns') &= \exists m. |(ds, ns) \rightsquigarrow (ds', ns') \wedge ns = (n, v, m) : ns' \wedge ds' = ds ++ [(n, v, m)]|_a^i
\end{aligned}$$

Figure 8.10: Proxies and witnesses for the Michael-Scott queue.

to describe that node with address x proceeds the node with address y in the list vs .

We define auxiliary notation for referring to proxies and witnesses for the enqueue and dequeue operations in Figure 8.10.

In order to describe the set of possible proxies and witnesses when the **dequeue** operation observes an empty queue, but has not yet decided if that is the atomic update, we define the following set:

$$\begin{aligned}
D_a(\emptyset) &\stackrel{\text{def}}{=} \{\text{True}\} \\
D_a(X) &\stackrel{\text{def}}{=} \{y * x \mid i \in X, x \in \{\text{PD}(i), \text{WD1}(i)\}, y \in D_a(X \setminus \{i\})\}
\end{aligned}$$

The **MSQueue** region type has the following interpretation:

$$I(\text{MSQueue}_a(x, (ds, ns))) = \exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X))$$

The abstract predicate $\text{Queue}(s, x, vs, (ds, ns))$ combines the region and the guard:

$$\begin{aligned}
\mathbb{T}_1 &\triangleq \text{Rld} \\
\mathbb{T}_2 &\triangleq \mathbb{L}_2 \times \mathbb{L}_3 \\
\text{Queue}(a, x, vs, (ds, ns)) &\triangleq \text{MSQueue}_a(x, (ds, ns)) * [\text{G}]_a * vs = \text{proj}_2(ns)
\end{aligned}$$

The proof of the **dequeue** operation is given in Figure 8.11.

When the queue is empty, we speculatively execute the atomic update, and keep the guard $\text{EMPTY}(i)$ to track it. Note that the node pointed by the **head**, must always be of the ds list. This means that if its next pointer is 0, then it must be the case that the queue is empty. If the following read value of the head pointer matches the previously read, we collapse the speculative state and retrieve the witness. Otherwise, we retrieve the original proxy from the speculative component and try to update the **tail** of the queue. Updating the tail does not affect the abstract state. If the queue was not empty, then we attempt to update the **head** to the next node using a compare-and-set. If successful, the atomic update occurs and we hold the guard $\text{VALUE}(n, v)$ to prevent other threads from nullifying the original value in the node. We then read the value from the node and set its value to 0 before returning.

The proof of the **enqueue** operation is omitted, as it is very similar to the proof for **dequeue**.

$$\begin{array}{l}
\mathbb{W}(vs, t) \in \mathbb{L}_1 \times \mathbb{T}_2. \\
\langle \text{Queue}(s, x, vs, t) \rangle \\
\mathbb{W}(ds, ns) \in \mathbb{L}_2 \times \mathbb{L}_3. \\
\langle \text{MSQueue}_a(x, (ds, ns)) * [G]_a * vs = \text{proj}_2(ns) \rangle \\
a : i \vdash \\
\text{while (true) } \{ \\
\quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) \} \\
\quad h := [x.\text{head}]; t := [x.\text{tail}]; \\
\quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) * h \in \text{proj}_1(ds) * (h = t \vee \text{after}(h, t, \text{proj}_1(ds++ns))) \} \\
\quad \mathbb{W}(ds, ns) \in \mathbb{L}_2 \times \mathbb{L}_3. \\
\quad \langle \exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X)) * \text{PD}(i) * h \in \text{proj}_1(ds) \rangle \\
\quad \langle * (h = t \vee \text{after}(h, t, \text{proj}_1(ds++ns))) \rangle \\
\quad \text{first} := [h.\text{next}]; \\
\quad \text{UPDATEREGION} \left\langle \begin{array}{l} \text{if first} = 0 \text{ then } \begin{array}{l} \overset{a}{\diamond} \left(\exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X \uplus \{i\})]_a * \mathbb{S}_a(D_a(X \uplus \{i\})) \right) \\ * [\text{EMPTY}(i)] * h \in \text{proj}_1(ds) * h = t \end{array} \\ \text{else } \begin{array}{l} \overset{a}{\diamond} \left(\exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X)) * \text{PD}(i) * h \in \text{proj}_1(ds) \right) \\ * (h = t \vee \text{after}(h, t, \text{proj}_1(ds++ns))) * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \end{array} \end{array} \right\rangle \\
\quad \left\{ \begin{array}{l} \text{if first} = 0 \text{ then } \left(\exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * [\text{EMPTY}(i)] * h \in \text{proj}_1(ds) * h = t \right) \\ \text{else } \left(\begin{array}{l} \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) * h \in \text{proj}_1(ds) \\ * (h = t \vee \text{after}(h, t, \text{proj}_1(ds++ns))) * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \end{array} \right) \end{array} \right\} \\
\quad h' := [x.\text{head}]; \\
\quad \text{if } (h = h') \{ \\
\quad \quad \text{if } (h = t) \{ \\
\quad \quad \quad \left\{ \begin{array}{l} \text{if first} = 0 \text{ then } \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * [\text{EMPTY}(i)] * h \in \text{proj}_1(ds) * h = t \\ \text{else } \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) * h \in \text{proj}_1(ds) * h = t * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \end{array} \right\} \\
\quad \quad \quad \text{if } (\text{first} = 0) \{ \\
\quad \quad \quad \quad \{ \text{WD1}(i) \} \text{ // open region and swap } [\text{EMPTY}(i)] \text{ by the witness.} \\
\quad \quad \quad \quad \text{return 0;} \\
\quad \quad \quad \text{else } \{ \\
\quad \quad \quad \quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) * h = t * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \} \\
\quad \quad \quad \quad \text{CAS}(x.\text{tail}, t, \text{first}); \\
\quad \quad \quad \quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) \} \\
\quad \quad \quad \} \\
\quad \quad \quad \} \text{ else } \{ \\
\quad \quad \quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \} \\
\quad \quad \quad \text{UPDATEREGION} \left\langle \begin{array}{l} \mathbb{W}(ds, ns) \in \mathbb{L}_2 \times \mathbb{L}_3. \\ \langle \exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X)) * \text{PD}(i) * \text{after}(h, \text{first}, \text{proj}_1(ds++ns)) \rangle \\ b := \text{CAS}(x.\text{head}, h, \text{first}); \\ \left\langle \begin{array}{l} \text{if } b = 1 \text{ then } \exists ds', ns', v. \overset{a}{\diamond} \left(\begin{array}{l} \exists X. \text{queue}(a, x, ds', ns') * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X)) \\ * \text{WD2}(i, \text{first}, v, ds, ns, ds', ns') * [\text{VALUE}(\text{first}, v)]_a \end{array} \right) \\ \text{else } \overset{a}{\diamond} \left(\exists X. \text{queue}(a, x, ds, ns) * [\text{DEQ}(X)]_a * \mathbb{S}_a(D_a(X)) * \text{PD}(i) \right) \end{array} \right\rangle \\ \left\{ \begin{array}{l} \text{if } b = 1 \text{ then } \left(\begin{array}{l} \exists ds, ns, ds', ns', ds'', ns'', v. \text{MSQueue}_a(x, (ds, ns)) \\ * \text{WD2}(i, \text{first}, v, ds', ns', ds'', ns'') * [\text{VALUE}(\text{first}, v)]_a \end{array} \right) \\ \text{else } \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) \end{array} \right\} \\ \text{if } (b = 1) \{ \\
\quad \quad \quad \left\{ \begin{array}{l} \exists ds, ns, ds', ns', ds'', ns'', v. \text{MSQueue}_a(x, (ds, ns)) \\ * \text{WD2}(i, \text{first}, v, ds', ns', ds'', ns'') * [\text{VALUE}(\text{first}, v)]_a \end{array} \right\} \\
\quad \quad \quad v := [\text{first}.\text{value}]; [\text{first}.\text{value}] := 0; \\
\quad \quad \quad \{ \exists ds', ns', ds'', ns''. \text{WD2}(i, \text{first}, v, ds', ns', ds'', ns'') \} \\
\quad \quad \quad \text{return } v; \\
\quad \quad \quad \} \\
\quad \quad \quad \{ \exists ds, ns. \text{MSQueue}_a(x, (ds, ns)) * \text{PD}(i) \} \\
\quad \quad \quad \} \} \} \\
\quad \left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{MSQueue}_a(x, (ds, ns)) * [G]_a * \text{ret} = 0 \\ \text{else } \left(\begin{array}{l} \exists vs', ds', ns', n, v, m. \text{MSQueue}_a(x, (ds', ns')) * [G]_a * vs' = \text{proj}_2(ns') * ns = (n, v, m) : ns' \\ * ds' = ds++[(n, v, m)] \end{array} \right) \end{array} \right\rangle \\
\langle \text{if } vs = [] \text{ then } \text{Queue}(s, x, vs, t) * \text{ret} = 0 \text{ else } \exists vs', t'. \text{Queue}(s, x, vs', t') * vs = \text{ret} : vs' \rangle
\end{array}$$

Figure 8.11: Proof of correctness of dequeue operation.

9 Conclusions

We have introduced a program logic, TaDA, which uses *atomic triples* for specifying abstract atomicity, as well as separation-style Hoare triples for specifying abstract disjointness. With the combination of abstract atomicity and abstract disjointness that TaDA provides, we can specify and verify modules which feature both atomic and non-atomic operations, possibly at different levels of abstraction. Additionally, TaDA allows us to easily extend modules with new operations, as well as build new modules on top of existing ones.

We have shown how TaDA can be extended to prove termination of non-blocking programs. Using our abstract specifications, clients can reason about total correctness without needing to know any of the details of the underlying implementation. Different implementations that satisfy the same specification may have different termination arguments, but these arguments are never exposed to the clients. By using ordinals to bound interference, our specifications can express traditional non-blocking properties. Moreover, they can also capture the notion of *non-impedance*: that one operation does not disrupt the progress of another.

Finally, we have proposed a potential extension to the logic to handle more advanced patterns of concurrency. Using the extension, we were able to give a novel atomic specification to a ticket lock module and to introduce a two-step counter that exhibits speculation, despite its simplicity. Additionally, we applied the technique to a larger example and shown that the approach scales. It remains an open question what the semantic model that justifies such an extension should be.

9.1 Future Work

9.1.1 Tool Support

So far, all of the soundness proofs and proofs shown in this thesis using TaDA or Total-TaDA have been done by hand. One of our goals is to formalise the logic in a proof assistant, such as Coq, and prove its soundness. Moreover, we have implemented CAPER [14], a semi-automated verification tool, on a fragment of TaDA. We intend to extend the tool to be able to prove all of the the examples shown.

9.1.2 Helping/Speculation

We have presented a technique that supports reasoning with about abstract atomicity when implementations make use of helping and speculation chapter 8. It seems clear that we require a different semantic model. However, it is not clear what changes must be made in order to support the extension. There are approaches based on contextual refinement that support helping and speculation [59, 39], but are restricted to linearisability and suffer from the problems described in §2.7. Other approaches use higher-order to support helping such as [32, 54, 35], but have no support for speculation.

9.1.3 Higher-order Support

Recently, Iris [35] combined TaDA with iCAP [54] into a new higher order program logic, which encodes TaDA’s proof rules in logic. This gives it the expressive power to handle higher-order programs and reentrancy. It would be interesting to explore if we can encode the extensions presented here to prove termination and speculation.

9.1.4 Weak Memory

Burkhardt *et al.* [5] have extended the concept of linearisability to the total store order (TSO) memory model [48]. Additionally, in recent years, there have been a series of program logics that adapt some of the program logics described in the background chapter to several weak memory models [62, 60, 36, 53]. An interesting research direction would be to investigate extensions of TaDA that can specify and verify programs that make use of weak memory models, such as TSO.

9.1.5 Liveness

Blocking

Many concurrent modules make use of *blocking*, for example by using locks or monitors. Properties such as starvation-freedom can be expressed in terms of termination, but require the assumption of a fair scheduler. Some aspects of our approach are likely to be applicable here as well. However, it is also necessary to constrain future behaviours, for instance, to specify that a lock that has been acquired will be released in finite time. Liang and Feng [38] have developed Lili, a program logic based on rely-guarantee and refinement that can verify linearisability and progress properties for concurrent modules under fair scheduling. Unfortunately, their approach suffers from the same problems common to linearisability, such as the lack of ownership transfer and not being able to handle modules with operations such as the `wkIncr`. Understanding how we could extend our termination extension to blocking algorithms and overcome the limitations of linearisability-based approaches would be another interesting research question.

Non-termination

Some programs, such as operating systems, are designed not to terminate. Nonetheless, such programs should still exhibit progress. It would be interesting to extend Total-TaDA to specify and verify progress properties of non-terminating systems. Progress can be seen as localised termination, so the same reasoning techniques should apply. However, it is not clear what kind of specifications will be necessary to express and verify these properties.

9.1.6 Fault Tolerance

An underlying assumption in our logic is that the machine where the program is running does not fail. In practice, machines can fail unpredictably for various reasons, e.g. power loss, corrupting resources. Critical software, e.g. file systems, employ recovery methods to mitigate these effects. In future, we want to combine abstract atomicity from concurrency with host failure atomicity in the style of [46].

Bibliography

- [1] Martín Abadi and Leslie Lamport. The existence of refinement mappings. In *Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK, July 5-8, 1988*, pages 165–175. IEEE Computer Society, 1988. doi: 10.1109/LICS.1988.5115. URL <http://dx.doi.org/10.1109/LICS.1988.5115>. (Cited on page 133.)
- [2] Richard Bornat, Cristiano Calcagno, Peter W. O'Hearn, and Matthew J. Parkinson. Permission accounting in separation logic. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 259–270, 2005. doi: 10.1145/1040305.1040327. URL <http://doi.acm.org/10.1145/1040305.1040327>. (Cited on page 32.)
- [3] Pontus Boström and Peter Müller. Modular Verification of Finite Blocking in Non-terminating Programs. In John Tang Boyland, editor, *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic*, volume 37 of *LIPICs*, pages 639–663. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. doi: 10.4230/LIPICs.ECOOP.2015.639. URL <http://dx.doi.org/10.4230/LIPICs.ECOOP.2015.639>. (Cited on page 132.)
- [4] John Boyland. Checking Interference with Fractional Permissions. In *Static Analysis, 10th International Symposium, SAS 2003, San Diego, CA, USA, June 11-13, 2003, Proceedings*, pages 55–72, 2003. doi: 10.1007/3-540-44898-5_4. URL http://dx.doi.org/10.1007/3-540-44898-5_4. (Cited on pages 32 and 114.)
- [5] Sebastian Burckhardt, Alexey Gotsman, Madanlal Musuvathi, and Hongseok Yang. Concurrent library correctness on the TSO memory model. In Helmut Seidl, editor, *Programming Languages and Systems - 21st European Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings*, volume 7211 of *Lecture Notes in Computer Science*, pages 87–107. Springer, 2012. doi: 10.1007/978-3-642-28869-2_5. URL http://dx.doi.org/10.1007/978-3-642-28869-2_5. (Cited on page 150.)
- [6] Georg Cantor. Beiträge zur begründung der transfiniten mengenlehre. *Mathematische Annalen*, 49(2):207–246, 1897. ISSN 0025-5831. doi: 10.1007/BF01444205. URL <http://dx.doi.org/10.1007/BF01444205>. (Cited on page 109.)
- [7] Pedro da Rocha Pinto, Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, and Mark J. Wheelhouse. A simple abstraction for complex concurrent indexes. In Cristina Videira Lopes and Kathleen Fisher, editors, *Proceedings of the 26th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2011, part of SPLASH 2011, Portland, OR, USA, October 22 - 27, 2011*, pages 845–864. ACM, 2011. doi:

10.1145/2048066.2048131. URL <http://doi.acm.org/10.1145/2048066.2048131>. (Cited on page 43.)

- [8] Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. TaDA: A Logic for Time and Data Abstraction. In Richard Jones, editor, *ECOOP 2014 - Object-Oriented Programming - 28th European Conference, Uppsala, Sweden, July 28 - August 1, 2014. Proceedings*, volume 8586 of *Lecture Notes in Computer Science*, pages 207–231. Springer, 2014. doi: 10.1007/978-3-662-44202-9_9. URL http://dx.doi.org/10.1007/978-3-662-44202-9_9. (Cited on page 21.)
- [9] Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. Steps in Modular Specifications for Concurrent Modules (Invited Tutorial Paper). *Electr. Notes Theor. Comput. Sci.*, 319:3–18, 2015. doi: 10.1016/j.entcs.2015.12.002. URL <http://dx.doi.org/10.1016/j.entcs.2015.12.002>. (Cited on page 21.)
- [10] Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. Modular Termination Verification for Non-blocking Concurrency. In Peter Thiemann, editor, *Programming Languages and Systems - 25th European Symposium on Programming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, volume 9632 of *Lecture Notes in Computer Science*, pages 176–201. Springer, 2016. doi: 10.1007/978-3-662-49498-1_8. URL http://dx.doi.org/10.1007/978-3-662-49498-1_8. (Cited on page 21.)
- [11] Markus Dahlweid, Michal Moskal, Thomas Santen, Stephan Tobies, and Wolfram Schulte. VCC: Contract-based modular verification of concurrent C. In *31st International Conference on Software Engineering, ICSE 2009, May 16-24, 2009, Vancouver, Canada, Companion Volume*, pages 429–430. IEEE, 2009. doi: 10.1109/ICSE-COMPANION.2009.5071046. URL <http://dx.doi.org/10.1109/ICSE-COMPANION.2009.5071046>. (Cited on page 38.)
- [12] Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. Concurrent Abstract Predicates. In Theo D’Hondt, editor, *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings*, volume 6183 of *Lecture Notes in Computer Science*, pages 504–528. Springer, 2010. doi: 10.1007/978-3-642-14107-2_24. URL http://dx.doi.org/10.1007/978-3-642-14107-2_24. (Cited on pages 19, 31, 33, 38, 40, 43, 90, 92, and 107.)
- [13] Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. Views: compositional reasoning for concurrent programs. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’13, Rome, Italy - January 23 - 25, 2013*, pages 287–300. ACM, 2013. doi: 10.1145/2429069.2429104. URL <http://doi.acm.org/10.1145/2429069.2429104>. (Cited on pages 47, 62, 65, and 68.)
- [14] Thomas Dinsdale-Young, Pedro da Rocha Pinto, Kristoffer Just Andersen, and Lars Birkedal. Caper: Automatic Verification for Fine-grained Concurrency. In *Programming Languages and*

- [15] Mike Dodds, Xinyu Feng, Matthew J. Parkinson, and Viktor Vafeiadis. Deny-Guarantee Reasoning. In *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, pages 363–377, 2009. doi: 10.1007/978-3-642-00590-9_26. URL http://dx.doi.org/10.1007/978-3-642-00590-9_26. (Cited on pages 31 and 107.)
- [16] Simon Doherty, David Detlefs, Lindsay Groves, Christine H. Flood, Victor Luchangco, Paul Alan Martin, Mark Moir, Nir Shavit, and Guy L. Steele Jr. DCAS is not a silver bullet for nonblocking algorithm design. In Phillip B. Gibbons and Micah Adler, editors, *SPAA 2004: Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, June 27-30, 2004, Barcelona, Spain*, pages 216–224. ACM, 2004. doi: 10.1145/1007912.1007945. URL <http://doi.acm.org/10.1145/1007912.1007945>. (Cited on page 102.)
- [17] Derek Dreyer, Georg Neis, and Lars Birkedal. The impact of higher-order state and control effects on local relational reasoning. In Paul Hudak and Stephanie Weirich, editors, *Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010*, pages 143–156. ACM, 2010. doi: 10.1145/1863543.1863566. URL <http://doi.acm.org/10.1145/1863543.1863566>. (Cited on page 107.)
- [18] Xinyu Feng. Local rely-guarantee reasoning. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 315–327, 2009. doi: 10.1145/1480881.1480922. URL <http://doi.acm.org/10.1145/1480881.1480922>. (Cited on pages 31 and 38.)
- [19] Ivana Filipovic, Peter W. O’Hearn, Noam Rinetzkky, and Hongseok Yang. Abstraction for concurrent objects. In Giuseppe Castagna, editor, *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5502 of *Lecture Notes in Computer Science*, pages 252–266. Springer, 2009. doi: 10.1007/978-3-642-00590-9_19. URL http://dx.doi.org/10.1007/978-3-642-00590-9_19. (Cited on pages 34 and 107.)
- [20] R. W. Floyd. Assigning Meanings to Programs. *Proceedings of the American Mathematical Society Symposia on Applied Mathematics*, 19:19–31, 1967. (Cited on pages 108 and 109.)
- [21] Brian Goetz, Tim Peierls, Joshua Bloch, Joseph Bowbeer, David Holmes, and Doug Lea. *Java Concurrency in Practice*. Addison-Wesley, 2006. ISBN 978-0-321-34960-6. (Cited on page 141.)
- [22] Alexey Gotsman and Hongseok Yang. Linearizability with Ownership Transfer. In Maciej Koutny and Irek Ulidowski, editors, *CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012, Newcastle upon Tyne, UK, September 4-7, 2012. Proceedings*, volume 7454 of *Lecture Notes in Computer Science*, pages 256–271. Springer, 2012. doi: 10.1007/

- 978-3-642-32940-1_19. URL http://dx.doi.org/10.1007/978-3-642-32940-1_19. (Cited on pages 101 and 107.)
- [23] Alexey Gotsman, Byron Cook, Matthew J. Parkinson, and Viktor Vafeiadis. Proving that non-blocking algorithms don't block. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 16–28. ACM, 2009. doi: 10.1145/1480881.1480886. URL <http://doi.acm.org/10.1145/1480881.1480886>. (Cited on pages 128 and 132.)
 - [24] Steve Heller, Maurice Herlihy, Victor Luchangco, Mark Moir, William N. Scherer III, and Nir Shavit. A lazy concurrent list-based set algorithm. *Parallel Processing Letters*, 17(4):411–424, 2007. doi: 10.1142/S0129626407003125. URL <http://dx.doi.org/10.1142/S0129626407003125>. (Cited on page 138.)
 - [25] Maurice Herlihy. Wait-free synchronization. *ACM Trans. Program. Lang. Syst.*, 13(1):124–149, 1991. doi: 10.1145/114005.102808. URL <http://doi.acm.org/10.1145/114005.102808>. (Cited on page 109.)
 - [26] Maurice Herlihy and Jeannette M. Wing. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990. doi: 10.1145/78969.78972. URL <http://doi.acm.org/10.1145/78969.78972>. (Cited on pages 19, 22, and 33.)
 - [27] Maurice Herlihy, Victor Luchangco, and Mark Moir. Obstruction-free synchronization: Double-ended queues as an example. In *23rd International Conference on Distributed Computing Systems (ICDCS 2003), 19-22 May 2003, Providence, RI, USA*, pages 522–529. IEEE Computer Society, 2003. doi: 10.1109/ICDCS.2003.1203503. URL <http://dx.doi.org/10.1109/ICDCS.2003.1203503>. (Cited on page 109.)
 - [28] Gerhard Hessenberg. *Grundbegriffe der Mengenlehre*. Abhandlungen der Fries'schen Schule / Neue Folge. Vandenhoeck & Ruprecht, 1906. (Cited on page 114.)
 - [29] C. A. R. Hoare. An Axiomatic Basis for Computer Programming. *Commun. ACM*, 12(10):576–580, 1969. doi: 10.1145/363235.363259. URL <http://doi.acm.org/10.1145/363235.363259>. (Cited on page 24.)
 - [30] Jan Hoffmann, Michael Marmar, and Zhong Shao. Quantitative Reasoning for Proving Lock-Freedom. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 124–133. IEEE Computer Society, 2013. doi: 10.1109/LICS.2013.18. URL <http://dx.doi.org/10.1109/LICS.2013.18>. (Cited on pages 17, 128, and 132.)
 - [31] Samin S. Ishtiaq and Peter W. O'Hearn. BI as an assertion language for mutable data structures. In Chris Hankin and Dave Schmidt, editors, *Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, London, UK, January 17-19, 2001*, pages 14–26. ACM, 2001. (Cited on page 31.)
 - [32] Bart Jacobs and Frank Piessens. Expressive modular fine-grained concurrency specification. In Thomas Ball and Mooly Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT*

- Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 271–282. ACM, 2011. doi: 10.1145/1926385.1926417. URL <http://doi.acm.org/10.1145/1926385.1926417>. (Cited on pages 107 and 149.)
- [33] Xiao Jia, Wei Li, and Viktor Vafeiadis. Proving lock-freedom easily and automatically. In *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP '15*, pages 119–127, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3296-5. doi: 10.1145/2676724.2693179. URL <http://doi.acm.org/10.1145/2676724.2693179>. (Cited on page 132.)
- [34] Cliff B. Jones. Specification and Design of (Parallel) Programs. In *IFIP Congress*, pages 321–332, 1983. (Cited on pages 22 and 29.)
- [35] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 637–650. ACM, 2015. doi: 10.1145/2676726.2676980. URL <http://doi.acm.org/10.1145/2676726.2676980>. (Cited on pages 33, 43, 107, 149, and 150.)
- [36] Ori Lahav and Viktor Vafeiadis. Owicki-gries reasoning for weak memory models. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 311–323. Springer, 2015. doi: 10.1007/978-3-662-47666-6_25. URL http://dx.doi.org/10.1007/978-3-662-47666-6_25. (Cited on page 150.)
- [37] Ruy Ley-Wild and Aleksandar Nanevski. Subjective auxiliary state for coarse-grained concurrency. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 561–574. ACM, 2013. doi: 10.1145/2429069.2429134. URL <http://doi.acm.org/10.1145/2429069.2429134>. (Cited on pages 33 and 107.)
- [38] Hongjin Liang and Xinyu Feng. A program logic for concurrent objects under fair scheduling. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 385–399. ACM, 2016. doi: 10.1145/2837614.2837635. URL <http://doi.acm.org/10.1145/2837614.2837635>. (Cited on page 150.)
- [39] Hongjin Liang, Xinyu Feng, and Ming Fu. A rely-guarantee-based simulation for verifying concurrent program transformations. In John Field and Michael Hicks, editors, *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, pages 455–468. ACM, 2012. doi: 10.1145/2103656.2103711. URL <http://doi.acm.org/10.1145/2103656.2103711>. (Cited on page 149.)
- [40] Hongjin Liang, Xinyu Feng, and Zhong Shao. Compositional verification of termination-preserving refinement of concurrent programs. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting*

- of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014, pages 65:1–65:10. ACM, 2014. doi: 10.1145/2603088.2603123. URL <http://doi.acm.org/10.1145/2603088.2603123>. (Cited on page 132.)
- [41] Henry Massalin and Calton Pu. A lock-free multiprocessor OS kernel (abstract). *Operating Systems Review*, 26(2):8, 1992. doi: 10.1145/142111.993246. URL <http://doi.acm.org/10.1145/142111.993246>. (Cited on page 109.)
- [42] John M. Mellor-Crummey and Michael L. Scott. Algorithms for Scalable Synchronization on Shared-Memory Multiprocessors. *ACM Trans. Comput. Syst.*, 9(1):21–65, 1991. doi: 10.1145/103727.103729. URL <http://doi.acm.org/10.1145/103727.103729>. (Cited on page 23.)
- [43] Maged M. Michael and Michael L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In James E. Burns and Yoram Moses, editors, *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing, Philadelphia, Pennsylvania, USA, May 23-26, 1996*, pages 267–275. ACM, 1996. doi: 10.1145/248052.248106. URL <http://doi.acm.org/10.1145/248052.248106>. (Cited on pages 138 and 141.)
- [44] Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. Communicating state transition systems for fine-grained concurrent resources. In Zhong Shao, editor, *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8410 of *Lecture Notes in Computer Science*, pages 290–310. Springer, 2014. doi: 10.1007/978-3-642-54833-8_16. URL http://dx.doi.org/10.1007/978-3-642-54833-8_16. (Cited on page 33.)
- [45] Gian Ntzik. *Reasoning about POSIX File Systems*. PhD thesis, Imperial College London, 2016. (Cited on page 107.)
- [46] Gian Ntzik, Pedro da Rocha Pinto, and Philippa Gardner. Fault-Tolerant Resource Reasoning. In Xinyu Feng and Sungwoo Park, editors, *Programming Languages and Systems - 13th Asian Symposium, APLAS 2015, Pohang, South Korea, November 30 - December 2, 2015, Proceedings*, volume 9458 of *Lecture Notes in Computer Science*, pages 169–188. Springer, 2015. doi: 10.1007/978-3-319-26529-2_10. URL http://dx.doi.org/10.1007/978-3-319-26529-2_10. (Cited on page 150.)
- [47] Peter W. O’Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3): 271–307, 2007. doi: 10.1016/j.tcs.2006.12.035. URL <http://dx.doi.org/10.1016/j.tcs.2006.12.035>. (Cited on pages 22, 31, and 107.)
- [48] Scott Owens, Susmit Sarkar, and Peter Sewell. A better x86 memory model: x86-tso. In Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel, editors, *Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009, Munich, Germany, August 17-20, 2009. Proceedings*, volume 5674 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2009. doi: 10.1007/978-3-642-03359-9_27. URL http://dx.doi.org/10.1007/978-3-642-03359-9_27. (Cited on page 150.)

- [49] Susan S. Owicki and David Gries. Verifying properties of parallel programs: An axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976. doi: 10.1145/360051.360224. URL <http://doi.acm.org/10.1145/360051.360224>. (Cited on pages 22 and 27.)
- [50] Matthew J. Parkinson and Gavin M. Bierman. Separation logic and abstraction. In Jens Palsberg and Martín Abadi, editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 247–258. ACM, 2005. doi: 10.1145/1040305.1040326. URL <http://doi.acm.org/10.1145/1040305.1040326>. (Cited on pages 36 and 108.)
- [51] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, pages 55–74. IEEE Computer Society, 2002. doi: 10.1109/LICS.2002.1029817. URL <http://dx.doi.org/10.1109/LICS.2002.1029817>. (Cited on page 31.)
- [52] Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. Specifying and verifying concurrent algorithms with histories and subjectivity. In Jan Vitek, editor, *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9032 of *Lecture Notes in Computer Science*, pages 333–358. Springer, 2015. doi: 10.1007/978-3-662-46669-8_14. URL http://dx.doi.org/10.1007/978-3-662-46669-8_14. (Cited on page 33.)
- [53] Filip Sieczkowski, Kasper Svendsen, Lars Birkedal, and Jean Pichon-Pharabod. A separation logic for fictional sequential consistency. In Jan Vitek, editor, *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9032 of *Lecture Notes in Computer Science*, pages 736–761. Springer, 2015. doi: 10.1007/978-3-662-46669-8_30. URL http://dx.doi.org/10.1007/978-3-662-46669-8_30. (Cited on page 150.)
- [54] Kasper Svendsen and Lars Birkedal. Impredicative Concurrent Abstract Predicates. In Zhong Shao, editor, *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8410 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2014. doi: 10.1007/978-3-642-54833-8_9. URL http://dx.doi.org/10.1007/978-3-642-54833-8_9. (Cited on pages 33, 38, 65, 92, 107, 149, and 150.)
- [55] Kasper Svendsen, Lars Birkedal, and Matthew J. Parkinson. Modular reasoning about separation of concurrent data structures. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7792 of *Lecture Notes in Computer Science*, pages 169–188. Springer, 2013. doi: 10.1007/978-3-642-37036-6_11. URL http://dx.doi.org/10.1007/978-3-642-37036-6_11. (Cited on pages 33 and 107.)

- [56] Alfred Tarski et al. A lattice-theoretical fixpoint theorem and its applications. *Pacific journal of Mathematics*, 5(2):285–309, 1955. (Cited on page 126.)
- [57] R Kent Treiber. Systems programming: Coping with parallelism. Technical Report RJ 5118, IBM Almaden Research Center, April 1986. (Cited on pages 120 and 121.)
- [58] Alan M. Turing. Checking a Large Routine. In *Report of a Conference on High Speed Automatic Calculating Machines*, pages 67–69, 1949. URL <http://www.turingarchive.org/browse.php/B/8>. (Cited on page 108.)
- [59] Aaron Turon, Derek Dreyer, and Lars Birkedal. Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. In Greg Morrisett and Tarmo Uustalu, editors, *ACM SIGPLAN International Conference on Functional Programming, ICFP’13, Boston, MA, USA - September 25 - 27, 2013*, pages 377–390. ACM, 2013. doi: 10.1145/2500365.2500600. URL <http://doi.acm.org/10.1145/2500365.2500600>. (Cited on pages 33, 34, 107, and 149.)
- [60] Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. GPS: navigating weak memory with ghosts, protocols, and separation. In Andrew P. Black and Todd D. Millstein, editors, *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20-24, 2014*, pages 691–707. ACM, 2014. doi: 10.1145/2660193.2660243. URL <http://doi.acm.org/10.1145/2660193.2660243>. (Cited on page 150.)
- [61] Aaron Joseph Turon, Jacob Thamsborg, Amal Ahmed, Lars Birkedal, and Derek Dreyer. Logical relations for fine-grained concurrency. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’13, Rome, Italy - January 23 - 25, 2013*, pages 343–356. ACM, 2013. doi: 10.1145/2429069.2429111. URL <http://doi.acm.org/10.1145/2429069.2429111>. (Cited on page 107.)
- [62] Viktor Vafeiadis and Chinmay Narayan. Relaxed separation logic: a program logic for C11 concurrency. In Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes, editors, *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 867–884. ACM, 2013. doi: 10.1145/2509136.2509532. URL <http://doi.acm.org/10.1145/2509136.2509532>. (Cited on page 150.)
- [63] Viktor Vafeiadis and Matthew J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings*, pages 256–271, 2007. doi: 10.1007/978-3-540-74407-8_18. URL http://dx.doi.org/10.1007/978-3-540-74407-8_18. (Cited on pages 31, 38, and 107.)
- [64] Shale Xiong, Pedro da Rocha Pinto, Gian Ntzik, and Philippa Gardner. Abstract Specifications for Concurrent Maps. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Lecture Notes in Computer Science*. Springer, 2017. (Cited on pages 21 and 90.)