

# Semi-Automated Modular Formal Verification of Critical Software

Liveness and Completeness Thresholds

**Tobias Reinhard**

Supervisors:

Prof. dr. B. Jacobs

Prof. dr. A. Timany

Dissertation presented in partial  
fulfillment of the requirements for the  
degree of Doctor of Engineering  
Science (PhD): Computer Science

February 2024



# **Semi-Automated Modular Formal Verification of Critical Software**

Liveness and Completeness Thresholds

**Tobias REINHARD**

Examination committee:

Prof. dr. ir. D. Vandermeulen, chair

Prof. dr. B. Jacobs, supervisor

Prof. dr. A. Timany, supervisor

Prof. dr. ir. F. Piessens

Prof. dr. ir. T. Schrijvers

Prof. dr. M. Huisman

(Universiteit Twente)

Prof. dr. P. Müller

(ETH Zürich)

Dissertation presented in partial  
fulfillment of the requirements for  
the degree of Doctor of Engineering  
Science (PhD): Computer Science

February 2024

© 2024 KU Leuven – Faculty of Engineering Science  
Uitgegeven in eigen beheer, Tobias Reinhard, Celestijnenlaan 200A box 2402, B-3001 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

# Preface

I am grateful to my advisor, Bart Jacobs, and my co-advisor, Amin Timany, for providing me with the opportunity to delve further into the realm of software verification and, especially, for introducing me to separation logic. Your steadfast encouragement and profound expertise have been deeply appreciated, along with all the insightful and interesting discussions we've had. The insights I gained from separation logic have significantly enhanced my understanding of software verification, an experience I deeply value. Furthermore, I appreciate the freedom Bart offered me to cultivate and pursue my own visions while venturing into new and uncharted territories.

I want to express my gratitude to my dissertation committee, Marieke Huisman, Peter Müller, Frank Piessens, and Tom Schrijvers, for their insightful evaluation of my dissertation. Your expertise was invaluable, and I sincerely appreciate the time and effort you dedicated to reviewing my work. Additionally, I thoroughly enjoyed the engaging and enriching discussion of my thesis during the defence. I would also like to thank Dirk Vandermeulen for chairing my dissertation committee.

Besides my work in academia, I had the invaluable opportunity to intern twice at Amazon's Automated Reasoning Group. This experience profoundly impacted not only my personal growth but also my development as a researcher. I want to extend my heartfelt thanks to my mentor, Nathan Chong, and manager, Mark Tuttle, for the engaging discussions and for their thoughtfulness that defined my time in Boston. Additionally, I am grateful to my other managers and colleagues, Mike Whalen, Jim Grundy, Saswat Padhi, Felipe Monteiro, Daniel Kroening, and Michael Tautschnig, for our fruitful collaboration.

I want to thank my friend and collaborator Joachim Kristensen for the joint project we recently started. Working with you has been a pleasure and I am sure that we have some exciting times ahead of us. I am particularly excited about the opportunity to visit you and your group in Oslo soon. Next, I want

to thank all the people who read my works and listened to my pitches numerous times. Note that all the feedback you gave me over the years and all the discussions we had are truly appreciated. While this list is not exhaustive, I want to give particular shout-outs to the following friends: Friedrich Weber for being my go-to model checking expert and for pointing me to the first work on completeness thresholds I ever read. Ramana Nagasamudram for your warm and constructive attitude, especially during those discussions that made me change my mind. Justus Fasse for being the personification of a related work library and for the fun I had during our collaborations. Next, I want to extend my gratitude to my SIGPLAN mentor Stavros Aronis. Thanks for sharing your experience with me and for the insightful perspectives you offered during our meetings.

I am happy to say that, during my time in Leuven, I found the most supportive and inclusive circle of friends and fellow PhD candidates one could wish for (in order of first appearance): Justus Fasse, Nima Rahimi Froushaani, Đorđe Marković, Vida Ranjbar, Niels Mommen, Denis Carnier, Weihong Wang. All of you are truly amazing and have made my experience in Leuven unforgettable. Thank you for your unconditional support and the kindness you showed me during the years. Your friendship has enriched my life in ways I'll always cherish and never forget.

Moreover, I thank my other friends and colleagues in the research community and the department: Hanneli Tavante, Robbe Van den Eede, Federico Quin, Gints Engelen, Fritz Alder, Shirin Kalantari, Pieter Luyten, Gerald Budigiri, Lesly-Ann Daniel, Ignace Bossuyt, Petra Asma, Katrien Janssens, Annick Vandijck, and An Makowski.

I am deeply thankful to my parents, Erwin and Karin, and my brother, Dominik Reinhard, for their unwavering support and enduring belief in my capabilities. Their steadfast encouragement has been invaluable, and their continuous support has played a crucial role in helping me pursue and achieve my goals throughout my journey.

Coming to an end, I just want to give a big shout-out to my partner, Carina. Your support over the years has been nothing short of incredible. Thank you for always being there to listen to my possibly overly detailed descriptions of random verification topics. Anyways, without you, none of this would have been possible! And now, I can't wait to see what adventures await us in the next chapter of our journey together! Schau'n mer mal was wird.

# Abstract

In this dissertation we describe two contributions to the state of the art in reasoning about liveness and safety, respectively.

Programs for multiprocessor machines commonly perform busy waiting for synchronization. We propose the first separation logic for modularly verifying termination of such programs under fair scheduling. Our logic requires the proof author to associate a ghost signal with each busy-waiting loop and allows such loops to iterate while their corresponding signal  $s$  is not set. The proof author further has to define a well-founded order on signals and to prove that if the looping thread holds an obligation to set a signal  $s'$ , then  $s'$  is ordered above  $s$ . By using conventional shared state invariants to associate the state of ghost signals with the state of data structures, programs busy-waiting for arbitrary conditions over arbitrary data structures can be verified.

Moreover, we present the first study of completeness thresholds for bounded memory safety proofs. Specifically, we consider heap-manipulating programs that iterate over arrays without allocating or freeing memory. In this setting, we present the first notion of completeness thresholds for program verification which reduce unbounded memory safety proofs to bounded ones. Furthermore, we demonstrate that we can characterise completeness thresholds for simple classes of array traversing programs. Finally, we suggest avenues of research to scale this technique theoretically, i.e., to larger classes of programs (heap manipulation, tree-like data structures), and practically by highlighting automation opportunities.





# Beknopte samenvatting

In dit proefschrift beschrijven we twee bijdragen aan de state-of-the-art in redeneren over respectievelijk liveness en safety.

Programma's voor multiprocessor machines maken gebruik van busy waiting voor synchronisatie. Wij stellen de eerste separation logic voor om modulair de beëindiging van zulke programma's onder fair scheduling te verifiëren. Onze logica vereist dat de auteur van het bewijs een ghost signal toekent aan elke busy waiting lus. Deze lussen mogen itereren zolang hun corresponderende signaal  $s$  niet is ingesteld. De auteur van het bewijs moet daarnaast een welgefundeerde orderrelatie op signalen definiëren, en moet daarbovenop bewijzen dat als de looping thread een obligatie van een signaal  $s'$  bezit, dat  $s'$  groter is dan  $s$ . Door klassieke gedeelde toestandsinvarianten te gebruiken om de toestand van ghost signals te associëren met de toestand van datastructuren, kunnen programma's met busy waiting geverifieerd worden, voor arbitraire condities over arbitraire datastructuren.

Daarnast presenteren we de eerste studie naar completeness thresholds voor bounded memory safety bewijzen. Specifiek beschouwen we heap-manipulerende programma's die itereren over arrays zonder het toekennen of vrijgeven van geheugen. In deze context presenteren we de eerste notie van completeness thresholds voor programmaverificatie die unbounded memory safety bewijzen reduceren tot bounded bewijzen. Tevens laten we zien dat we completeness thresholds kunnen karakteriseren voor eenvoudige klassen van programma's die arrays doorlopen. Tot slot suggereren we theoretische onderzoekspistes voor grotere klassen van programma's (heapmanipulatie, boomachtige datastructuren), en praktische mogelijkheden voor bewijsautomatisatie.



# Contents

<b>Abstract</b>	<b>iii</b>
<b>Beknopte samenvatting</b>	<b>v</b>
<b>Contents</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 State of the Art</b>	<b>9</b>
2.1 Deductive Liveness Verification of Concurrent Software . . . . .	9
2.2 Reducing Unbounded to Bounded Proofs . . . . .	12
<b>3 Ghost Signals: Verifying Termination of Busy Waiting</b>	<b>17</b>
3.1 Introduction . . . . .	17
3.2 A Guide on Verifying Termination of Busy Waiting . . . . .	19
3.2.1 Simplest Setting: Thread-Safe Physical Signals . . . . .	20
3.2.2 Non-Thread-Safe Physical Signals . . . . .	25
3.2.3 Arbitrary Data Structures . . . . .	32
3.2.4 Signal Erasure . . . . .	33
3.3 A Realistic Example . . . . .	35
3.4 Specifying Busy-Waiting Concurrent Objects . . . . .	39
3.5 Tool Support . . . . .	40
3.6 Integrating Higher-Order Features . . . . .	41
3.7 Related & Future Work . . . . .	41
3.8 Conclusion . . . . .	42
<b>4 Completeness Thresholds for Memory Safety of Array Traversing Programs</b>	<b>43</b>
4.1 Introduction . . . . .	43
4.2 Limitations of Bounded Proofs . . . . .	44
4.3 Completeness Thresholds . . . . .	46

4.3.1	Approximating CTs via Verification Conditions . . . . .	47
4.3.2	Modularity of Completeness Thresholds . . . . .	52
4.4	Conclusion . . . . .	54
4.5	Related Work . . . . .	55
4.6	Outlook . . . . .	55
<b>5</b>	<b>Conclusion and Future Work</b>	<b>57</b>
5.1	Liveness Verification . . . . .	57
5.2	Completeness Thresholds . . . . .	60
	<b>List of Contributions</b>	<b>65</b>
<b>A</b>	<b>Ghost Signals: Formalization &amp; Case Studies</b>	<b>69</b>
A.1	General . . . . .	70
A.2	Language . . . . .	71
A.3	Logic . . . . .	74
A.4	Soundness . . . . .	81
A.4.1	Annotated Semantics . . . . .	81
A.4.2	Hoare Triple Model Relation . . . . .	90
A.4.3	Soundness Proof . . . . .	92
A.5	Case Studies . . . . .	99
A.5.1	Verification of Realistic Example . . . . .	100
A.5.2	Case Study: Statically Unbounded Number of Communi- cating Parties . . . . .	107
<b>B</b>	<b>Ghost Signals: Generalised Logic</b>	<b>115</b>
B.1	Universe . . . . .	115
B.2	General . . . . .	116
B.3	Syntax . . . . .	117
B.4	Example . . . . .	118
B.5	Resources . . . . .	118
B.6	Semantics . . . . .	119
B.7	Assertions . . . . .	121
B.8	Proof Rules . . . . .	125
B.9	Annotated Semantics . . . . .	127
B.10	Hoare Triple Model Relation . . . . .	137
B.11	Soundness . . . . .	138
B.12	Verification Example . . . . .	147
B.12.1	Minimal Example . . . . .	148
B.12.2	Bounded FIFO . . . . .	148
<b>C</b>	<b>Completeness Thresholds: Formalization</b>	<b>163</b>
C.1	Introduction . . . . .	163

C.2	General Notation and Basic Definitions . . . . .	164
C.3	Syntax . . . . .	166
C.4	Dynamic Semantics . . . . .	167
C.5	Assertion Language . . . . .	169
C.6	Memory Safety . . . . .	172
C.7	Verification Conditions . . . . .	173
C.7.1	VC Generation . . . . .	174
C.8	Completeness Thresholds . . . . .	175
C.8.1	Iteratively Extracting Completeness Thresholds . . . . .	178
C.8.2	Iterating over Arrays . . . . .	179
<b>D</b>	<b>Generalising Completeness Thresholds</b>	<b>181</b>
D.1	Basic Definitions . . . . .	181
D.2	Generalised Completeness Thresholds . . . . .	182
	<b>Bibliography</b>	<b>189</b>



# Chapter 1

## Introduction

Critical software is everywhere and almost every piece of critical infrastructure depends on it, in one way or another. The electrical power we use every day is produced by power plants whose safety heavily depends on the correctness of a variety of software systems. For instance, in case of an emergency dedicated *shutdown systems* [113, 179] are responsible to automatically power down reactors and prevent disaster. But also a power plant's normal operation is heavily software controlled. Specifically, they rely on sophisticated networks of hardware sensors and software components, called *supervisory control and data acquisition* (SCADA) systems [95]. These collect data and make automatic decisions about the plant's operation or assist human decision makers. Their correctness is critical to any power plant's safety. Moreover, the complex networks that supply our homes every day with electricity, gas and water are also operated by SCADA systems [171]. They can typically be accessed remotely, which makes them vulnerable to cyber attacks exploiting bugs [171, 95].

To a similar extent, we critically rely on software when we travel. Many of the safety features that cars offer nowadays are provided by software [140]: ABS (anti-lock braking system) prevents wheels from locking during emergency braking. ACC (adaptive cruise control) ensures a safe distance to the vehicle ahead. ESC (electronic stability control) is yet another service helping the car to stay on track by controlled breaking. Just to name a few.

Aircrafts use onboard software to control about every aspect of their journey. They navigate via GPS software [160]. Dedicated software controls the slats and flaps – the parts of the wing used to alter or keep the altitude [31]. The most critical part of any flight, the landing, is also computer guided [160]. Faults in any of these programs can have critical consequences. The two groundings of

Boeing 737 MAX planes in October 2018 and March 2019 should serve as a sad reminder of that [89].<sup>1</sup>

**Correctness Properties** It is easy to agree that correctness is essential for critical systems. Though, what exactly it means for software to be correct differs from case to case. Two important classes of correctness properties are *safety* and *liveness* [111, 4].

**Safety** We call a program *safe* if something bad never happens. For instance, a program is memory safe if it never accesses invalid memory (e.g. unallocated or already freed memory). This is typically the most basic property that we want for every program, since we need it to ensure any other correctness property. A very common class of memory safety errors are array out-of-bounds accesses, where we read or write memory past the array bounds. This type of error is often exploited in cyber attacks [96, 48, 75].

Another safety property that we commonly care about in concurrent applications is data race freedom [11]. A data race occurs when a thread (non-atomically) writes a memory cell (or any other shared resource) concurrently with another thread writing or reading it. Whenever that happens the outcome is undefined or non-deterministic if the involved operations are atomic. In the worst case, a concurrent write can corrupt the memory cell. If the subsequent program depends on the memory cell's value, its behaviour becomes unpredictable. Even worse, depending on the concrete programming language used, the semantics of the entire program might become undefined.

**Liveness** Liveness is the counterpart to safety. We call a program *live* if something good eventually happens. For instance, we often want computations to terminate, because we care about their result. Non-terminating computations block the entire subsequent program. Hence, termination is a commonly desired liveness property, even though the name liveness suggests otherwise.

Imagine a distributed banking system processing monetary transactions. For each transaction, it would debit the specified amount from the sending account, send it to the recipient's institution and then add it to their account. One of the most important properties of such a system is that the transferred money eventually arrives on the recipient's account, which is also a liveness property.

---

<sup>1</sup>We should note that the software was only one of multiple flaws that ultimately led to the groundings.



**Approaches to Correctness** Over the years the community came up with different approaches to reduce the number of bugs and thereby improve the correctness of software systems. It remains, however, very difficult to reason about programs. Many interesting properties are undecidable, e.g., memory safety [157], deadlock freedom [80]<sup>2</sup> and termination [156]. In a nutshell this means that for sufficiently complex programs it is not possible to reason precisely and automatically about every possible execution.

**Testing** Test driven development [78], i.e., writing tests first and implementing the actual functionality later, has proven valuable in finding bugs early on. However, as Dijkstra said: “Program testing can be used to show the presence of bugs, but never to show their absence!” [60] Even for relatively simple programs, the number of tests we had to write to cover every possible scenario, would be infeasibly large. Approaches like property based testing [33] and fuzzing [129] try to tackle this issue by automatically generating test cases and thereby increasing coverage. However, due to the sheer number of possible inputs and runtime behaviours, even automatic testing is not sufficient to guarantee that we did not miss any bug. The situation gets even worse when we move from sequential, single-threaded software to concurrent, multi-threaded computations. In the latter case, we not only have to explore all possible inputs but also all possible thread interleavings to ensure that our software works as expected. By that the number of required test cases explodes exponentially. Meanwhile, the main advantage of testing is its simplicity.

**Bounded Model Checking** Model checking [37, 146] is an alternative approach to automatically find bugs by exhaustively simulating every possible program execution. Since programs do not necessarily terminate, neither does model checking them. Moreover, the number of different states that the program might be in and which we have to check grows exponentially with the complexity of the program, the number of possible inputs and the length of the executions. This is known as the state space explosion problem [27, 36, 142, 35] and it makes straight-forward model checking unfeasible, even for terminating programs.

Bounded model checking (BMC) [15] overcomes this issue by bounding both the length of executions and the input size. That is, when we bounded model check a program sorting lists of integers, we specify an input size bound  $S$  and a depth bound  $D$ . Then the model checker simulates all executions for lists with up to  $S$  elements and it stops the simulation after  $D$  program steps.

---

<sup>2</sup>The cited work concerns communicating finite state machines, but the work applies straightforwardly to deadlock freedom in concurrent programs.

Compared to testing, BMC is a much more exhaustive approach to ensure program correctness. Over the years, BMC has become an established technique in industry, though still far less widespread or well-known than testing. At this point we have industrial grade bounded model checkers like CBMC [39] and ESBMC [46] that are relatively easy to use. Writing a CBMC proof for instance requires little more than specifying the bounds  $S$  and  $D$  as well as which correctness properties the model checker should look for. The rest happens automatically. Since CBMC simulates concrete executions, it can report a counterexample for every found bug. Just like with testing, this makes it easy for the developer to understand and fix the problem.

The main selling point of bounded model checkers is that they manage to strike a compromise between usability and exhaustiveness. It achieves this by under-approximating the program behaviour, i.e., exploring a (small) subset of finite prefixes of concrete program executions. However, as with all under-approximating approaches it does not suffice to ensure that the program we checked is correct. Maybe we missed a bug that only shows up for bigger inputs or longer program executions. To make things worse, the state space explosion makes it usually only feasible to check small inputs.

**Deductive Verification** The approaches we discussed so far focus on finding bugs. With every bug we find and fix our program converges a bit more towards correctness. With these approaches we can, however, never be certain that no bug is hiding in a corner that we didn't explore. One way to overcome this limitation is deductive verification [173, 92]. The underlying idea is that we analyse the program and use mathematical reasoning to rule out every possible error. When this succeeds, we know for sure that the program is correct. Hence, it is the most powerful correctness approach among the ones discussed here.

So why do we bother with testing or model checking at all? For once, to validate the scenarios considered by our proof. Yet also because proofs are hard! In order to logically reason about the behaviour of a program we have to understand how it works and – even worse – why it should be correct. That is, we have to manually come up with a correctness argument in the form of specifications, invariants and inductive proofs. The deductive verifier then checks that every proof step we took is correct and that our argument holds. This requires not only a detailed understanding of the software's inner workings but also expertise in the logic the deductive verifier is built upon. This makes deductive verification hard to access for developers without any background in formal logic. In contrast, testing and BMC allow us to more or less treat the program as a black box.

Another difference is that a deductive verifier never executes the program. That

means that we never look at any concrete program state in an execution. Instead we deduce correctness from the program's static structure. Our correctness argument must be general enough to cover all possible executions. This often forces us to over-generalise in the sense that our argument also covers impossible executions, especially for loops and recursion. As a result, it is not possible to verify every correct program.

Nevertheless, deductive verification (complemented by validation testing) is the best approach that we have to truly ensure correctness of complex critical systems. For some applications failure is not an option and the stakes are high enough to warrant the tedious correctness proofs. Over the years, many research tools such as VeriFast [99], VerCors [18], Viper [132], Why3 [70], and VCC [43] have been developed. While their focus is primarily academic, some are also used in industrial projects such as the verification of FreeRTOS [153, 55] or cloud infrastructure [44]. Moreover, a small number of industrial grade tools such as Dafny [116] and Adacore Spark [125] exist with a long track record in critical areas like cryptography and security [44, 158, 1, 86, 68]. Meanwhile, works like loop invariant synthesis [145, 117] and – recently – the application of large language models [185, 183] try to make deductive verification simpler and thereby more accessible to the masses.

**Advantages and Disadvantages** Testing is a good way to find bugs early in the development cycle. It requires little expertise and the results are easy to understand. But given the very limited guarantees it entails, testing alone is clearly insufficient to guarantee correctness of any critical system. Yet, it complements deductive verification well and serves as an easy way to validate proof assumptions.

In contrast, BMC provides much stronger guarantees. Nonetheless, it succeeds in preserving the main advantages of testing: It is fairly easy to apply. Especially because it allows us to treat most of the target program as a black box. And it yields real, i.e., non-spurious, counterexamples when it fails. The combination makes it easy to find and fix bugs. Its main drawback is that BMC proofs are unsound, because they only explore a subset of all inputs and only bounded executions. Consequently, no matter how big we set the bounds, we can never be certain that no bug hides in some dark corner we did not explore.

Deductive verification, in contrast, is sound and yields unbounded guarantees. This makes it an appropriate approach for critical systems, where we cannot tolerate any failure. Yet, it should be complemented by validation, e.g., testing, to ensure the proof considers the right scenarios. Deductive proofs are hard, though, and require significant expertise. At the same time, the approach is incomplete. That is, many programs are impossible to verify even though they

are fully correct. Hence, choosing the right correctness approach is a trade-off between the degree of guarantees we desire and resources we are willing to spend.

**Termination of Busy Waiting** In concurrent programs responsibilities are often distributed across threads. The same holds for dependencies. It is very common for a thread  $A$  to depend on data produced by another thread  $B$ . Both threads can run concurrently until  $A$  arrives at a point where it needs to access said data. Whenever that happens  $A$  has to wait until the data arrives. That is, the thread is blocked until  $B$  finishes the computation it is responsible for. We call this pattern *blocking concurrency*. One possible implementation that we focus on in this dissertation is for  $A$  to loop until the awaited data arrives and to check the status in every iteration. We call this *busy waiting* [128].

Such a system can easily get stuck if one thread fails to adhere to its responsibility or if we accidentally end up with cyclic dependencies. Ensuring that this not the case, i.e., that the system is *live*, is far from trivial. Even more so to formally prove it live. Every such proof has to show that  $A$ 's wait loop eventually terminates. To achieve that, we can first tie  $A$ 's busy waiting to the progress of  $B$ . In a next step, we can prove that  $B$  eventually produces the data  $A$  is waiting for.

We see that reasoning about whether or not  $A$ 's busy waiting eventually terminates, forces us to take  $B$ 's progress into account. Hence, reasoning about such liveness properties and dependencies is difficult to achieve modularly, i.e., thread-locally. In § 3 we present the first work that allows us verify termination of busy waiting modularly. We present a separation logic [155, 138] and a new type of ghost resource [103] that we call *ghost signal*.

Ghost resources are objects that do not exist during the runtime of a program. Instead we only introduce them for the sake of our proof. Naturally, we have to make sure that they can be introduced and erased without affecting the program's semantics or control flow. Intuitively, a ghost signal is a witness of some other thread promising to do something. For instance, suppose thread  $A$  possesses a ghost signal tied to the data  $A$  is waiting for. Then it knows that another thread promised to produced said data. Hence, we can use it in a thread-local proof to justify that  $A$ 's busy waiting eventually terminates.

**Completeness Thresholds** BMC is unsound by construction since there's always a chance that we chose our bounds too small to find all bugs. While this is obviously a big disadvantage, it is also where BMC's usability comes from. In general, bounded proofs such as the one that we see in BMC are easy to automate

since they only explore a finite set of finite, concrete program executions. Also for that reason, failed proofs yield non-spurious counterexamples. The latter help developers by showing them precisely under which circumstance the bugs occur. The same does not hold for unbounded proofs like the ones we get with deductive verification.

BMC proofs are often justified by the intuition that if there is a bug in a program, then it also tends to show up for small inputs and early in an execution. The problem is that we do not know when this intuition is true and when it is not. If we knew that it was correct for certain programs, it would allow us to turn bounded proofs into unbounded ones.

Though, we know at least that this intuition holds for bounded proofs of specific properties  $\phi$  over finite state transition systems  $T$  [15, 40, 28]. To formalise it, the model checking community came up with the concept of *completeness thresholds* (CTs) [40]: Intuitively, such a threshold is a number  $k$  that is large enough such that if  $T$  violates  $\phi$ , then there is also a violation that can be reached with at most  $k$  steps. In other words, if there is a bug, we can find it by exploring all finite path prefixes up to the length  $k$ .

Any bounded proof that exceeds the threshold is sound. From a verification point of view it would make more sense to call it a *soundness threshold*. Unfortunately, the model checking and verification communities use the terms *sound* and *complete* in opposite ways. For the verification community, soundness means that any proven property indeed holds. The model checking community calls this completeness, because the proof did not miss any bugs.

In § 4 we present the first notion of CTs for program verification. In particular, we define them in the context of memory safety proofs and show that we can extract them for array traversing programs. The programs we study do not correspond to finite transition systems. Hence, we cannot reuse the existing results or approaches to CTs from the model checking literature. Though, to highlight the connection to previous CT research in the model checking community, we chose to continue the terminology instead of calling them soundness thresholds. Yet, apart from the term CT we continue to use *sound* and *complete* in the verification sense.

For us, a CT is not a number  $k$ . Instead it is a subdomain  $Q \subseteq X$  of an input domain  $X$  that is big enough to ensure that a bounded proof only exploring  $Q$  covers all relevant cases. Suppose we want to prove memory safety of a program that iterates over an array of size  $s \in \mathbb{N}$ . How can we know whether it is sufficient to check memory safety for arrays of size  $s < 10$ ? In § 4 we show that we can extract that knowledge from the program's verification condition and thereby reduce unbounded memory safety proofs to bounded ones.



# Chapter 2

## State of the Art

### 2.1 Deductive Liveness Verification of Concurrent Software

We can broadly classify major deductive approaches to reason about concurrent software along two dimensions: (i) The abstraction level and (ii) the foundation.

**Abstraction Level** Most popular approaches [135, 139, 184, 176, 177, 104, 122], like the one we propose in § 3, work directly with the target program and reason about the effects of every single instruction. A less widespread alternative is contextual refinement. The idea is to reason about a more abstract model instead of the actual program we want to verify. In this context, a refinement is a program transformation that preserves the properties we are interested in and at the same time abstracts implementation details away. Refinement-based works include Turon et al. [175], Gu et al. [82], Boussabbeh et al. [20], de Carvalho Gomes et al. [51]. Works that employ a mixture of both approaches include: CaReSL [174], LiLi [120, 121], Conditional Contextual Refinement [166].

**Logical Foundation** The second dimension is the foundation we choose to base our reasoning on. While the landscape of logical foundations is vast, most fall into one of two broad categories: (i) Type-based reasoning and (ii) program logics à la Floyd-Hoare.

(i) The foundation for all type-based approaches is the Curry-Howard correspondence which equates logical propositions  $P$  with types  $T$  and consequently proofs of  $P$  with values of type  $T$ . Dependent type theories like the calculus of constructions [45] form the basis of popular state-of-the-art proof assistants like Coq [54] and Lean [57, 52]. Due to their expressiveness, they mostly serve as a trusted foundation to mechanise and check higher-level reasoning. For instance, Iris [104] and Verifiable C/VST-Floyd [8, 29] are both higher-order separation logics mechanised in Coq. The expressiveness of dependent types comes at the cost of automation. Meanwhile refinement types are sufficiently restricted such that assertions/typings can be automatically checked by SMT solvers. One of the most notable projects using refinement types for automatic program verification is Liquid Haskell [178]. More recent works include Flux [114] to verify Rust code and CN [144] which combines refinement types and separation logic to verify C programs.

(ii) The central idea of Hoare-logic is that we specify the behaviour of a program in the form of pre- and post conditions. This simple paradigm proved very effective to reason about imperative programs. Since its introduction, many successors have been proposed [135, 139, 184, 176, 177, 104, 122]. One notable step along the evolution of Hoare-style logics is the Owicki-Gries method [139], which introduced some of the concepts we still use today, such as ghost variables, critical sections and resource invariants. Another notable, though not very wide-spread, type of Hoare-style logics is rely-guarantee reasoning [184] and its successor deny-guarantee reasoning [62].

Most notably, separation logic [155, 138] became the de-facto standard to reason about heap-manipulating programs. Dynamic frames [105] and implicit dynamic frames [165] are conceptually similar and used in verifiers such as Dafny [116], though certainly less wide-spread than separation logic. In particular, concurrent separation logic [155, 138] is now the de-facto standard to reason about concurrent, imperative programs. The key idea is fine-grained reasoning about ownership of resources. This is also the foundation for the approach we present in § 3. Over the years, separation logic has been extended by many concepts to address the challenges that arise when we try to reason about concurrent programs. Higher-order ghost state [103] makes very expressive invariants and specifications possible. Fractional permissions [22, 26] allow us to distribute permissions among threads in a fine-grained manner that for instance differentiates between read and write permissions.

**Proving Liveness** Hoare-style logics – and in particular separation logic – are by construction safety logics. There are, however, several extensions that allow us to prove liveness properties by reducing them to safety ones. Obligations [107]



allow us to prove finite blocking by reasoning about responsibility distribution among threads. Obligations are a very versatile concept that can be used to reason about many different forms of blocking, e.g., acquisition of locks [19, 118], reading from channels [19, 118], waiting for a heap cell to be set (§ 3). Call permissions [100] support reasoning about termination of loops and function calls by forming a decreasing measure.

**Non-Blocking Concurrency** We can divide the landscape of concurrent programs into blocking and non-blocking ones. In case of the former, threads' progress depends on that of other threads. Non-blocking ones are in contrast programs where different threads cannot block each other's progress. Several previous works such as Moir and Anderson [130], Hurault and Quéinnec [94], Gotsman et al. [79], da Rocha Pinto et al. [50] focus on verifying liveness properties of non-blocking algorithms. Notably, Total-TaDA [50] is a total correctness logic that supports proving termination of concurrent, non-blocking programs. Moreover, it also allows us to prove non-impedance, i.e., that threads cannot prevent each other's progress.

**Blocking Concurrency** Boström and Müller [19] propose a technique to verify blocking behaviour in non-terminating, concurrent programs. They consider blocking behaviour that stems from primitive blocking operations: Reading from built-in channels, acquiring built-in locks, joining another thread, i.e., waiting until it terminates. Their approach is based on several techniques that we also reuse for the work we present in § 3. In particular, they use obligations, decreasing measures and levels to prove that waiting terminates.

A non-primitive form of blocking behaviour is busy-waiting, where a thread waits by looping until the desired event occurs. Liang and Feng [120, 121] propose LiLi, a separation logic to verify liveness of blocking constructs implemented via busy waiting. Their proofs rely on contextual, liveness preserving refinement. Ghost signals (§ 3) allow the sound generation of fuel to justify indefinite waiting in a context that provably leads to termination. Tada Live [64], a separation logic published in parallel to our ghost signal paper, also allows to prove termination of busy waiting loops. We discuss it in more depth in § 5.

Condition variables are another way to implement blocking behaviour. Threads can register to wait for access to a specific resource and are woken up when the resource becomes available. In their paper de Carvalho Gomes et al. [51] propose technique to verify liveness of Java programs that synchronise via condition variables. The liveness property they consider is that each thread which enters a critical section eventually leaves it, which is essentially a termination property. They reduce termination to reachability in coloured Petri Nets. The latter can

be checked by existing Petri Net analysis tools. Hamin and Jacobs [84] show that we can also use obligations to reason about the liveness of code that uses condition variables for synchronisation.

## 2.2 Reducing Unbounded to Bounded Proofs

The fundamental idea behind model checking is to express a problem as property of a transition system and then exhaustively check that it holds for all paths. This is often unfeasible for two reasons: (i) The system contains infinite paths that can never be fully checked. (ii) The state space explosion problem [27, 36, 142, 35]. Transition systems grow very fast with the complexity of the original problem we care about. Since our computational resources are limited, we can often only check a small prefix of the entire system.

The model checking community came up with two fundamentally different approaches to overcome this limitation: (i) Heavily reducing the state space via abstractions and (ii) approximating unbounded proofs by bounded ones.

**Simplification by Abstraction** *Predicate abstraction* [81] over-approximates concrete states in a finite transition system by abstract states expressed in terms of hand-crafted predicates. Coming up with suitable predicates is hard and leads to spurious counterexamples [12]. To counter this issue, Clarke et al. [34] proposed *counterexample-guided abstraction refinement* (CEGAR). Each found counterexample is checked against the concrete transition system to determine whether it is spurious. When that is the case, we automatically refine our abstractions to ensure that the refined abstract transition system does not permit the spurious counterexample anymore. A special form of this approach is *lazy abstraction* [88] where we limit the refinement to states where the previous abstraction yields a spurious counterexample. This reduces the size increase inherent to abstraction refinement. Another abstraction refinement approach worth mentioning is *property-directed reachability* (PDR) [24, 66, 93]. While both CEGAR and PDR use counterexamples to refine state abstractions, a major difference is that PDR preserves a precise transition relation [93].

**Approximating Proofs** Bounded model checking (BMC), i.e., approximating unbounded proofs by bounded ones, was originally introduced by Biere et al. [15] and is now a wide-spread technique. State-of-the-art software-model checkers such as CBMC [39] and ESBMC [46] use this paradigm. Instead of aiming for an exhaustive check, they only focus on a small prefix of the transition system. One way is to keep the transition system as is, but to set a depth bound that

limits the length up to which path prefixes are explored [15]. Checkers typically also allow for more fine-grained bounds, such as separate bounds for specific input sizes and loop depths. Fine-grained bounds essentially allow the user to trim specific parts of the transition system.

**Completeness Thresholds** While this approach is very popular, it only gives us bounded guarantees that can not straight-forwardly be extrapolated to the entire transition system. However, in the very paper introducing BMC Biere et al. [15] already observed that finite, lasso-shaped prefixes (i.e., prefixes ending in a backloop) can represent the behaviour of infinite lasso-shaped paths. Thus they can serve as finite witnesses of infinite behaviour. This might sound profane at first glance, but it is actually a very fundamental insight. As Biere et al. showed, it allows us to extract unbounded guarantees from bounded proofs, given the right circumstances. The concept was later popularised by Clarke et al. [40] under the term *completeness threshold*. Consider a finite transition system  $T$  and a property of interest  $\phi$ . We call a number  $k \in \mathbb{N}$  a *completeness threshold* (CT) for  $T$  and  $\phi$  if the following holds: If  $\phi$  holds for all path prefixes up to length  $k$ , then  $\phi$  holds for all paths in  $T$ , i.e.,  $T \models_k \phi \Rightarrow T \models \phi$ .

Notice that the term CT is used inconsistently in the literature. While Clarke et al. [40] use the definition above, Kroening and Strichman [109] refer to a CT as the minimal  $k$  with  $T \models_k \phi \Rightarrow T \models \phi$ . We stick to the original definition.

Since their introduction, CTs have been studied extensively for finite state transition systems [15, 40, 109, 28, 2, 87, 10, 126, 108]. Finding minimal CTs is in general just as hard as checking the unbounded system [110]. Moreover, the question when we can soundly extrapolate a bounded proof to the entire system highly depends on the property  $\phi$  we are trying to prove. Hence, the research focuses on finding over-approximations specific to certain types of formulas. Existing works characterise CTs in terms of key properties of the transition system  $T$ . For instance, the diameter is a CT for the simple LTL safety properties  $Gp$  (i.e.  $\Box p$ ), where  $p$  is local [15, 110]. Meanwhile the recurrence diameter is a CT for simple LTL liveness properties  $Fp$  (i.e.  $\Diamond p$ ) [110]. In this context, the diameter of  $T$  is the maximal distance between two nodes in  $T$  and the recurrence diameter is the length of the longest loop-free path in  $T$ . That is, whenever we want to model check a simple safety or liveness property, it suffices to run a bounded proof up to the diameter or recurrence diameter. That way, we can get sound unbounded guarantees from a bounded proof.

That sounds great, but we should note that both CTs pose very large bounds, which are typically too big to explore in practice. Both are often exponential in the number of state variables in  $T$  [108]. Moreover, since both are hard to compute, they are typically over-approximated as well [108, 141, 17], leading

to even larger CTs. They are also worse-case tight in terms of the chosen key property [110]. Hence, we cannot hope to find smaller CTs for  $Gp$  or  $Fp$  that are expressible purely in terms of  $T$ 's diameter or recurrence diameter. To make things worse, CTs for more complex LTL formulas have much larger CTs that can be quadratic, super-polynomial and even exponential [28, 110].

An interesting combination of CTs and CEGAR is [Sinha et al.](#)'s *Short-Small-World* ( $S^2W$ ) work [164]. They use counterexamples to determine which parts of their state should be represented precisely. The rest can be abstracted away, which makes the system smaller and easier to model check. Finally, they compute an over-approximation of the reduced system's reachability diameter and use it as CT to bounded model check their system. If BMC fails, they collect the found counterexample and return to the CEGAR phase. The CT guarantees  $S^2W$ 's soundness. In addition, it also guarantees "short" counterexamples (relatively speaking), making them easier to process in the CEGAR phase.

All existing, non-abstracting works on CTs focus on finite state transition systems. Software that processes arbitrary large data structures, e.g., arrays or lists, corresponds, however, in general to infinite state transition systems. Hence, the key properties we typically use to describe CTs, like the (recurrence) diameter, are in general infinite. The same holds thereby for the existing CT approximations. In other words, we cannot reuse existing results to reason about CTs for software that processes arbitrary large data. In § 4 we present the first notion of CTs for memory safety proofs and show how to extract them from the program's verification condition. In this work we consider array-traversing programs and focus on CTs for input variables, in particular on CTs for the size of the traversed array. Hence, for us, CTs are not depth bounds  $k \in \mathbb{N}$  but subdomains. Let  $\{A\} c \{B\}$  be a program specification and let  $x$  be an input variable with domain  $X$ . A subdomain  $Q \subseteq X$  is a CT for  $x$  in  $\{A\} c \{B\}$  if correctness of the specification on  $Q$  implies the unbounded specification, i.e.,

$$\models \forall x \in Q. \{A\} c \{B\} \quad \Rightarrow \quad \models \forall x \in X. \{A\} c \{B\}.$$

In § 4 we illustrate the concept by characterising CTs for a simple array traversal pattern. Moreover, in § 4 and § 5 we also sketch plans to characterise CTs for complex programs that traverse a data structure (array, list, etc.) in a memory-layout-preserving way. We also sketch plans our approach to arbitrary correctness properties.

**Reasoning about Coverage** Intuitively, a subdomain  $Q \subseteq X$  is a CT if it contains the right inputs to reach (and refute) all potentially problematic behaviours of our program. In this sense, reasoning about CTs means reasoning about a specific form of coverage. Coverage also plays a big role in traditional

testing techniques, like fuzzing [129] and property-based testing [33]. Compared to hand-written tests, automated approaches make it easier to reach a higher coverage due to the bigger number of tests we can generate. Just as in the context of BMC and CTs, the higher the coverage the stronger the guarantees that our tests convey. The various notions of code coverage that we consider in testing, e.g., method coverage, statement coverage, path coverage [6, 91], are much weaker than what we need to reason about CTs. Furthermore, existing testing works do not care about providing formal guarantees for the coverage they achieve. Recently, Zhou et al. [186] proposed *coverage types* to reason formally and automatically about coverage of test generators. Their approach combines refinement types [178] and incorrectness logic [137] and supports sophisticated structural notions of coverage.

**Cut-Offs, Small Worlds, Small Models** All of these terms are equivalent to CTs but are mostly used in a different model checking subcommunity than CTs. An exception that proves the rule is the *Small-Short-World* work by Sinha et al. [164]. Many distributed protocols can be represented as a parameterised system  $P^n = P_1 \parallel \dots \parallel P_n$  consisting of  $n$  identical and parallel processes  $P_i$  that each can be represented by the same finite state transition system. Intuitively, a distributed protocol should not behave wildly different for every possible number of processes  $n$ . In other words, there should be a finite bound  $K$  big enough such that  $K$  running processes  $P_1 \parallel \dots \parallel P_K$  exhibit all interesting properties [134]. Emerson and Namjoshi [67] formalised this intuition in the context of token rings [134]. Finding a small cut-off  $K$  then allows us to represent the finite system  $P^K = P_1 \parallel \dots \parallel P_K$  in one finite transition system and use classic model checking to reason about it. Note that this setting is fundamentally different from the one considered by classic CTs à la Biere et al. [15] and Clarke et al. [40]. In particular, it differs fundamentally from the setting we consider for our CT work in § 4.

The technique is also known as *small model* or *small world* approach [164]. Since its introduction, it has proven a valuable technique to simplify reasoning about a range of topologies, e.g., token rings, hybrid automata networks [102], parameterised networks of timed automata [167], general composed network topologies [41, 5]. Furthermore, it has since been applied to other domains where systems can be represented by parameterised transition systems. For instance, several works use cut-offs to simplify the verification of security properties in hypervisors [73, 74, 164].

Franklin et al. [73] consider the verification of security properties concerning page table accesses. They focus on a restricted setting that yields nice cutoffs by construction: A guarded programming language and a restricted LTL for the

specification of security properties  $\phi(n)$ , where  $n$  refers to the number of page table entries. Page table accesses are modelled via array accesses. Neither their programming nor specification language support programs or specifications  $\phi(n)$  that differentiate between array accesses at different indices. They show that they can reduce any expressible security property  $\phi(n)$  to  $\phi(1)$ . Impressively, even though they consider such a restricted setting, they manage to verify two realistic case studies: The access control mechanisms in the hypervisors SecVisor [161] and sHype [159]. Franklin et al. [74] extend this work to address separation properties over multi-level data structures such as nested page tables.

***k*-induction** Another approach that has gained popularity since its introduction is *k*-induction. Originally, Sheeran et al. [162] proposed *k*-induction as a SAT-based technique to prove safety properties over finite transition systems. Since then, the approach has been adopted to program verification [63, 83, 3, 76] and integrated in the C model checker ESBMC [77]. As the name suggests, *k*-induction is a generalised form of proofs by induction. The approach's special peculiarity is that we replace loop invariants by *k* unrolled loop iterations. This eliminates the difficult task of manually finding or generating loop invariants. Once a *k* has been set, the loop has been unrolled and once the right assertions have been set to represent the inductive step, we can use standard BMC to carry out the proof. We see that *k*-induction does not reduce unbounded proofs to bounded ones. However, it reduces the task of finding a traditional loop invariant to finding a big enough *k* such that *k* unrolled iterations carry enough information to be representative of the loop's general behaviour. Both approaches are thus similar in spirit.

## Chapter 3

# Ghost Signals: Verifying Termination of Busy Waiting

This chapter was previously published as:

**T. Reinhard**, and B. Jacobs. *Ghost signals: Verifying Termination of Busy Waiting*. In Proceedings of 33rd International Conference on Computer-Aided Verification (CAV) (2021), vol. 12760 of Lecture Notes in Computer Science, Springer, pp. 27–50.  
[doi.org/10.1007/978-3-030-81688-9\\_2](https://doi.org/10.1007/978-3-030-81688-9_2)

### 3.1 Introduction

Programs for multiprocessor machines commonly perform busy waiting for synchronization [133, 128]. In this chapter, we propose a separation logic [155, 138] to modularly verify termination of such programs under fair scheduling. Specifically, we consider programs where some threads busy-wait for a certain condition  $C$  over a shared data structure to hold, e.g., a memory flag being set by other threads. By modularly, we mean that we reason about each thread and each function in isolation. That is, we do not reason about thread scheduling or interleavings. We only consider these issues when proving the soundness of our logic. Assuming fair scheduling is necessary since busy-waiting for a condition  $C$  only terminates if the thread responsible for establishing the condition is sufficiently often scheduled to establish  $C$ .

Busy waiting is an example of *blocking* behaviour, where a thread’s progress *requires interference* from other threads. This is not to be confused with *non-blocking* concurrency, where a thread’s progress does not rely on—and may in fact be *impeded* by—interference from other threads. Existing proposed approaches for verifying termination of concurrent programs consider only programs that only involve non-blocking concurrent objects [50], or *primitive blocking constructs* of the programming language, such as acquiring built-in mutexes, receiving from built-in channels, joining threads, or waiting for built-in monitor condition variables [118, 19, 84], or both [100]. Existing techniques that do support busy waiting are not Hoare logics; instead, they verify termination-preserving *contextual refinements* between more concrete and more abstract implementations of busy-waiting concurrent objects [121, 106]. In contrast, we here propose the first conventional program logic for modular verification of termination of programs involving busy waiting, using Hoare triples as module specifications.

In order to prove that a busy-waiting loop terminates, we have to prove that it performs only finitely many iterations. To do this we introduce a special form of *ghost resources* [103] which we call *ghost signals*. As ghost resources they only exist on the verification level and hence do not affect the program’s runtime behaviour. Signals are initially unset and come with an obligation to set them. Setting a signal does not by definition correspond to any runtime condition. So, in order to use a signal  $s$  effectively, anyone using our approach has to prove an invariant stating that  $s$  is set if and only if the condition of interest holds. Further, the proof author must prove that every thread discharges all its obligations by performing the corresponding actions, e.g., by setting a signal and establishing the corresponding condition by setting the memory flag.

In our verification approach we tie every busy-waiting loop to a finite set of ghost signals  $S$  that correspond to the set of conditions the loop is waiting for. Every iteration that does not terminate the loop must be justified by the proof author proving that some signal  $s \in S$  has indeed not been set, yet. This way, we reduce proving termination to proving that no signal is waited for infinitely often.

Our approach ensures that no thread directly or indirectly waits for itself by requiring the proof author (i) to choose a well-founded and partially ordered set of levels  $\mathcal{L}evs$  and (ii) to assign a level to every signal and by (iii) only allowing a thread to wait for a signal if the signal’s level is lower than the level of each held obligation. This guarantees that every signal is waited for only finitely often and hence that every busy-waiting loop terminates. We use this to prove that every program that is verified using our approach indeed terminates.

We start by gradually introducing the intuition behind our verification approach



and the concepts we use. In § 3.2.1 and § 3.2.2 we present the main aspects of using signals to verify termination. We start by treating them as physical thread-safe resources and only consider busy waiting for a signal to be set. Then, we drop thread-safety and explain how to prove data-race- and deadlock-freedom. In § 3.2.3 and § 3.2.4 we generalize our approach to busy waiting for arbitrary conditions over arbitrary data structures and then lift signals to the verification level by introducing ghost signals.

In § 3.3 we sketch the verification of a realistic producer-consumer example involving a bounded FIFO to demonstrate our approach’s usability and address fine-grained concurrency in § 3.4. Further, we describe the available tool support in § 3.5 and discuss integrating higher-order features in § 3.6. We conclude by comparing our approach to related work and reflecting on it in § 3.7 and § 3.8.

We formally define our logic and prove its soundness in appendix § A. To keep the presentation in this chapter simple, we assume busy-waiting loops to have a certain syntactical form. In our technical report [151], included as appendix § B, we present a generalised version of our logic and its soundness proof. Further, we verify the realistic example presented in § 3.3 in full detail in appendix § A and in the technical report (appendix § B), using the respective version of our logic. We used our tool support to verify C versions of the bounded FIFO example and the CLH lock. The tool we used and the annotated .c files can be found at [98, 149, 148].

## 3.2 A Guide on Verifying Termination of Busy Waiting

When we try to verify termination of busy-waiting programs, multiple challenges arise. Throughout this section, we describe these challenges and our approach to overcome them. In § 3.2.1 we start by discussing the core ideas of our logic. In order to simplify the presentation we initially consider a simple language with built-in thread-safe *signals* and a corresponding minimal example where one thread busy-waits for such a signal. Signals are heap cells containing Boolean values that are specially marked as being solely used for busy waiting. Throughout this section, we generalize our setting as well as our example towards one that allows to verify programs with busy waiting for arbitrary conditions over arbitrary shared data structures. In § 3.2.2 we present the concepts necessary to verify data-race-, deadlock-freedom and termination in the presence of built-in signals that are not thread safe. In § 3.2.3 we explain how to use these non-thread-safe signals to verify programs that wait for arbitrary conditions over shared data structures. We illustrate this by an example waiting

```

let sig := new_signal in
fork await is_set(sig);
set_signal(sig)

```

Figure 3.1: Minimal example with two threads communicating via a physical thread-safe signal.

for a shared heap cell to be set. In § 3.2.4 we erase the signals from our program and lift them to the verification level in the form of a concept we call *ghost signals*.

### 3.2.1 Simplest Setting: Thread-Safe Physical Signals

We want to verify programs that busy-wait for arbitrary conditions over arbitrary shared data structures. As a first step towards achieving this, we first consider programs that busy-wait for simple Boolean flags, specially marked as being used for the purpose of busy waiting. We call these flags *signals*. For now, we assume that read and write operations on signals are thread-safe. Consider a simple programming language with built-in signals and with the following commands: (i) **new\_signal** for creating a new unset signal, (ii) **set\_signal**( $x$ ) for setting  $x$  and (iii) **await is\_set**( $x$ ) for busy-waiting until  $x$  is set. The main thread creates the signal **sig** and forks a new thread that busy-waits for **sig** to be set. Then, the main thread sets the signal. As we assume signal operations to be thread-safe in this example, we do not have to care about potential data races. Notice that like all busy-waiting programs, this program is guaranteed to terminate only under fair thread scheduling: Indeed, it does not terminate if the main thread is never scheduled after it forks the new thread. In this chapter we verify termination under fair scheduling.

## Augmented Semantics

**Obligations** The only construct in our language that can lead to non-termination are busy-waiting loops of the form `await is_set(sig)`. In order to prove that programs terminate it is therefore sufficient to prove that all created signals are eventually set. We use so-called *obligations* [85, 84, 118, 107] to ensure this. These are *ghost resources* [103], i.e., resources that do not exist during runtime and can hence not influence a program’s runtime behaviour. They carry, however, information relevant to the program’s verification. Generally, holding an obligation requires a thread to discharge it by performing a certain action. For instance, when the main thread in our example creates signal `sig`, it simultaneously creates an obligation to set it. The only way to discharge this obligation is to set `sig`.

We denote thread IDs by  $\theta$  and describe which obligations a thread  $\theta$  holds by bundling them into an obligations chunk  $\theta.\text{obs}(O)$ , where  $O$  is a multiset of signals. We denote multisets by double braces  $\{\dots\}$  and multiset union by  $\uplus$ . Each occurrence of a signal  $s$  in  $O$  corresponds to an obligation by thread  $\theta$  to set  $s$ . Consequently,  $\theta.\text{obs}(\emptyset)$  asserts that thread  $\theta$  does not hold any obligations.

**Augmented Semantics** In the *real* semantics of the programming language we consider here, ghost resources such as obligations do not exist during runtime. To prove termination, we consider an *augmented* version of it that keeps track of ghost resources during runtime. In this semantics, we maintain the invariant that every thread holds exactly one `obs` chunk. That is, for every running thread  $\theta$ , our heap contains a unique heap cell  $\theta.\text{obs}$  that stores the thread’s bag of obligations. Further, we let a thread get stuck if it tries to finish while it still holds undischarged obligations. Note that we use the term *finish* to refer to thread-local behaviour while we write *termination* to refer to program-global behaviour, i.e., meaning that every thread finishes. For every augmented execution there trivially exists a corresponding execution in the real semantics.

Fig. 3.2 presents some of the reduction rules we use to define the augmented semantics. We use  $\hat{h}$  to refer to augmented heaps, i.e., heaps that can contain ghost resources. A reduction step has the form  $\hat{h}, c \xrightarrow{\theta}_{\text{aug}} \hat{h}', c', T$  expresses that thread  $\theta$  reduces heap  $\hat{h}$  (which is shared by all threads) and command  $c$  to heap  $\hat{h}'$  and command  $c'$ . Further,  $T$  represents the set of threads forked during this step. It is either empty or a singleton containing the new thread’s ID and the command it is going to execute, i.e.,  $\{(\theta_f, c_f)\}$ . We omit it whenever it is clear from the context that no thread is forked. Further, we denote disjoint union of sets by  $\sqcup$ .

AUG-RED-NEWSIGNAL

$$\frac{id \notin \text{ids}(\widehat{h}) \quad L \in \mathcal{Levs}}{\widehat{h} \sqcup \{\theta.\text{obs}(O)\}, \text{new\_signal} \xrightarrow{\theta}_{\text{aug}} \widehat{h} \sqcup \{\theta.\text{obs}(O \uplus \{(id, L)\})\}, \text{signal}(\{(id, L)\}), id}$$

AUG-RED-SETSIGNAL

$$\widehat{h} \sqcup \{\theta.\text{obs}(O \uplus \{s\})\}, \text{set\_signal}(s.id) \xrightarrow{\theta}_{\text{aug}} \widehat{h} \sqcup \{\theta.\text{obs}(O), \text{signalSet}(s)\}, \text{tt}$$

AUG-RED-FORK

$$\frac{\theta_f \notin \text{thlds}(\widehat{h})}{\widehat{h} \sqcup \{\theta.\text{obs}(O \uplus O_f)\}, \text{fork } c \xrightarrow{\theta}_{\text{aug}} \widehat{h} \sqcup \{\theta.\text{obs}(O), \theta_f.\text{obs}(O_f)\}, \text{tt}, \{(\theta_f, c)\}}$$

AUG-RED-AWAIT

$$\frac{\theta.\text{obs}(O) \in \widehat{h} \quad \text{signal}(s) \in \widehat{h} \quad \text{signalSet}(s) \notin \widehat{h} \quad s.\text{lev} \prec_L O}{\widehat{h}, \text{await is\_set}(s.id) \xrightarrow{\theta}_{\text{aug}} \widehat{h}, \text{await is\_set}(s.id)}$$

Figure 3.2: Reduction rules for augmented semantics.

Our reduction rules comply with the intuition behind obligations we outlined above. **AUG-RED-NEWSIGNAL** creates a new signal and simultaneously a corresponding obligation. The only way to discharge it is by setting the signal using **AUG-RED-SETSIGNAL**.

**Forking** Whenever a thread forks a new thread, it can pass some of its obligations to the newly forked thread, cf. **AUG-RED-FORK**. Forking a new thread with ID  $\theta_f$  also allocates a new heap cell  $\theta_f.\text{obs}$  to store its bag of obligations. Since this is the only way to allocate a new **obs** heap cell, we will never run into a heap  $\widehat{h} \sqcup \{\theta.\text{obs}(O)\} \sqcup \{\theta.\text{obs}'(O')\}$  that contains multiple obligations chunks belonging to the same thread  $\theta$ . Remember that threads cannot finish while holding obligations. This prevents them from dropping obligations via dummy forks.

**Levels** In order to prove that a busy-waiting loop **await is\_set(sig)** terminates, we must ensure that the waiting thread does not directly or indirectly wait for itself. We could just check that it does not hold an obligation for the signal it is waiting for, but that is not sufficient as the following example demonstrates: Consider a program with two signals  $\text{sig}_1, \text{sig}_2$  and two threads. Let one thread hold the obligation for  $\text{sig}_2$  and execute **await is\_set(sig<sub>1</sub>); set\_signal(sig<sub>2</sub>)**. Likewise, let the other thread hold the obligation for  $\text{sig}_1$  and let it execute **await is\_set(sig<sub>2</sub>); set\_signal(sig<sub>1</sub>)**.

To prevent such *wait cycles* modularly, we apply the usual approach [118, 21, 72]. For every program that we want to execute in our augmented semantics, we

choose a partially ordered set of levels  $\mathcal{Levs}$ . Further, during every reduction step in the augmented semantics that creates a signal  $s$ , we pick a level  $L \in \mathcal{Levs}$  and associate it with  $s$ . Note that much like obligations, levels do not exist during runtime in the real semantics. Signal chunks in the augmented semantics have the form  $\text{signal}((id, L))$  where  $id$  is the unique signal identifier returned by **new\_signal**. The level assigned to any signal can be chosen freely, cf. AUG-RED-NEWSIGNAL. In practice, determining levels boils down to solving a set of constraints that reflect the dependencies. In our example, however, the choice is trivial as it only involves a single signal. We choose  $\mathcal{Levs} = \{0\}$  and 0 as level for **sig** and thereby get  $\text{signal}((\text{sig}, 0))$ . Generally, we denote signal tuples by  $s = (id, L)$ . Now we can rule out cyclic wait dependencies by only allowing a thread to busy-wait for a signal  $s$  if its level  $s.\text{lev}$  is smaller than the level of each held obligation, cf. AUG-RED-AWAIT<sup>1</sup>. Given a bag of obligations  $O$ , we denote this by  $s.\text{lev} \prec_L O$ .

**Proving Termination** As we will explain below, the augmented semantics has no fair infinite executions. We can use this as follows to prove that a program  $c$  terminates under fair scheduling: For every fair infinite execution of  $c$ , show that we can construct a corresponding augmented execution. (This requires that each step's side conditions in the augmented semantics are satisfied. Note that we thereby prove certain properties for the real execution, like absence of cyclic wait dependencies.) As there are no fair infinite executions in the augmented semantics, we get a contradiction. It follows that  $c$  has no fair infinite executions in the real semantics.

**Soundness** In order to prove soundness of our approach, we must prove that there indeed are no fair infinite executions in the augmented semantics. This boils down to proving that no signal can be waited for infinitely often. Consider any program and any fair augmented execution of it. Consider the execution's *program order graph*, (i) whose nodes are the execution steps and (ii) which has an edge from a step to the next step of the same thread and to the first step of the forked thread, if it is a fork step. Notice that for each obligation created during the execution, the set of nodes corresponding to a step made by a thread while that thread holds the obligation constitutes a path that ends when the obligation is discharged. We say that this path *carries* the obligation.

It is not possible that a signal is waited for infinitely often. Indeed, suppose some signals  $S^\infty$  are. Take  $s_{\min} \in S^\infty$  with minimal level. Since  $s_{\min}$  is never

---

<sup>1</sup>For simplicity, our augmented semantics assumes that the level order and the level associated with any object remains fixed for the entire execution. However, following the approach presented in [115], it would be sound to add a step rule that allows a thread to change the level of an object it has exclusive access to (cf. § 3.2.2).

set, the path in the program order graph that carries the obligation must be infinite as well. Indeed, suppose it is finite. The final node  $N$  of the path cannot discharge the obligation without setting the signal, so it must pass the obligation on either to the next step of the same thread or to a newly forked thread. By fairness of the scheduler, both of these threads will eventually be scheduled. This contradicts  $N$  being the final node of the path.

The path carrying the obligation for  $s_{\min}$  waits only for signals that are waited for finitely often. (Remember that AUG-RED-AWAIT requires the signal waited for to be of a lower level than all held obligations, i.e., a lower level than that of  $s_{\min}$ .) It is therefore a finite path. A contradiction.

Notice that the above argument relies on the property that every non-empty set of levels has a minimal element. For this reason, for termination verification we require that  $\mathcal{L}evs$  is not just partially ordered, but also well-founded.

## Program Logic

Directly using the augmented semantics to prove that our example program terminates is cumbersome. In the following, we present a separation logic that simplifies this task.

**Safety** We call a program  $c$  *safe* under a (partial) heap  $\hat{h}$  if it provides all the resources necessary such that both  $c$  and any threads it forks can execute without getting stuck in the augmented semantics. (This depends on the angelic choices.) We denote this by  $\mathit{safe}(\hat{h}, c)$  [176]<sup>2</sup>.

Consider a program  $c$  that is safe under an augmented heap  $\hat{h}$ . Let  $h$  be the real heap that matches  $\hat{h}$  apart from the ghost resources. Then, for every real execution that starts with  $h$  we can construct a corresponding augmented execution.

**Specifications** We use Hoare triples  $\{A\} c \{\lambda r. B(r)\}$  [92] to specify the behaviour of a program  $c$ . Such a triple expresses the following: Consider any evaluation context  $E$ , such that for every return value  $v$ , running  $E[v]$  from a state that satisfies  $B(v)$  is safe. Then, running  $E[c]$  from a state that satisfies  $A$  is safe.

---

<sup>2</sup>For a formal definition see appendix § A.4.2 and the technical report [151] (appendix § B.10).

**Proof System** We define a proof relation  $\vdash$  which ensures that whenever we can prove  $\vdash \{A\} c \{\lambda r. B(r)\}$ , then  $c$  complies with the specification  $\{A\} c \{\lambda r. B(r)\}$ . Fig. 3.3b presents some of the proof rules we use to define  $\vdash$ . As we evolve our setting throughout this section, we also adapt our proof rules. Rules that will be changed later are marked with a prime in their name. The full set of rules is presented in the appendix (cf. Fig. A.7 and A.8). Our proof rules PR-SETSIGNAL' and PR-AWAIT' are similar to the rules for sending and receiving on a channel presented in [118].

Notice how the proof rules enforce the side-conditions of the augmented semantics. Hence, all we have to do to prove that a program  $c$  terminates is to prove that every thread eventually discharges all its obligations. That is, we have to prove  $\vdash \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$ . Fig. 3.3a illustrates how we can apply our rules to verify that our minimal example terminates.

### 3.2.2 Non-Thread-Safe Physical Signals

As a step towards supporting waiting for arbitrary conditions over shared data structures, including non-thread-safe ones, we now move to non-thread-safe signals. For simplicity, in this chapter we consider programs that use mutexes to synchronize concurrent accesses to shared data structures. (Our ideas apply equally to programs that use other constructs, such as atomic machine instructions.) Fig. 3.4 presents our updated example.

As signal `sig` is no longer thread-safe, the two threads can no longer use it directly to communicate. Instead, we have to synchronize accesses to avoid data races. Hence, we protect the signal by a mutex `mut` created by the main thread. In each iteration, the forked thread acquires the mutex, checks whether `sig` has been set and releases it again. After forking, the main thread acquires the mutex, sets the signal and releases it again.

**Exposing Signal Values** Signals are specially marked heap cells storing Boolean values. We make this explicit by extending our signal chunks from `signal(s)` to `signal(s, b)` where  $b$  is the current value of  $s$  and by updating our proof rules accordingly. Upon creation, signals are unset. Hence, creating a signal `sig` now spawns an *unset* signal chunk `signal((sig, L), False)` for some freely chosen level  $L$  and an obligation for `(sig, L)`, cf. PR-NEWSIGNAL''. We present our new proof rules in Fig. 3.6 and demonstrate their application in Fig. 3.5.

**Data Races** As read and write operations on signals are no longer thread-safe, our logic has to ensure that two threads never try to access `sig` at the same

<pre> {obs(∅)} let sig := new_signal in {obs({(sig, 0)}) * signal((sig, 0))} fork ({obs(∅) * signal(s)}       await is_set(sig)       {obs(∅) * signal(s)}); {obs({s})} set_signal(sig) {obs(∅)} </pre>	<pre> PR-NEWSIGNAL' with L = 0 s := (sig, 0) s.lev = 0 &lt;_L ∅ </pre>
---	--

(a) Proof outline for program from Fig. 3.1. Applied proof rule marked in purple. Abbreviation marked in red. General hint marked in grey.

$\frac{\text{PR-NEWSIGNAL}' \quad L \in \mathcal{L}evs}{\vdash \{\text{obs}(O)\} \text{ new\_signal } \{\lambda r. \text{obs}(O \uplus \{(r, L)\}) * \text{signal}((r, L))\}}$
$\frac{\text{PR-SET SIGNAL}'}{\vdash \{\text{obs}(O \uplus \{s\})\} \text{ set\_signal}(s.\text{id}) \{\text{obs}(O)\}}$
$\frac{\text{PR-FORK}' \quad \vdash \{\text{obs}(O_f) * A\} c \{\text{obs}(\emptyset) * B\}}{\vdash \{\text{obs}(O_m \uplus O_f) * A\} \text{ fork } c \{\text{obs}(O_m)\}}$
$\frac{\text{PR-AWAIT}' \quad s.\text{lev} <_L O}{\vdash \{\text{obs}(O) * \text{signal}(s)\} \text{ await is\_set}(s.\text{id}) \{\text{obs}(O) * \text{signal}(s)\}}$
$\frac{\text{PR-LET} \quad \vdash \{A\} c \{\lambda r. C(r)\} \quad \forall v. \vdash \{C(v)\} c'[x \mapsto v] \{B\}}{\vdash \{A\} \text{ let } x := c \text{ in } c' \{B\}}$

(b) Proof rules. Rules only used in this section marked with '.

Figure 3.3: Verifying termination of minimal example with physical thread-safe signal.



```

let sig := new_signal in
let mut := new_mutex in
fork with mut await is_set(sig);
acquire mut;
set_signal(sig);
release mut

```

(a) Code.

```

with mut await c := (while acquire mut;
                    let r := c in
                    release mut;
                    ¬r
                    do skip)

```

(b) Syntactic sugar.  $r$  not free in  $\text{mut}$ .

Figure 3.4: Minimal example with two threads communicating via a physical non-thread-safe signal protected by a mutex.

time. Hence, in our logic possession of a signal chunk  $\text{signal}(s, b)$  expresses (temporary) *exclusive ownership* of  $s$ . Further, our logic requires threads to own any signal they are trying to access. Specifically, when a thread wants to set  $\text{sig}$ , it must hold a chunk of the form  $\text{signal}((\text{sig}, L), b)$ , cf.  $\text{PR-SET-SIGNAL''}$ . The same holds for reading a signal's value, cf.  $\text{PR-IS-SIGNAL-SET''}$ . Note that signal chunks are not duplicable and only created upon creation of the signal they refer to. Therefore, holding a signal chunk for  $\text{sig}$  indeed guarantees that the holding thread has the exclusive right to access  $\text{sig}$  (while holding the signal chunk).

**Synchronization & Lock Invariants** After the main thread creates  $\text{sig}$ , it exclusively owns the signal. The main thread can transfer ownership of this resource during forking, cf.  $\text{PR-FORK'}$ , and thereby allow the forked thread to busy-wait for  $\text{sig}$ . This would, however, leave the main thread without any permission to set the signal and thereby discharge its obligation.

We use mutexes to let multiple threads share ownership of a common set of resources in a synchronized fashion. Every mutex is associated with a *lock invariant*  $P$ , an assertion chosen by the proof author that specifies which resources the mutex protects. In our example, we want both threads to share  $\text{sig}$ . To reflect the fact that the signal's value changes over time, we

<pre> {obs(∅)} let sig := new_signal in {obs({(sig, 1)}) * signal((sig, 1), False)} {obs({(sig, 1)}) * ∃b. signal((sig, 1), b)} let mut := new_mutex in {obs({s}) * mutex(m, P)} {obs({s}) * mutex(m, P) * mutex(m, P)} fork ({obs(∅) * mutex(m, P)}       with m await         {obs({m}) * P}         ∀b. {obs({m}) * signal(s, b)}         is_set(sig)           {             λr.obs({m})               * signal(s, b)               ∧ r = b             λr.obs({m}) *               if r then P               else signal(s, False)           }         {obs(∅) * mutex(m, P)}         {obs(∅)}); {obs({s}) * mutex(m, P)} acquire mut; {   obs({s, m}) * locked(m, P)   * ∃b. signal(s, b) } ∀b. {   obs({s, m}) * locked(m, P)   * signal(s, b) } set_signal(sig); {   obs({m}) * locked(m, P)   * signal(s, True) } {obs({m}) * locked(m, P) * P} release mut {obs(∅) * mutex(m, P)} {obs(∅)} </pre>	<pre> PR-NEWSIGNAL'' with L = 1 PR-VIEWSHIFT &amp; VS-SEMIMP s := (sig, 1), P := ∃b. signal(s, b) PR-NEWMUTEX'' with L = 0 PR-VIEWSHIFT &amp; VS-CLONEMUT''  m.lev, s.lev &lt;_L ∅ PR-EXISTS  PR-VIEWSHIFT &amp; VS-SEMIMP  PR-VIEWSHIFT &amp; VS-SEMIMP  m.lev = 0 &lt; 1 = s.lev PR-EXISTS  PR-VIEWSHIFT &amp; VS-SEMIMP  PR-VIEWSHIFT &amp; VS-SEMIMP </pre>
---	---

Figure 3.5: Proof outline for program 3.4, verifying termination with mutexes & non-thread safe signals. Applied proof and view shift rules marked in purple. Abbreviations marked in red. General hints marked in grey.

$$\frac{\text{PR-NEWSIGNAL''}}{L \in \mathcal{L}evs} \\ \vdash \{ \text{obs}(O) \} \text{ new\_signal } \{ \lambda id. \text{ obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False}) \}$$

$$\text{PR-SET SIGNAL''} \\ \vdash \{ \text{obs}(O \uplus \{s\}) * \text{signal}(s, \_) \} \text{ set\_signal}(s.id) \{ \text{obs}(O) * \text{signal}(s, \text{True}) \}$$

$$\text{PR-IS SIGNAL SET''} \\ \vdash \{ \text{signal}(s, b) \} \text{ is\_set}(s.id) \{ \lambda r. \text{signal}(s, b) \wedge r = b \}$$

$$\text{PR-AWAIT''} \\ \frac{m.\text{lev}, s.\text{lev} \prec_L O \quad \text{signal}(s, \text{False}) * R \Rightarrow P}{\vdash \{ \text{obs}(O \uplus \{m\}) * P \} c \{ \lambda r. \text{obs}(O \uplus \{m\}) * \text{if } r \text{ then } P \text{ else } \text{signal}(s, \text{False}) * R \}} \\ \vdash \{ \text{obs}(O) * \text{mutex}(m, P) \} \text{ with } m.\text{loc} \text{ await } c \{ \text{obs}(O) * \text{mutex}(m, P) \}$$

(a) Signals & busy waiting.

$$\frac{\text{PR-NEWMUTEX''}}{L \in \mathcal{L}evs} \\ \vdash \{ P \} \text{ new\_mutex } \{ \lambda \ell. \text{mutex}((\ell, L), P) \}$$

$$\begin{array}{ll} \text{PR-ACQUIRE''} & \text{PR-RELEASE''} \\ \vdash \{ \text{obs}(O) * \text{mutex}(m, P) \wedge m.\text{lev} \prec_L O \} & \vdash \{ \text{obs}(O \uplus \{m\}) * \text{locked}(m, P) * P \} \\ \vdash \text{acquire } m.\text{loc} & \vdash \text{release } m.\text{loc} \\ \vdash \{ \text{obs}(O \uplus \{m\}) * \text{locked}(m, P) * P \} & \vdash \{ \text{obs}(O) * \text{mutex}(m, P) \} \end{array}$$

(b) Mutexes.

$$\begin{array}{ll} \text{PR-FRAME} & \text{PR-EXISTS} \\ \frac{\vdash \{ A \} c \{ B \}}{\vdash \{ A * F \} c \{ B * F \}} & \frac{\forall a \in A. \vdash \{ a \} c \{ B \}}{\vdash \{ \bigvee A \} c \{ B \}} \end{array}$$

$$\begin{array}{ll} \text{PR-FORK} & \text{PR-VIEWSHIFT} \\ \frac{\vdash \{ \text{obs}(O_f) * A \} c \{ \text{obs}(\emptyset) \}}{\vdash \{ \text{obs}(O_m \uplus O_f) * A \} \text{ fork } c \{ \text{obs}(O_m) \}} & \frac{A \Rightarrow A' \quad \vdash \{ A' \} c \{ B' \} \quad B' \Rightarrow B}{\vdash \{ A \} c \{ B \}} \end{array}$$

(c) Standard rules.

$$\begin{array}{ll} \text{VS-SEMIMP} & \text{VS-TRANS} \\ \frac{\forall H. \text{consistent}_{\text{th}}(H) \wedge H \vDash_A A \Rightarrow H \vDash_A B}{A \Rightarrow B} & \frac{A \Rightarrow C \quad C \Rightarrow B}{A \Rightarrow B} \end{array}$$

$$\text{VS-CLONEMUT''} \\ \text{mutex}(m, P) \Rightarrow \text{mutex}(m, P) * \text{mutex}(m, P)$$

(d) View shifts.

Figure 3.6: Proof rules & view shift rules for mutexes & non-thread safe signals. Rules only used in this section marked with ''.

choose a lock invariant that abstracts over its concrete value. We choose  $P := \exists b. \text{signal}((\text{sig}, L), b)$ . Let us ignore the chosen signal level  $L$  for now. Creating the mutex `mut` consumes this lock invariant and binds it to `mut` by creating a mutex chunk  $\text{mutex}((\text{mut}, \dots), P)$ , cf. PR-NEWMUTEX''. Thereby, the main thread loses access to `sig`. The only way to regain access is by acquiring `mut`, cf. PR-ACQUIRE''. Once the thread releases `mut`, it again loses access to all resources protected by the mutex, cf. PR-RELEASE''.

**Deadlocks** We have to ensure that any acquired mutex is eventually released, again. Hence, acquiring a mutex spawns a release obligation for this mutex and the only way to discharge this obligation is indeed by releasing it, cf. PR-ACQUIRE'' and PR-RELEASE''.

Any attempt to acquire a mutex will block until the mutex becomes available. In order to prove that our program terminates, we have to prove that it does not get stuck during an acquisition attempt. To prevent wait cycles involving mutexes, we require the proof author to associate every mutex as well (just like signals) with a level  $L$ . This level can be freely chosen during the mutex' creation, cf. PR-NEWMUTEX''. Mutex chunks therefore have the form  $\text{mutex}((\ell, L), P)$  where  $\ell$  is the heap location the mutex is stored at. Their only purpose is to record the level and lock invariant a mutex is associated with. Hence, these chunks can be freely duplicated as we will see later. Generally, we denote mutex tuples by  $m = (\ell, L)$ . We only allow to acquire a mutex if its level is lower than the level of each held obligation, cf. PR-ACQUIRE''. This also prevents any thread from attempting to acquire mutexes twice, e.g., **acquire mut; acquire mut** or **with mut await acquire mut**.

**View Shifts** When verifying a program, it can be necessary to reformulate the proof state and to draw semantic conclusions. To allow this we introduce a so-called *view shift* relation  $\Rightarrow$  [104]. By applying proof rule PR-VIEWSHIFT and VS-SEMIMP we can strengthen the precondition and weaken the postcondition. In our example, we use this to convert the unset signal chunk into the lock invariant which abstracts over the signal's value, i.e.,  $\text{signal}(s, \text{False}) \Rightarrow \exists b. \text{signal}(s, b)$ .

The logic we present in this work is an intuitionistic separation logic that allows us to drop chunks. <sup>3</sup> This allows us to simplify the postcondition of our fork

---

<sup>3</sup>This allows a thread to drop its obligations chunk  $\text{obs}(O)$ . Note, however, that by dropping this chunk the thread does not drop its obligations, but only its ability to show what its obligations are. In particular the thread would be unable to present an empty obligations chunk upon termination.

proof rule's premise from  $\text{obs}(\emptyset) * B$  to  $\text{obs}(\emptyset)$ , cf. PR-FORK, and drop all unneeded chunks via a semantic implication  $\text{obs}(\emptyset) * B \Rightarrow \text{obs}(\emptyset)$ .

We also allow to clone mutex chunks via view shifts, cf. VS-CLONEMUT''. In our example, this is necessary to inform both threads which level and lock invariant mutex `mut` is associated with. That is, the main thread clones the mutex chunk  $\text{mutex}(m, P)$  and passes one chunk on when it forks the busy-waiting thread.

In § 3.2.4 we extend our view shift relation and revisit our interpretation of what a view shift expresses. The full set of rules we use to define  $\Rightarrow$  is presented in the appendix (cf. Fig. A.6).

**Busy Waiting** In the approach presented in this chapter, for simplicity we only support busy-waiting loops of the form **with** `mut` **await**  $c$ , which is syntactic sugar for **while** **acquire** `mut`; **let**  $r := c$  **in** **release** `mut`;  $\neg r$  **do** **skip** where  $r$  denotes a fresh variable.<sup>4</sup> In each iteration, the loop tries to acquire `mut`, executes  $c$ , releases `mut` again and lets the result returned by  $c$  determine whether the loop continues. Such loops can fail to terminate for two reasons: (i) Acquiring `mut` can get stuck and (ii) the loop could diverge.

We prevent the loop from getting stuck by requiring `mut`'s level to be lower than the level of each held obligation, cf. PR-AWAIT''. Further, we enforce termination by requiring the loop to wait for a signal. That is, when verifying a busy-waiting loop using our approach, the proof author must choose a fixed signal and prove that this signal remains unset at the end of every non-finishing iteration. This way, we can prove that the loop terminates by proving that every signal is eventually set, just as in § 3.2.1. And just as before, our logic requires the level of the waited-for signal to be lower than the level of each held obligation.

Acquiring the mutex in every iteration makes the lock invariant available during the verification of the loop body  $c$ . This lock invariant has to be restored at the end of the iteration such that it can be consumed during the mutex's release. PR-AWAIT'' allows for an additional view shift to restore the invariant. In our example, we end our busy-waiting loop's non-finishing iterations with the assertion  $\text{signal}(s, \text{False})$ . We use a semantic implication view shift to convert the signal chunk into the mutex invariant  $\exists b. \text{signal}(s, b)$ .

**Choosing Levels** In our example, we have to assign levels to the mutex `mut` and to the signal `sig`. Our proof rules for mutex acquisition and busy waiting

---

<sup>4</sup>As we discuss in § 3.5, in the technical report accompanying this chapter (appendix B) we present a more general logic that imposes no such syntactic restrictions.

```

let x := cons(0) in
let mut := new_mutex in
fork with mut await [x] = 1;
acquire mut;
[x] := 1;
release mut

```

(a) Example program with busy waiting for heap cell  $x$  to be set.

```

let x := cons(0) in
let sig := new_signal in
let mut := new_mutex in
fork with mut await [x] = 1;
acquire mut;
[x] := 1;
set_signal(sig);
release mut

```

(b) Example program 3.7a with additional signal  $\text{sig}$  inserted, marked in green.  $\text{sig}$  and  $x$  are kept in sync.

$$[e] = e' \quad := \quad (\text{let } r := [e] \text{ in } r = e')$$

(c) Syntactic sugar.  $r$  free in  $e'$ .

Figure 3.7: Minimal example illustrating busy waiting for condition over heap cell.

impose some restrictions on the levels of the involved mutexes and signals. By analysing the corresponding rule applications that occur in our proof, we can derive which constraints our level choice must comply with. Our example's verification involves one application of  $\text{PR-ACQUIRE}''$  and one application of  $\text{PR-AWAIT}''$ : (i) Our main thread tries to acquire  $\text{mut}$  while holding an obligation to set  $\text{sig}$ . (ii) The forked thread busy-waits for  $\text{sig}$  while not holding any obligations. Our assignment of levels must therefore satisfy the single constraint  $m.\text{lev} <_{\perp} s.\text{lev}$ . So, we choose  $\mathcal{L}evs = \{0, 1\}$ ,  $m.\text{lev} = 0$  and  $s.\text{lev} = 1$ .

### 3.2.3 Arbitrary Data Structures

The proof rules we introduced in § 3.2.2 allow us to verify programs busy-waiting for arbitrary conditions over arbitrary shared data structures as follows: For every condition  $C$  the program waits for, the proof author inserts a signal  $s$  into the program. They ensure that  $s$  is set at the same time the program establishes  $C$  and prove an invariant stating that the signal's value expresses whether  $C$  holds. Then, the waiting thread can use  $s$  to wait for  $C$ . We illustrate this here for the simplest case of setting a single heap cell in Fig. 3.7a.

The program involves three new non-thread-safe commands: (i)  $\text{cons}(v)$  for allocating a new heap cell and initializing it with value  $v$ , (ii)  $[\ell] := v$  for assigning value  $v$  to heap location  $\ell$ , (iii)  $[\ell]$  for reading the value stored in heap location  $\ell$ . We use  $[\ell] = v$  as syntactic sugar for  $\text{let } r := [e] \text{ in } r = e'$ .

In our example, the main thread allocates  $x$ , initializes it with the value 0 and protects it using mutex  $mut$ . It forks a new thread busy-waiting for  $x$  to be set. Afterwards, the main thread sets  $x$ . As explained above, we verify the program by inserting a signal  $sig$  that reflects whether  $x$  has been set, yet. Fig. 3.7b presents the resulting code. The main thread creates the signal and sets it when it sets  $x$ .

**Heap Cells** Verifying this example does not conceptually differ from the example we presented in § 3.2.2. Fig. 3.8b presents the new proof rules we need and Fig. 3.8a sketches our example’s verification. As with non-thread-safe signals, we have to prevent multiple threads from trying to access  $x$  at the same time in order to prevent data races. For this we use so-called *points-to* chunks [155, 138]. They have the form  $\ell \mapsto v$  and express that heap location  $\ell$  stores the value  $v$ . When a thread holds such a chunk, it exclusively owns the right to access heap location  $\ell$ .

Heap locations are unique and the only way to create a new points-to chunk is to allocate and initialize a new heap cell via  $\mathbf{cons}(v)$ , cf. PR-CONS. Hence, there will never be two points-to chunks involving the same heap location. In order to read or write a heap cell via  $[\ell]$  or  $[\ell] := e$ , the acting thread must first acquire possession of the corresponding points-to chunk, cf. PR-ASSIGNTOHEAP and PR-READHEAPLOC'''.

**Relating Signals to Conditions** In our example, the forked thread busy-waits for  $x$  to be set while our proof rules require us to justify each iteration by showing an unset signal. That is, we must prove an invariant stating that the value of  $x$  matches  $sig$ . As this invariant must be shared between both threads, we encode it in the lock invariant:  $P := \exists v. x \mapsto v * \mathbf{signal}(s, v = 1)$ . This does not only allow both threads to share the heap cell and the signal but it also automatically enforces that they maintain the invariant whenever they acquire and release the mutex.

### 3.2.4 Signal Erasure

In the program from Fig. 3.7b signal  $sig$  is never read and does hence not influence the waiting thread’s runtime behaviour. Therefore, we can verify the original program presented in Fig. 3.7a by erasing the physical signal and treating it as ghost code.

```

{obs(∅)}
let x := cons(0) in
{obs(∅) * x ↦ 0}
let sig := new_signal in
let mut := new_mutex in
s := (sig, 1), m := (mut, 0)
P := ∃v. x ↦ v * signal(s, v = 1)
{obs({s}) * mutex(m, P) * mutex(m, P)}
fork ({obs(∅) * mutex(m, P)}
  with m await
    {obs({m}) * P}
    ∀v. { obs({m}) * x ↦ v
        * signal(s, v = 1)
        [x] = 1
        { λr. obs({m})
          * if r then P
            else x ↦ v ∧ v ≠ 1
            * signal(s, False)
          }
        }
    {obs(∅)});
{obs({s}) * mutex(m, P)}
acquire mut;
∀v. { obs({s, m}) * locked(m, P) * x ↦ v
    * signal(s, v = 1)
    [x] := 1;
    { obs({s, m}) * locked(m, P) * x ↦ 1
      * signal(s, v = 1)
    }
    set_signal(sig);
    { obs({m}) * locked(m, P) * x ↦ 1
      * signal(s, True)
    }
    release mut
    {obs(∅)}

```

PR-NEWSIGNAL'' with  $L = 1$   
 PR-NEWMUTEX'' with  $L = 0$

$m.\text{lev}, s.\text{lev} \prec_L \emptyset$

$m.\text{lev} = 0 < 1 = s.\text{lev}$

(a) Proof outline for program 3.7b. Applied proof rules marked in purple. Abbreviations marked in red. General hints marked in grey.

<p>PR-CONS</p> $\vdash \{\text{True}\} \text{cons}(v) \{\lambda\ell. \ell \mapsto v\}$	<p>PR-ASSIGNTOHEAP</p> $\vdash \{\ell \mapsto \_ \} [\ell] := v \{\ell \mapsto v\}$
<p>PR-READHEAPLOC'''</p> $\vdash \{\ell \mapsto v\} [\ell] \{\lambda r. r = v * \ell \mapsto v\}$	<p>PR-EXP</p> $\frac{\llbracket e \rrbracket \in \text{Values}}{\vdash \{\text{True}\} e \{\lambda r. r = \llbracket e \rrbracket\}}$

(b) Proof rules. Evaluation function  $\llbracket \cdot \rrbracket$ . Rules only used in this section marked with ''.

Figure 3.8: Verifying termination of busy waiting for condition over heap cell.



**Ghost Signals** Central aspects of the proof sketch we presented in Fig. 3.8a are that (i) the main thread was obliged to set `sig` and that (ii) the value of `sig` reflected whether `x` was already set. *Ghost signals* allow us to keep this information but at the same time to remove the physical signals from the code. Ghost signals are essentially identical to the physical non-thread-safe signals we used so far. However, as ghost resources they cannot influence the program’s runtime behaviour. They merely carry information we can use during the verification process.

**View Shifts Revisited** We implement ghost signals by extending our view shift relation. In particular, we introduce two new view shift rules: VS-NEWSIGNAL and VS-SET SIGNAL presented in Fig. 3.9b. The former creates a new unset signal and simultaneously spawns an obligation to set it. The latter can be used to set a signal and thereby discharge a corresponding obligation. We say that these rules change the *ghost state* and therefore call their application a *ghost proof step*. With this extension, a view shift  $A \Rightarrow B$  expresses that we can reach postcondition  $B$  from precondition  $A$  by (i) drawing semantic conclusions or by (ii) manipulating the ghost state. In Fig. 3.9a we use ghost signals to verify the program from 3.7a.

Note that lifting signals to the verification level does not affect the soundness of our approach. The argument we presented in § 3.2.1 still holds. We formalize our logic and provide a formal soundness proof in appendix § A and in the technical report [151] (appendix § B). The latter contains a more general version of the presented logic that (i) is not restricted to busy-waiting loops of the form **with mut await**  $c$  and that (ii) is easier to integrate into existing tools like VeriFast [99], as explained in § 3.5.

### 3.3 A Realistic Example

To demonstrate the expressiveness of the presented verification approach, we verified the termination of the program presented in Fig. 3.10a. It involves two threads, a consumer and a producer, communicating via a shared bounded FIFO with a maximal capacity of 10. The producer enqueues numbers 100, ..., 1 into the FIFO and the consumer dequeues those. Whenever the queue is full, the producer busy-waits for the consumer to dequeue an element. Likewise, whenever the queue is empty, the consumer busy-waits for the producer to enqueue the next element. Each thread’s finishing depends on the other thread’s productivity. This is, however, no cyclic dependency. For instance, in order to prove that the producer eventually pushes number  $i$  into the queue, we only

<pre> {obs(∅)} let x := cons(0) in {obs(∅) * x ↦ 0} new_ghost_signal; { ∃sig. obs(⟦(sig, 1)⟧) * x ↦ 0 } { * signal((sig, 1), False) } ∀sig. {obs(⟦s⟧) * x ↦ 0 * signal(s, False)}  let mut := new_mutex in { obs(⟦s⟧) * mutex((mut, 0), P) } { * mutex((mut, 0), P) } fork ( {obs(∅) * mutex(m, P)} ) with m await   {obs(⟦m⟧) * P}   ∀v. { obs(⟦m⟧) * x ↦ v }       { * signal(s, v = 1) }       [x] = 1       {         λr. obs(⟦m⟧) *           if r then P           else x ↦ v ∧ v ≠ 1           * signal(s, False)       }   {obs(∅)}; {obs(⟦s⟧) * mutex(m, P)} acquire mut; ∀v. { obs(⟦s, m⟧) * locked(m, P) }     { * x ↦ v * signal(s, v = 1) }     [x] := 1;     set_ghost_signal(s);     { obs(⟦m⟧) * locked(m, P) }     { * x ↦ 1 * signal(s, True) } release mut {obs(∅)} </pre>	<p style="color: purple;">VS-NEWSIGNAL with <math>L = 1</math></p> <p style="color: red;"><math>s := (\text{sig}, 1)</math></p> <p style="color: red;"><math>P := \exists v. x \mapsto v</math>  <math>\quad * \text{signal}(s, v = 1)</math></p> <p style="color: purple;">PR-NEWMUTEX'' with <math>L = 0</math></p> <p style="color: red;"><math>m := (\text{mut}, 0)</math></p> <p style="color: red;"><math>m.\text{lev}, s.\text{lev} \prec_L \emptyset</math></p> <p style="color: red;"><math>m.\text{lev} = 0 &lt; 1 = s.\text{lev}</math></p>
<p>(a) Proof outline for the program presented in Fig. 3.7a. Auxiliary commands hinting at view shifts and general hints marked in <i>grey (italic)</i>. Applied proof and view shift rules marked in <b>purple</b>. Abbreviations marked in <b>red</b>.</p>	<p style="text-align: center;">VS-NEWSIGNAL</p> $\frac{L \in \mathcal{Levs}}{\text{obs}(O) \Rightarrow \exists id. \text{obs}(O \uplus \llbracket(id, L)\rrbracket) * \text{signal}((id, L), \text{False})}$ <p style="text-align: center;">VS-SET SIGNAL</p> $\text{obs}(O \uplus \llbracket s \rrbracket) * \text{signal}(s, \_) \Rightarrow \text{obs}(O) * \text{signal}(s, \text{True})$

(b) Proof rules.

Figure 3.9: Verifying termination with ghost signals.

need to rely on the consumer to pop  $i + 10$ . A similar property holds for the consumer.

**Fine-Tuning Signal Creation** To simplify complex proofs involving many signals we refine the process of creating a new ghost signal. For simplicity, we combined the allocation of a new signal ID and its association with a level and a Boolean in one step. For some proofs, such as the one we outline in this section, it can be helpful to fix the IDs of all signals that will be created throughout the proof already at the beginning. To realize this, we replace view shift rule VS-NEWSIGNAL by the rules presented in Fig. 3.10b and adapt our signal chunks accordingly. With these more fine-grained view shifts, we start by allocating a signal ID, cf. VS-ALLOCSIGID. Thereby we obtain an *uninitialized* signal  $\text{uninitSig}(id)$  that is not associated with any level or Boolean, yet. Also, allocating a signal ID does not create any obligation because threads can only wait for *initialized* (and unset) signals. When we initialize a signal, we bind its already allocated ID to a level of our choice and associate the signal with **False**, cf. VS-SIGINIT. This creates an obligation to set the signal.

**Loops & Signals** In our program, both threads have a local counter initially set to 100 and run a nested loop. The outer loops are controlled by their thread's counter, which is decreased in each iteration until it reaches 0 and the loop stops. For such loops, we introduce a conventional proof rule for total correctness of loops, cf. Fig. A.7 in the appendix. Verifying termination of the inner loops is a bit more tricky and requires the use of ghost signals.

So far, we had to fix a single signal for the verification of every **await** loop. We can relax this restriction to considering a finite set of signals the loop may wait for, cf. PR-AWAIT presented in Fig. A.7 in the appendix. Apart from being a generalisation, this rule does not differ from PR-AWAIT'' introduced in § 3.2.2.

Initially, we allocate 200 signal IDs  $id_{\text{push}}^{100}, \dots, id_{\text{push}}^1, id_{\text{pop}}^{100}, \dots, id_{\text{pop}}^1$ . We are going to ensure that always at most one push signal and at most one pop signal are initialized and unset. The producer and consumer are going to hold the obligation for the push and pop signal, respectively. The producer will hold the obligation for  $s_{\text{push}}^i$  while  $i$  is the next number to be pushed into the FIFO and it will set  $s_{\text{push}}^i$  when it pushes the number  $i$  into the FIFO. Meanwhile, the consumer will use  $s_{\text{push}}^i$  to wait for the number  $i$  to arrive in the queue when it is empty. Similarly, the consumer will hold the obligation for  $s_{\text{pop}}^i$  while number  $i$  is the next number to be popped from the FIFO and will set  $s_{\text{pop}}^i$  when it pops the number  $i$ . The producer uses  $s_{\text{pop}}^i$  to wait for the consumer to pop  $i$  from the queue when it is full. At any time, we let the mutex **mut** protect the two active signals and thereby make them accessible to both threads.

```

alloc_ghost_signal_IDs( $id_{pop}^i, id_{push}^i$ ) for  $1 \leq i \leq 100$ ;
 $L_{pop}^i := 102 - i, L_{push}^i := 101 - i, s_x^i := (id_x^i, L_x^i)$  for  $1 \leq i \leq 100$ 
init_ghost_signals( $s_{pop}^{100}, s_{push}^{100}$ );
{obs( $\{s_{pop}^{100}, s_{push}^{100}\}$ ) * ...}
let  $fif_{10} := \text{cons}(\text{nil})$  in
let  $mut := \text{new\_mutex}$  in
let  $c_p := \text{cons}(100)$  in
let  $c_c := \text{cons}(100)$  in
fork ( while (
  with  $mut$  await (
    {obs( $\{s_{push}^{c_p}, (mut, 0)\}$ ) * ...}
    let  $f := [fif_{10}]$  in
    if  $\text{size}(f) < 10$  then (
      let  $c := [c_p]$  in
       $[fif_{10}] := f \cdot \langle c \rangle$ ;
       $[c_p] := c - 1$ ;
      set_ghost_signal( $s_{push}^c$ );
      if  $c - 1 \neq 0$  then
        init_ghost_signal( $s_{push}^{c-1}$ );
       $\text{size}(f) \neq 10$  );
       $[c_p] \neq 0$ 
    do skip);
  while (
    with  $mut$  await (
      {obs( $\{s_{pop}^{c_c}, (mut, 0)\}$ ) * ...}
      let  $f := [fif_{10}]$  in
      if  $\text{size}(f) > 0$  then (
        let  $c := [c_c]$  in
         $[fif_{10}] := \text{tail}(f)$ ;
         $[c_c] := c - 1$ ;
        set_ghost_signal( $s_{pop}^c$ );
        if  $c - 1 \neq 0$  then
          init_ghost_signal( $s_{pop}^{c-1}$ );
         $\text{size}(f) > 0$  );
         $[c_c] \neq 0$ 
      do skip);
  while (
    if  $\text{size}(f) = 10$  then wait for  $s_{pop}^{c_p+10}$ 
     $L_{pop}^{c_p+10} = 92 - c_p < 101 - c_p = L_{push}^{c_p}$ 
     $c_p$  decreases in each iteration.
    Busy-wait for  $fif_{10}$  not being full.
    → Wait for consumer to pop.
    If  $fif_{10}$  not full, push next element.
     $c_c$  decreases in each iteration.
    Busy-wait for  $fif_{10}$  not being empty.
    → Wait for producer to push.
    If  $fif_{10}$  not empty, pop next element.
    if  $\text{size}(f) = 0$  then wait for  $s_{push}^{c_c}$ 
     $L_{push}^{c_c} = 101 - c_c < 102 - c_c = L_{pop}^{c_c}$ 
  )
do skip);

```

(a) Example program with two threads communicating via a shared bounded FIFO with maximal size 10. Auxiliary commands hinting at view shifts & general hints marked in *grey (italic)*. Abbreviations marked in **red**. Hints on proof state marked in **blue**.

VS-ALLOCSIGID $\text{True} \Rightarrow \exists id. \text{uninitSig}(id)$	VS-SIGINIT $\text{obs}(O) * \text{uninitSig}(id)$ $\Rightarrow \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})$
---	---

(b) Fine-grained view shift rules for signal creation.

Figure 3.10: Realistic example program.

**Choosing the Levels** Note that we ignored the levels so far. The producer and the consumer both acquire the mutex while holding an obligation for a signal. Hence, we choose  $\mathcal{L}evs = \mathbb{N}$ ,  $m.le\!v = 0$  and  $s.le\!v > 0$  for every signal  $s$ . Both threads will justify iterations of their respective **await** loop by using an unset signal at the end of such an iteration. Our proof rules allow us to ignore the mutex obligation during this step. Hence, the mutex level does not interfere with the level of the unset signal. Whenever the queue is full, the producer waits for the consumer to pop an element and whenever the queue is empty, the consumer waits for the producer to push. That is, the producer waits for  $s_{pop}^{i+10}$  while holding an obligation for  $s_{push}^i$  and the consumer waits for  $s_{push}^i$  while holding an obligation for  $s_{pop}^i$ . So, we have to choose the signal levels such that  $s_{pop}^{i+10}.lev < s_{push}^i.lev$  and  $s_{push}^i.lev < s_{pop}^i.lev$  hold. We solve this by choosing  $s_{pop}^i.lev = 102 - i$  and  $s_{push}^i.lev = 101 - i$ .

**Verifying Termination** This setup suffices to verify the example program. Via the lock invariant, each thread has access to both active signals. Whenever the producer pushes a number  $i$  into the queue, it sets  $s_{push}^i$  which discharges the held obligation and decreases its counter. Afterwards, if  $i > 1$ , it uses the uninitialized signal chunk `uninitSig( $id_{pop}^{i-1}$ )` to initialize  $s_{push}^{i-1} = (id_{push}^{i-1}, 101 - (i - 1))$  and replaces  $s_{push}^i$  in the lock invariant by  $s_{push}^{i-1}$  before it releases the lock. If  $i = 1$ , the counter reached 0 and the loop ends. In this case, the producer holds no obligation. The consumer behaves similarly. Since we proved that each thread discharged all its obligations, we proved that the program terminates. Fig. 3.10a illustrates the most important proof steps. We present the program’s verification in full detail in appendix § A.5.1 and in the technical report [151] (appendix B.12.2). Furthermore, we encoded [149] the proof in VeriFast [99].

The number of threads in this program is fixed. However, our approach also supports the verification of programs where the number of threads is not even statically bounded. In the appendix in § A.5.2 we present and verify such a program. It involves  $N$  producer and  $N$  consumer threads that communicate via a shared buffer of size 1, for a random number  $N > 0$  determined during runtime.

### 3.4 Specifying Busy-Waiting Concurrent Objects

Our approach can be used to verify busy-waiting concurrent objects with respect to abstract specifications. For example, we have verified [148] the CLH lock [90] against a specification that is very similar to our proof rules for built-in mutexes shown in Fig. 3.6. The main difference is that it is slightly more abstract: when

a lock is initialized, it is associated with a *bounded infinite set* of levels rather than with a single particular level. (To make this possible, an appropriate universe of levels should be used, such as the set of lists of natural numbers, ordered lexicographically.) To acquire a lock, the levels of the obligations held by the thread must be above the elements of the set; the new obligation’s level is an element of the set.

### 3.5 Tool Support

We have extended the VeriFast tool [98] for separation logic-based modular verification of C and Java programs so that it supports verifying termination of busy-waiting C or Java programs. When verifying termination, VeriFast consumes a *call permission* at each recursive call or loop iteration. In the technical report [151] (appendix § B) we define a generalised version of our logic that instead of providing a special proof rule for busy-waiting loops, provides *wait permissions* and a *wait view shift*. A call permission of a *degree*  $\delta$  can be turned into a wait permission of a degree  $\delta' < \delta$  for a given signal  $s$ . A wait view shift for an unset signal  $s$  for which a wait permission of degree  $\delta$  exists produces a call permission of degree  $\delta$ , which can be used to fuel a busy-waiting loop. When busy-waiting for some signal  $s$ , we can generate new permissions to justify each iteration as long as  $s$  remains unset.

VeriFast allows threads to freely exchange permissions. This is useful to verify termination of non-blocking algorithms involving compare-and-swap loops [100]. However, we must be careful to prevent self-fueling busy-waiting loops. Hence, we restrict where a permission can be consumed based on the *thread phase* it was created in. The main thread’s initial phase is  $\epsilon$ . When a thread in phase  $p$  forks a new thread, its phase changes to  $p.\text{Forker}$  and the new thread starts in phase  $p.\text{Forkee}$ . We allow a thread in phase  $p$  to consume a permission only if it was produced in an *ancestor thread phase*  $p' \sqsubseteq p$ .

The only change we had to make to VeriFast’s symbolic execution engine was to enforce the thread phase rule. We encoded the other aspects of the logic simply as axioms in a *trusted header file*. We used this tool support to verify the bounded FIFO (§ 3.3) and the CLH lock (§ 3.4). The bounded FIFO proof [149] contains 160 lines of proof annotations for 37 lines of code (an annotation overhead of 435%) and takes 0.08s to verify. The CLH lock proof [148] contains 343 lines of annotations for 49 lines of code (an overhead of 700%) and takes 0.1s to verify.

## 3.6 Integrating Higher-Order Features

The logic we presented in this chapter does not support higher-order features such as assertions that quantify over assertions, or storing assertions in the (logical) heap as the values of ghost cells. While we did not need such features to carry out our example proofs, they are generally useful to verify higher-order program modules against abstract specifications. The typical way to support such features in a program logic is by applying *step indexing* [7, 112], where the domain of logical heaps is indexed by the number of execution steps left in the (partial) program trace under consideration. Assertions stored in a logical heap at index  $n + 1$  talk about logical heaps at index  $n$ ; i.e., they are meaningful only *later*, after at least one more execution step has been performed.

It follows that such logics apply directly only to *partial* correctness properties. Fortunately, we can reduce a termination property to a safety property by writing our program in a programming language *instrumented* with runtime checks that guarantee termination. Specifically, we can write our program in a programming language that fulfils the following criteria: It tracks signals, obligations and permissions at runtime and has constructs for signal creation, waiting and setting a signal. The **fork** command takes as an extra operand the list of obligations to be transferred to the new thread (and the other constructs similarly take sufficient operands to eliminate any need for angelic choice). Threads get stuck when these constructs' preconditions are not satisfied, such as when a thread waits for a signal while holding the obligation for that signal. We can then use a step-indexing-based higher-order logic such as Iris [104] to verify that no thread in our program ever gets stuck. Once we established this, we know none of the instrumentation has any effect and can be safely *erased* from the program.

## 3.7 Related & Future Work

In recent work [154] we propose a separation logic to verify termination of programs where threads busy-wait to be abruptly terminated. We generalize this work to support busy waiting for arbitrary conditions.

In [100] we propose an approach based on *call permissions* to verify termination of single- and multithreaded programs that involve loops and recursion. However, that work does not consider busy-waiting loops. In the technical report [151] (appendix § B), we present a generalised logic that uses call permissions and allows busy waiting to be implemented using arbitrary looping and/or recursion.

Furthermore, the use of call permissions allowed us to encode our case studies in our VeriFast tool which also uses call permissions for termination verification.

Liang and Feng [120, 121] propose LiLi, a separation logic to verify liveness of blocking constructs implemented via busy waiting. In contrast to our verification approach, theirs is based on the idea of contextual refinement. In their approach, client code involving calls of blocking methods of the concurrent object is verified by first applying the contextual refinement result to replace these calls by code involving primitive blocking operations and then verifying the resulting client code using some other approach. In contrast, specifications in our approach are regular Hoare-style triples and proofs are regular Hoare-style proofs.

In [97] we propose a Hoare logic to verify liveness properties of the I/O behaviour of programs that do not perform busy waiting. By combining that approach with the one we proposed in this chapter, we expect to be able to verify I/O liveness of realistic concurrent programs involving both I/O and busy waiting, such as a server where one thread receives requests and enqueues them into a bounded FIFO, and another one dequeues them and responds. To support this claim, we encoded the combined logic in VeriFast and verified a simple server application where the receiver and responder thread communicate via a shared buffer [150].

### 3.8 Conclusion

We propose what is to the best of our knowledge the first separation logic for verifying termination of programs with busy waiting. We offer a soundness proof of the system presented in this chapter in appendix § A, and of a more general system in the technical report [151] (appendix § B). Further, we demonstrated its usability by verifying a realistic example. We encoded our logic and the realistic example in VeriFast [149] and used this encoding also to verify the CLH lock [148]. Moreover, we expect that our approach can be integrated into other existing concurrent separation logics such as Iris [104].



## Chapter 4

# Completeness Thresholds for Memory Safety of Array Traversing Programs

This chapter was previously published as:

**T. Reinhard**, J. Fasse, and B. Jacobs. *Completeness Thresholds for Memory Safety of Array Traversing Programs*. In Proceedings of the 12th ACM SIGPLAN International Workshop on the State Of the Art in Program Analysis (SOAP) (2023), ACM, pp. 47–54. [doi.org/10.1145/3589250.3596143](https://doi.org/10.1145/3589250.3596143)

To streamline the notation, we replaced some constraints by constraint sets.

### 4.1 Introduction

**Unbounded vs Bounded Proofs** Many techniques have been developed to convince ourselves of the trustworthiness of software. A fundamental pillar for any higher-level property is memory safety. In memory-unsafe languages the burden of proof lies with the programmer. Yet, it remains hard to prove and in general requires us to write tedious, inductive proofs. One way to automate the verification process is to settle on bounded proofs and accept bounded guarantees.

Consider a program  $c$  that searches through an array of size  $s$ . An unbounded memory safety proof for  $c$  would yield that the program is safe for any possible input, in particular for any array size, i.e.,  $\forall s. \text{memsafe}(c(s))$ . A bounded proof that only considers input sizes up to ten would only guarantee that the program is safe for any such bounded array, i.e.,  $\forall s \leq 10. \text{memsafe}(c(s))$ .

**Completeness Thresholds** Approximating unbounded proofs by bounded ones is a technique often used in model checking. Hence, the relationship between bounded and unbounded proofs about finite state transition systems has been studied extensively [15, 40, 109, 28, 2, 87, 10, 126]. For a finite transition system  $T$  and a property of interest  $\phi$ , a *completeness threshold* is any number  $k$  such that we can prove  $\phi$  by only examining path prefixes of length  $k$  in  $T$ , i.e.,  $T \models_k \phi \Rightarrow T \models \phi$  [40]<sup>1</sup>. Over the years, various works characterised over-approximations of least completeness thresholds for different types of properties  $\phi$ . These over-approximations are typically described in terms of key attributes of the transition system  $T$ , such as the *recurrence diameter* (longest loop-free path) [109]. For instance, consider the class of global safety properties of the form  $Gp$  for finite transition systems  $T$ , where  $p$  is a local property. We know that the smallest completeness threshold for this class expressible solely in terms of  $T$ 's diameter is exactly  $\text{diam}(T)$  [15, 110]. As safety property of the form  $Gp$ , this also applies to memory safety of finite transition systems.

In general, heap-manipulating programs' state space can be infinite. That is because the program's input data can be arbitrarily large and because executions can be arbitrarily long. Therefore, the key attributes described above will generally be infinite as well. This vast structural difference between the programs we are interested in and the transition systems for which completeness thresholds have been studied prevents us from reusing any of the existing definitions or results.

## 4.2 Limitations of Bounded Proofs

**Bounded Model Checking** Generally, if we want unbounded memory safety guarantees, we have to consider all possible input sizes and all possible executions. This is often hard and requires us to write tedious inductive proofs. An alternative is to give up on the idea of unbounded guarantees and to settle for

---

<sup>1</sup>Note that the term completeness threshold is used inconsistently in literature. Some papers such as [40] use the definition above, according to which completeness thresholds are not unique. Others like [109] define them as the minimal number  $k$  such that  $T \models_k \phi \Rightarrow T \models \phi$ , which makes them unique.

$$\begin{array}{l}
 1 \text{ for } i \text{ in } [L : s-R] \text{ do} \\
 2 \quad !a[i+Z]
 \end{array}
 \left. \vphantom{\begin{array}{l} 1 \\ 2 \end{array}} \right\} =: \text{trav}_{L,R}^Z(a, s)$$

$$\text{Trav} := \{ \text{trav}_{L,R}^Z \mid L, R, Z \in \mathbb{Z} \}$$

Figure 4.1: Class *Trav* of programs  $\text{trav}_{L,R}^Z$  traversing an array  $a$  of size  $s$ , attempting to read elements.  $L, R, Z$  are constants.

bounded ones. One approach that has proven useful during development of critical software is bounded model checking (BMC) [32].

The underlying idea is to approximate the original verification problem by a finite model that we can check automatically. With this approach, we choose a size bound  $S$  and only consider inputs with sizes up to  $S$ . Further, we also only check finite execution prefixes. A common approach is to unwind loops and recursion up to a certain depth.

The intuition behind this approach is that if the program contains errors, they likely already occur for small input sizes and early loop iterations. As long as BMC does not perform abstraction [35], all reported counterexamples are real bugs. However, we should be careful not to forget that this way we only obtain a bounded proof yielding bounded guarantees.

**Array Traversal Pattern** Consider the class of programs *Trav* presented in Fig. 4.1 in a WHILE language with pointer arithmetic. Given a pointer  $a$  and a variable  $s$ , such that  $a$  points to an array of size  $s$ , each program iterates through the array and attempts to read elements. The class models a basic programming pattern and common off-by-n errors [47]. We use upper case letters for constants and lower case letters for (program) variables. A program  $\text{trav}_{L,R}^Z$  from this class iterates from  $i = L$  to  $i = s - R$  (bounds incl.) and attempts to read the array at index  $i + Z$ . We use  $!x$  to express accesses to a heap location  $x$ . Whether memory errors occur for a concrete instance  $\text{trav}_{L,R}^Z$  depends on how the constants  $L, R, Z$  are chosen. We use it as minimal example throughout this chapter.

**What Could Go Wrong with Bounded Proofs?** To illustrate the issue, let us use BMC to check various instances of the array traversal pattern: (i) traversal of the entire array:  $\text{trav}_{0,1}^0$ , (ii) traversal of the array with accesses offset by two from the index:  $\text{trav}_{0,1}^2$  and (iii) an additional reduction of the index variable's upper bound by one:  $\text{trav}_{0,2}^2$ . It is easy to see that (i) is memory-safe while (ii) and (iii) are not. However, before we run a model checking algorithm we have

to choose appropriate bounds. The pattern we are looking at is quite simple. So, we choose size bound  $S = 1$  and unwinding depth  $D = 1$  for the BMC procedure. Note that the latter effectively means: we do not restrict the loop depth for the input sizes we chose.

For the standard variant (i) we cannot find any errors within the bounds. This is fine because the program is safe. In variant (ii) array accesses  $a[i + 2]$  are incorrectly shifted to the right. This already leads to an out-of-bounds error for arrays of size 1. This size falls within our chosen bounds, so BMC reports this error and we can correct it. Finally, (iii)'s reduction of the index variable's range to  $[0, s - 2]$  means that the program only performs loop iterations for arrays of size  $s \geq 2$ . Consequently, it is trivially safe for the sizes 0 and 1. These are the sizes our bounded proof explores. Hence, BMC does not report any errors and leads us to wrongly believe that  $trav_{0,2}^Z$  is safe.

### 4.3 Completeness Thresholds

As illustrated above, bounded proofs are in general unsound approximations of unbounded proofs. A concrete approximation is sound iff we choose the bounds large enough, such that we can be sure that we do not miss any errors. We focus on bounding input sizes (in our examples array sizes), ignoring loop bounds that do not depend on these parameters.

Recall from § 4.1 that completeness thresholds are a concept from model checking of finite transition systems [40]. We borrow this terminology and apply it to memory safety verification. Hence, for a program  $c(x)$  with input parameter  $x \in X$ , we call any subdomain  $Q \subseteq X$  a *completeness threshold* (CT) for  $x$  in  $c$  if we can prove memory safety of  $c$  by only considering inputs from  $Q$ , i.e.,  $\forall x \in Q. \text{ memsafe}(c(x)) \Rightarrow \forall x \in X. \text{ memsafe}(c(x))$ .

**Intuitive CT Extraction** Returning to our example class of programs  $Trav$  implementing the array traversal pattern. Some of these are memory safe, some are not. So, let us try to compute completeness thresholds for these programs. First, let's take a look at the errors that might occur. This gives us an idea which sizes a sound bounded proof must cover. Any instance  $trav_{L,R}^Z(a, s)$  iterates ascendingly  $i = L, \dots, s - R$  and accesses  $a[i + Z]$ . For sizes  $s$  that cause the ascending range  $L, \dots, s - R$  to be empty, we do not execute the loop at all. Any such run is trivially memory safe. Therefore, any meaningful bounded proof of  $trav_{L,R}^Z(a, s)$  must include sizes  $s$  with  $\{L, \dots, s - R\} \neq \emptyset$ , i.e.,  $s \geq L + R$ .

Suppose  $s \geq L + R$ . An error occurs if the index  $i + Z$  violates the array bounds, i.e., if  $i + Z < 0$  or  $i + Z \geq s$ . Taking the index range into account, we see that we get an error if  $L + Z < 0$  or  $s - R + Z \geq s$  holds. We can simplify the latter to  $Z - R \geq 0$ .

Note that neither  $L + Z < 0$  nor  $Z - R \geq 0$  depend on the array size  $s$ . This means that as long as we focus on sizes above the threshold  $s \geq L + R$ , the concrete choice of  $s$  does not influence whether an error occurs or not. In other words, it suffices for our bounded proof to only check a single (arbitrarily chosen) size  $q \geq L + R$  and then we can extrapolate the result, i.e.,

$$\forall a. \text{memsafe}(trav_{L,R}^Z(a, q)) \Rightarrow \forall s. \forall a. \text{memsafe}(trav_{L,R}^Z(a, s))$$

Hence, any set  $\{q\}$  for  $q \geq L + R$  is a CT for the array size parameter  $s$  in  $trav_{L,R}^Z$ . We just found a uniform characterization of CTs for the entire class  $Trav$ . Note that  $\{q\}$  is not necessarily the smallest CT. For safe instances such as  $trav_{0,1}^0$ , the empty set  $\emptyset$  is a valid CT as well.

**Our Approach** We study CTs for  $x$  in  $c(x)$  by studying its *verification condition* (VC). The latter is an automatically generated logical formula of the form  $\forall x \in X. vc(x)$  and proving it entails memory safety of  $c(x)$  for all choices of  $x$ . Next, we currently simplify  $vc(x)$  by hand until it becomes clear how the choice of  $x$  affects the validity of  $vc(x)$ . Knowing this allows us to partition the domain into  $X = \bigcup Q_i$ . For each subdomain we get  $vc_i(x) = \forall x \in Q_i. vc(x)$ . If possible, we simplify each  $vc_i(x)$  into  $vc'_i(x)$  based on the restricted subdomain  $Q_i$  with the goal to eliminate occurrences of  $x$ . If  $vc'_i$  does not mention  $x$  we pick any element of  $Q_i$  as representative  $Q'_i$ . Otherwise,  $Q'_i = Q_i$ . Hence,  $\bigcup Q'_i$  is a CT for  $x$  in  $c(x)$ . In the following we elaborate this in more detail.

### 4.3.1 Approximating CTs via Verification Conditions

Now that we have an intuition for the CTs of  $Trav$ , let's turn our informal argument from above into a formal one. Formal definitions of the language and logic we consider and proofs for the presented lemmas can be found in the technical report [147].

**Hoare Triples** We use Hoare triples [92] to express program specifications. A triple  $\{A\} c \{B\}$  expresses that the following properties hold for every execution that starts in a state which satisfies precondition  $A$ : Firstly, the

execution does not encounter any runtime errors. Secondly, it either (i) does not terminate or (ii) it terminates in a state complying with postcondition  $B$ .

In this work, we study the memory safety of programs that do not change the shape of the data structures they process. Hence, we choose preconditions that merely describe the memory layout of the data structures which our programs receive as input. For the array traversal program, we choose the predicate  $\text{array}(a, s)$  as precondition, which expresses that  $a$  points to a contiguous memory chunk of size  $s$ . Given that our target programs do not change the memory layout, specifications simplify to  $\{A\} c \{B\}$ . For the array traversal we get  $\{\text{array}(a, s)\} \text{trav}_{L,R}^Z \{\text{array}(a, s)\}$ .

**Definition 4.3.1** (Completeness Thresholds for Programs). *Let  $\{A\} c \{B\}$  be a program specification containing a free variable  $x$  with domain  $X$ . We call a subdomain  $Q \subseteq X$  a completeness threshold for  $x$  in  $\{A\} c \{B\}$  if*

$$\models \forall x \in Q. \{A\} c \{B\} \quad \Rightarrow \quad \models \forall x \in X. \{A\} c \{B\}$$

We omit spelling out the pre- and postconditions when they are clear from the context. Instead we say that  $Q$  is a completeness threshold for  $x$  in program  $c$ .

**Separation Logic** We use a first-order affine/intuitionistic separation logic with recursion predicates [136, 138, 155] to describe memory. Since we focus on heap-manipulating programs, we use assertions to describe heaps. Separation logic comes with a few special operators: (i) The points-to chunk  $x \mapsto v$  describes a heap containing a location  $x$  which holds the value  $v$ . We write  $x \mapsto \_$  to express that we do not care about the value stored in the heap cell. (ii) The separating conjunction  $a_1 * a_2$  expresses that  $a_1$  and  $a_2$  describe disjoint heaps. Hence,  $x \mapsto \_ * y \mapsto \_$  implies that  $x \neq y$ . (iii) The separating implication  $a_1 \multimap a_2$  can be read as  $a_2$  *without*  $a_1$ . That is, combining the heap described by  $a_1 \multimap a_2$  with a disjoint heap described by  $a_1$  yields a heap compliant with  $a_2$ . (iii) In our logic, the persistence modality  $\Box a$  means that  $a$  does not describe resources and hence holds under the empty heap (cf. [104]).

We assume that  $\text{array}$  denotes a (recursively defined) predicate, such that for every fixed size  $s$ , we can express it as iterated separating conjunction:

$$\text{array}(a, s) \equiv \bigstar_{0 \leq k < s} a[k] \mapsto \_.$$

**Verification Conditions** A common way to verify programs is via verification conditions [71, 143]. For any specification  $\{A\} c \{B\}$ , a *verification condition* (VC) is any logical formula  $vc$ , such that we can verify  $\{A\} c \{B\}$  by proving  $vc$ , i.e.,  $\models vc \Rightarrow \models \{A\} c \{B\}$ .

**Definition 4.3.2** (Verification Condition). *We call an assertion  $a$  a verification condition for  $\{A\} c \{B\}$  if*

$$\models a \quad \Rightarrow \quad \models \{A\} c \{B\}.$$

**Definition 4.3.3** (Completeness Thresholds for Assertions). *Let  $a$  be an assertion with a free variable  $x$  of domain  $X$ . We call a subdomain  $Q \subseteq X$  a completeness threshold for  $x$  in  $a$  if*

$$\models \forall x \in Q. a \quad \Rightarrow \quad \models \forall x \in X. a.$$

Consider a specification  $\{A\} c \{B\}$  with a free variable  $x \in X$  and a corresponding VC  $\forall x \in X. vc$ . Suppose we get a completeness threshold  $Q$  for  $x$  in  $vc$ . Knowing this threshold reduces correctness of the specification to the bounded VC, i.e.,  $\models \forall x \in Q. vc \Rightarrow \{A\} c \{B\}$ . That is, we can derive unbounded guarantees from a bounded proof. We usually omit quantification domains when they are clear from the context.

**Weakest Liberal Preconditions** A common way to generate VCs is via weakest liberal preconditions [71, 61]. For any program  $c$  and postcondition  $B$ , the *weakest liberal precondition*  $wlp(c, \lambda r. B)$  is an assertion for which

$$\forall A. ( A \models wlp(c, \lambda r. B) \Rightarrow \models \{A\} c \{B\} )$$

holds. That is, if the weakest liberal precondition holds for the starting state, then  $c$  does either not terminate or it terminates in a state complying with postcondition  $B$ . In particular, no memory error occurs during the execution. The canonical VC for  $\{A\} c \{B\}$  is  $\forall \bar{x}. A \rightarrow wlp(c, \lambda r. B)$  where  $\bar{x}$  is the tuple of variables occurring freely in  $A$ ,  $c$  and  $B$ .

**Limitations of CTs** In general, VCs are over-approximations. Hence, CTs derived from a VC do not always apply to the corresponding program. Consider the specification  $\{\text{array}(a, s)\} \text{trav}_{0,2}^2 \{\text{array}(a, s)\}$  and any unsatisfiable assertion  $vc_{\text{false}} \equiv \text{False}$ . Then  $\forall s \in \mathbb{N}. vc_{\text{false}}$  is an over-approximating VC, since  $\text{False} \Rightarrow \models \{\text{array}(a, s)\} \text{trav}_{0,2}^2 \{\text{array}(a, s)\}$ . As discussed in § 4.2, a bounded proof that only covers sizes 0, 1 does not discover the errors in  $\text{trav}_{0,2}^2$ . Hence, the set  $\{0\}$  is not a CT for  $x$  in  $\{\text{array}(a, s)\} \text{trav}_{0,2}^2 \{\text{array}(a, s)\}$ . However,  $\{0\}$  is a CT for  $s$  in  $vc_{\text{false}}$ , since  $\forall s \in \{0\}. vc_{\text{false}} \equiv \text{False}$  and  $\text{False} \Rightarrow \models \forall s \in \mathbb{N}. vc_{\text{false}}$ . We see that we can only transfer a CT  $Q$  for some variable  $x$  derived from a VC  $vc$  to the corresponding program, if  $vc$  does not over-approximate with regard to  $x$ . (We are currently studying this connection.) Note that this is, however, the case for the examples we discuss in this chapter. Though, even if  $vc$  does

over-approximate, CT  $Q$  still applies to any proof considering a property at least as strong as  $vc$ .

As the name suggests, VCs derived from weakest preconditions yield very weak properties. It is reasonable to assume that often a bounded proof would imply a bounded version of the wlp-based VC. Therefore, it is reasonable to use them during our study of CTs.

**Extracting CTs via VCs** Applying the above approach to our specification  $\{\text{array}(a, s)\} \text{trav}_{L,R}^Z(a, s) \{\text{array}(a, s)\}$ , we get a VC  $\forall a. \forall s. vc_{\text{trav}}$ , where  $vc_{\text{trav}}$  is as follows <sup>2</sup>:

$$\begin{aligned}
 vc_{\text{trav}} := & \text{array}(a, s) \rightarrow & \text{\#1} \\
 & \text{array}(a, s) & \text{\#2} \\
 & * \square(\forall i. (L \leq i \leq s - R) \wedge \text{array}(a, s) \multimap & \text{\#3} \\
 & \quad \exists v. a[i + Z] \mapsto v \wedge \text{array}(a, s)) \\
 & * (\text{array}(a, s) \multimap \text{array}(a, s)) & \text{\#4}
 \end{aligned}$$

Here, **\#1** is the precondition describing our memory layout. The separating conjunction **\#2**-**\#4** is the weakest precondition derived from our specification and  $vc_{\text{trav}}$  says that it should follow from precondition **\#1**.

The weakest precondition states that the memory layout must stay invariant under the loop execution. **\#2** says that it should hold before the loop starts. **\#3** demands that every loop iteration preserves the layout, i.e., that the layout description is a loop invariant. **\#4** states this invariant implies the postcondition from our specification, which, again, is the unchanged memory layout.

Remember that we use an affine separation logic. Clearly, this VC contains many trivially obsolete parts. We can simplify  $vc_{\text{trav}}$  to  $vc_1$ :

---

<sup>2</sup>The weakest precondition calculus requires us to annotate loops with loop invariants. In the setting we study, the initial memory layout is invariant under the program's execution. The preconditions we consider describe exactly the initial memory layout, nothing else. Hence, we can reuse preconditions as loop invariants during the wlp computation.



$$\begin{aligned}
vc_{\text{trav}} &\equiv // \text{ Eliminate (\#2), (\#4)} \\
&\text{array}(a, s) \rightarrow \\
&\quad \square(\forall i. (L \leq i \leq s - R) \wedge \text{array}(a, s) \text{ -*} \\
&\quad\quad \exists v. a[i + Z] \mapsto v \wedge \text{array}(a, s)) \\
&\equiv // \text{ Persistency makes pre. obsolete} \\
&\quad \square(\forall i. (L \leq i \leq s - R) \wedge \text{array}(a, s) \text{ -*} \\
&\quad\quad \exists v. a[i + Z] \mapsto v \wedge \text{array}(a, s)) \\
&\equiv // \text{ array}(a, s) \text{ equiv. to } \bigstar_{0 \leq k < s} a[k] \mapsto \_ \\
&\quad \forall i. (L \leq i \leq s - R) \rightarrow (0 \leq i + Z < s) \\
&=: vc_1
\end{aligned}$$

This equivalent VC  $\forall a. \forall s. vc_1$  does reflect the intuition we developed when analysing the program informally: For sizes  $s < L + R$ , the program does not perform any loop iterations and hence it is trivially memory safe. For bigger arrays, a memory error occurs iff index  $i + Z$  violates the array bounds.

We can justify this intuition by partitioning the domain of  $s$  into  $\mathbb{N} = \{0, \dots, L + (R - 1)\} \cup \{L + R, \dots\}$ . Let's analyse  $vc_1$  for both subdomains separately. For a size  $s_- < L + R$ , we get

$$vc_1(s_-) \equiv \forall i. \text{False} \rightarrow (0 \leq i + Z < s_-) \equiv \text{True}$$

So, we do not have to bother checking sizes  $s < L + R$ . For bigger sizes  $s_+ \geq L + R$ , we get

$$\begin{aligned}
vc_1(s_+) &\equiv \forall i. (L \leq i \leq s_+ - R) \rightarrow (0 \leq i + Z < s_+) \\
&\equiv \forall i. (L \leq i \rightarrow 0 \leq i + Z) \wedge (i \leq s_+ - R \rightarrow i + Z < s_+) \\
&\equiv \forall i. (L \leq i \rightarrow 0 \leq i + Z) \wedge (i \leq -R \rightarrow i + Z < 0) \\
&=: vc_2.
\end{aligned}$$

Since  $s_+$  does not occur freely in  $vc_2$ , the truth of  $vc_1(s_+)$  does not depend on the choice of  $s_+$ . Remember that we have,  $vc_{\text{trav}}(s_-) \equiv \text{True}$  and  $vc_{\text{trav}}(s_+) \equiv vc_2$ . Hence,

$$\models \forall a. \forall s. vc_{\text{trav}} \Leftrightarrow \models \forall a. vc_{\text{trav}}(s_+)$$

We see that it suffices to check the original VC  $vc_{\text{trav}}$  for any size  $s_+ \geq L - R$  to prove memory safety of our array traversing program  $\text{trav}_{L,R}^Z$ . That is,  $\{s_+\}$  is a CT.

**Characterizing CTs via Constraints** Note that the constraint  $s \geq L - R$  we just derived is a uniform representation for the CTs of the entire class  $\text{Trav}$ . We often

use constraint sets  $\{k_1, \dots, k_n\}$  to concisely characterize CTs. Each  $k_i$  describes a property that some element in our CT must cover to reach a specific potential error. If  $k_i$  is unsatisfiable, it means that the error is unreachable. Hence, a set  $Q$  models the constraint set if it contains models  $k_1, \dots, k_n$  for each satisfiable constraint, i.e.,  $\not\models \neg k_i \Rightarrow \models k_i[s \mapsto k_i]$ . For singleton constraint sets  $\{k\}$ , we drop the set notation and just write  $k$ .

Overall, a constraint set formulates a property that is sufficiently strong so that every subdomain covering it is a CT. For every program  $trav_{L,R}^Z \in Trav$ , every subdomain  $Q' \subseteq \mathbb{N}$  is a CT for the array size  $s$  if  $Q' \cap \{s \in \mathbb{N} \mid s \geq L - R\} \neq \emptyset$ .

### 4.3.2 Modularity of Completeness Thresholds

**Unrelated Data Structures** Consider the program  $sum_{L,R}^Z(a, s, n)$  that attempts to sum up all elements of array  $a$  and writes the result to heap location  $n$ :

$$\left. \begin{array}{l} 1 \text{ for } i \text{ in } [L : s-R] \text{ do} \\ 2 \quad !n := !n + !a[i+Z] \end{array} \right\} =: sum_{L,R}^Z(a, s, n)$$

Analogously to  $trav_{L,R}^Z$  it iterates through the array and attempts to read array elements. Additionally, it uses the read value to update the sum stored at heap location  $n$ . Let us assume that the array  $a$  and the result variable  $n$  do not alias. We get the specification  $\{A\} sum_{L,R}^Z \{A\}$  for  $A := \text{array}(a, s) * n \mapsto \_$ . We assume structured memory. Therefore, it is not possible to access  $n$  via an array access  $a[\dots]$ .

Intuitively, it is clear that the array size does not affect the memory accesses to heap location  $n$ . Hence, the CTs for  $s$  in  $sum_{L,R}^Z$  should be the same as the ones for  $trav_{L,R}^Z$ . In fact, analysing the wlp-based VC for  $sum_{L,R}^Z$  confirms this intuition. It has the form  $\forall a. \forall s. \forall r. vc_{\text{sum}}$  and we can rewrite  $vc_{\text{sum}}$  into

$$vc_{\text{sum}} \equiv vc_{\text{trav}} * (n \mapsto \_ \rightarrow A)$$

where  $vc_{\text{trav}}$  is the VC from § 4.3.1 for the array traversing program  $trav_{L,R}^Z$ . Moreover,  $\text{freeVars}(A) = \{n\}$ . Since  $s$  does not occur freely in  $n \mapsto \_ \rightarrow A$ , we can ignore it while searching for a CT for  $s$  in  $vc_{\text{sum}}$ .

**Lemma 4.3.4** (VC Slicing). *Let  $a, a_x, a_y$  be assertions with  $x \in \text{freeVars}(a_x)$  and  $x \notin \text{freeVars}(a_y)$  and  $a \equiv a_x * a_y$ . Let  $Q \subseteq X$  be a CT for  $x$  in  $a_x$ . Then,  $Q$  is also a CT for  $x$  in  $a$ , i.e.,*

$$\models \forall x \in Q. \forall \bar{y} \in \bar{Y}. a \Rightarrow \models \forall x \in X. \forall \bar{y} \in \bar{Y}. a$$

$$\left. \begin{array}{l}
 1 \text{ for } i \text{ in } [L : s-R] \text{ do} \\
 \quad ( \\
 2 \quad n := a[i+Z]; \\
 3 \quad \text{complex\_fct}(n, y, k) \\
 4 \quad )
 \end{array} \right\} =: \text{comp}_{L,R}^Z(a, s, y, k)$$

$$\begin{array}{l}
 M(a, s) \quad := \quad \text{array}(a, s) * \text{complex\_data}(y, k) \\
 \text{Comp} \quad := \quad \{ \{M\} \text{comp}_{L,R}^Z \{M\} \mid L, R, Z \in \mathbb{Z} \}
 \end{array}$$

Figure 4.2: Class of programs involving a complex data structure and computation that do not depend on the array size  $s$ .

We can extrapolate what we saw in the  $\text{sum}_{L,R}^Z$  example to more complex classes of programs. Consider the class  $\text{Comp}$  presented in Fig. 4.2. A program  $\text{comp}_{L,R}^Z \in \text{Comp}$  receives two non-aliasing data structures: an array  $a$  of size  $s$  and a complex data structure  $y$  of size  $k$ , described by the predicate  $\text{complex\_data}(y, k)$ .  $\text{comp}_{L,R}^Z$  reads elements from array  $a$ , stores the result in a local variable  $n$  and then calls a complex function  $\text{complex\_fct}(n, y, k)$  which does neither depend on  $a$  nor  $s$ . The VC  $vc_{\text{comp}}$  will reflect this. That is, analogous to the example above, it should be expressible as  $vc_{\text{comp}} \equiv vc_{\text{trav}} * (\text{complex\_data}(y, k) \rightarrow \dots)$  where the right conjunct does not depend on  $s$ .

VC Slicing lemma 4.3.4 tells us that whenever we want to characterize a CT for a specific parameter, we can ignore all separated VC conjuncts that do not involve this parameter. Effectively, this means that we can ignore all the complex parts of  $\text{comp}_{L,R}^Z$  that are not related to the array size while searching for a CT for  $s$ . This allows us to reduce the search to the CTs of  $\text{Trav}$ .

**Compositionality** We can describe the CTs of complex programs in terms of the CTs of their building blocks. Consider the program  $c_1; c_2$  and suppose that  $c_1$  and  $c_2$  are instances of patterns we studied before. So we know that each  $c_i$  corresponds to a VC  $\forall x. vc_i$  with a CT  $Q_i$  for  $x$ . Let  $vc_{1;2}$  be the VC for  $c_1; c_2$  that we want to prove. Suppose it can be rewritten into  $\forall x. vc_1 \wedge vc_2$ . Then, we know that  $Q_1 \cup Q_2$  is a CT for  $vc_{1;2}$ . Therefore, our approach to studying CTs is to study patterns and combinators.

**Basic Patterns** We view basic patterns such as the array traversal pattern discussed above as the basic building blocks. They tend to occur frequently in programs and they are sufficiently concise to extract CTs by studying their VCs. In particular, we focus on traversal and access patterns that preserve the

memory layout. For now, we focus on arrays, but we are going to generalize it to arbitrary inductive data structures.

**Managing Complexity** One of our main goals is to describe CTs for interesting classes of programs. VCs tend to become very complex very fast as a program gets more complex. Hence, we need a way to deal with this complexity and to break the CT analysis down into simpler problems. Following the structure of the program we want to reason about is a natural approach.

**Combinator Patterns** In order to exploit the program structure while analysing CTs, we need to study how control structures affect CTs. For instance, as described above, we can characterize the CT of a sequence  $c_1; c_2$  as the union of the CTs derived from  $c_1$  and  $c_2$ . Further, consider the command `if  $e$  then  $c_1$  else  $c_2$`  and suppose that we can describe CTs for a size  $s$  in each  $c_i$  via a constraint set  $K_i$ . Then, we can describe the CT for the entire command via the constraint set  $(e \wedge K_1) \cup (\neg e \wedge K_2)$ .

## 4.4 Conclusion

Past approaches to program verification either targeted unbounded guarantees and relied on unbounded, often inductive, proofs or they targeted bounded guarantees and tried to approximate the program behaviour using techniques like bounded model checking. We have, however, seen little interaction between the two communities.

In this work we propose a new perspective on memory safety proofs that connects unbounded and bounded proofs. We show that we can reduce unbounded memory safety proofs to bounded ones for certain programs that traverse arrays and preserve the memory layout. For any such program considering a few select array sizes yields the same guarantees as considering arrays of all possible sizes. We call this concept *completeness thresholds* in reference to a similar concept from model checking of finite transition systems. Moreover, we show that studying verification conditions are an adequate way to study completeness thresholds.

## 4.5 Related Work

Completeness thresholds were first introduced by Kroening and Strichman [109]. So far, the study of CTs has been limited to finite state systems. Indeed, the well-known CTs for classes of LTL properties are defined with respect to the (recurrence) diameter of the finite state system in question (e.g. [15, 110]). Determining the worst-case execution time of a program and discovering upper bounds on loops by iterative unrolling can also be used to determine CTs [65, 39]. For a possibly infinite state system those CTs can naturally be infinite as well. By specializing in just one property, memory safety, we are able to characterize and possibly find useful CTs for these systems as well. Model checking for parameterized network topologies of identical (e.g. bisimilar [27] or isomorphic [67]) processes features a related concept to completeness thresholds called *cutoff*. That is, model-checking up to the cutoff implies correctness of scaling the topology up to infinitely many processes. Positive results exist for properties of such token rings [27, 67] but also other topologies [41, 5].

The model checking literature (cf. [42]) boasts a wealth of alternative approaches to obtain unbounded guarantees on finite state systems, e.g., k-induction [16, 162], Craig interpolation [126] and property-directed reachability [24, 66] with adaptations to the software verification setting (e.g. [63, 14, 13]).

Array-manipulating programs are well-studied across different domains [23, 25, 101, 30]. However, we consider our main contribution to be a novel approach to connect unbounded and bounded proofs about memory safety. Ultimately, as discussed in the outlook, we aim to generalize and automate our approach to tree-like data structures. In that regard, Mathur et al. [124] consider the special case of proving memory safety of heap-manipulating programs as well. They prove that memory safety is decidable if the initial heap is forest-like and the program only performs a single-pass over the data-structure (see also § 4.6). They do not cover arrays and buffer overflows, but they support (de-)allocation.

## 4.6 Outlook

**CTs for Programs** Consider any program  $c$  with a free variable  $x$  and a corresponding VC  $vc$ . Suppose we derived a CT  $Q$  for  $x$  in  $vc$ . In general, this does not allow us to conclude that  $Q$  is also a CT for  $x$  in  $c$ . Intuitively, this is only true if  $vc$  does not over-approximate with regard to  $x$ . This holds for the programs and VCs studied in this chapter. While CTs for VCs still tell us something about proofs targeting these VCs, our ultimate goal is to derive thresholds for programs. Hence, we are currently studying this connection.

**Scalability** In this work, we focus on a restricted array traversal pattern to illustrate CTs. Our goal is to scale this approach to complex programs. Therefore, we are currently studying more array traversal and access patterns and combinators. This will allow us to better understand how the structure of programs affects the relation between bounded and unbounded proofs. Knowing this will allow us to characterise CTs for complex classes of programs that cover errors besides off-by-n errors.

Afterwards, we are going to extend our approach to include arbitrary inductive data types. In particular, we plan to describe CTs for a class of programs *Sort* that includes (safe and unsafe) implementations of in-place sorting algorithms involving nested loops. Once we managed that, we are going to investigate how allocation and deallocation affect CTs.

**Decidability** The memory accesses we observe in sorting functions typically mainly depend on the size of the sorted data structure, the traversal strategy and a comparison relation  $<$ . Suppose we managed to derive finite CTs for the class *Sort*. We conjecture that this will be sufficient to conclude decidability of memory safety for *Sort*. Note that this targeted result would escape the scope of the related work by Mathur et al. [124]. The latter showed decidability of memory safety for a certain form of single-pass programs, i.e., programs that traverse a datastructure exactly once.

**Improving BMC Guarantees** Bounded model checking suffers from the state space explosion problem [36, 142, 35] with respect to the chosen bounds. It is often only practical to check very small bounds on each parameter to keep the verification time practical. We plan to automate the approach introduced in this chapter and to leverage existing static analysis techniques (e.g. [180, 123, 9, 131]) to simplify the generated VCs. Once we are able to automatically compute a small CT for one parameter, say the traversed data structure’s size, we can adjust the corresponding bound. We can be sure that we didn’t miss to check any size that could lead to memory errors. Hence, it would strengthen the guarantees we get from our bounded proof. Ultimately, we would like to integrate with the industrial-strength BMC tool CBMC [38, 39].

Further, it would speed up the verification time and free resources that can be spend on exploring other parts of the program. That is, the lowered size bound might allow us to increase other bounds that seem more important. We would be able to do this while keeping the verification time stable and without sacrificing guarantees.

# Chapter 5

## Conclusion and Future Work

In this thesis we described two contributions to the state of the art in reasoning about liveness and safety, respectively. The following sections 5.1 and 5.2 reflect on the contributions we presented in § 3 and § 4, respectively, and outline directions for future work.

### 5.1 Liveness Verification

**Ghost Signals** In § 3 we presented the first separation logic to modularly verify termination of concurrent programs involving busy waiting. Standard separation logic only allows us to prove safety and partial correctness properties. Safety properties are much better understood and easier to verify than liveness properties. This is particularly true due the wealth of Hoare-style [92] logics that have been developed for safety verification [135, 139, 184, 176, 177, 104, 122]. Hoare triples  $\{A\} c \{B\}$  like the ones used in separation logic allow us to reason locally and intuitively. The logic we presented in § 3 inherits these advantages by extending separation logic. Moreover, extending a standard separation logic made it easy to implement our approach in VeriFast, a separation logic based program verifier for C.

Our approach to verifying termination is to reduce it to safety. We do this by introducing a new type of ghost resource called *ghost signal*. These signals only exist on the proof level and each is tied to an event a thread might wait for. Creating a ghost signal also spawns an obligation – another ghost resource – to perform said event. Our logic ensures that no thread terminates with

undischarged obligations. It also takes care to rule out cyclic wait dependencies that could lead to non-terminating wait loops. Since the latter two are safety properties, we indeed reduced liveness to safety.

Busy waiting verification is an active area of research. In the remainder of this section, we briefly discuss TaDA Live, the main new development in this area since our work appeared. We also briefly discuss an ongoing project of our research group continuing our ghost signal work. Afterwards, we briefly sketch another active research area: Reasoning about liveness in step-indexed logics.

**TaDA Live** D’Oswaldo et al. propose TaDA Live [64], a separation logic to verify termination of programs that involve blocking, fine-grained concurrency, e.g., busy waiting for a heap cell to be set. The paper was submitted in parallel to our ghost signal paper (cf. § 3) and published shortly after it. Both approaches share similarities such as the use of obligations and levels (or layers in TaDA Live terms) to reason about dependencies. Those aside, their logic differs widely from ours. TaDA Live offers a notion of liveness invariants  $\Box\Diamond P$ , expressed in terms of subjective obligations and guards [62]. Their obligations are subjective in the sense introduced by Ley-Wild and Nanevski [119]. That is, they support two different perspectives: Local and environment. This makes knowledge about the environment duplicable and a first class citizen in TaDA Live. (The term *duplicable* should not be confused with Iris’ notion of duplicable assertions [104]. In Iris, duplicable assertions are assertions that remain true forever. This is not the case here, since in TaDA Live duplicable assertions are not necessarily stable.)

TaDA Live offers a notion of abstract atomic specifications to specify blocking behaviour. This allows them to assign reusable and provable specifications to, e.g., spin and CLH locks. Moreover, TaDA Live’s liveness invariants offer a great deal of flexibility and control to the client over when subjective obligations are spawned and discharged. In general, the client is responsible for proving that it enforces the liveness assumptions required by the lock specifications. This allows them to verify programs where a lock is acquired in one thread but released in another one. Our ghost signal work does not concern the question of how to construct optimally abstract yet expressive specifications for busy-waiting modules.

The expressiveness that TaDA Live offers comes at the price of a proof system that is significantly more complex than ours. Hence, it is also harder to mechanize. In particular, TaDA Live is an unstable separation logic. That is, not all assertions are stable in the sense that assertions can generally be invalidated by other threads. Meanwhile, assertions used as pre- or postconditions in Hoare triples require stability since they describe locally owned resources. That is, if



a relevant shared region protocol permits the environment to take a transition, then we have to prove that this transition does not invalidate the assertion. In contrast, we implemented our ghost signals and our CLH lock case study in VeriFast [99]. Since VeriFast is a stable separation logic, it is difficult or impossible to use it to implement TaDA Live. For the same reason, it is hard or impossible to implement TaDA Live in the popular mechanised separation logic Iris [104]. There is, however, a proof checker called Voila [182, 181] for TaDA [49] a predecessor of TaDA Live.

Since the publication of our ghost signal paper, we focused on different research areas. Meanwhile, our research group, specifically Fasse and Jacobs [69], continues this line of work.

TaDA Live also supports specifying the abstract blocking behaviour of spin and CLH lock [90] implementations as well as a flexible handling of obligations. The ongoing work of Fasse and Jacobs [69] extends our ghost signal logic to support more flexible, modular lock specifications for various lock implementations such as spin, CLH, ticket [127] and cohort [58, 59] locks. In order to verify the latter one, they need to treat the creation of obligations more flexibly than it is currently the case in our ghost signal work. They achieve this flexibility by letting the client create the obligations instead of the module. In the example lock specifications we developed, every acquire call spawns an obligation to release the lock. Since we do not allow threads to pass obligations to arbitrary other threads, this means that we force the acquiring thread to eventually release it. In a cohort lock, however, an acquired lock can be passed on to another thread within the same cohort without releasing it. Our logic does not support these *hand-offs*. By passing on the lock, the original owner loses their chance to discharge their release obligation. Moreover, Fasse and Jacobs are working on mechanising their logic in Iris.

**Liveness and Step-Indexed Logics** Reasoning about liveness in step-indexed logics is an active research topic [170, 168, 172]. The main challenge is to overcome the restriction that by default, step-indexed logics only allow reasoning about finite execution prefixes.

Tassarotti et al. [170] focus on verifying liveness in concurrent settings that only have bounded non-determinism (i.e. languages where each state has only finitely many successors). This restriction allows them to extend Iris to reason about fair termination-preserving refinement. Spies et al. [168] propose Transfinite Iris which uses ordinals instead of natural numbers as step indices. Using these transfinite step indices instead of finite ones allows them to reason about termination and termination-preserving refinement in non-concurrent programs. Timany et al. [172] propose Trillium, a separation logic framework to prove

intensional refinement relations between a program and an abstract model given as transition system. In particular, they consider liveness-preserving refinements concerning concurrent programs under fair scheduling.

## 5.2 Completeness Thresholds

In § 4 we presented the first notion of completeness thresholds (CT) for software verification, which allows to obtain unbounded guarantees from bounded proofs. To be precise, we study CTs for memory safety proofs of array traversing programs. Bounded verification approaches like bounded model checking (BMC) offer a high degree of automation and convenience. Yet this comes at the expense of soundness, since they only explore a finite number of finite execution prefixes. Bounded model checkers require the user to set two types of bounds upfront: (i) A bound for input data and (ii) a depth bound limiting the size of the explored execution prefixes (or alternatively distinct depth bounds for each loop and recursive function call). In this work, we focus on bounding input data. In practice, coming up with meaningful bounds is very difficult and requires the user to take both the target program into account and also the class of errors they want to find.

**Extracting CTs from VCs** CTs are a concept originally developed in the context of model checking finite state transition systems [15, 40]. In this context CTs are natural numbers that bound the length of path prefixes a model checker needs to explore to produce a sound proof. The finite nature of these transition systems allows for the characterisation of CTs in terms of key properties of the transition system. Programs processing arbitrarily large data structures, however, correspond to infinite transition systems where said key properties are often infinite. Naively applying existing approaches would hence lead to infinite and thus unusable CTs. Instead we study a program’s verification condition (VC). The VC reflects all potential memory errors that might arise during the program execution. Most importantly, it shows how different classes of concrete inputs influence the correctness of the program. In our work, we show that we can extract a CT  $Q$  for a variable  $x$  by extracting a validity preserving subdomain  $Y \subseteq X$  from a VC. We prove that  $Y$  is validity preserving by restricting  $x$  in the VC to  $Y$  and then rewriting the VC until we eliminated  $x$ . If this succeeds, the concrete choice of  $x \in Y$  does not affect whether or not a memory error occurs in our program. That is, we can collapse the entire subdomain  $Y$  to an arbitrary representative  $y \in Y$ . Hence, we found our CT  $Q = X \setminus Y \cup \{y\}$ .

**Increasing Trust in BMC** The state space explosion problem and resource constraints make it often only practical to explore small bounds. Especially traversing over data structures drives up the computational costs. Obtaining small CTs for the size of the traversed structure does not only allow us to obtain unbounded guarantees with respect to one dimension of the proof. In many cases it will also allow us to reduce the explored size bound and thereby significantly speed up the proof. We can then reinvest the freed up computation resources into checking more critical bounds. That is, we can increase another bound while maintaining the proof's overall performance. Hence, ultimately we obtain a bounded proof we can trust significantly more.

As this is an ongoing project, there are several angles that we are going to study further.

**Limitations of CTs** In general, VCs over-approximate a program's correctness behaviour. That is, it is possible that the VC does not hold, even though the program is correct. We have to be careful about this when relating CTs that we extracted from a VC back to the original program. We are currently investigating this relationship and found that CTs for a variable  $x$  extracted from a VC hold for the program iff the VC does not over-approximate the influence of  $x$  on the correctness of the program. Note that this is a significantly weaker restriction than requiring that the VC as a whole does not over-approximate. In appendix § C we formalise this relationship and call such VCs *precise* in  $x$  (cf. Def. C.7.2). Moreover, we prove a soundness theorem (cf. Theo. C.8.3) stating that precision is a sufficiently strong requirement to ensure soundness of any extracted CT with respect to the original program specification.

Note that this applies to the array traversal pattern and its VC studied in § 4. Loops are a common cause for over-approximations in VCs. Meanwhile, iteration over data structures is a common cause for exponential blow ups in model checking. Therefore, we are going to focus on common traversal patterns for inductive data structures. We are going to study whether their VCs over-approximate with respect to the size of the traversed data structures. In case we find that they do not, we plan to manually compute CTs for the most common patterns. Model checkers could then automatically search for these patterns and directly apply our CTs to speed up the model checking process.

**Scaling** Ultimately, we want to extract CTs for realistic programs that are much more complex than the array traversal pattern studied in § 4. To achieve this, we are going to look into several different directions:

(i) We are going to study the combinatorial properties of memory safety CTs. As already sketched in the outlook of § 4, we can exploit the structure of a program when extracting CTs. We are going to investigate how control structures allow us to extract a CT by combining CTs from their respective sub-ASTs. Once we know sufficiently many CT combinators, we will be a big step closer to extracting CTs from realistic programs.

(ii) Consider a program that traverses an array, reads elements and uses the read values for some further computation  $f$  not involving the array nor its size  $s$ . Intuitively,  $f$  should not affect the CT for  $s$ . Yet, VCs tend to get very complex very fast. Hence, proving that we can ignore the sub-VC corresponding to  $f$  proves tricky on the VC level. Meanwhile, program slicing [180] is an established technique that allows us to do exactly that: Slicing off the parts of the program not affected by  $s$ . We are going to look into this. It seems promising that (i) and (ii) together will allow us to compute array-size CTs for realistic programs traversing arrays.

(iii) We are also going to investigate CTs for programs traversing inductive data structures besides arrays, e.g., lists and trees. The mechanism of extracting a CT remains the same regardless of the data structure. The VC, however, and also the traversal patterns we need to study grow in complexity with the traversed structure. We are confident that extending our results to lists will be straight-forward. Though, we expect it to be more difficult to do the same for tree structures. Mainly because of the complexity of common traversal patterns like breadth-first-search.

**Case Study** Once we are able to extract CTs for realistic programs, we plan to empirically evaluate their impact on the runtime of industrial-grade BMC proofs. Our evaluation target will be FreeRTOS's TCP/IP stack [56]. FreeRTOS [55] is an open source, real-time operating system used by many commercial IoT vendors. The FreeRTOS team heavily relies on CBMC – an industrial-strength, open source bounded model checker – to ensure the correctness of their system. In particular, the repository of their TCP stack contains many CBMC proofs for functions traversing over TCP buffer arrays. We plan to select a number of these functions, which we then manually transcribe to the WHILE language we study in our CT work. Once we extracted CTs from those transcribed programs, we are going to apply them to the existing CBMC proofs and measure the speed up. We are going to run each proof with two sets of bounds: Once with its original bounds and once with minimal bounds according to the extracted CTs. The runtime difference, will then be the speed up we can attribute to our CTs.

We are aware that CTs extracted from transcribed programs are not a perfect fit for the original C programs. There's always a risk that the transcription

misses some intricacies of  $C$  semantics. However, we are confident that these experiments are sufficiently precise to convey insights into the magnitude of speed up CTs allow for.

**Generalise CTs** At last, we are currently working on generalising CTs to arbitrary correctness properties. We found that this is straight-forward for the theoretical foundations of our work, i.e., the definition of CTs and all theorems presented in § 4 and appendix § C.

In appendix § D we present a generalised CT theory that considers an abstract notion of program correctness. Thereby, it covers both safety and liveness properties such as memory safety and termination, but also more complex functional correctness properties. In particular, it subsumes the memory safety CT theory presented in appendix § C. We present a generalised notion of precision (cf. Def. D.2.4) and a soundness theorem (cf. Theo D.2.11) for the generalised theory. The latter states that (analogous to the special case, cf. Theo. C.8.3) precision is a sufficiently strong requirement to ensure soundness of any extracted CT with respect to the original program specification. Moreover, we show that – following the intuition behind CTs – CTs are subdomains that cover error equivalence classes (cf. Theo. D.2.20). This justifies extracting CTs by factoring out potential error equivalence classes.

Most of the time, it is sufficient to treat correctness as an arbitrary predicate. The same can, however, not be said for their combinatorial properties. Those heavily depend on the studied correctness property. Memory safety CTs are fairly easy to combine, because memory safety is essentially a local property, as long as we do not change the memory layout relevant for the CT. We are planning to look into more correctness properties and see whether we can from that extract a locality requirement that allows for CT combinators similar to the memory safety ones.



# List of Contributions

## Published Papers

- **T. Reinhard**, J. Fasse, and B. Jacobs. *Completeness Thresholds for Memory Safety of Array Traversing Programs*. In Proceedings of the 12th ACM SIGPLAN International Workshop on the State of the Art in Program Analysis (SOAP) (2023), ACM, pp. 47–54. [doi.org/10.1145/3589250.3596143](https://doi.org/10.1145/3589250.3596143)
- **T. Reinhard**, and B. Jacobs. *Ghost Signals: Verifying Termination of Busy Waiting*. In Proceedings of 33rd International Conference on Computer-Aided Verification (CAV) (2021), vol. 12760 of Lecture Notes in Computer Science, Springer, pp. 27–50. [doi.org/10.1007/978-3-030-81688-9\\_2](https://doi.org/10.1007/978-3-030-81688-9_2)
- **T. Reinhard**, A. Timany, and B. Jacobs. *A Separation Logic to Verify Termination of Busy Waiting for Abrupt Program Exit*. In Proceedings of the 22nd ACM SIGPLAN International Workshop on Formal Techniques for Java-Like Programs (FTfJP) (2020), ACM, pp. 26–32. [doi.org/10.1145/3427761.3428345](https://doi.org/10.1145/3427761.3428345)
- P. Weisenburger, **T. Reinhard**, G. Salvaneschi. *Static Latency Tracking with Placement Types*. In Companion Proceedings for the ISSTA/ECOOP 2018 Workshops. (2018) ACM, pp. 34–36. [doi.org/10.1145/3236454.3236486](https://doi.org/10.1145/3236454.3236486)

## Posters and Competition Contributions

- **T. Reinhard**, J. Fasse, and B. Jacobs. *Completeness Thresholds for Memory Safety: Unbounded Guarantees Via Bounded Proofs*. ACM SIGPLAN International Conference on Systems, Programming, Languages and Applications: Software for Humanity (SPLASH) (2023), Poster Track. Zenodo [doi.org/10.5281/zenodo.10066197](https://doi.org/10.5281/zenodo.10066197)
- **T. Reinhard**. *A Core Calculus for Static Latency Tracking with Placement Types*. ACM SIGPLAN Student Research Competition at Symposium on Principles of Programming Languages (SRC@POPL) (2019), Extended Abstract Round. [doi.org/10.48550/arXiv.2007.15617](https://doi.org/10.48550/arXiv.2007.15617)
- **T. Reinhard**. *Type-Level Latency Tracking with Placement Types*. ACM SIGPLAN Student Research Competition at Symposium on Principles of Programming Languages (SRC@POPL) (2019), Poster Round. [doi.org/10.5281/zenodo.10160638](https://doi.org/10.5281/zenodo.10160638)

## Abstracts and Technical Reports

- **T. Reinhard**, J. Fasse, and B. Jacobs. *Completeness Thresholds for Memory Safety: Unbounded Guarantees via Bounded Proofs (Extended Abstract)*. (2023) [arxiv.org/abs/2309.09731](https://arxiv.org/abs/2309.09731)
- **T. Reinhard**. *Completeness Thresholds for Memory Safety of Array Traversing Programs: Early Technical Report*. (2023) [arxiv.org/abs/2211.11885](https://arxiv.org/abs/2211.11885)
- **T. Reinhard** and B. Jacobs. *Ghost Signals: Verifying Termination of Busy Waiting (Technical Report)*. Zenodo. (2021) [doi.org/10.5281/zenodo.4775181](https://doi.org/10.5281/zenodo.4775181)
- **T. Reinhard** and B. Jacobs. *A Separation Logic to Verify Termination of Busy-Waiting for Abrupt Program Exit: Technical Report*. (2020) [arxiv.org/abs/2007.10215](https://arxiv.org/abs/2007.10215)



## Invited Talks

- *Completeness Thresholds for Memory Safety: Unbounded Guarantees Via Bounded Proofs.* Programming Languages and Verification Seminar, Portland State University, 2023
- *Completeness Thresholds for Memory Safety: Unbounded Guarantees Via Bounded Proofs.* Programming Group, University of St. Gallen (HSG), 2023

## Research Internships

- *Verification of the Multicore Implementation of the FreeRTOS Scheduler.* Low-level Code Reasoning Team, Automated Reasoning Group, Amazon Web Services (2022)
- *Strategies to Progress From Bounded to Unbounded Verification.* Proof Platforms Team, Automated Reasoning Group, Amazon Web Services (2021)



# Appendix A

## Ghost Signals: Formalization & Case Studies

This chapter formalizes the work described in chapter 3 and presents additional case studies. The contents of this chapter were previously published on arXiv in the form of an appendix of the extended version of the ghost signal paper:

**T. Reinhard** and Bart Jacobs. *Ghost Signals: Verifying Termination of Busy Waiting (Extended Version)*. arXiv, CoRR, abs/2010.11762, 2021. <https://doi.org/10.48550/arXiv.2010.11762>

In this appendix, we formalize our approach and prove its soundness. We start in § A.1 by defining the notations we use. In § A.2 we define the simple programming language we consider in this work. In § A.3 we define our logic and in particular provide a full overview of all the proof and view shift rules we use and state a soundness theorem. In § A.4 we prove our approach sound. Our proof utilizes an annotated semantics that keeps track of ghost resources and thereby connects the runtime and the verification level. We define it in § A.4.1. Afterwards, in § A.4.2 we define a model for Hoare triples and prove that every triple that we can derive with our proof rules also holds in our model. In § A.4.3 we use this model and the annotated semantics to prove our soundness theorem. At last, in § A.5 we present two verification case studies. In § A.5.1 we provide a detailed proof outline for the realistic example presented in § 3.3. In § A.5.2 we present and verify a similar program but with a statically unbounded number of producer and consumer threads.

## A.1 General

**Definition A.1.1** (Projections). For any Cartesian product  $C = \prod_{i \in I} A_i$  and any index  $k \in I$ , we denote the  $k^{\text{th}}$  projection by  $\pi_k^C : \prod_{i \in I} A_i \rightarrow A_k$ . We define

$$\pi_k^C((a_i)_{i \in I}) := a_k.$$

In case the domain  $C$  is clear from the context, we write  $\pi_k$  instead of  $\pi_k^C$ .

**Definition A.1.2** (Disjoint Union). Let  $A, B$  be sets. We define their disjoint union as

$$A \sqcup B := A \cup B$$

if  $A \cap B = \emptyset$  and leave it undefined otherwise.

**Definition A.1.3** (Bags). For any set  $X$  we define the set of bags  $\text{Bags}(X)$  and the set of finite bags  $\text{Bags}_{\text{fin}}(X)$  over  $X$  as

$$\begin{aligned} \text{Bags}(X) &:= X \rightarrow \mathbb{N}, \\ \text{Bags}_{\text{fin}}(X) &:= \{B \in \text{Bags}(X) \mid \{x \in B \mid B(x) > 0\} \text{ finite}\}. \end{aligned}$$

We define union and subtraction of bags as

$$\begin{aligned} (B_1 \uplus B_2)(x) &:= B_1(x) + B_2(x), \\ (B_1 \setminus B_2)(x) &:= \max(0, B_1(x) - B_2(x)). \end{aligned}$$

For finite bags where the domain is clear from the context, we define the following set-like notation:

$$\begin{aligned} \emptyset &:= x \mapsto 0, \\ \{x\} &:= \begin{cases} x &\mapsto 1 \\ y &\mapsto 0 \end{cases} \text{ for } y \neq x, \\ \{x_1, \dots, x_n\} &:= \bigoplus_{i=1}^n \{x_i\}. \end{aligned}$$

Further, we define for any  $n \in \mathbb{N}$  the notation:

$$n \cdot \{x\} := \underbrace{\{x, \dots, x\}}_{n \text{ times}}.$$

We define the following set-like notations for element and subset relationship:

$$\begin{aligned} x \in B &\Leftrightarrow B(x) > 0, \\ B_1 \subseteq B_2 &\Leftrightarrow \forall x \in B_1. B_1(x) \leq B_2(x), \\ B_1 \subset B_2 &\Leftrightarrow \exists C \subseteq B_1. C \neq \emptyset \wedge B_1 = B_2 \setminus C. \end{aligned}$$

For any bag  $B \in \text{Bags}(X)$  and predicate  $P \subseteq X$  we define the following refinement:

$$\llbracket x \in B \mid P(x) \rrbracket := \begin{cases} x \mapsto B(x) & \text{if } P(x), \\ x \mapsto 0 & \text{otherwise.} \end{cases}$$

**Definition A.1.4** (Disjoint Union). *Let  $A, B$  be sets. We define their disjoint union as*

$$A \sqcup B := A \cup B$$

*if  $A \cap B = \emptyset$  and leave it undefined otherwise.*

## A.2 Language

In this section and the next, we present our approach formally. In this section, we define the programming language; in § A.3 we define the proof system.

We consider a simple imperative programming language with support for multi-threading, shared memory and synchronization via mutexes. For its definition we assume (i) an infinite set of program variables  $x \in \mathcal{X}$ , (ii) an infinite set of heap locations  $\ell \in \mathcal{Locs}$ , (iii) a set of values  $v \in \mathcal{Values}$  which includes heap locations, Booleans  $\mathbb{B} = \{\text{True}, \text{False}\}$  and the unit value  $\text{tt}$ , (iv) a set of operations  $op \in \mathcal{Ops}$  and (v) an infinite, totally ordered and well-founded set of thread IDs  $\theta \in \Theta$ .

**Definition A.2.1** (Syntax). *We define the sets of commands  $\mathit{Cmds}$  and expressions  $\mathit{Exps}$  according to the syntax presented in Fig. A.1.*

The language contains standard sets of pure expressions  $\mathit{Exps}$  and (potentially) side-effectful commands  $\mathit{Cmds}$ . The latter includes commands for heap allocation and manipulation, forking and loops. We define *physical heaps* [100] (as opposed to *logical heaps* [100] presented in the next section) as a finite set of *physical resource chunks*. A points-to chunk  $\ell \mapsto v$  expresses that heap location  $\ell$  points to value  $v$  [155, 100]. Moreover, we have chunks to represent unlocked and locked mutexes.

**Definition A.2.2** (Physical Resources). *We define the set of physical resources  $\mathcal{R}^{\text{phys}}$  syntactically as follows:*

$$\begin{aligned} r^{\text{p}} \in \mathcal{R}^{\text{phys}} & ::= \ell \mapsto v \mid \text{unlocked}_{\text{pRes}}(\ell) \mid \text{locked}_{\text{pRes}}(\ell) \\ & \ell \in \mathcal{Locs} \qquad v \in \mathcal{Values} \end{aligned}$$

$$x \in \mathcal{X} : \text{Program variables} \quad \ell \in \mathcal{Locs} : \text{Heap locations}$$

$$v \in \text{Values} \supseteq \{\text{tt}\} \cup \mathbb{B} \cup \mathcal{Locs} : \text{Values} \quad op \in \text{Ops} : \text{Operations}$$

(a) Assumed sets and variables.  $\mathcal{X}$  and  $\mathcal{Locs}$  infinite.

$$e \in \text{Exps} \quad ::= \quad x \mid v \mid e = e \mid \neg e \mid op(\bar{e})$$

$$c \in \text{Cmds} \quad ::= \quad e \mid \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \mid \mathbf{fork} \ c \mid \mathbf{let} \ x := c \ \mathbf{in} \ c \mid$$

$$\mathbf{if} \ c \ \mathbf{then} \ c \mid \mathbf{cons}(e) \mid [e] \mid [e] := e \mid \mathbf{new\_mutex} \mid$$

$$\mathbf{acquire} \ e \mid \mathbf{release} \ e$$

$$E \in \text{EvalCtxts} \quad ::= \quad \mathbf{if} \ \square \ \mathbf{then} \ c \mid \mathbf{let} \ x := \square \ \mathbf{in} \ c$$

(b) Expressions and commands.

$$\mathbf{with} \ e \ \mathbf{await} \ c \quad ::= \quad (\mathbf{while} \ \mathbf{acquire} \ e; \mathbf{let} \ r := c \ \mathbf{in} \ \mathbf{release} \ e; \neg r \ \mathbf{do} \ \mathbf{skip})$$

$$c; c' \quad ::= \quad (\mathbf{let} \ r := c \ \mathbf{in} \ c')$$

$$e \neq e' \quad ::= \quad \neg(e = e')$$

$$e''([e]) \quad ::= \quad (\mathbf{let} \ r := [e] \ \mathbf{in} \ e''(r))$$

(c) Syntactic sugar.  $e''(y)$  expression with free variable  $y$ .  $r \in \mathcal{X}$  not free in  $e, e'', c'$ .

Figure A.1: Syntax.

**Definition A.2.3** (Physical Heaps). *We define the set of physical heaps as*

$$\text{Heaps}^{\text{phys}} := \mathcal{P}_{\text{fin}}(\mathcal{R}^{\text{phys}})$$

*and the function  $\text{locs}_{\text{pRes}} : \text{Heaps}^{\text{phys}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$  mapping physical heaps to the sets of allocated heap locations as*

$$\text{locs}_{\text{pRes}}(h) \quad ::= \quad \{\ell \in \mathcal{Locs} \mid \text{unlocked}_{\text{pRes}}(\ell) \in h \vee \text{locked}_{\text{pRes}}(\ell) \in h \vee$$

$$\exists v \in \text{Values}. \ell \mapsto v \in h\}.$$

*We denote physical heaps by  $h$ .*

We represent a program state by a physical heap and a *thread pool*, which we define as a partial function mapping a finite number of thread IDs to threads. Thread IDs are unique and never reused. Hence, we represent running threads by commands and terminated ones by **term** instead of removing threads from the pool. For the following definition, remember that we assume the set of thread IDs  $\Theta$  to be infinite and well-founded.

**Definition A.2.4** (Thread Pools). *We define the set of thread pools  $\mathcal{TP}$  as the set of finite partial functions mapping thread IDs to threads:*

$$\mathcal{TP} := \Theta \rightarrow_{\text{fin}} (\text{Cmds} \cup \{\text{term}\}).$$

The symbol `term` represents a terminated thread. We denote thread pools by  $P$ , thread IDs by  $\theta$  and the empty thread pool by  $\emptyset_{\text{tp}}$ , i.e.,

$$\begin{aligned} \emptyset_{\text{tp}} &: \Theta \rightarrow_{\text{fin}} (\text{Cmds} \cup \{\text{term}\}), \\ \text{dom}(\emptyset_{\text{tp}}) &= \emptyset. \end{aligned}$$

We define the operation  $+_{\text{tp}} : \mathcal{TP} \times \{C \subset \text{Cmds} \mid |C| \leq 1\} \rightarrow \mathcal{TP}$  as follows:

$$\begin{aligned} P +_{\text{tp}} \emptyset &:= P, \\ P +_{\text{tp}} \{c\} &:= P[\theta_{\text{new}} := c] \quad \text{for } \theta_{\text{new}} := \min(\Theta \setminus \text{dom}(P)). \end{aligned}$$

We define the operational semantics of our language in terms of two small-step reduction relations:  $\rightsquigarrow_{\text{st}}$  for single threads and  $\rightsquigarrow_{\text{tp}}$  for thread pools. Since expressions are pure and their evaluation is deterministic we identify closed expressions with their ascribed value. (i)  $h, c \rightsquigarrow_{\text{st}} h', c', T$  expresses that heap  $h$  and command  $c$  are reduced in a single step to  $h'$  and  $c'$  and that this thread forks a set of threads  $T$ . This set is either empty or a singleton as no step forks more than one thread. (ii)  $h, P \xrightarrow{\theta}_{\text{tp}} h', P'$  expresses that heap  $h$  and thread pool  $P$  are reduced in a single step to  $h'$  and  $P'$ . ID  $\theta$  identifies the thread reduced in this step.

**Definition A.2.5** (Evaluation of Closed Expressions). *We define a partial evaluation function  $\llbracket \cdot \rrbracket : \text{Exps} \rightarrow \text{Values}$  on expressions by recursion on the structure of expressions as follows:*

$$\begin{aligned} \llbracket v \rrbracket &:= v && \text{if } v \in \text{Values}, \\ \llbracket e = e' \rrbracket &:= \text{True} && \text{if } \llbracket e \rrbracket = \llbracket e' \rrbracket \neq \perp, \\ \llbracket e = e' \rrbracket &:= \text{False} && \text{if } \llbracket e \rrbracket \neq \llbracket e' \rrbracket \wedge \llbracket e \rrbracket \neq \perp \wedge \llbracket e' \rrbracket \neq \perp, \\ \llbracket \neg e \rrbracket &:= \text{False} && \text{if } \llbracket e \rrbracket = \text{True}, \\ \llbracket \neg e \rrbracket &:= \text{True} && \text{if } \llbracket e \rrbracket = \text{False}, \\ \llbracket e \rrbracket &:= \perp && \text{otherwise.} \end{aligned}$$

We identify closed expressions  $e$  with their ascribed value  $\llbracket e \rrbracket$ .

**Definition A.2.6** (Evaluation Context). *We define the set of evaluation contexts  $\text{EvalCtxts}$  syntactically as follows:*

$$\begin{aligned} E \in \text{EvalCtxts} &::= \text{if } \square \text{ then } c \mid \text{let } x := \square \text{ in } c \\ c \in \text{Cmds} & \quad x \in \mathcal{X} \end{aligned}$$

For any  $c \in \text{Cmds}$  and  $E \in \text{EvalCtxts}$ , we define  $E[c] := E[\square \mapsto c]$ .

Note that for every  $c \in \text{Cmds}$ ,  $E \in \text{EvalCtxts}$ , we have  $E[c] \in \text{Cmds}$ .

**Definition A.2.7** (Single Thread Reduction Relation). *We define a reduction relation  $\rightsquigarrow_{\text{st}}$  for single threads according to the rules presented in Figure A.2. A reduction step has the form*

$$h, c \rightsquigarrow_{\text{st}} h', c', T$$

for a set of forked threads  $T \subset \text{Cmds}$  with  $|T| \leq 1$ .

For simplicity of notation, we omit  $T$  if it is clear from the context that no thread is forked and  $T = \emptyset$ .

**Definition A.2.8** (Thread Pool Reduction Relation). *We define a thread pool reduction relation  $\rightsquigarrow_{\text{tp}}$  according to the rules presented in Figure A.3. A reduction step has the form*

$$h, P \xrightarrow{\theta}_{\text{tp}} h', P'.$$

As thread scheduling is non-deterministic, so is our thread pool reduction relation  $\rightsquigarrow_{\text{tp}}$ . Consider the minimal example we presented in Fig. 3.7a in § 3.2.3 to illustrate busy waiting for a shared memory flag to be set. It does not terminate if the main thread is never scheduled after the new thread was forked. Hence, our verification approach relies on the assumption of fair scheduling. That is, we assume that every thread is always eventually scheduled while it remains running. Further, we represent program executions by sequences of reduction steps. As we primarily consider infinite sequences in this work, we define reduction sequences to be infinite to simplify our terminology.

**Definition A.2.9** (Reduction Sequence). *Let  $(h_i)_{i \in \mathbb{N}}$  and  $(P_i)_{i \in \mathbb{N}}$  be infinite sequences of physical heaps and thread pools, respectively. We call  $(h_i, P_i)_{i \in \mathbb{N}}$  a reduction sequence if there exists a sequence of thread IDs  $(\theta_i)_{i \in \mathbb{N}}$  such that  $h_i, P_i \xrightarrow{\theta_i}_{\text{tp}} h_{i+1}, P_{i+1}$  holds for every  $i \in \mathbb{N}$ .*

**Definition A.2.10** (Fairness). *We call a reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  fair iff for all  $i \in \mathbb{N}$  and  $\theta \in \text{dom}(P_i)$  with  $P_i(\theta) \neq \text{term}$  there exists some  $k \geq i$  with*

$$h_k, P_k \xrightarrow{\theta}_{\text{tp}} h_{k+1}, P_{k+1}.$$

## A.3 Logic

In this section we formalize the logic we sketched in § 3.2. For the definition we assume (i) an infinite set of ghost signal IDs  $id \in \mathcal{ID}$  and (ii) an infinite, partially ordered and well-founded set of levels  $L \in \mathcal{Levs}$ . We denote the level order relation by  $<_{\mathcal{L}}$ .



$$\begin{array}{c}
\text{ST-RED-EVALCTXT} \\
\frac{h, c \rightsquigarrow_{\text{st}} h', c', T}{h, E[c] \rightsquigarrow_{\text{st}} h', E[c'], T} \\
\\
\text{ST-RED-FORK} \qquad \text{ST-RED-IFTRUE} \\
h, \mathbf{fork} \ c \rightsquigarrow_{\text{st}} h, \mathbf{tt}, \{c\} \qquad h, \mathbf{if} \ \mathbf{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, c \\
\\
\text{ST-RED-IFFALSE} \qquad \text{ST-RED-LET} \\
h, \mathbf{if} \ \mathbf{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, \mathbf{tt} \qquad h, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{st}} h, c[x \mapsto v] \\
\\
\text{ST-RED-WHILE} \\
h, \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rightsquigarrow_{\text{st}} h, \mathbf{if} \ c \ \mathbf{then} \ \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \\
\\
\text{ST-RED-CONS} \qquad \text{ST-RED-READHEAPLOC} \\
\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{cons}(v) \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \ell} \qquad \frac{\ell \mapsto v \in h}{h, [\ell] \rightsquigarrow_{\text{st}} h, v} \\
\\
\text{ST-RED-ASSIGN} \\
h \sqcup \{\ell \mapsto v'\}, [\ell] := v \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \mathbf{tt} \\
\\
\text{ST-RED-NEWMUTEX} \\
\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{new\_mutex} \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \ell} \\
\\
\text{ST-RED-ACQUIRE} \\
h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \mathbf{acquire} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \mathbf{tt} \\
\\
\text{ST-RED-RELEASE} \\
h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \mathbf{release} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \mathbf{tt}
\end{array}$$

Figure A.2: Single thread reduction rules.

$$\begin{array}{c}
\text{TP-RED-LIFT} \qquad \text{TP-RED-TERM} \\
\frac{P(\theta) = c \quad h, c \rightsquigarrow_{\text{st}} h', c', T}{h, P \xrightarrow{\theta}_{\text{tp}} h', P[\theta := c'] +_{\text{tp}} T} \qquad \frac{P(\theta) = v}{h, P \xrightarrow{\theta}_{\text{tp}} h, P[\theta := \text{term}]}
\end{array}$$

Figure A.3: Thread pool reduction rules

**Definition A.3.1** (Fractions). *We define the set of fractions as*

$$\mathcal{F} := \{f \in \mathbb{Q} \mid 0 < f \leq 1\}$$

and denote fractions by  $f$ .

**Definition A.3.2** (Obligations, Signals & Mutexes). *We define the set of obligations  $\mathcal{O}$  and signals  $\mathcal{S}$  as*

$$\begin{aligned} \mathcal{O} &:= (\mathcal{Locs} \cup \mathcal{ID}) \times \mathcal{Levs}, \\ \mathcal{S} &:= \mathcal{ID} \times \mathcal{Levs}. \end{aligned}$$

We denote bags of obligations by  $O \in \text{Bags}(\mathcal{O})$ , signals by  $s \in \mathcal{S}$  and mutexes by  $m \in \mathcal{Locs} \times \mathcal{Levs}$ . For convenience of notation, we define the following selector functions:

$$\begin{aligned} (id, \_).id &:= id && \text{for signals,} \\ (\ell, \_).loc &:= \ell && \text{for mutexes,} \\ (\_, L).lev &:= L && \text{for signals, mutexes and obligations.} \end{aligned}$$

**Definition A.3.3** (Assertions). *We define the set of assertions  $\mathcal{A}$  according to the syntax presented in Figure A.4.<sup>1</sup> We omit the index set  $I$  in quantifications when its choice becomes clear from the context and write  $\exists i. a(i)$  and  $\forall i. a(i)$  instead of  $\exists i \in I. a(i)$  and  $\forall i \in I. a(i)$ , respectively.*

In § A.2, we used physical resources and heaps to model a program's state. We use assertions to capture which fraction of a physical resource and which ghost resources a thread owns. Therefore, we have to interpret assertions in an extended notion of state. We define *logical heaps* and *logical resources* [100] which correspond to physical ones but additionally encompass ghost resources and ownership. For instance, logical resources include signal chunks and initialized mutex chunks are associated with a lock invariant. Rather than being a set of resources, logical heaps map logical resources to fractions. This allows us to express which portion of a resource a thread owns.

**Definition A.3.4** (Logical Resources). *We define the set of logical resources  $\mathcal{R}^{\text{log}}$  syntactically as follows:*

$$r^l \in \mathcal{R}^{\text{log}} ::= \ell \mapsto v \mid \text{signal}_{\text{Res}}((id, L), b) \mid \text{uninit}_{\text{Res}}(\ell) \mid \text{mutex}_{\text{Res}}((\ell, L), a) \mid \text{locked}_{\text{Res}}((\ell, L), a, f) \mid \text{obs}_{\text{Res}}(O)$$

<sup>1</sup>That is, we define  $\mathcal{A}$  as the least fixpoint of  $F$  where  $F(A) = \{\text{True}, \text{False}\} \cup \{\neg a \mid a \in A\} \cup \{a_1 \wedge a_2 \mid a_1, a_2 \in A\} \cup \dots \cup \{\bigvee A' \mid A' \subseteq A\} \cup \dots$ . Since  $F$  is a monotonic function over a complete lattice, it has a least fixpoint according to the Knaster-Tarski theorem [169].

$$\begin{array}{l}
 A \subseteq \mathcal{A} \quad \text{Index set } I \\
 a \in \mathcal{A} \quad ::= \text{ True } \mid \text{ False } \mid \neg a \mid a \wedge a \mid a \vee a \mid a * a \mid [f]\ell \mapsto v \\
 \quad \quad \quad \mid \bigvee A \\
 \quad \quad \quad \mid \text{ uninit}(\ell) \mid [f]\text{mutex}((\ell, L), a) \mid \text{locked}((\ell, L), a, f) \\
 \quad \quad \quad \mid \text{signal}((id, L), b) \mid \text{obs}(O)
 \end{array}$$

(a) Assertion syntax. We omit quantification domain  $I$  when it is clear from the context.

$$\begin{array}{l}
 a_1 \rightarrow a_2 \quad ::= \neg a_1 \vee a_2 \\
 a_1 \leftrightarrow a_2 \quad ::= (a_1 \rightarrow a_2) \wedge (a_2 \rightarrow a_1) \\
 \exists i \in I. a(i) \quad ::= \bigvee \{a(i) \mid i \in I\} \\
 \forall i \in I. a(i) \quad ::= \neg \exists i \in I. \neg a(i)
 \end{array}$$

(b) Syntactic sugar.

Figure A.4: Assertions.

Further, we define the function  $\text{getHLocs}_{\text{IRes}} : \mathcal{R}^{\text{log}} \rightarrow \mathcal{Locs}$  mapping logical resources to their respective (either empty or singleton) set of involved heap locations as

$$\begin{array}{l}
 \text{getHLocs}_{\text{IRes}}(\ell \mapsto v) \quad \quad \quad ::= \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{uninit}_{\text{IRes}}(\ell)) \quad \quad \quad ::= \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{mutex}_{\text{IRes}}((\ell, L), a)) \quad \quad \quad ::= \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{locked}_{\text{IRes}}((\ell, L), a, f)) \quad \quad \quad ::= \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\_) \quad \quad \quad ::= \emptyset \quad \text{otherwise.}
 \end{array}$$

**Definition A.3.5** (Logical Heaps). We define the set of logical heaps as

$$H \in \text{Heaps}^{\text{log}} := \mathcal{R}^{\text{log}} \rightarrow \{q \in \mathbb{Q} \mid q \geq 0\}.$$

We define the empty logical heap  $\emptyset_{\text{log}}$  as the constant zero function

$$\emptyset_{\text{log}} : r^{\text{l}} \mapsto 0.$$

We denote logical heaps by  $H$ , point-wise addition by  $+$  and multiplication with non-negative rationals by  $\cdot$ , i.e.,

$$\begin{array}{l}
 (H_1 + H_2)(r^{\text{l}}) \quad ::= H_1(r^{\text{l}}) + H_2(r^{\text{l}}), \\
 (q \cdot H)(r^{\text{l}}) \quad ::= q \cdot (H(r^{\text{l}}))
 \end{array}$$

for  $q \in \mathbb{Q}$  with  $q \geq 0$ . We give  $\cdot$  a higher precedence than  $+$ . For convenience of notation we represent logical heaps containing finitely many resources by sets of resources as follows

$$\{r_1^{\text{l}}, \dots, r_n^{\text{l}}\} \quad ::= \begin{cases} r_i^{\text{l}} & \mapsto 1 \\ x & \mapsto 0 \quad \text{if } x \notin \{r_1^{\text{l}}, \dots, r_n^{\text{l}}\}. \end{cases}$$

**Definition A.3.6** (Logical Heap Predicates). *Let  $H$  be a logical heap. We call  $H$  complete and write  $\text{complete}_{\text{lh}}(H)$  if it contains exactly one obligations chunk, i.e., if there exists a bag of obligations  $O$  with  $H(\text{obs}_{\text{S|Res}}(O)) = 1$  and if there does not exist any bag of obligations  $O'$  with  $O \neq O'$  and  $H(\text{obs}_{\text{S|Res}}(O')) > 0$ .*

*We call  $H$  finite and write  $\text{finite}_{\text{lh}}(H)$  if it contains only finitely many resources, i.e., if the set  $\{r^l \in \mathcal{R}^{\text{log}} \mid H(r^l) > 0\}$  is finite.*

*We call  $H$  consistent and write  $\text{consistent}_{\text{lh}}(H)$  if (i) it contains only full obligations chunks, i.e., if*

$$H(\text{obs}_{\text{S|Res}}(O)) \in \mathbb{N}$$

*holds for all  $O \in \text{Bags}(\mathcal{O})$  and if (ii) heap locations are unique in  $H$ , i.e., if there are no  $r_1^l, r_2^l \in \mathcal{R}^{\text{log}}$  with  $r_1^l \neq r_2^l$ ,  $H(r_1^l) > 0$ ,  $H(r_2^l) > 0$  and with  $\text{getHLocs}_{\text{S|Res}}(r_1^l) \cap \text{getHLocs}_{\text{S|Res}}(r_2^l) \neq \emptyset$ .*

We interpret assertions in terms of a model relation.  $H \models_{\mathcal{A}} a$  expresses that assertion  $a$  holds with respect to logical heap  $H$ . Further, we define the view shift relation  $\Rightarrow$  and the proof relation  $\vdash$  we sketched in § 3.2.

**Definition A.3.7** (Assertion Model Relation). *We define a model relation  $\models_{\mathcal{A}} \subset \text{Heaps}^{\text{log}} \times \mathcal{A}$  for assertions by recursion on the structure of assertions according to the rules presented in Figure A.5. We write  $H \models_{\mathcal{A}} a$  to express that logical heap  $H$  models assertion  $a$  and  $H \not\models_{\mathcal{A}} a$  to express that  $H \models_{\mathcal{A}} a$  does not hold.*

**Definition A.3.8.** *Let  $L \in \mathcal{L}e\text{vs}$ ,  $O \in \text{Bags}(\mathcal{O})$ . We define  $\prec_{\mathcal{L}} \subset \mathcal{L}e\text{vs} \times \mathcal{O}$  as*

$$L \prec_{\mathcal{L}} O \iff \forall o \in O. L \prec_{\mathcal{L}} o.\text{lev.}$$

**Definition A.3.9** (View Shift). *We define a view shift relation  $\Rightarrow \subset \mathcal{A} \times \mathcal{A}$  according to the rules presented in Fig. A.6.*

**Definition A.3.10** (Proof Relation). *We define a proof relation  $\vdash \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$  according to the rules presented in Fig. A.7 and A.8. We state the provability of a Hoare triple in the form of  $\vdash \{A\} c \{\lambda r. B(r)\}$  where  $r$  captures the value returned by  $c$ . To simplify the notation, we omit the result value if it is clear from the context or irrelevant.*

Following the intuition provided in § 3.2, we can prove that a program terminates by proving that it discharges all obligations. The following theorem states that this approach is sound.

**Theorem A.3.11** (Soundness). *Let  $\vdash \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$  hold. There exists no fair, infinite reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  with  $h_0 = \emptyset$  and  $P_0 = \{(\theta_0, c)\}$  for any choice of  $\theta_0$ .*

$H \vDash_{\mathbf{A}} \text{True}$	
$H \not\vDash_{\mathbf{A}} \text{False}$	
$H \vDash_{\mathbf{A}} \neg a$	if $H \not\vDash_{\mathbf{A}} a$
$H \vDash_{\mathbf{A}} a_1 \wedge a_2$	if $H \vDash_{\mathbf{A}} a_1 \wedge H \vDash_{\mathbf{A}} a_2$
$H \vDash_{\mathbf{A}} a_1 \vee a_2$	if $H \vDash_{\mathbf{A}} a_1 \vee H \vDash_{\mathbf{A}} a_2$
$H \vDash_{\mathbf{A}} a_1 * a_2$	if $\exists H_1, H_2 \in \text{Heaps}^{\text{log}}. H = H_1 + H_2 \wedge$ $H_1 \vDash_{\mathbf{A}} a_1 \wedge H_2 \vDash_{\mathbf{A}} a_2$
$H \vDash_{\mathbf{A}} [f]\ell \mapsto v$	if $H(\ell \mapsto v) \geq f$
$H \vDash_{\mathbf{A}} \bigvee A$	if $\exists a \in A. H \vDash_{\mathbf{A}} a$
$H \vDash_{\mathbf{A}} [f]\text{uninit}(\ell)$	if $H(\text{uninit}_{\text{IRes}}(\ell)) \geq f$
$H \vDash_{\mathbf{A}} [f]\text{mutex}(m, P)$	if $H(\text{mutex}_{\text{IRes}}(m, P)) \geq f$
$H \vDash_{\mathbf{A}} [f]\text{locked}(m, P, f_u)$	if $H(\text{locked}_{\text{IRes}}(m, P, f_u)) \geq f$
$H \vDash_{\mathbf{A}} [f]\text{signal}(s, b)$	if $H(\text{signal}_{\text{IRes}}(s, b)) \geq f$
$H \vDash_{\mathbf{A}} \text{obs}(O)$	if $H(\text{obs}_{\text{IRes}}(O)) \geq 1$

Figure A.5: Assertion model relation.  $H \not\vDash_{\mathbf{A}} a$  expresses that  $H \vDash_{\mathbf{A}} a$  does not hold.

$\frac{\text{VS-SEMIMP} \quad \forall H. \text{consistent}_{\text{lh}}(H) \wedge H \vDash_{\mathbf{A}} A \Rightarrow H \vDash_{\mathbf{A}} B}{A \Rightarrow B}$	$\frac{\text{VS-TRANS} \quad A \Rightarrow C \quad C \Rightarrow B}{A \Rightarrow B}$
$\frac{\text{VS-OR} \quad A_1 \Rightarrow B \quad A_2 \Rightarrow B}{A_1 \vee A_2 \Rightarrow B}$	
$\frac{\text{VS-NEWSIGNAL} \quad L \in \mathcal{L}evs}{\text{obs}(O) \Rightarrow \exists id. \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})}$	
$\text{VS-SET SIGNAL} \quad \text{obs}(O \uplus \{s\}) * \text{signal}(s, \_) \Rightarrow \text{obs}(O) * \text{signal}(s, \text{True})$	
$\frac{\text{VS-MUTINIT} \quad L \in \mathcal{L}evs}{\text{uninit}(\ell) * P \Rightarrow \text{mutex}((\ell, L), P)}$	$\frac{\text{VS-GHOSTLOOP} \quad \forall n'. (I(n') \wedge n' > 0 \Rightarrow I(n' - 1))}{I(n) \wedge n \geq 0 \Rightarrow I(0)}$

Figure A.6: View shift rules.

$$\frac{\text{PR-FRAME} \quad \vdash \{A\} c \{B\}}{\vdash \{A * F\} c \{B * F\}} \quad \frac{\text{PR-VIEWSHIFT} \quad A \Rightarrow A' \quad \vdash \{A'\} c \{B'\} \quad B' \Rightarrow B}{\vdash \{A\} c \{B\}}$$

$$\frac{\text{PR-EXP} \quad \llbracket e \rrbracket \in \text{Values}}{\vdash \{\text{True}\} e \{\lambda r. r = \llbracket e \rrbracket\}} \quad \frac{\text{PR-EXISTS} \quad \forall a \in A. \vdash \{a\} c \{B\}}{\vdash \{\bigvee A\} c \{B\}}$$

$$\frac{\text{PR-FORK} \quad \vdash \{\text{obs}(O_f) * A\} c \{\text{obs}(\emptyset)\}}{\vdash \{\text{obs}(O_m \uplus O_f) * A\} \text{fork } c \{\lambda r. \text{obs}(O_m) * r = \text{tt}\}}$$

(a) Basic Proof Rules.

$$\frac{\text{PR-IF} \quad \begin{array}{l} \vdash \{A\} c_b \{\lambda b. C(b) \wedge (b = \text{True} \vee b = \text{False})\} \\ \vdash \{C(\text{True})\} c_t \{B\} \quad C(\text{False}) \Rightarrow B \end{array}}{\vdash \{A\} \text{if } c_b \text{ then } c_t \{B\}}$$

$$\frac{\text{PR-AWAIT} \quad \begin{array}{l} S \subset_{\text{fin}} \mathcal{S} \quad \forall O''. (\exists s. \text{signal}(s, \text{False}) * R(s, O'') \Rightarrow I(O'') * P) \\ \{\text{obs}(O' \uplus \{m\}) * I(O') * P\} \\ \forall O'. \vdash c \left\{ \begin{array}{l} \lambda b. \exists O''. \text{obs}(O'' \uplus \{m\}) \\ * \text{if } b \\ \text{then } P * B(O'') \\ \text{else } \exists s \in S. \text{signal}(s, \text{False}) * R(s, O'') * s.\text{lev} \prec_L O'' \\ * m.\text{lev} \prec_L O'' \end{array} \right\} \end{array}}{\vdash \begin{array}{l} \{\text{obs}(O) * [f]\text{mutex}(m, P) * I(O) * m.\text{lev} \prec_L O\} \\ \text{with } m \text{ await } c \\ \{\exists O'. \text{obs}(O') * [f]\text{mutex}(m, P) * B(O')\} \end{array}}$$

$$\frac{\text{PR-WHILE-DEC} \quad \forall n'. \vdash \{I(n')\} c_b \{\lambda b. \text{if } b \text{ then } \exists n''. 0 \leq n'' < n' \wedge I(n'') \text{ else } B\}}{\vdash \{I(n) \wedge n \geq 0\} \text{while } c_b \text{ do skip } \{B\}}$$

$$\frac{\text{PR-LET} \quad \vdash \{A\} c \{\lambda r. C(r)\} \quad \forall v. \vdash \{C(v)\} c'[x \mapsto v] \{B\}}{\vdash \{A\} \text{let } x := c \text{ in } c' \{B\}}$$

(b) Control Structures.

Figure A.7: Proof rules (part 1).

$$\begin{array}{c}
\text{PR-ACQUIRE} \\
\frac{m.\text{lev} \prec_L O}{\{ \text{obs}(O) * [f]\text{mutex}(m, P) \}} \\
\vdash \text{acquire } m.\text{loc} \\
\{ \lambda r. r = \text{tt} * \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * P \} \\
\\
\text{PR-RELEASE} \\
\{ \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * A \} \\
\vdash \text{release } m.\text{loc} \\
\{ \text{obs}(O) * [f]\text{mutex}(m, P) * B \} \\
\\
\text{PR-NEWMUTEX} \\
\vdash \{ \text{True} \} \text{new\_mutex } \{ \lambda \ell. \text{uninit}(\ell) \}
\end{array}$$

(a) Mutexes.

$$\begin{array}{c}
\text{PR-CONS} \qquad \qquad \qquad \text{PR-READHEAPLOC} \\
\vdash \{ \text{True} \} \text{cons}(v) \{ \lambda \ell. \ell \mapsto v \} \quad \vdash \{ [f]\ell \mapsto v \} [\ell] \{ \lambda r. r = v * [f]\ell \mapsto v \} \\
\\
\text{PR-ASSIGNTOHEAP} \\
\vdash \{ \ell \mapsto \_ \} [\ell] := v \{ \lambda r. r = \text{tt} * \ell \mapsto v \}
\end{array}$$

(b) Heap Access.

Figure A.8: Proof rules (part 2).

## A.4 Soundness

We already sketched the high-level intuition behind our soundness argument in § 3.2.1. In this section, we prove it formally. In § A.4.1 we define an annotated reduction semantics that tracks which resources threads own (including ghost resources). In § A.4.2 we define a model relation for Hoare triples and prove that every specification we can derive with our proof rules also holds in our model. In § A.4.3 we use the annotated semantics and the model relation to prove the soundness theorem we stated above.

### A.4.1 Annotated Semantics

During runtime, all threads share one physical heap where every thread is free to access every resource. This does not reflect the notions of ownership and

$$\begin{array}{l}
c \in \mathit{Cmds}^+ \quad ::= \quad e \mid \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \mid \mathbf{fork} \ c \mid \\
\mathbf{let} \ x := c \ \mathbf{in} \ c \mid \mathbf{if} \ c \ \mathbf{then} \ c \mid \mathbf{cons}(e) \mid \\
[e] \mid [e] := e \mid \\
\mathbf{new\_mutex} \mid \mathbf{acquire} \ e \mid \mathbf{release} \ e \mid \\
\mathbf{with} \ S, \ell \ \mathbf{await} \ c \mid \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip} \\
n \in \mathbb{N} \\
S \subseteq_{\text{fin}} \mathcal{S} \quad \text{finite set of signals}
\end{array}$$

Figure A.9: Extended set of commands for intermediate representation.

lock invariants which we maintain on the verification level. It also does not allow us to restrict actions based on levels, e.g., only allowing the acquisition of a mutex if its level is lower than the level of each held obligation. Hence:

- (i) We annotate every thread by a logical heap to express which resources it owns (including ghost resources) and thereby obtain an *annotated thread pool*.
- (ii) We represent the program state by an *annotated heap* that keeps track of lock invariants and levels. In particular, we associate unlocked mutexes with logical heaps to represent the resources they protect. Since annotated heaps keep track of levels, they also keep track of signals.

**Definition A.4.1** (Intermediate Representation). *We define an extended set of commands  $\mathit{Cmds}^+$  according to the syntax presented in Figure A.9.*

For the rest of this appendix, commands  $c$  refer to the extended set of commands, i.e.,  $c \in \mathit{Cmds}^+$ .

**Definition A.4.2** (Annotated Resources). *We define the set of annotated resources  $\mathit{AnnoRes}$  as*

$$\begin{array}{l}
r^a \in \mathit{AnnoRes} \quad ::= \quad \ell \mapsto v \mid \mathbf{uninit}_{\text{aRes}}(\ell) \mid \\
\mathbf{unlocked}_{\text{aRes}}((\ell, L), a, H) \mid \mathbf{locked}_{\text{aRes}}((\ell, L), a, f) \mid \\
\mathbf{signal}_{\text{aRes}}((id, L), b)
\end{array}$$

where  $H$  does not contain any obligations chunks.

**Definition A.4.3** (Annotated Heaps). *We define the set of annotated heaps as*

$$\mathit{Heaps}^{\text{annot}} := \mathcal{P}_{\text{fin}}(\mathit{AnnoRes}),$$



the function  $\text{locs}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$  mapping annotated heaps to the sets of allocated heap locations as

$$\begin{aligned} \text{locs}_{\text{ah}}(h^a) \quad := \quad & \{ \ell \in \mathcal{Locs} \mid \exists v \in \text{Values}. \exists L \in \mathcal{Levs}. \exists a \in \mathcal{A}. \\ & \exists H \in \text{Heaps}^{\text{log}}. \exists f \in \mathcal{F}. \\ & \ell \mapsto v \in h^a \vee \text{uninit}_{\text{aRes}}(\ell) \in h^a \vee \\ & \text{unlocked}_{\text{aRes}}((\ell, L), a, H) \in h^a \vee \\ & \text{locked}_{\text{aRes}}((\ell, L), a, f) \in h^a \} \end{aligned}$$

and the function  $\text{ids}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{ID})$  mapping annotated heaps to sets of allocated signal IDs as

$$\text{ids}_{\text{ah}}(h^a) \quad := \quad \{ id \in \mathcal{ID} \mid \exists L \in \mathcal{Levs}. \exists b \in \mathbb{B}. \text{signal}_{\text{aRes}}((id, L), b) \in h^a \}.$$

We denote annotated heaps by  $h^a$ .

We call an annotated heap  $h^a$  finite and write  $\text{finite}_{\text{ah}}(h^a)$  if there exists no chunk  $\text{unlocked}_{\text{aRes}}((\ell, L), a, H) \in h^a$  for which  $\text{finite}_{\text{lh}}(H)$  does not hold.

Physical, annotated and logical heaps represent program states on different abstraction levels. While each level focuses on different aspects of the program state, they also share information. For instance, all three kinds of heaps use points-to chunks. Also, annotated and logical heaps both have signal chunks. Therefore, when referring to multiple kinds of heaps to talk about the same program state, we have to ensure that their contents match and do not contradict each other. We call such heaps *compatible*. The following definitions make this precise.

**Definition A.4.4** (Compatibility of Annotated and Physical Heaps). *We inductively define a relation  $\text{ah} \sim_{\text{ph}} \subset \text{Heaps}^{\text{annot}} \times \mathcal{R}^{\text{phys}}$  between annotated and physical heaps such that the following holds:*

$$\begin{array}{ll} \emptyset & \text{ah} \sim_{\text{ph}} \emptyset, \\ \ell \mapsto v \cup h^a & \text{ah} \sim_{\text{ph}} \ell \mapsto v \cup h, \\ \text{uninit}_{\text{aRes}}(\ell) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{unlocked}_{\text{aRes}}((\ell, L), P, H_P) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{locked}_{\text{aRes}}((\ell, L), P, f) \cup h^a & \text{ah} \sim_{\text{ph}} \text{locked}_{\text{pRes}}(\ell) \cup h, \\ \text{signal}_{\text{aRes}}(s, b) \cup h^a & \text{ah} \sim_{\text{ph}} h, \end{array}$$

where  $h^a \in \text{Heaps}^{\text{annot}}$  and  $h \in \text{Heaps}^{\text{phys}}$  are annotated and physical heaps with  $h^a \sim_{\text{ph}} h$ .

**Definition A.4.5** (Compatibility of Annotated and Logical Heaps). *We inductively define a relation  $\text{ah} \sim_{\text{lh}} \subset \text{Heaps}^{\text{annot}} \times \text{Heaps}^{\text{log}}$  between annotated*

and logical heaps such that the following holds:

$$\begin{array}{ll}
\emptyset & \text{ah} \sim_{\text{lh}} \emptyset_{\text{log}}, \\
h^a \cup \{\ell \mapsto v\} & \text{ah} \sim_{\text{lh}} H + \{\ell \mapsto v\}, \\
h^a \cup \{\text{uninit}_{\text{aRes}}(\ell)\} & \text{ah} \sim_{\text{lh}} H + \{\text{uninit}_{\text{lRes}}(\ell)\}, \\
h^a \cup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\} & \text{ah} \sim_{\text{lh}} H + \{\text{mutex}_{\text{lRes}}(m, P)\} + H_P, \\
h^a \cup \{\text{locked}_{\text{aRes}}(m, P, f)\} & \text{ah} \sim_{\text{lh}} H + \{\text{locked}_{\text{lRes}}(m, P, f)\} \\
& \quad + (1 - f) \cdot \{\text{mutex}_{\text{lRes}}(m, P)\}, \\
h^a \cup \{\text{signal}_{\text{aRes}}(s, b)\} & \text{ah} \sim_{\text{lh}} H + \{\text{signal}_{\text{lRes}}(s, b)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{obs}_{\text{lRes}}(O)\},
\end{array}$$

where  $h^a \in \text{Heaps}^{\text{annot}}$  and  $H \in \text{Heaps}^{\text{log}}$  are annotated and logical heaps with  $\ell, m.\text{loc} \notin \text{locs}_{\text{ah}}(h^a)$ ,  $s.\text{id} \notin \text{ids}_{\text{ah}}(h^a)$  and  $h^a \text{ ah} \sim_{\text{lh}} H$ .

We define annotated versions  $\rightsquigarrow_{\text{atp}}$  and  $\rightsquigarrow_{\text{ast}}$  of the relations  $\rightsquigarrow_{\text{tp}}$  and  $\rightsquigarrow_{\text{st}}$ , respectively. The annotated reduction semantics we thereby obtain needs to reflect ghost proof steps implemented by view shifts. Hence, we define  $\rightsquigarrow_{\text{atp}}$  in terms of two relations: (i)  $\rightsquigarrow_{\text{ghost}}$  for ghost steps and (i)  $\rightsquigarrow_{\text{real}}$  for actual program execution steps. The annotated semantics ensure that a reduction gets stuck if a thread violates any of the restrictions formulated by our proof rules.

**Definition A.4.6** (Annotated Single Thread Reduction Relation). *We define a reduction relation  $\rightsquigarrow_{\text{ast}}$  for annotated threads according to the rules presented in Fig. A.10 and A.11. A reduction step has the form*

$$h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T^a$$

for a set of annotated forked threads  $T^a \subset \text{Heaps}^{\text{log}} \times \text{Cmds}$  with  $|T^a| \leq 1$ .

It indicates that given annotated heap  $h^a$  and a logical heap  $H$ , command  $c$  can be reduced to annotated heap  $h^{a'}$ , logical heap  $H'$  and command  $c'$ . The either empty or singleton set  $T^a$  represents whether a new thread is forked in this step.

For simplicity of notation we omit  $T^a$  if it is clear from the context that no thread is forked and  $T^a = \emptyset$ .

**Definition A.4.7** (Annotated Thread Pools). *We define the set of annotated thread pools  $\mathcal{TP}^a$  as the set of finite partial functions mapping thread IDs to annotated threads:*

$$\mathcal{TP}^a := \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds}^+ \cup \{\text{term}\}).$$

We denote annotated thread pools by  $P^a$  and the empty thread pool by  $\emptyset_{\text{atp}}$ , i.e.,

$$\begin{array}{l}
\emptyset_{\text{atp}} : \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds}^+ \cup \{\text{term}\}), \\
\text{dom}(\emptyset_{\text{atp}}) = \emptyset.
\end{array}$$

$$\frac{\text{AST-RED-EVALCTXT} \quad h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T}{h^a, H, E[c] \rightsquigarrow_{\text{ast}} h^{a'}, H', E[c'], T}$$

$$\begin{array}{l} \text{AST-RED-FORK} \\ h^a, H_m + \{\text{obs}_{\text{IRes}}(O_m \uplus O_f)\} + H_f, \mathbf{fork} \ c \rightsquigarrow_{\text{ast}} \\ h^a, H_m + \{\text{obs}_{\text{IRes}}(O_m)\}, \mathbf{tt}, \{(\{\text{obs}_{\text{IRes}}(O_f)\} + H_f), c\} \end{array}$$

(a) Basic constructs.

$$\begin{array}{ll} \text{AST-RED-IFTRUE} & \text{AST-RED-IFFALSE} \\ h^a, H, \mathbf{if} \ \text{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, c & h^a, H, \mathbf{if} \ \text{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt} \end{array}$$

$$\begin{array}{l} \text{AST-RED-LET} \\ h^a, H, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{ast}} h^a, H, c[x \mapsto v] \end{array}$$

(b) Control structures.

$$\begin{array}{ll} \text{AST-RED-CONS} & \text{AST-RED-READHEAPLOC} \\ \frac{\ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \mathbf{cons}(v)} & \frac{\ell \mapsto v \in h^a}{h^a, H, [\ell] \rightsquigarrow_{\text{ast}} h^a, H, v} \\ \rightsquigarrow_{\text{ast}} h^a \sqcup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, \ell & \end{array}$$

$$\begin{array}{l} \text{AST-RED-ASSIGN} \\ h \sqcup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, [\ell] := v \rightsquigarrow_{\text{ast}} h \sqcup \{\ell \mapsto v'\}, H + \{\ell \mapsto v'\}, \mathbf{tt} \end{array}$$

(c) Heap access.

$$\begin{array}{l} \text{AST-RED-NEWMUTEX} \\ \frac{\ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \mathbf{new\_mutex} \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{uninit}_{\text{aRes}}(\ell)\}, H + \{\text{uninit}_{\text{IRes}}(\ell)\}, \ell} \end{array}$$

$$\begin{array}{l} \text{AST-RED-ACQUIRE} \\ \frac{f \in \mathcal{F} \quad m.\text{lev} \prec_{\text{L}} O}{h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, a, H_P)\}, \\ H + \{\text{obs}_{\text{IRes}}(O)\} + f \cdot \{\text{mutex}_{\text{IRes}}(m, P)\}, \\ \mathbf{acquire} \ m.\text{loc} \\ \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, \\ H + \{\text{obs}_{\text{IRes}}(O \uplus \{\!| m |\!\}\}, \text{locked}_{\text{IRes}}(m, P, f)\} + H_P, \\ \mathbf{tt}} \end{array}$$

$$\begin{array}{l} \text{AST-RED-RELEASE} \\ \frac{H_P \vDash_A P \quad \text{consistent}_{\text{lh}}(H_P) \quad \exists O. H(\text{obs}_{\text{IRes}}(O)) \geq 1}{h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, H + \{\text{obs}_{\text{IRes}}(O \uplus \{\!| m |\!\}\}, \text{locked}_{\text{IRes}}(m, P, f)\} + H_P, \\ \mathbf{release} \ m.\text{loc} \\ \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\}, H + \{\text{obs}_{\text{IRes}}(O)\} + f \cdot \{\text{mutex}_{\text{IRes}}(m, P)\}, \\ \mathbf{tt}} \end{array}$$

(d) Mutexes.

Figure A.10: Annotated single thread reduction rules (part 1).

$$\begin{array}{c}
\text{AST-RED-WHILEDEC-INIT} \\
\frac{n \geq 0}{h^a, H, \text{ while } c \text{ do skip} \rightsquigarrow_{\text{ast}} h^a, H, \text{ while } n, c \text{ do skip}} \\
\\
\text{AST-RED-WHILEDEC} \\
\frac{n \geq 0}{h^a, H, \text{ while } n, c \text{ do skip} \rightsquigarrow_{\text{ast}} h^a, H, \text{ if } c \text{ then while } n-1, c \text{ do skip}} \\
\\
\text{AST-RED-AWAIT-INIT} \\
\frac{S \subset_{\text{fin}} \mathcal{S} \quad r \in \mathcal{X}}{h^a, H, \text{ with } \ell \text{ await } c \rightsquigarrow_{\text{ast}} h^a, H, \text{ if acquire } \ell; \text{ let } r := c \text{ in release } \ell; \neg r \\ \text{ then with } S, \ell \text{ await } c} \\
\\
\text{AST-RED-AWAIT} \\
\frac{\begin{array}{c} r \in \mathcal{X} \quad H(\text{obs}_{|\text{Res}}(O)) \geq 1 \\ (id, L) \in S \quad L \prec_L O \quad H(\text{signal}_{|\text{Res}}((id, L), \text{False})) \geq 1 \end{array}}{h^a, H, \text{ with } S, \ell \text{ await } c \rightsquigarrow_{\text{ast}} h^a, H, \text{ if acquire } \ell; \text{ let } r := c \text{ in release } \ell; \neg r \\ \text{ then with } S, \ell \text{ await } c}
\end{array}$$

(a) Loops and auxiliary commands.

Figure A.11: Annotated single thread reduction rules (part 2).

We define the extension operation  $+_{\text{atp}}$  analogously to  $+_{\text{tp}}$ , cf. Definition A.2.4.

For convenience of notation we define selector functions for annotated threads as

$$\begin{aligned}
(H, c).\text{heap} &:= H, \\
(H, c).\text{cmd} &:= c.
\end{aligned}$$

**Definition A.4.8** (Ghost Reduction Relation). *We define a thread pool reduction relation  $\rightsquigarrow_{\text{ghost}}$  according to the rules presented in Fig. A.12 to express ghost steps. A ghost reduction step has the form*

$$h^a, P^a \overset{\theta}{\rightsquigarrow}_{\text{ghost}} h^{a'}, P^{a'}.$$

We denote its reflexive transitive closure by  $\rightsquigarrow_{\text{ghost}}^*$ .

**Definition A.4.9** (Non-ghost Thread Pool Reduction Relation). *We define a thread pool reduction relation  $\rightsquigarrow_{\text{real}}$  according to the rules presented in Fig. A.13 to express real (i.e. non-ghost) reduction steps. A reduction step has the form*

$$h^a, P^a \overset{\theta}{\rightsquigarrow}_{\text{real}} h^{a'}, P^{a'}.$$

$$\begin{array}{c}
\text{GTP-RED-NEWSIGNAL} \\
\frac{P^a(\theta) = (H + \{\text{obs}_{\text{IRes}}(O)\}, c) \quad id \notin \text{ids}_{\text{ah}}(h^a) \quad H' = H + \{\text{signal}_{\text{IRes}}((id, L), \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{id, L\}\})\}}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}((id, L), \text{False})\}, P^a[\theta := (H', c)]} \\
\\
\text{GTP-RED-SETSIGNAL} \\
\frac{P^a(\theta) = (H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{s\}\})\}, c) \quad H' = H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O)\}}{h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{False})\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{True})\}, P^a[\theta := (H', c)]} \\
\\
\text{GTP-RED-MUTINIT} \\
\frac{P^a(\theta) = (H + \{\text{uninit}_{\text{IRes}}(\ell)\} + H_P, c) \quad H' = H + \{\text{mutex}_{\text{IRes}}((\ell, L), H_P)\} \quad H_P \Vdash_A P \quad \text{consistent}_{\text{lh}}(H_P) \quad \exists O. H(\text{obs}_{\text{IRes}}(O)) \geq 1}{h^a \sqcup \{\text{uninit}_{\text{aRes}}(\ell)\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}((\ell, L), a, H_P)\}, P^a[\theta := (H', c)]}
\end{array}$$

Figure A.12: Ghost thread pool reduction rules.

$$\begin{array}{c}
\text{RTP-RED-LIFT} \\
\frac{\theta_f = \min(\Theta \setminus \text{dom}(P^a)) \quad P^a(\theta) = (H, c) \quad h^a, H, c \xrightarrow{\text{ast}} h^{a'}, H', c', T^a}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^{a'}, P^a[\theta := (H', c')] +_{\text{atp}} T^a} \\
\\
\text{RTP-RED-TERM} \\
\frac{P^a(\theta) = (H, v) \quad H.\text{obs} = \emptyset}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^a, P^a -_{\text{atp}} \theta}
\end{array}$$

Figure A.13: Non-ghost thread pool reduction rules.

**Definition A.4.10** (Annotated Thread Pool Reduction Relation). *We define the annotated thread pool reduction relation  $\xrightarrow{\text{atp}}$  as*

$$\xrightarrow{\text{atp}} := \xrightarrow{\text{ghost}} \cup \xrightarrow{\text{real}} .$$

Note that the reduction relation  $\xrightarrow{\text{atp}}$  indeed reflects the restrictions our proof rules impose. For instance, proof rule PR-AWAIT only allows a thread to wait for a signal if the signal's level is below the level of all held obligations. AST-RED-AWAIT ensures that any thread that does not comply with this restriction gets stuck.

**Definition A.4.11** (Annotated Reduction Sequence). *Let  $(h^a_i)_{i \in \mathbb{N}}$  and  $(P^a_i)_{i \in \mathbb{N}}$  be infinite sequences of annotated heaps and annotated thread pools, respectively. Let  $\text{sig} : \mathbb{N} \rightarrow \mathcal{S}$  be a partial function mapping indices to signals.*

*We call  $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$  an annotated reduction sequence if there exists a sequence of thread IDs  $(\theta_i)_{i \in \mathbb{N}}$  such that the following holds for every  $i \in \mathbb{N}$ :*

- $h^a_i, P^a_i \xrightarrow{\theta_i}_{\text{atp}} h^a_{i+1}, P^a_{i+1}$
- *If this reduction step does not result from an application of RTP-RED-LIFT in combination with AST-RED-AWAIT, then  $\text{sig}(i) = \perp$ . If AST-RED-AWAIT is applied to some signal  $s$ , then  $\text{sig}(i) = s$ .*

*In case the signal annotation  $\text{sig}$  is clear from the context or not relevant, we omit it and write  $(h^a_i, P^a_i)_{i \in \mathbb{N}}$  instead of  $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$ .*

*We call  $(h^a_i, P^a_i)$  an annotated machine configuration.*

**Definition A.4.12** (Fairness of Annotated Reduction Sequences). *We call an annotated reduction sequence  $(h^a_i, P^a_i)_{i \in \mathbb{N}}$  fair iff for all  $i \in \mathbb{N}$  and  $\theta \in \text{dom}(P^a_i)$  with  $P^a_i(\theta).\text{cmd} \neq \text{term}$  there exists some  $k \geq i$  with*

$$h^a_k, P^a_k \xrightarrow{\theta}_{\text{real}} h^a_{k+1}, P^a_{k+1}.$$

**Lemma A.4.13** (Preservation of Finiteness). *Let  $(h^a_i, P^a_i)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{finite}_{\text{ah}}(h^a_0)$  and  $\text{finite}_{\text{lh}}(P^a_0(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P^a_0)$ .*

*Then,  $\text{finite}_{\text{lh}}(P^a_i(\theta).\text{heap})$  holds for all  $i \in \mathbb{N}$  and all  $\theta \in \text{dom}(P^a_i)$ .*

*Proof.* Proof by induction on  $i$ . □

**Lemma A.4.14** (Preservation of Completeness). *Let  $(h^a_i, P^a_i)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{complete}_{\text{lh}}(P^a_0(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P^a_0)$ . Then,  $\text{complete}_{\text{lh}}(P^a_i(\theta).\text{heap})$  holds for every  $i \in \mathbb{N}$  and every  $\theta \in \text{dom}(P^a_i)$ .*

*Proof.* Proof by induction on  $i$ . □

Every thread of an annotated thread pool is annotated by a thread-local logical heap that expresses which resources are owned by this thread. In the following we define a function to extract the logical heap expressing which resources are owned by threads of a thread pool (i.e. the sum of all thread-local logical heaps).

**Definition A.4.15.** We define the function  $\text{ownedResHeap}_{\text{atp}} : \mathcal{TP}^a \rightarrow \text{Heaps}^{\text{log}}$  mapping annotated thread pools to logical heaps as follows:

$$P^a \mapsto \sum_{\theta \in \text{dom}(P^a)} P^a(\theta).\text{heap}$$

Annotated resources representing unlocked locks, i.e.,  $\text{unlocked}_{\text{aRes}}(m, a, H_a)$ , contain a logical heap  $H_a$  that expresses which resources are protected by this lock. In the following, we define a function that extracts a logical heap from an annotated heap  $h^a$  expressing which resources are protected by unlocked locks in  $h^a$ .

**Definition A.4.16.** We define the function  $\text{protectedResHeap}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \text{Heaps}^{\text{log}}$  mapping annotated heaps to logical heaps as follows:

For any annotated heap  $h^a$  let

$$\text{LockInvs}(h^a) := \{ \{ H_P \in \text{Heaps}^{\text{log}} \mid \exists m \in \mathcal{Locs} \times \mathcal{Levs}. \exists P \in \mathcal{A}. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \} \}$$

be the auxiliary bag aggregating all logical heaps corresponding to lock invariants of unlocked locks stored in  $h^a$ . We define  $\text{protectedResHeap}_{\text{ah}}$  as

$$h^a \mapsto \sum_{H_P \in \text{LockInvs}(h^a)} H_P.$$

We consider a machine configuration  $(h^a, P^a)$  to be *consistent* if it fulfils three criteria: (i) Every thread-local logical heap is consistent, i.e., for all used thread IDs  $\theta$ ,  $P^a(\theta).\text{heap}$  only stores full obligations chunks. (ii) Every logical heap protected by an unlocked lock in  $h^a$  is consistent. (iii)  $h^a$  is compatible with the logical heap that represents (a) the resources owned by threads in  $P^a$  and (b) the resources protected by unlocked locks stored in  $h^a$ .

**Definition A.4.17** (Consistency of Annotated Machine Configurations). We call an annotated machine configuration  $(h^a, P^a)$  consistent and write  $\text{consistent}_{\text{conf}}(h^a, P^a)$  if all of the following hold:

- $\text{consistent}_{\text{lh}}(P^a(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P^a)$ ,
- $\forall m. \forall P. \forall H_P. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \rightarrow \text{consistent}_{\text{lh}}(H_P)$ ,
- $h^a \underset{\text{ah} \sim_{\text{lh}}}{\text{ownedResHeap}_{\text{atp}}(P^a) + \text{protectedResHeap}_{\text{ah}}(h^a)}$ .

**Lemma A.4.18** (Preservation of Consistency). Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{consistent}_{\text{conf}}(h_0^a, P_0^a)$ . Then,  $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$  holds for every  $i \in \mathbb{N}$ .

*Proof.* Proof by induction on  $i$ . □

## A.4.2 Hoare Triple Model Relation

We interpret program specifications  $\{A\} c \{B\}$  in terms of a model relation  $\models_{\text{H}} \{A\} c \{B\}$  and an auxiliary safety relation  $\text{safe}(H, c)$ .

In the annotated semantics we annotate threads by local logical heaps that express which resources they own (including ghost resources) and use an extended intermediate representation for commands. We say that an annotated thread pool  $P^a$  is an *annotation* of a plain thread pool  $P$  if they represent matching information. That is, they must store matching threads under the same thread IDs.

**Definition A.4.19** (Command Annotation). *We define the predicate  $\text{annot}_{\text{cmd}} \subseteq \text{Cmds}^+ \times \text{Cmds}$  such that  $\text{annot}_{\text{cmd}}(c^+, c)$  holds iff  $c$  results from  $c^+$  by replacing all subcommands of the form (i) **while**  $n, c'$  **do skip** by **while**  $c'$  **do skip** and (ii) **with**  $S, e$  **await**  $c'$  by **with**  $e$  **await**  $c'$ .*

**Definition A.4.20** (Thread Pool Annotation). *We define a predicate  $\text{annot}_{\text{tp}} \subseteq \mathcal{TP}^a \times \mathcal{TP}$  such that:*

$$\begin{aligned} & \text{annot}_{\text{tp}}(P^a, P) \\ & \iff \\ \text{dom}(P^a) = \text{dom}(P) \wedge \forall \theta \in \text{dom}(P). \text{annot}_{\text{cmd}}(P^a(\theta).\text{cmd}, P(\theta)) \end{aligned}$$

Intuitively, a command  $c$  is *safe* under a logical heap  $H$  if  $H$  provides the necessary resources so that for every execution of  $c$ , there is a corresponding annotated execution of  $c$  that does not get stuck. That is, if we consider a reduction step from  $c$  to  $c'$  in the plain operational semantics, then the resources that thread  $c$  owns according to  $H$ , allow us to perform a corresponding sequence of annotated reduction steps leading to  $c'$ . Specifically, we can perform a finite number of ghost steps to manipulate ghost resources followed by one real step to reduce  $c$  to  $c'$ . Furthermore, safety requires  $c'$  to be safe under its respective local logical heap. The same must hold for any potentially forked thread.

**Definition A.4.21** (Safety). *We define the safety predicate  $\text{safe} \subseteq \text{Heaps}^{\text{log}} \times \text{Cmds}$  coinductively as the greatest solution (with respect to  $\sqsubseteq$ ) of the following*



equation:

$$\begin{aligned}
& \text{safe}(H, c) \\
& \iff \\
& \text{complete}_{\text{lh}}(H) \rightarrow \\
& \forall P, P'. \forall \theta \in \text{dom}(P). \forall h, h'. \forall P^a. \forall h^a. \\
& \text{consistent}_{\text{conf}}(h^a, P^a) \wedge h^a \sim_{\text{ph}} h \wedge \\
& P(\theta) = c \wedge P^a(\theta) = (H, c) \wedge \text{annot}_{\text{tp}}(P^a, P) \wedge h, P \xrightarrow{\theta}_{\text{tp}} h', P' \rightarrow \\
& \exists P^G, P^{a'}. \exists h^G, h^{a'}. \\
& h^a, P^a \xrightarrow{\theta}_{\text{ghost}}^* h^G, P^G \wedge h^G, P^G \xrightarrow{\theta}_{\text{real}} h^{a'}, P^{a'} \wedge \text{annot}_{\text{tp}}(P^{a'}, P') \wedge \\
& h^{a'} \sim_{\text{ph}} h' \wedge \\
& \forall (H_f, c_f) \in \text{range}(P^{a'}) \setminus \text{range}(P^a). \text{safe}(H_f, c_f).
\end{aligned}$$

We interpret Hoare triples  $\{A\} c \{B\}$  in terms of safety as follows: Let  $E$  be an evaluation context that (when instantiated) is safe under any heap  $H_B$  which fulfills postcondition  $B$ , i.e., let  $E$  be a context for which  $\forall v. \forall H_B. H_B \vDash_A B(v) \rightarrow \text{safe}(H_B, E[v])$  holds. Then, any heap  $H_A$  that satisfies precondition  $A$  provides all resources necessary to safely run  $E[c]$ , i.e.,  $\forall H_A. H_A \vDash_A A \rightarrow \text{safe}(H_A, E[c])$ . Thereby, if the reduction of  $c$  under  $H_A$  finishes and returns a value  $v$ , then postcondition  $B(v)$  holds in the final state. In the following definition, we also allow  $H_A$  and  $H_B$  to be embedded into arbitrary heap frames  $H_F$ .

**Definition A.4.22** (Hoare Triple Model Relation). *We define the model relation for Hoare triples  $\vDash_{\text{H}} \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$  such that:*

$$\begin{aligned}
& \vDash_{\text{H}} \{A\} c \{\lambda r. B(r)\} \\
& \iff \\
& \forall H_F. \forall E. (\forall v. \forall H_B. H_B \vDash_A B(v) \rightarrow \text{safe}(H_B + H_F, E[v])) \\
& \rightarrow \forall H_A. H_A \vDash_A A \rightarrow \text{safe}(H_A + H_F, E[c])
\end{aligned}$$

We can instantiate context  $E$  in above definition to **let**  $x := \square$  **in** **tt**, which yields the consequent  $\text{safe}(H_A + H_F, \text{let } x := c \text{ in tt})$ . Note that this implies  $\text{safe}(H_A + H_F, c)$ . Also note that compliance with the frame rule directly follows from above definition, i.e.,  $\vDash_{\text{H}} \{A\} c \{B\}$  implies  $\vDash_{\text{H}} \{A * F\} c \{B * F\}$  for any frame  $F$ . Further, every specification we can derive in our proof system also holds in our model.

**Lemma A.4.23** (Hoare Triple Soundness). *Let  $\vdash \{A\} c \{B\}$  hold, then also  $\vDash_{\text{H}} \{A\} c \{B\}$  holds.*

*Proof.* Proof by induction on the derivation of  $\vdash \{A\} c \{B\}$ . □

### A.4.3 Soundness Proof

In this section we prove our soundness theorem.

**Constructing Annotated Executions** Given a command  $c$  which provably discharges all obligations and a fair reduction sequence for  $c$ , we can construct a corresponding annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ . This is a useful tool to analyse program executions as  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  carries much more information than the original sequence, e.g., which obligations a thread holds. Note that our definition of fairness forbids  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  to perform ghost steps forever and that we use  $h^a_{\text{ah}} \sim_{\text{lh}} H$  to express that  $h^a$  is compatible with  $H$ .

**Lemma A.4.24** (Construction of Annotated Reduction Sequences). *Suppose we can prove  $\vDash_{\text{H}} \{A\} c \{\text{obs}(\emptyset)\}$ . Let  $H_A$  be a logical heap with  $H_A \vDash_A A$  and  $\text{complete}_{\text{lh}}(H_A)$  and  $h_0^a$  an annotated heap with  $h_0^a_{\text{ah}} \sim_{\text{lh}} H_A$ . Let  $(h_i, P_i)_{i \in \mathbb{N}}$  be a fair plain reduction sequence with  $h_0^a_{\text{ah}} \sim_{\text{ph}} h_0$  and  $P_0 = \{(\theta_0, c)\}$  for some thread ID  $\theta_0$  and command  $c$ .*

*Then, there exists a fair annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  with  $P^a = \{(\theta_0, (H_A, c))\}$  and  $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$  for all  $i \in \mathbb{N}$ .*

*Proof.* We can construct the annotated reduction sequence inductively from the plain reduction sequence.  $\square$

**Program Order Graph** In the remainder of this section, we prove that programs where every thread discharges all obligations terminate. For this, we need to analyse each thread's control flow, i.e., the subsequence of execution steps belonging to the thread. We do this by taking a sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  representing a program execution, constructing an annotation  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  carrying additional information and then analysing its *program order graph*  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  which represents the execution's control flow.

**Definition A.4.25** (Program Order Graph). *Let  $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$  be an annotated reduction sequence. Let  $N^r$  be the set of names referring to reduction rules defining the relations  $\rightsquigarrow_{\text{real}}$ ,  $\rightsquigarrow_{\text{ghost}}$  and  $\rightsquigarrow_{\text{ast}}$ . We define the set of annotated reduction rule names  $N^a$  where *AST-RED-AWAIT* is annotated by signals as*

$$N^a := (N^r \setminus \{\text{AST-RED-AWAIT}\}) \cup (\{\text{AST-RED-AWAIT}\} \times \mathcal{S})$$

*We define the program order graph  $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})) = (\mathbb{N}, E)$  with root 0 where  $E \subset \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$ .*

A node  $a \in \mathbb{N}$  corresponds to the sequence's  $a^{\text{th}}$  reduction step, i.e., to the step  $h_a^a, P_a^a \xrightarrow{\theta}_{\text{atp}} h_{a+1}^a, P_{a+1}^a$  for some  $\theta \in \text{dom}(P_a^a)$ . An edge from node  $a$  to node  $b$  expresses that the  $b^{\text{th}}$  reduction step continues the control flow of step  $a$ . For any  $\ell \in \mathbb{N}$ , let  $\theta_\ell$  denote the ID of the thread reduced in step  $\ell$ . Furthermore, let  $n_\ell^a$  denote the name of the reduction rule applied in the  $\ell^{\text{th}}$  step, in the following sense:

- If  $h_\ell^a, P_\ell^a \xrightarrow{\theta}_{\text{atp}} h_{\ell+1}^a, P_{\ell+1}^a$  results from an application of RTP-RED-LIFT in combination with single-thread reduction rule  $n^{\text{st}}$  other than AST-RED-AWAIT, then  $n_\ell^a = n^{\text{st}}$ .
- If  $h_\ell^a, P_\ell^a \xrightarrow{\theta}_{\text{atp}} h_{\ell+1}^a, P_{\ell+1}^a$  results from an application of RTP-RED-LIFT in combination with AST-RED-AWAIT, then  $n_\ell^a = (\text{AST-RED-AWAIT}, \text{sig}(\ell))$ .
- Otherwise,  $n^a$  denotes the applied (real or ghost) thread pool reduction rule.

An edge  $(a, \theta, n^a, b) \in \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$  is contained in  $E$  if  $n^a = n_a^a$  and one of the following conditions applies:

- $\theta = \theta_a = \theta_b$  and  $b = \min(\{k > a \mid h_k^a, P_k^a \xrightarrow{\theta_a}_{\text{atp}} h_{k+1}^a, P_{k+1}^a\})$ .  
In this case, the edge expresses that step  $b$  marks the first time that thread  $\theta_a$  is rescheduled for reduction (after step  $a$ ).
- $\text{dom}(P_{a+1}^a) \setminus \text{dom}(P_a^a) = \{\theta\}$  and  
 $b = \min\{k \in \mathbb{N} \mid h_k^a, P_k^a \xrightarrow{\theta}_{\text{atp}} h_{k+1}^a, P_{k+1}^a\}$ .  
In this case,  $\theta$  identifies the thread forked in step  $a$ . The edge expresses that step  $b$  marks the first reduction of the forked thread.

In case the choice of reduction sequence  $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$  is clear from the context, we write  $G$  instead of  $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig}))$ .

**Observation A.4.26.** Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $|\text{dom}(P_0^a)| = 1$ . The sequence's program order graph  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  is a binary tree.

For any reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ , the paths in its program order graph  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  represent the sequence's control flow paths. Hence, we are going to use program order graphs to analyse reduction sequences' control flows. We refer to a program order graph's edges by the kind of reduction step they represent. For instance, we call any edge  $(a, \theta, R, b)$  a *loop edge* where  $R$  refers to one of the rules related to the execution loops, i.e., AST-RED-WHILEDEC-INIT,

AST-RED-WHILEDEC, AST-RED-AWAIT-INIT or AST-RED-AWAIT. In the following, we prove that any path in a program order graph that does not involve a loop edge is finite. This follows from the fact that the size of the command reduced along this path decreases with each non-ghost non-loop step.

**Lemma A.4.27.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence. Let  $p = (V, E)$  be a path in  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ . Let*

$$R = \{ \text{AST-RED-WHILEDEC-INIT}, \text{AST-RED-WHILEDEC}, \\ \text{AST-RED-AWAIT-INIT} \} \cup (\{ \text{AST-RED-AWAIT} \} \times \mathcal{S})$$

*be the set of names of single-thread reduction rules related to loops and let  $L = \{e \in E \mid \pi_3(e) \in R\}$  be the set of loop edges contained in  $p$ . Then,  $p$  is infinite if and only if  $L$  is infinite.*

*Proof.* If  $L$  is infinite,  $p$  is obviously infinite as well. So, suppose  $L$  is finite.

For any command, we consider its size to be the number of nodes contained in its abstract syntax tree. By structural induction over the set of commands, it follows that the size of a command  $c = P^a(\theta).\text{cmd}$  decreases in every non-ghost reduction step  $h^a, P^a \xrightarrow{\theta}_{\text{atp}} h^{a'}, P^{a'}$  that is not an application of RTP-RED-LIFT in combination with some  $r \in R$ .

Since  $L$  is finite, there exists a node  $x$  such that the suffix  $p_{\geq x}$  starting at node  $x$  does not contain any loop edges. By fairness of  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ , every non-empty suffix of  $p_{\geq x}$  contains an edge corresponding to a non-ghost reduction step. For any edge  $e = (i, \theta, n, j)$  consider the command  $c_e = P_i^a(\theta).\text{cmd}$  reduced in this edge. The size of these commands decreases along  $p_{\geq x}$ . So,  $p_{\geq x}$  must be finite and thus  $p$  must be finite as well.  $\square$

**Definition A.4.28.** *Let  $G = (V, E)$  be a subgraph of some program order graph. We define the function  $\text{awaitEdges}_G : \mathcal{S} \rightarrow \mathcal{P}(E)$  mapping any signal  $s$  to the set of await edges in  $G$  concerning  $s$  as:*

$$\text{awaitEdges}_G(s) := \{(a, \theta, (\text{AST-RED-AWAIT}, s'), b) \in E \mid s' = s\}.$$

*Furthermore, we define the set  $\mathcal{S}_G \subseteq \mathcal{S}$  of signals being waited for in  $G$  and its subset  $\mathcal{S}_G^\infty \subseteq \mathcal{S}_G$  of signals waited-for infinitely often in  $G$  as follows:*

$$\begin{aligned} \mathcal{S}_G &:= \{s \in \mathcal{S} \mid \text{awaitEdges}_G(s) \neq \emptyset\}, \\ \mathcal{S}_G^\infty &:= \{s^\infty \in \mathcal{S}_G \mid \text{awaitEdges}_G(s^\infty) \text{ infinite}\}. \end{aligned}$$

**Definition A.4.29** (Partial Order on Finite Bags). *Let  $X$  be a set and  $<_X \subseteq X \times X$  a partial order on  $X$ . We define the partial order  $\prec_X \subseteq \text{Bags}_{\text{fin}}(X) \times$*

$Bags_{\text{fin}}(X)$  on finite bags over  $X$  as the Dershowitz-Manna ordering [53] induced by  $<_X$ :

$$\begin{aligned} A \prec_X B &\iff \exists C, D \in Bags_{\text{fin}}(X). \quad \emptyset \neq C \subseteq B \\ &\quad \wedge A = (B \setminus C) \uplus D \\ &\quad \wedge \forall d \in D. \exists c \in C. d <_X c. \end{aligned}$$

We define  $\preceq_X \subset Bags_{\text{fin}}(X) \times Bags_{\text{fin}}(X)$  such that

$$A \preceq_X B \iff A = B \vee A \prec_X B$$

holds.

**Corollary A.4.30.** *The partial order  $\prec_{\mathbb{N}} \subset Bags_{\text{fin}}(\mathbb{N}) \times Bags_{\text{fin}}(\mathbb{N})$  is well-founded.*

*Proof.* Follows from [53] and well-foundedness of  $<_{\mathbb{N}}$ .  $\square$

Below we define a metric on commands in a graph that allows us to prove that every control flow path in which no signal is waited-for infinitely often is finite. We construct the metric in two steps. (i) We consider the maximal degree for which the command contains a nested loop. Here, we consider an uninitialized loop **while**  $c$  **do skip** as a doubly nested one since its reduction in the annotated semantics involves its conversion into either a decrease loop or into an await loop, cf. AST-RED-WHILEDEC-INIT and AST-RED-AWAIT-INIT. (ii) We use the extracted degree to construct a rank (in form of a finite bag of degrees) that intuitively spoken measures the number of loop iterations the command's execution causes in the graph we consider.

**Definition A.4.31** (Degree Extraction). *We define the function  $\langle \cdot \rangle : Ccmds \rightarrow \mathbb{N}$  mapping augmented commands to degrees by recursion on the structure of commands as follows:*

$$\begin{aligned} \langle \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rangle &:= \langle c \rangle + 2, \\ \langle \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip} \rangle &:= \langle c \rangle + 1, \\ \langle \mathbf{with} \ S, m \ \mathbf{await} \ c \rangle &:= \langle c \rangle + 1, \\ \langle \mathbf{fork} \ c \rangle &:= \langle c \rangle, \\ \langle \mathbf{let} \ x := c \ \mathbf{in} \ c' \rangle &:= \max(\langle c \rangle, \langle c' \rangle), \\ \langle \mathbf{if} \ c \ \mathbf{then} \ c' \rangle &:= \max(\langle c \rangle, \langle c' \rangle), \\ \langle c \rangle &:= 0 \quad \textit{otherwise.} \end{aligned}$$

**Definition A.4.32** (Rank Extraction). *For any subgraph  $G$  of a program order graph with  $S_G^\infty = \emptyset$ , we define the function  $\langle \langle \cdot \rangle \rangle_G : Ccmds \rightarrow Bags_{\text{fin}}(\mathbb{N})$  mapping*

commands to finite bags of degrees by recursion on the structure of commands as follows:

$$\begin{aligned}
\langle\langle \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rangle\rangle_G &:= \{\langle \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rangle\}, \\
\langle\langle \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip} \rangle\rangle_G &:= n \cdot \{\langle \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip} \rangle\}, \\
\langle\langle \mathbf{with} \ S, m \ \mathbf{await} \ c \rangle\rangle_G &:= \left| \bigcup_{s \in S} \mathbf{awaitEdges}_G(s) \right| \cdot \{\langle \mathbf{with} \ S, m \ \mathbf{await} \ c \rangle\}, \\
\langle\langle \mathbf{fork} \ c \rangle\rangle_G &:= \langle\langle c \rangle\rangle_G, \\
\langle\langle \mathbf{let} \ x := c \ \mathbf{in} \ c' \rangle\rangle_G &:= \langle\langle c \rangle\rangle_G \uplus \langle\langle c' \rangle\rangle_G, \\
\langle\langle \mathbf{if} \ c \ \mathbf{then} \ c' \rangle\rangle_G &:= \langle\langle c \rangle\rangle_G \uplus \langle\langle c' \rangle\rangle_G, \\
\langle\langle c \rangle\rangle_G &:= \emptyset \qquad \text{otherwise.}
\end{aligned}$$

We view paths in a program order graph as single-branched subgraphs. This allows us to apply above definition on graphs to paths. In the proof of the following lemma, we see that in any control flow path where every signal is waited-for only finitely often, the rank of the reduced command decreases step by step.

**Lemma A.4.33.** *Let  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  be a program order graph and let  $p = (V, E)$  be a path in  $G$  with  $S_p^\infty = \emptyset$ . For every  $\theta \in \text{dom}(P_0^a)$  let  $P_0^a(\theta).\text{heap}$  be finite and complete. Then,  $p$  is finite.*

Assume  $p$  is infinite. We prove a contradiction by assigning a decreasing metric to every node along the path. For every  $i \in V$ , let  $p_{\geq i}$  be the suffix of  $p$  starting at node  $i$ . Every node  $i \in V$  corresponds to a reduction step  $h_i^a, P_i^a \xrightarrow{\theta_i}_{\text{atp}} h_j^a, P_j^a$ . In the following, we let  $c_i := P_i^a(\theta_i)$  denote the command reduced in this step. Consider the function  $f : V \rightarrow \text{Bags}_{\text{fin}}(\mathbb{N})$ ,  $i \mapsto \langle\langle c_i \rangle\rangle_{p_{\geq i}}$  mapping every node on the path to the rank of the command whose reduction the node represents.

Consider the sequence  $(f(i))_{i \in V}$ . Since every element is a finite bag of natural numbers, we can order it by  $\prec_{\mathbb{N}}$ . We are going to prove a contradiction by proving that the sequence is an infinitely descending chain.

Consider any edge  $e = (i, \theta, r, j) \in E$ . There are only 4 cases in which  $f(i) \neq f(j)$  holds.

- $r = \text{AST-RED-WHILEDEC-INIT}$ :  
In this case we have  $c_i = \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip}$  and  $c_j = \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip}$  for some  $c, n$ . We get

$$f(j) = n \cdot \{\langle c \rangle + 1\} \prec_{\mathbb{N}} \{\langle c \rangle + 2\} = f(i).$$

- $r = \text{AST-RED-WHILEDEC}$ :  
In this case we have  $c_i = \mathbf{while} \ n, c \ \mathbf{do} \ \mathbf{skip}$  and  $c_j = \mathbf{if} \ c \ \mathbf{then} \ \mathbf{while} \ n -$

1,  $c$  **do skip** for some  $n, c$ . We get

$$f(j) = \{\langle c \rangle\} \uplus (n-1) \cdot \{\langle c \rangle + 1\} \prec_{\mathbb{N}} n \cdot \{\langle c \rangle + 1\} = f(i).$$

- $r = \text{AST-RED-AWAIT-INIT}$ :

In this case we have  $c_i = \mathbf{with } m \mathbf{ await } c$  and

$c_j = \mathbf{if acquire } m; \mathbf{let } x := c \mathbf{ in release } m; \neg x \mathbf{ then with } S, m \mathbf{ await } c$   
for some  $m, x, c, S$ . We get

$$f(j) = \{\langle c \rangle\} \uplus |\dots| \cdot \{\langle c \rangle + 1\} \prec_{\mathbb{N}} \{\langle c \rangle + 2\} = f(i).$$

- $r = (\text{AST-RED-AWAIT}, s')$  for some  $s'$ :

In this case we have  $c_i = \mathbf{with } S, m \mathbf{ await } c$  and

$c_j = \mathbf{if acquire } m; \mathbf{let } x := c \mathbf{ in release } m; x \mathbf{ then with } S, m \mathbf{ await } c$   
for some  $m, x, c, S$ . We get

$$\begin{aligned} f(j) &= \{\langle c \rangle\} \uplus \left| \bigcup_{s \in S} \text{awaitEdges}_{p \geq j}(s) \right| \cdot \{\langle c \rangle + 1\} \\ &\prec_{\mathbb{N}} \left( 1 + \left| \bigcup_{s \in S} \text{awaitEdges}_{p \geq j}(s) \right| \right) \cdot \{\langle c \rangle + 1\} \\ &= \left| \bigcup_{s \in S} \text{awaitEdges}_{p \geq i}(s) \right| \cdot \{\langle c \rangle + 1\} = f(i) \end{aligned}$$

By application of Lemma A.4.27 we get that  $p$  contains infinitely many of the loop edges listed above. Hence,  $(f(i))_{i \in V}$  is an infinitely decreasing chain. By Corollary A.4.30,  $\prec_{\mathbb{N}}$  is well-founded. A contradiction. So,  $p$  is finite.

We proceed to prove that no signals are waited for infinitely often.

**Lemma A.4.34.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence with  $P_0^a = \{(\theta_0, (H_0, c))\}$ ,  $\text{finite}_{\text{ah}}(h_0^a)$ ,  $\text{complete}_{\text{lh}}(H_0)$ ,  $\text{finite}_{\text{lh}}(H_0)$  and also  $\text{consistent}_{\text{conf}}(h_0^a, P_0^a)$ . Let  $H_0$  contain no signal chunks. Further, let  $h_0^a$  contain no chunks of the form  $\text{unlocked}_{\text{aRes}}(m, P, H_P)$  where  $H_P$  contains any signal chunks. Let  $G$  be the program order graph of  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ . Then,  $S_G^\infty = \emptyset$ .*

*Proof.* Suppose  $S_G^\infty \neq \emptyset$ . Since  $\mathcal{L}evs$  is well-founded, the same holds for the set  $\{s.\text{lev} \mid s \in S^\infty\}$ . Hence, there is some  $s_{\min} \in S^\infty$  for which no  $z \in S^\infty$  with  $z.\text{lev} <_{\mathbb{L}} s_{\min}.\text{lev}$  exists.

Since neither the initial logical heap  $H_0$  nor any unlocked lock invariant stored in  $h_0^a$  does contain any signals,  $s_{\min}$  must be created during the reduction sequence. The reduction step creating signal  $s_{\min}$  is an application of GTP-RED-NEWSIGNAL, which simultaneously creates an obligation to set  $s_{\min}$ . By preservation of completeness, Lemma A.4.14, every thread-local logical heap  $P_i^a(\theta).\text{heap}$  annotating some thread  $\theta$  in some step  $i$

is complete. According to reduction rule AST-RED-AWAIT, every await edge  $(a, \theta, (\text{AST-RED-AWAIT}, s_{\min}), b)$  implies together with completeness that in step  $a$  (i) thread  $\theta$  does not hold any obligation for  $s_{\min}$  (i.e.  $P_a^a(\theta).\text{heap}(\text{obs}_{\text{Res}}(O)) = 1$  for some bag of obligations  $O$  with  $s_{\min} \notin O$ ) and (ii)  $s_{\min}$  has not been set, yet (i.e.  $\text{signal}_{\text{aRes}}(s_{\min}, \text{False}) \in h_a^a$ ). Hence, in step  $a$  another thread  $\theta_{\text{ob}} \neq \theta$  must hold the obligation for  $s_{\min}$  (i.e.  $P_a^a(\theta_{\text{ob}}).\text{heap}(\text{obs}_{\text{Res}}(O)) = 1$  for some bag of obligations  $O$  with  $s_{\min} \in O$ ). Since there are infinitely many await edges concerning  $s_{\min}$  in  $G$ , the signal is never set.

By fairness, for every await edge as above, there must be a non-ghost reduction step  $h_k^a, P_k^a \xrightarrow{\theta_{\text{ob}}}_{\text{atp}} h_{k+1}^a, P_{k+1}^a$  of the thread  $\theta_{\text{ob}}$  holding the obligation for  $s_{\min}$  with  $k \geq a$ . Hence, there exists an infinite path  $p_{\text{ob}}$  in  $G$  where each edge  $(e, \theta_{\text{ob}}, n, f) \in \text{edges}(p_{\text{ob}})$  concerns some thread  $\theta_{\text{ob}}$  holding the obligation for  $s_{\min}$ . (Note that this thread ID does not have to be constant along the path, since the obligation can be passed on during fork steps.)

The path  $p_{\text{ob}}$  does not contain await edges  $(e, \theta_{\text{ob}}, (\text{AST-RED-AWAIT}, s^\infty), f)$  for any  $s^\infty \in S^\infty$ , since reduction rule AST-RED-AWAIT would (together with completeness of  $P_e^a(\theta_{\text{ob}}).\text{heap}$ ) require  $s^\infty$  to be of a lower level than all held obligations. This restriction implies  $s^\infty.\text{lev} <_{\text{L}} s_{\min}.\text{lev}$  and would hence contradict the minimality of  $s_{\min}$ . That is,  $S_{p_{\text{ob}}}^\infty = \emptyset$ .

By preservation of finiteness, Lemma A.4.13, we get that every logical heap associated with the root of  $p_{\text{ob}}$  is finite. This allows us to apply Lemma A.4.33, by which we get that  $p_{\text{ob}}$  is finite. A contradiction.  $\square$

Finally, we got everything we need to prove that any program that discharges all its obligations terminates.

**Lemma A.4.35.** *Let  $\models_{\text{H}} \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$  hold. There exists no fair, infinite annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  with  $P_0^a = \{(\theta_0, (H_0, c))\}$ ,  $h_0^a = \emptyset$  and  $H_0 = \{\text{obs}_{\text{Res}}(\emptyset)\}$ .*

*Proof.* Suppose a reduction sequence as described above exists. We are going to prove a contradiction by considering its infinite program order graph  $G$ .

Since  $P_0^a$  contains only a single thread,  $G$  is a binary tree with an infinite set of vertices. By the Weak König's Lemma [163]  $G$  has an infinite branch, i.e. an infinite path  $p$  starting at root 0.

The initial logical heap  $H_0$  is complete and finite and the initial annotated machine configuration  $(h_0^a, P_0^a)$  is consistent. By Lemma A.4.34 we know



that  $S_G^\infty = \emptyset$ . Since  $S_p^\infty \subseteq S_G^\infty$ , we get  $S_p^\infty = \emptyset$ . This allows us to apply Lemma A.4.33, by which we get that  $p$  is finite, which is a contradiction.  $\square$

**Theorem B.10.7** (Soundness). *Let  $\vdash \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$  hold. There exists no fair, infinite reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  with  $h_0 = \emptyset$  and  $P_0 = \{(\theta_0, c)\}$  for any choice of  $\theta_0$ .*

*Proof.* Assume that such a reduction sequence exists. By Hoare triple soundness, Lemma A.4.23, we get  $\models_{\text{H}} \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$  from  $\vdash \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$ . Consider the logical heap  $H_0 = \{\text{obs}_{|\text{Res}}(\emptyset)\}$  and the annotated heap  $h_0^a = \emptyset$ . It holds  $H_0 \models_{\text{A}} \text{obs}(\emptyset)$ ,  $h_0^a \text{ah} \sim_{\text{lh}} H_0$  (since  $H_0$  does not contain any logical resources with an annotated counterpart) and  $h_0^a \text{ah} \sim_{\text{ph}} h_0$  (since both heaps are empty). This allows us to apply Lemma A.4.24, by which we can construct a corresponding fair annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  that starts with  $h_0^a = \emptyset$  and  $P_0^a = \{(\theta_0, (H_0, c))\}$ . By Lemma A.4.35  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  does not exist. A contradiction.  $\square$

## A.5 Case Studies

In this section, we verify two example programs in detail. In § A.5.1 we verify the realistic example program presented in § 3.3. It involves two threads communicating via a shared bounded FIFO. In § A.5.2 we present and verify a program similar to the one from § 3.3 but with an unbounded number of producer and consumer threads. In order to lower the visual burden on the reader, we use the following colour coding:

As in the paper body, we present the proof state in **blue**, applied proof and view shift rules in **purple**, and abbreviations and definition in **red**. Since the verification outlines we present in this section span multiple figures, we include hints and explanations concerning other figures (e.g. a hint pointing to the figure where an invariant was defined). To set these remarks and in general any hints off from the essential parts of the proof, we present them in a *grey and italic font*. We occasionally remind the reader of earlier proof steps performed in a previous figure by repeating them in *the same grey font* at the beginning of the current figure.

Further, we highlight how our proof steps effect the proof state as follows: Consider a proof state of the form  $\{A * \ell \mapsto v\}$ . (i) When a proof step adds a new chunk  $C$ , we highlight it in **green**, i.e.  $\{A * \ell \mapsto v * C\}$ . (ii) When a proof step removes the chunk  $A$ , we highlight this change by underlying the removed part of the assertion with a **dark grey** background,

<pre> {obs(0)} let fifo10 := cons(nil) in let mut := new_mutex in ∀ℓ<sub>fifo10</sub>, ℓ<sub>mut</sub>. {obs(0) * ℓ<sub>fifo10</sub> ↦ nil * uninit(ℓ<sub>mut</sub>) } ∀id<sup>1</sup><sub>pop</sub>, ..., id<sup>100</sup><sub>pop</sub>, id<sup>1</sup><sub>push</sub>, ..., id<sup>100</sup><sub>push</sub>. {   *<sub>i=1, ..., 100</sub> uninitSig(id<sup>i</sup><sub>pop</sub>) * *<sub>i=1, ..., 100</sub> uninitSig(id<sup>i</sup><sub>push</sub>) * ... } L<sup>i</sup><sub>pop</sub> := 102 - i, L<sup>i</sup><sub>push</sub> := 101 - i for 1 ≤ i ≤ 100 (Later L<sup>i+10</sup><sub>pop</sub> &lt; L<sup>i</sup><sub>push</sub> and L<sup>i</sup><sub>push</sub> &lt; L<sup>i</sup><sub>pop</sub> must hold, cf. Figures A.19 and A.21.) s<sup>i</sup><sub>push</sub> := (id<sup>i</sup><sub>push</sub>, L<sup>i</sup><sub>push</sub>), s<sup>i</sup><sub>pop</sub> := (id<sup>i</sup><sub>pop</sub>, L<sup>i</sup><sub>pop</sub>) for 1 ≤ i ≤ 100 {   uninitSig(id<sup>100</sup><sub>pop</sub>) * uninitSig(id<sup>100</sup><sub>push</sub>) signal(s<sup>100</sup><sub>pop</sub>, False) * signal(s<sup>100</sup><sub>push</sub>, False)   * obs({ s<sup>100</sup><sub>pop</sub>, s<sup>100</sup><sub>push</sub> }) * ... } let c<sub>p</sub> := cons(100) in let c<sub>c</sub> := cons(100) in ∀ℓ<sub>c<sub>p</sub></sub>, ℓ<sub>c<sub>c</sub></sub>. { ℓ<sub>c<sub>p</sub></sub> ↦ 100 * ℓ<sub>c<sub>c</sub></sub> ↦ 100 * ... } {   obs({ s<sup>100</sup><sub>push</sub>, s<sup>100</sup><sub>pop</sub> }) * [½]ℓ<sub>c<sub>p</sub></sub> ↦ 100 * [½]ℓ<sub>c<sub>c</sub></sub> ↦ 100 * uninit(ℓ<sub>mut</sub>)   * P<sub>mut</sub>   * *<sub>i=1, ..., 99</sub> (uninitSig(id<sup>i</sup><sub>push</sub>) * uninitSig(id<sup>i</sup><sub>pop</sub>)) } m := (ℓ<sub>mut</sub>, 0) (Later m.lev &lt; L<sup>i</sup><sub>push</sub> and m.lev &lt; L<sup>i</sup><sub>pop</sub> must hold for all 1 ≤ i ≤ 100, cf. Figures A.19 and A.21.) { uninit(ℓ<sub>mut</sub>) * P<sub>mut</sub> mutex(m, P<sub>mut</sub>) * ... } {   obs({ s<sup>100</sup><sub>push</sub>, s<sup>100</sup><sub>pop</sub> }) * [½]ℓ<sub>c<sub>p</sub></sub> ↦ 100 * [½]ℓ<sub>c<sub>c</sub></sub> ↦ 100   * *<sub>i=1, ..., 99</sub> (uninitSig(id<sup>i</sup><sub>push</sub>) * uninitSig(id<sup>i</sup><sub>pop</sub>))   * [½]mutex(m, P<sub>mut</sub>) * [½]mutex(m, P<sub>mut</sub>) } ... </pre>	<p>PR-LET (2x) &amp; PR-CONS &amp; PR-NEWMUTEX PR-VIEWSHIFT &amp; VS-ALLOC SIGID &amp; PR-EXISTS (200x)</p> <p>PR-VIEWSHIFT &amp; VS-SIGINT</p> <p>PR-LET &amp; PR-CONS (2x)</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>For definition of lock invariant P<sub>mut</sub> cf. Fig. A.15. PR-VIEWSHIFT &amp; VS-MUTINT</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>Continued in Fig. A.18.</p>
---	---

Figure A.14: Verification of realistic example 3.10a: Initialization.

i.e.,  $\{A * \ell \mapsto v\}$ . Note that in this case, the greyed out  $A *$  is not a part of the proof state anymore. (iii) When a proof step changes only part of a chunk, we highlight this change in yellow. For instance, if the step changes the value of heap location  $\ell$  from  $v$  to  $v'$ , we highlight it in the resulting state as  $\{A * \ell \mapsto v'\}$ .

## A.5.1 Verification of Realistic Example

In this section, we verify the realistic example program from § 3.3 presented in Fig. 3.10a. We present the full proof outline in Fig. A.14 – A.22. We scaled the proof outlines down to fit them into the page layout. You can find the original proof outlines in the extended version of the ghost signal paper [152] and in the technical report [151].

$$\begin{aligned}
P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) &:= \exists v_{\text{cp}}^{\text{mut}}, v_{\text{cc}}^{\text{mut}}. \\
&[\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}}^{\text{mut}} * [\frac{1}{2}] \ell_{\text{cc}} \mapsto v_{\text{cc}}^{\text{mut}} * 0 \leq v_{\text{cp}}^{\text{mut}} \leq 100 * 0 \leq v_{\text{cc}}^{\text{mut}} \leq 100 \\
&* \ell_{\text{fifo10}} \mapsto v_{\text{fifo10}}^{\text{mut}} * v_{\text{cc}}^{\text{mut}} = v_{\text{cp}}^{\text{mut}} + \text{size}(v_{\text{fifo10}}^{\text{mut}}) * 0 \leq \text{size}(v_{\text{fifo10}}^{\text{mut}}) \leq 10 \\
&* v_{\text{fifo10}}^{\text{mut}} = (v_{\text{cp}} + \text{size}(v_{\text{fifo10}}^{\text{mut}})) :: \dots :: (v_{\text{cp}} + 1) :: \text{nil} \\
&* \left( v_{\text{cp}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{push}}^{\text{cp}}, L_{\text{push}}^{\text{cp}}, \text{False}) \right) \\
&* \left( v_{\text{cc}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{pop}}^{\text{cc}}, L_{\text{pop}}^{\text{cc}}, \text{False}) \right)
\end{aligned}$$

*Producer & consumer counters.*

*Bounded FIFO & its relationship to counters*

*Signal set by producer.*

*Signal set by consumer.*

$$\begin{aligned}
P_{\text{mut}} &:= \exists v_{\text{fifo10}}^{\text{mut}}. P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\
P_{\text{mut}}^{\text{no:s push}} &:= \exists v_{\text{cp}}^{\text{mut}}, v_{\text{cc}}^{\text{mut}}. \\
&[\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}}^{\text{mut}} * [\frac{1}{2}] \ell_{\text{cc}} \mapsto v_{\text{cc}}^{\text{mut}} * 0 \leq v_{\text{cp}}^{\text{mut}} \leq 100 * 0 \leq v_{\text{cc}}^{\text{mut}} \leq 100 \\
&* \ell_{\text{fifo10}} \mapsto v_{\text{fifo10}}^{\text{mut}} * v_{\text{cc}}^{\text{mut}} = v_{\text{cp}}^{\text{mut}} + \text{size}(v_{\text{fifo10}}^{\text{mut}}) * 0 \leq \text{size}(v_{\text{fifo10}}^{\text{mut}}) \leq 10 \\
&* v_{\text{fifo10}}^{\text{mut}} = (v_{\text{cp}} + \text{size}(v_{\text{fifo10}}^{\text{mut}})) :: \dots :: (v_{\text{cp}} + 1) :: \text{nil} \\
&* \left( v_{\text{cp}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{push}}^{\text{cp}}, L_{\text{push}}^{\text{cp}}, \text{False}) \right) \\
&* \left( v_{\text{cc}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{pop}}^{\text{cc}}, L_{\text{pop}}^{\text{cc}}, \text{False}) \right)
\end{aligned}$$

*Shorthand for lock invariant without push-signal chunk.*

$$\begin{aligned}
P_{\text{mut}}^{\text{no:s pop}} &:= \exists v_{\text{cp}}^{\text{mut}}, v_{\text{cc}}^{\text{mut}}. \\
&[\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}}^{\text{mut}} * [\frac{1}{2}] \ell_{\text{cc}} \mapsto v_{\text{cc}}^{\text{mut}} * 0 \leq v_{\text{cp}}^{\text{mut}} \leq 100 * 0 \leq v_{\text{cc}}^{\text{mut}} \leq 100 \\
&* \ell_{\text{fifo10}} \mapsto v_{\text{fifo10}}^{\text{mut}} * v_{\text{cc}}^{\text{mut}} = v_{\text{cp}}^{\text{mut}} + \text{size}(v_{\text{fifo10}}^{\text{mut}}) * 0 \leq \text{size}(v_{\text{fifo10}}^{\text{mut}}) \leq 10 \\
&* v_{\text{fifo10}}^{\text{mut}} = (v_{\text{cp}} + \text{size}(v_{\text{fifo10}}^{\text{mut}})) :: \dots :: (v_{\text{cp}} + 1) :: \text{nil} \\
&* \left( v_{\text{cp}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{push}}^{\text{cp}}, L_{\text{push}}^{\text{cp}}, \text{False}) \right) \\
&* \left( v_{\text{cc}}^{\text{mut}} > 0 \rightarrow \text{signal}(\text{id}_{\text{pop}}^{\text{cc}}, L_{\text{pop}}^{\text{cc}}, \text{False}) \right)
\end{aligned}$$

*Shorthand for lock invariant without pop-signal chunk.*

Figure A.15: Verification of realistic example 3.10a: Lock invariant.

$$\begin{aligned}
L_p(n, O_p) &:= 0 < n \leq 100 * \bigstar_{i=1, \dots, n-1} \text{uninitSig}(\text{id}_{\text{push}}^i) \\
&* (n > 0 \leftrightarrow O_p = \{\text{s}_{\text{push}}^n\}) * (n = 0 \leftrightarrow O_p = \emptyset)
\end{aligned}$$

*Loop invariant of producer.*

Figure A.16: Verification of realistic example 3.10a: Producer's loop invariant.

$$\begin{aligned}
L_c(n, O_c) &:= 0 < n \leq 100 * \bigstar_{i=1, \dots, n-1} \text{uninitSig}(\text{id}_{\text{pop}}^i) \\
&* (n > 0 \leftrightarrow O_c = \{\text{s}_{\text{pop}}^n\}) * (n = 0 \leftrightarrow O_c = \emptyset)
\end{aligned}$$

*Loop invariant of consumer.*

Figure A.17: Verification of realistic example 3.10a: Consumer's loop invariant.




Figure A.18: Verification of realistic example 3.10a: Main thread (consumer) forks producer thread.

$\forall \ell_{\text{fifo10}}, \ell_{\text{mut}}, \ell_{\text{cp}}, \ell_{\text{c}}.$ $\dots$ $\{\text{obs}(\{s_{\text{push}}^{100}\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto 100 * L_P(100, \{s_{\text{push}}^{100}\}) * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}})\}$ <b>while</b> ( $\forall v_{\text{cp}}.$ $\{\text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * L_P(v_{\text{cp}}, \{s_{\text{push}}^{v_{\text{cp}}}\}) * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}})\}$ $m.\text{lev} = 0 < 101 - v_{\text{cp}} = s_{\text{push}}^{v_{\text{cp}}}. \text{lev}$ <b>with mut await</b> ( $\forall O_P.$ $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}\} \ O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * L_P(v_{\text{cp}}, \{s_{\text{push}}^{v_{\text{cp}}}\} \ O_P) \\ * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}}) \ P_{\text{mut}} \end{array} \right\}$ $\forall v_{\text{fifo10}}^{\text{mut}}.$ $\{P_{\text{mut}} \ P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) * \dots\}$ <b>let</b> $f := [\text{fifo10}]$ <b>in</b> $\{L_P(v_{\text{cp}}, O_P) * \text{obs}(O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}})\}$ <b>if</b> $\text{size}(f) < 10$ <b>then</b> ( $\left\{ \begin{array}{l} \text{size}(v_{\text{fifo10}}^{\text{mut}}) < 10 * L_P(v_{\text{cp}}, O_P) * \text{obs}(O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} \\ * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \dots \text{A.20} \dots \\ \exists O'_P. \text{obs}(O'_P \ \{m\}) \\ * \text{if } \text{size}(v_{\text{fifo10}}^{\text{mut}}) \neq 10 \\ \text{then } [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 * L_P(v_{\text{cp}} - 1, O'_P) * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \text{else } \text{signal}(s_{\text{pop}}^{v_{\text{cp}}+10}, \text{False}) * L_{\text{pop}}^{v_{\text{cp}}+10} < L_{\text{push}}^{v_{\text{cp}}} * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} \\ * P_{\text{mut}}^{\text{no:pop}}(v_{\text{fifo10}}^{\text{mut}}) * L_P(v_{\text{cp}}, O'_P) \end{array} \right\}$ $\text{) else}$ ( $\left\{ \begin{array}{l} \text{size}(v_{\text{fifo10}}^{\text{mut}}) = 10 * L_P(v_{\text{cp}}, O_P) * \text{obs}(O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} \\ * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \text{size}(v_{\text{fifo10}}^{\text{mut}}) = 10 * L_P(v_{\text{cp}}, O_P) * \text{obs}(O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} \\ * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \text{Postlf}_p \end{array} \right\}$ $\text{);}$ $\{L_P(v_{\text{cp}}, O_P) * \text{obs}(O_P \ \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \ \text{Postlf}_p \}$ $\text{size}(\text{fifo10}) \neq 10$ $\{\dots * \text{if } \dots \text{ then } \dots \ P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \ P_{\text{mut}} \ \text{else } \dots\}$ $\text{);}$ $\left\{ \begin{array}{l} \exists O'_P. \text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}\} \ O'_P) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 \\ * L_P(v_{\text{cp}} - 1, \{s_{\text{push}}^{v_{\text{cp}}}\} \ O'_P) \\ * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}}) \end{array} \right\}$ $[\text{cp}] \neq 0$ <i>Remember that command is syntactic sugar.</i> $\left\{ \begin{array}{l} \text{if } v_{\text{cp}} - 1 \neq 0 \\ \text{then } \text{obs}(\{s_{\text{push}}^{v_{\text{cp}}-1}\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 * L_P(v_{\text{cp}} - 1, \{s_{\text{push}}^{v_{\text{cp}}-1}\}) \\ * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}}) \\ \text{else } \text{obs}(\emptyset) \end{array} \right\}$ $\text{)$ $\{\text{obs}(\{s_{\text{push}}^{100}\} \ \emptyset) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto 100 * L_P(100, \{s_{\text{push}}^{100}\}) * [\frac{1}{2}]\text{mutex}(m, P_{\text{mut}})\}$ $\dots$	<i>Continuation of Fig. A.18.</i>  PR-WHILE-DEC  <i>Justification for application of:</i> PR-AWAIT  PR-EXISTS  PR-LET & PR-READHEAPLOC  PR-IF  $\text{=: Postlf}_p$  PR-VIEWSHIFT & VS-SEMIMP  PR-EXP & PR-VIEWSHIFT & VS-SEMIMP  PR-LET & PR-READHEAPLOC & PR-EXP & PR-VIEWSHIFT & VS-SEMIMP  <i>Continued in Fig. A.18.</i>
--	--

Figure A.19: Verification of realistic example 3.10a: Producer loops.

$\forall \ell_{\text{fifo10}}, \ell_{\text{mut}}, \ell_{\text{cp}}, \ell_{\text{cc}}, v_{\text{cp}}, O_p, v_{\text{fifo10}}^{\text{mut}}.$   
 $\dots$   
 For definition of  $P_{\text{mut}}, L_p(n, O)$  and variations cf. Fig. A.15 and A.16.  
 $\{\text{size}(v_{\text{fifo10}}^{\text{mut}}) < 10 * L_p(v_{\text{cp}}, O_p) * \text{obs}(O_p \uplus \{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}})\}$

$\left\{ \begin{array}{l} [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}}^{\text{mut}} \quad \ell_{\text{cp}} \mapsto v_{\text{cp}} * v_{\text{cp}} = v_{\text{cp}}^{\text{mut}} \\ * (v_{\text{cp}} > 0 \mapsto \text{signal}(s_{\text{push}}^{\text{mut}}, \text{False})) \quad \text{signal}(s_{\text{push}}, \text{False}) \\ * (v_{\text{cp}} > 0 \leftrightarrow O_p = \{s_{\text{push}}^{\text{cp}}\}) * (v_{\text{cp}} = 0 \leftrightarrow O_p = \emptyset) \quad O_p = \{s_{\text{push}}^{\text{cp}}\} \\ * \dots \end{array} \right\}$	Continuation of Fig. A.19. PR-VIEWSHIFT & VS-SEMIMP
<p> <b>let</b> <math>c := [cp]</math> <b>in</b>  <math>[fifo10] := f \cdot (c :: \text{nil}); [cp] := c - 1</math>  <math>\{\ell_{\text{fifo10}} \mapsto v_{\text{fifo10}}^{\text{mut}} \cdot (v_{\text{cp}} :: \text{nil}) * \ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 * \dots\}</math>  <math>\{\text{obs}(\{s_{\text{push}}^{\text{cp}}, m\}) * \text{signal}(s_{\text{push}}^{\text{cp}}, \text{True}) * \dots\}</math>  <math>\{(v_{\text{cp}} - 1 = 0 \vee v_{\text{cp}} &gt; 0) * \dots\}</math> </p> <p> <i>case:</i> <math>v_{\text{cp}} - 1 = 0</math> <span style="float: right;"><i>Last iteration, nothing left to do.</i></span>  <math>\{\text{obs}(\{m\}) * [\frac{1}{2}]\ell_{\text{cp}} \mapsto 0 * P_{\text{mut}}\}</math> </p> $\left\{ \begin{array}{l} \exists O'_p. \text{obs}(O'_p \uplus \{m\}) \\ * \text{if } \text{size}(v_{\text{fifo10}}^{\text{mut}}) \neq 10 \\ \text{then } [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 * L_p(v_{\text{cp}} - 1, O'_p) * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \text{else } \text{signal}(s_{\text{pop}}^{\text{cp}}, \text{False}) * L_{\text{pop}}^{\text{cp}+10} < L_{\text{push}}^{\text{cp}} * [\frac{1}{2}]\ell_{\text{cp}} \mapsto v_{\text{cp}} \\ * P_{\text{mut}}^{\text{no};s} \text{pop}(v_{\text{fifo10}}^{\text{mut}}) * L_p(v_{\text{cp}}, O'_p) \end{array} \right\}$	PR-LET & PR-READHEAPLOC PR-ASSIGNTOHEAP (2x) PR-VIEWSHIFT & VS-SET SIGNAL PR-VIEWSHIFT & VS-SEMIMP PR-VIEWSHIFT & VS-OR PR-VIEWSHIFT & VS-SEMIMP PR-VIEWSHIFT & VS-SEMIMP = $\text{Postlf}_p$ For definition of $\text{Postlf}_p$ cf. Fig. A.19.
<p> <i>case:</i> <math>v_{\text{cp}} - 1 &gt; 0</math> <span style="float: right;"><i>Must create signal for next iteration.</i></span> </p> $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push}}^{\text{cp}-1}, m\}) * \text{uninitSig}(id_{\text{push}}^i) \\ * \text{uninitSig}(id_{\text{push}}^{\text{cp}-1}) \quad \text{signal}(s_{\text{push}}^{\text{cp}-1}, \text{False}) * \dots \end{array} \right\}$ <p> <math>\{\text{Postlf}_p\}</math>  <math>\{\text{Postlf}_p\}</math>  <math>\dots</math> </p>	PR-VIEWSHIFT & VS-SIGINIT PR-VIEWSHIFT & VS-SEMIMP Continued in Fig. A.19.

Figure A.20: Verification of realistic example 3.10a: Producer thread's production step.

$\forall \ell_{\text{fifo10}}, \ell_{\text{mut}}, \ell_{c_p}, \ell_{c_c}.$ $\dots$ $\{ \text{obs}(\{s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \ell_{c_c} \mapsto 100 * L_C(100, \{s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \}$ $\text{while (}$ $\quad \forall v_{c_c}. \left\{ \text{obs}(\{s_{\text{pop}}^{v_{c_c}}\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} * L_C(v_{c_c}, \{s_{\text{pop}}^{v_{c_c}}\}) * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \right\}$ $m.\text{lev} = 0 < 102 - v_{c_c} = s_{\text{pop}}^{v_{c_c}}.\text{lev}$ $\text{with mut await (}$ $\quad \forall O_C. \left\{ \begin{array}{l} \text{obs}(\{s_{\text{pop}}^{v_{c_c}}\} \ O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} * L_C(v_{c_c}, \{s_{\text{pop}}^{v_{c_c}}\} \ O_C) \\ * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \ P_{\text{mut}} \end{array} \right\}$ $\quad \forall v_{\text{mut}}^{\text{fifo10}}. \left\{ P_{\text{mut}} \ P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) * \dots \right\}$ $\text{let } f := [\text{fifo10}] \text{ in}$ $\left\{ L_C(v_{c_c}, O_C) * \text{obs}(O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \right\}$ $\text{if size}(f) > 0 \text{ then (}$ $\quad \left\{ \begin{array}{l} \text{size}(v_{\text{mut}}^{\text{fifo10}}) > 0 * L_C(v_{c_c}, O_C) * \text{obs}(O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} \\ * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \\ \dots \text{A.22} \dots \end{array} \right\}$ $\quad \text{Consumption step presented in Fig. A.22.}$ $\quad \left\{ \begin{array}{l} \exists O'_C. \text{obs}(O'_C \ \{m\}) \\ * \text{if size}(v_{\text{mut}}^{\text{fifo10}}) \neq 0 \\ \text{then } [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} - 1 * L_C(v_{c_c} - 1, O'_C) * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \\ \text{else signal}(s_{\text{push}}^{v_{c_c}}, \text{False}) * L_{\text{push}}^{v_{c_c}} < L_{\text{pop}}^{v_{c_c}} * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} \\ * P_{\text{mut}}^{\text{no};s_{\text{push}}^{v_{c_c}}} (v_{\text{mut}}^{\text{fifo10}}) * L_C(v_{c_c}, O'_C) \end{array} \right\}$ $\quad \text{Postf}_c$ $\text{) else (}$ $\quad \left\{ \begin{array}{l} \text{size}(v_{\text{mut}}^{\text{fifo10}}) = 0 * L_C(v_{c_c}, O_C) * \text{obs}(O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} \\ * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \\ \text{size}(v_{\text{mut}}^{\text{fifo10}}) = 10 * L_C(v_{c_c}, O_C) * \text{obs}(O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} \\ * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \\ \text{Postf}_c \end{array} \right\}$ $\text{);}$ $\left\{ L_C(v_{c_c}, O_C) * \text{obs}(O_C \ \{m\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} * P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \ \text{Postf}_c \right\}$ $\text{size}(\text{fifo10}) > 0$ $\left\{ \dots * \text{if} \dots \text{ then} \dots P'_{\text{mut}}(v_{\text{mut}}^{\text{fifo10}}) \ P_{\text{mut}} \ \text{else} \dots \right\}$ $\text{);}$ $\left\{ \begin{array}{l} \exists O'_C. \text{obs}(\{s_{\text{pop}}^{v_{c_c}}\} \ O'_C) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} - 1 \\ * L_C(v_{c_c} - 1, \{s_{\text{pop}}^{v_{c_c}}\} \ O'_C) \\ * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \end{array} \right\}$ $[c] \neq 0 \quad \text{Remember that command is syntactic sugar.}$ $\left\{ \begin{array}{l} \text{if } v_{c_c} - 1 \neq 0 \\ \text{then } \text{obs}(\{s_{\text{pop}}^{v_{c_c}-1}\}) * [\frac{1}{2}] \ell_{c_c} \mapsto v_{c_c} - 1 * L_C(v_{c_c} - 1, \{s_{\text{pop}}^{v_{c_c}-1}\}) \\ * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \\ \text{else } \text{obs}(\emptyset) \end{array} \right\}$ $\text{)}$ $\left\{ \text{obs}(\{s_{\text{pop}}^{100}\} \ \emptyset) * [\frac{1}{2}] \ell_{c_c} \mapsto 100 * L_C(100, \{s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \text{mutex}(m, P_{\text{mut}}) \right\}$ $\dots$	<p>Continuation of Fig. A.18.</p> <p>PR-WHILE-DEC</p> <p>Justification for application of: PR-AWAIT</p> <p>PR-EXISTS</p> <p>PR-LET &amp; PR-READHEAPLOC</p> <p>PR-IF</p> <p>=: Postf<sub>c</sub></p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-EXP &amp; PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-LET &amp; PR-READHEAPLOC &amp; PR-EXP &amp; PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>Continued in Fig. A.18.</p>
---	---

Figure A.21: Verification of realistic example 3.10a: Consumer loops.

$\forall \ell_{\text{fifo10}}, \ell_{\text{mut}}, \ell_{\text{cp}}, \ell_{\text{c}}, v_{\text{c}}, O_{\text{c}}, v_{\text{fifo10}}^{\text{mut}}.$   
 $\dots$   
 For definition of  $P_{\text{mut}}, L_{\text{c}}(n, O)$  and variations cf. Figures A.15 and A.17.

$\{\text{size}(v_{\text{fifo10}}^{\text{mut}}) > 0 * L_{\text{c}}(v_{\text{c}}, O_{\text{c}}) * \text{obs}(O_{\text{c}} \sqcup \{m\}) * [\frac{1}{2}] \ell_{\text{c}} \mapsto v_{\text{c}} * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}})\}$

$$\left\{ \begin{array}{l} [\frac{1}{2}] \ell_{\text{c}} \mapsto v_{\text{c}} * [\frac{1}{2}] \ell_{\text{c}} \mapsto v_{\text{c}}^{\text{mut}} \quad \ell_{\text{c}} \mapsto v_{\text{c}} * v_{\text{c}} = v_{\text{c}}^{\text{mut}} \\ * (v_{\text{c}} > 0 \mapsto \text{signal}(s_{\text{pop}}^{\text{mut}}, \text{False})) \quad \text{signal}(s_{\text{pop}}^{\text{v}_{\text{c}}}, \text{False}) \\ * (v_{\text{c}} > 0 \leftrightarrow O_{\text{c}} = \{s_{\text{pop}}^{\text{v}_{\text{c}}}\}) * (v_{\text{c}} = 0 \leftrightarrow O_{\text{c}} = \emptyset) \quad O_{\text{c}} = \{s_{\text{pop}}^{\text{v}_{\text{c}}}\} \\ * \dots \end{array} \right\}$$

**let**  $c := [c_{\text{c}}]$  **in**  
 $\text{[fifo10]} := \text{tail}(f): [c_{\text{c}}] := c - 1$   
 $\{\ell_{\text{fifo10}} \mapsto \text{tail}(v_{\text{fifo10}}^{\text{mut}}) * \ell_{\text{c}} \mapsto v_{\text{c}} - 1 * \dots\}$   
 $\{\text{obs}(\{s_{\text{pop}}^{\text{v}_{\text{c}}}, m\}) * \text{signal}(s_{\text{pop}}^{\text{v}_{\text{c}}}, \text{True}) * \dots\}$   
 $\{(v_{\text{c}} - 1 = 0 \vee v_{\text{c}} > 0) * \dots\}$

**case:**  $v_{\text{c}} - 1 = 0$  *Last iteration, nothing left to do.*

$\{\text{obs}(\{m\}) * [\frac{1}{2}] \ell_{\text{c}} \mapsto 0 * P_{\text{mut}}\}$

$$\left\{ \begin{array}{l} \exists O'_{\text{c}}. \text{obs}(O'_{\text{c}} \sqcup \{m\}) \\ * \text{if } \text{size}(v_{\text{fifo10}}^{\text{mut}}) \neq 0 \\ \text{then } [\frac{1}{2}] \ell_{\text{c}} \mapsto v_{\text{c}} - 1 * L_{\text{c}}(v_{\text{c}} - 1, O'_{\text{c}}) * P'_{\text{mut}}(v_{\text{fifo10}}^{\text{mut}}) \\ \text{else } \text{signal}(s_{\text{pop}}^{\text{v}_{\text{c}}}, \text{False}) * L_{\text{push}}^{\text{v}_{\text{c}}} < L_{\text{pop}}^{\text{v}_{\text{c}}} * [\frac{1}{2}] \ell_{\text{c}} \mapsto v_{\text{c}} \\ * P_{\text{mut}}^{\text{push}}(v_{\text{fifo10}}^{\text{mut}}) * L_{\text{c}}(v_{\text{c}}, O'_{\text{c}}) \end{array} \right\}$$

**case:**  $v_{\text{c}} - 1 > 0$  *Must create signal for next iteration.*

$$\left\{ \begin{array}{l} \text{obs}(\{s_{\text{pop}}^{\text{v}_{\text{c}} - 1}, m\}) * \bigstar_{i=1, \dots, v_{\text{c}} - 2} \text{uninitSig}(id_{\text{pop}}^i) \\ * \text{uninitSig}(id_{\text{pop}}^{\text{v}_{\text{c}} - 1}) \quad \text{signal}(s_{\text{pop}}^{\text{v}_{\text{c}} - 1}, \text{False}) * \dots \end{array} \right\}$$

$\{\text{Postf}_{\text{c}}\}$   
 $\{\text{Postf}_{\text{c}}\}$   
 $\dots$

*Continuation of Fig. A.21.*  
 PR-VIEWSHIFT & VS-SEMIMP  
 PR-LET & PR-READHEAPLOC  
 PR-ASIGNTOHEAP (2x)  
 PR-VIEWSHIFT & VS-SETIGNAL  
 PR-VIEWSHIFT & VS-SEMIMP  
 PR-VIEWSHIFT & VS-OR  
 PR-VIEWSHIFT & VS-SEMIMP  
 PR-VIEWSHIFT & VS-SEMIMP  
 = Postf<sub>c</sub>  
 For definition of Postf<sub>c</sub>  
 cf. Fig. A.21.  
 PR-VIEWSHIFT & VS-SIGNAL  
 PR-VIEWSHIFT & VS-SEMIMP  
 Continued in Fig. A.21.

Figure A.22: Verification of realistic example 3.10a: Consumer thread's consumption step.



## A.5.2 Case Study: Statically Unbounded Number of Communicating Parties

In this section we present and verify a program similar to the one from the previous section, but where the number of producer and consumer threads is not statically bounded. Fig. A.23 presents this program. It involves three parties: the main thread, producer threads and consumer threads. The main thread creates a shared buffer of size 1, generates a random number  $N > 0$  and spawns  $N$  producer and  $N$  consumer threads., which communicate via the shared buffer. Each producer tries to push a single random number into the buffer. In case the buffer is full, the producer busy-waits for it to become empty, i.e., it busy-waits for a consumer to pop the number currently stored in the buffer. After it pushed, the producer terminates. Each consumer tries to pop a single number from the buffer. In case the buffer is empty, it waits for some producer to push a number into the buffer. After it popped, the consumer terminates.

**Threads Racing for Buffer** Since the program is conceptionally similar to the one from the previous section, the same holds for its termination proof. There is, however, one fundamental difference. The previous program involved one producer and one consumer. So, it was clear which thread would push the  $k^{\text{th}}$  element, set the  $k^{\text{th}}$  push signal and discharge the corresponding obligation. The same was true for the consumer and the pop signals. Now,  $N$  producers and  $N$  consumers race for the buffer. We cannot statically determine which ones will win. Therefore, we cannot statically decide which threads we should delegate the push and pop obligations to.

**Signals** To solve this, we create  $N$  variations of every push and pop signal. That is, we create signals  $s_{\text{push},i}^k$  and  $s_{\text{pop},i}^k$  for  $1 \leq i, k \leq N$  where  $i$  refers to the  $i$ 's producer or consumer thread and  $k$  to the  $k^{\text{th}}$  element pushed to or popped from the buffer. We delegate the obligations for  $s_{\text{push},i}^1, \dots, s_{\text{push},i}^N$  to the  $i^{\text{th}}$  producer and the obligations for  $s_{\text{pop},i}^1, \dots, s_{\text{pop},i}^N$  to the  $i^{\text{th}}$  consumer.

When the  $i^{\text{th}}$  producer waits because the buffer is full and contains the  $k^{\text{th}}$  number, it sets all its remaining unset signals up to the  $k^{\text{th}}$  one. That is, signals  $s_{\text{push},i}^k, \dots, s_{\text{push},i}^N$  remain unset. (Reducing the number of obligations that a producer holds while waiting for a consumer to pop, reduces the potential for level conflicts and simplifies ordering the signals.) When the  $i^{\text{th}}$  producer pushes any number to the buffer, it sets all its remaining signals, before it terminates. A consumer can use any unset  $s_{\text{push},j}^k$  for any  $j$  to wait for the  $k^{\text{th}}$  element to arrive in the buffer. We proceed analogously with the consumers and their pop signals. Similarly to the proof from the previous section, we include our signals

<b>let</b> N := <b>randomNat</b> + 1 <b>in</b>	Random number of threads.
<b>let</b> buf := <b>cons</b> (-1) <b>in</b>	Shared buffer. Negative value $\leftrightarrow$ empty.
<b>let</b> mut := <b>new_mutex</b> <b>in</b>	
<b>let</b> c := <b>cons</b> (N) <b>in</b>	
<b>while</b> (	Main thread forks N producer and N consumer threads.
<b>fork</b> (	Producer
<b>with</b> mut <b>await</b> (	Producer busy-waits for empty buffer.
<b>let</b> b := [buf] <b>in</b>	
<b>if</b> b < 0 <b>then</b>	
[buf] := <b>randomNat</b> ;	Push random number to buffer.
b < 0	
)	
);	
<b>fork</b> (	Consumer
<b>with</b> mut <b>await</b> (	Consumer busy-waits for non-empty buffer.
<b>let</b> b := [buf] <b>in</b>	
<b>if</b> b $\geq$ 0 <b>then</b>	
[buf] := -1;	Pop buffer value.
b $\geq$ 0	
)	
);	
[c] := [c] - 1;	
c > 0	
) <b>do skip</b>	

Figure A.23: Program with statically unbounded number of threads communicating via a shared buffer of size 1.

$$\begin{array}{ll}
\text{VS-NEWGCELL} & \text{VS-SETGCELL} \\
\text{True} \Rightarrow \exists \widehat{\ell}. \widehat{\ell} \mapsto \widehat{v} & \widehat{\ell} \mapsto \widehat{v} \Rightarrow \widehat{\ell} \mapsto \widehat{v}'
\end{array}$$

Figure A.24: View shift rules for ghost variables.

in the lock invariant associated with the mutex that protects the shared buffer. Thereby, we allow all threads to share the signals in a synchronised fashion.

**Levels** The  $i^{\text{th}}$  producer uses some unset  $s_{\text{pop},j}^k$  to wait for the  $k^{\text{th}}$  element to be popped from the buffer. Meanwhile it holds obligations  $s_{\text{push},i}^k, \dots, s_{\text{push},i}^N$ . Analogously, the  $i^{\text{th}}$  consumer holds obligations  $s_{\text{pop},i}^k, \dots, s_{\text{pop},i}^N$  while using some unset  $s_{\text{push},j}^k$  to wait for the  $k^{\text{th}}$  element to be pushed. We assign one common level  $L_{\text{push}}^k$  to all  $s_{\text{push},1}^k, \dots, s_{\text{push},N}^k$  and one common level  $L_{\text{pop}}^k$  to all  $s_{\text{pop},1}^k, \dots, s_{\text{pop},N}^k$ . It must hold  $L_{\text{pop}}^k < L_{\text{push}}^j$  for all  $j > k$  and  $L_{\text{push}}^k < L_{\text{pop}}^h$  for all  $h \geq k$ . Hence, we chose  $L_{\text{push}}^k = 2 \cdot k - 1$  and  $L_{\text{pop}}^k = 2 \cdot k$ .

**Ghost Variables** In order to simplify the verification of this program, we introduce *ghost variables*, which are standard [139, 103, 100]. Ghost variables are ghost resources that behave like regular program variables, but since they only exist on the verification level, they cannot affect the runtime behaviour. We assume an infinite set of ghost locations  $\text{Locs}^G$  with  $\text{Locs}^G \cap \text{Locs} = \emptyset$  and a set of ghost values  $\text{Values}^G$ . We implement ghost variables in the form of *ghost heap cells* and represent them by *ghost points-to chunks*  $\widehat{\ell} \mapsto \widehat{v}$ . Fig. A.24 presents two new view shift rules for ghost variables. Rule VS-NEWGCELL allows to create a new ghost variables and VS-SETGCELL allows to change the value of an existing one. Adding ghost variables does not affect the soundness of our verification approach. In fact, the generalised logic we present and prove sound in the technical report [151] (appendix § B) includes ghost variables. In our termination proof, we introduce ghost variables  $w_i$  and  $r_i$ . Variable  $w_i$  tracks whether the  $i^{\text{th}}$  producer has already pushed ("written") an element to the buffer and  $r_i$  tracks whether the  $i^{\text{th}}$  consumer has already popped ("read") an element from the buffer. These variables conveniently allow us to refer to the number of elements written to and read from the buffer in the lock invariant. In Fig. A.25 – A.32, we present the full proof outline for our program from Fig. A.23. We scaled the proof outlines down to fit them into the page layout. You can find the original proof outlines in the extended version of the ghost signal paper [152] and in the technical report [151].

$\{\text{obs}(\emptyset)\}$

let  $N := \text{randomNat} + 1$  in let  $\text{buf} := \text{cons}(-1)$  in

let  $\text{mut} := \text{new\_mutex}$  in let  $c := \text{cons}(N)$  in

$\forall \ell_{\text{buf}}, \ell_{\text{mut}}, \ell_c.$

$\{\text{obs}(\emptyset) * N > 0 * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N\}$

$L_{\text{push}}^k := 2 \cdot k - 1, L_{\text{pop}}^k := 2 \cdot k$  for  $1 \leq k \leq N$

Later  $L_{\text{pop}}^k < L_{\text{push}}^j$  and  $L_{\text{push}}^k < L_{\text{pop}}^h$  must hold for  $j > k$  and  $h \geq k$ .

$$\left\{ \begin{array}{l} \exists id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N \\ \text{obs}(\{(id_{\text{push},1}^1, L_{\text{push}}^1), (id_{\text{pop},1}^1, L_{\text{pop}}^1), \dots, (id_{\text{push},N}^N, L_{\text{push}}^N), (id_{\text{pop},N}^N, L_{\text{pop}}^N)\}) \\ * N > 0 * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N \\ * \bigstar_{i,k=1,\dots,N} \text{signal}((id_{\text{push},i}^k, L_{\text{push}}^k), \text{False}) * \text{signal}((id_{\text{pop},i}^k, L_{\text{pop}}^k), \text{False}) \end{array} \right\}$$

$\forall id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N.$   
 $s_{\text{push},i}^k := (id_{\text{push},i}^k, L_{\text{push}}^k), s_{\text{pop},i}^k := (id_{\text{pop},i}^k, L_{\text{pop}}^k)$  for  $1 \leq i, k \leq N$

$$\left\{ \begin{array}{l} \exists id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N \text{ obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) \\ * N > 0 * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N \\ * \bigstar_{i,k=1,\dots,N} \text{signal}(s_{\text{push},i}^k, \text{False}) * \text{signal}(s_{\text{pop},i}^k, \text{False}) \end{array} \right\}$$

A consumer can use any unset  $s_{\text{push},i}^k$  to wait for the  $k^{\text{th}}$  number to be pushed.

When producer  $i$  pushes, it sets all  $s_{\text{push},i}^j$ . When it waits for a consumer to pop the  $k^{\text{th}}$  number, it sets all  $s_{\text{push},i}^j$  for  $j \leq k$  and keeps obligations for  $j > k$ .

So, if  $k < N$  numbers have been pushed, there exists some  $i$  s.t.  $s_{\text{push},i}^{k+1}$  is unset.

$$\left\{ \begin{array}{l} \exists \widehat{\ell}_1^w, \widehat{\ell}_1^r, \dots, \widehat{\ell}_N^w, \widehat{\ell}_N^r. \text{ obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) * N > 0 \\ * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N * \bigstar_{i=1,\dots,N} \widehat{\ell}_i^w \mapsto \text{False} * \widehat{\ell}_i^r \mapsto \text{False} \\ * \bigstar_{i,k=1,\dots,N} \text{signal}(s_{\text{push},i}^k, \text{False}) * \text{signal}(s_{\text{pop},i}^k, \text{False}) \end{array} \right\}$$

Ghost heap cell  $\widehat{\ell}_i^w / \widehat{\ell}_i^r$  records whether prod./cons.  $i$  has already pushed/popped.

$$\left\{ \begin{array}{l} \exists \widehat{\ell}_1^w, \widehat{\ell}_1^r, \dots, \widehat{\ell}_N^w, \widehat{\ell}_N^r. \text{ obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) * N > 0 \\ * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N * \bigstar_{i=1,\dots,N} \widehat{\ell}_i^w \mapsto \text{False} * \widehat{\ell}_i^r \mapsto \text{False} \\ * \bigstar_{i,k=1,\dots,N} \text{signal}(s_{\text{push},i}^k, \text{False}) * \text{signal}(s_{\text{pop},i}^k, \text{False}) \end{array} \right\}$$

$m := (\text{mut}, 0)$

Later  $m.\text{lev} < L_{\text{push}}^k$  and  $m.\text{lev} < L_{\text{pop}}^k$  must hold.

$$\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) * \text{mutex}(m, P) \\ N > 0 * \ell_{\text{buf}} \mapsto -1 * \text{uninit}(\ell_{\text{mut}}) * \ell_c \mapsto N * \bigstar_{i=1,\dots,N} \frac{1}{2} \widehat{\ell}_i^w \mapsto \text{False} * \frac{1}{2} \widehat{\ell}_i^r \mapsto \text{False} \\ * \bigstar_{i,k=1,\dots,N} \text{signal}(s_{\text{push},i}^k, \text{False}) * \text{signal}(s_{\text{pop},i}^k, \text{False}) \end{array} \right\}$$

...

PR-LET (2x)  
& PR-CONS

PR-LET (2x)  
& PR-NEWMUTEX  
& PR-CONS

PR-VIEWSHIFT  
& VS-GHOSTLOOP  
& VS-NEWSIGNAL

Creating  
 $N \cdot N$  push  
and  $N \cdot N$  pop  
signals.

cf. Fig. A.31  
and A.32.

PR-EXISTS

PR-VIEWSHIFT  
& VS-GHOSTLOOP  
& VS-NEWCCELL

PR-EXISTS

PR-VIEWSHIFT  
& VS-SEMIMP  
& VS-MUTINT

cf. Fig. A.31  
and A.32.

For definition of  
lock invariant  $P$   
cf. Fig. A.26.

Continued in  
Fig. A.30

Figure A.25: Verification of program A.23: Initialisation.

$$\begin{aligned}
P'(v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N) & \\
:= \ell_{\text{buf}} \mapsto v_{\text{buf}} \wedge v_{\text{buf}} \in \mathbb{Z} \wedge N > 0 & \\
& \\
* \bigstar_{i=1, \dots, N} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto w_i * W = |\{i \in \{1, \dots, N\} \mid w_i = \text{True}\}| & \\
& \\
* \bigstar_{i=1, \dots, N} [\frac{1}{2}] \widehat{\ell}_i^r \mapsto r_i * R = |\{i \in \{1, \dots, N\} \mid r_i = \text{True}\}| & \\
& \\
* (v_{\text{buf}} \geq 0 \rightarrow W > 0 \wedge R < N) & \\
* (v_{\text{buf}} < 0 \leftrightarrow R = W) * (v_{\text{buf}} \geq 0 \leftrightarrow R = W - 1) & \\
& \\
* \bigstar_{\substack{i=1, \dots, N \\ k=W+1, \dots, N}} \text{signal}(s_{\text{push}, i}^k, w_i) * \bigstar_{\substack{i=1, \dots, N \\ k=1, \dots, W}} \text{signal}(s_{\text{push}, i}^k, \_ ) & \\
& \\
* \bigstar_{\substack{i=1, \dots, N \\ k=R+1, \dots, N}} \text{signal}(s_{\text{pop}, i}^k, r_i) * \bigstar_{\substack{i=1, \dots, N \\ k=1, \dots, R}} \text{signal}(s_{\text{pop}, i}^k, \_ ) & \\
& \\
P := \exists v_{\text{buf}} \in \mathbb{Z}. \exists W, R \in \mathbb{N}. \exists w_1, \dots, w_N, r_1, \dots, r_N \in \mathbb{B}. & \\
P'(v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N) &
\end{aligned}$$

Shared buffer.

Negative  $\leftrightarrow$  empty.

Ghost value  $w_i$  tracks whether producer  $i$  already pushed.

$W$ : Number of elements written to buffer.

Ghost value  $r_i$  tracks whether consumer  $i$  already popped.

$R$ : Number of elements read from buffer.

Signals set by producers  $\mathcal{E}$  used by consumers to wait when buffer is empty.

Signals set by consumers  $\mathcal{E}$  used by producers to wait when buffer is empty.

Figure A.26: Verification of program A.23: Lock invariant.

$$\begin{aligned}
L_f(c) & := \ell_c \mapsto c \wedge 0 \leq c \leq N \wedge N > 0 \\
& \\
* \text{obs}(\{s_{\text{push}, i}^k, s_{\text{pop}, i}^k \mid 1 \leq i \leq c \wedge 1 \leq k \leq N\}) & \\
& \\
* \bigstar_{i=1, \dots, c} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} & \\
& \\
* [\frac{2-c}{2-N}] \text{mutex}(m, P) &
\end{aligned}$$

Decreasing counter.

Prods.  $\mathcal{E}$  cons.

$c + 1, \dots, N$  already forked.

Remaining obligations for threads that have not been forked, yet.

Ghost heap cells for unforked consumers.

Partial mutex chunk for unforked producers  $\mathcal{E}$  consumers.

Figure A.27: Verification of program A.23: Fork loop invariant.

$$\begin{aligned}
L_p^i(O) & := \exists n_{\text{push}}. 1 \leq n_{\text{push}} \leq N \\
& \\
* O = \{s_{\text{push}, i}^{n_{\text{push}}}, \dots, s_{\text{push}, i}^N\} & \\
& \\
* [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} &
\end{aligned}$$

Lower bound for index of next element that will be pushed by any thread.

Figure A.28: Verification of program A.23: Loop invariant for producer  $i$ .

$$\begin{aligned}
L_c^i(O) & := \exists n_{\text{pop}}. 1 \leq n_{\text{pop}} \leq N \\
& \\
* O = \{s_{\text{pop}, i}^{n_{\text{pop}}}, \dots, s_{\text{pop}, i}^N\} & \\
& \\
* [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} &
\end{aligned}$$

Lower bound for index of next element that will be popped by any thread.

Figure A.29: Verification of program A.23: Loop invariant for consumer  $i$ .

...

$$\forall \ell_{\text{buf}}, \ell_{\text{mut}}, \ell_c, id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N, \widehat{\ell}_1^w, \widehat{\ell}_1^r, \dots, \widehat{\ell}_N^w, \widehat{\ell}_N^r.$$

$$\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) * \text{mutex}(m, P) \\ * \ell_c \mapsto N * \bigstar_{i=1, \dots, N} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} \end{array} \right\}$$

$$\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},1}^1, s_{\text{pop},1}^1, \dots, s_{\text{push},N}^N, s_{\text{pop},N}^N\}) * \text{mutex}(m, P) \\ * \ell_c \mapsto N * \bigstar_{i=1, \dots, N} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} \quad L_f(N) \end{array} \right\}$$

**while** (  
 $\forall v_c.$   
 $\{L_f(v_c)\}$   
**fork** (  
 $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}) * [\frac{1}{N}] \text{mutex}(m, P) \\ * [\frac{1}{2}] \widehat{\ell}_{v_c}^w \mapsto \text{False} \end{array} \right\}$   
...  
 $\{ \text{obs}(\emptyset) * [\frac{1}{N}] \text{mutex}(m, P) \}$   
);  
 $\left\{ \begin{array}{l} \ell_c \mapsto v_c * \bigstar_{i=1, \dots, v_c-1} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * \bigstar_{i=1, \dots, v_c} [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} \\ * [\frac{2 \cdot v_c - 1}{2 \cdot N}] \text{mutex}(m, P) \\ * \text{obs} \left( \begin{array}{l} \{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\} \uplus \{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\} \\ \uplus \{s_{\text{push},i}^k, s_{\text{pop},i}^k \mid 1 \leq i \leq v_c - 1 \wedge 1 \leq k \leq N\} \end{array} \right) \end{array} \right\}$   
**fork** (  
 $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) * [\frac{1}{N}] \text{mutex}(m, P) \\ * [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto \text{False} \end{array} \right\}$   
...  
 $\{ \text{obs}(\emptyset) * [\frac{1}{N}] \text{mutex}(m, P) \}$   
);  
 $\left\{ \begin{array}{l} \ell_c \mapsto v_c * \bigstar_{i=1, \dots, v_c-1} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} \\ * [\frac{2 \cdot v_c - 2}{2 \cdot N}] \text{mutex}(m, P) \\ * \text{obs} \left( \begin{array}{l} \{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\} \uplus \\ \{s_{\text{push},i}^k, s_{\text{pop},i}^k \mid 1 \leq i \leq v_c - 1 \wedge 1 \leq k \leq N\} \end{array} \right) \end{array} \right\}$   
 $[c] := [c] - 1;$   
 $\left\{ \begin{array}{l} \ell_c \mapsto v_c - 1 * \bigstar_{i=1, \dots, v_c-1} [\frac{1}{2}] \widehat{\ell}_i^w \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_i^r \mapsto \text{False} \\ * [\frac{2 \cdot v_c - 2}{2 \cdot N}] \text{mutex}(m, P) \\ * \text{obs}(\{s_{\text{push},i}^k, s_{\text{pop},i}^k \mid 1 \leq i \leq v_c - 1 \wedge 1 \leq k \leq N\}) \end{array} \right\}$   
 $\{L_f(v_c - 1)\}$   
 $c > 0$   
 $\{L_f(v_c - 1)\}$   
if  $v_c - 1 > 0$  then  $L_f(v_c - 1)$  else  $\text{obs}(\emptyset)$   
**do skip**  
 $\{L_f(N) \quad \text{obs}(\emptyset)\}$

Continuation of Fig. A.25

PR-VIEWSHIFT & VS-SEMIMP

For definition of fork loop invariant cf. Fig. A.27.

PR-WHILE-DEC

PR-FORK

Producer loop on Fig. A.31.

PR-FORK

Consumer loop on Fig. A.32.

PR-LET & PR-READHEAPLOC & PR-ASSTOHEAP  
Remember that this command is syntactic sugar.

PR-EXP  
PR-VIEWSHIFT & VS-SEMIMP

Figure A.30: Verification of program A.23: Fork loop.

<p>...</p> $\forall \ell_{\text{buf}}, \ell_{\text{mut}}, \ell_c, id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N, \widehat{\ell}_1^w, \widehat{\ell}_1^r, \dots, \widehat{\ell}_N^w, \widehat{\ell}_N^r, v_c.$ $\{ \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}) * [\frac{1}{N}] \text{mutex}(m, P) * [\frac{1}{2}] \ell_{v_c}^w \mapsto \text{False} \}$ $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}) * L_P^{v_c}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}) \\ * [\frac{1}{N}] \text{mutex}(m, P) \end{array} \right\}$ <p><math>m.\text{lev} = 0 &lt; 2 \cdot k - 1 = s_{\text{push},v_c}^k.\text{lev}</math> for <math>1 \leq k \leq N</math></p> <p><b>with mut await</b> (</p> <p><math>\forall O.</math></p> $\{ \text{obs}(O \uplus \{m\}) * L_P^{v_c}(O) * P \}$ <p><math>\forall n_{\text{push}}, v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N.</math></p> $\widehat{P} := P'(v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N)$ $\{ \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}, m) * [\frac{1}{2}] \ell_{v_c}^w \mapsto \text{False} * \widehat{P} \}$ $\left\{ \begin{array}{l} \text{max}(n_{\text{push}}, W + 1) \\ \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}, m) * [\frac{1}{2}] \ell_{v_c}^w \mapsto \text{False} * \widehat{P} \end{array} \right\}$ <p><b>if</b> <math>b &lt; 0</math> <b>then</b></p> $\{ v_{\text{buf}} < 0 * \dots \}$ <p><math>[b_{\text{buf}}] := \text{randomNat}</math></p> $\left\{ \begin{array}{l} \exists v'_{\text{buf}}. \ell_{\text{buf}} \mapsto v'_{\text{buf}} \wedge v'_{\text{buf}} \geq 0 \\ * [\frac{1}{2}] \ell_{v_c}^w \mapsto \text{False} * [\frac{1}{2}] \ell_{v_c}^w \mapsto w_{v_c} \ell_{v_c}^w \mapsto \text{False} * \dots \end{array} \right\}$ $\{ \ell_{v_c}^w \mapsto \text{True} * \dots \}$ $\left\{ \begin{array}{l} \text{obs}(\{s_{\text{push},v_c}^{\text{max}(n_{\text{push}}, W + 1)}, \dots, s_{\text{push},v_c}^N\}, m) \\ * \bigstar_{k=\text{max}(n_{\text{push}}, W + 1), \dots, N} \text{signal}(s_{\text{push},v_c}^k, \text{True}) * \dots \end{array} \right\}$ $\left\{ \begin{array}{l} \exists O'. \text{obs}(O' \uplus \{m\}) \\ * \text{if } v_{\text{buf}} < 0 \\ \text{then } P * O = \emptyset \\ \text{else } L_P^{v_c}(O') * \exists i. \text{signal}(s_{\text{pop},i}^{R+1}, \text{False}) * P_P^{\text{rest}} \\ * L_{\text{pop}}^{R+1} \prec_L O' * m.\text{lev} \prec_L O' \end{array} \right\}$ <p><b>else</b></p> $\{ v_{\text{buf}} \geq 0 \wedge R = W - 1 * \text{obs}(\{s_{\text{push},v_c}^{\text{max}(n_{\text{push}}, W + 1)}, \dots, s_{\text{push},v_c}^N\}, m) * \dots \}$ $L_{\text{pop}}^{R+1} = 2 \cdot (R + 1) = 2 \cdot W$ $< 2 \cdot \text{max}(n_{\text{push}}, W + 1) - 1 \leq L_{\text{push}}^k$ <p>for <math>\text{max}(n_{\text{push}}, W + 1) \leq k \leq N.</math></p> <p><math>m.\text{lev} = 0 &lt; L_{\text{push}}^k</math> for any <math>k.</math></p> $R < N \wedge \widehat{P} \Rightarrow \exists i. r_i = \text{False} \wedge \bigstar_{j=1, \dots, N} \text{signal}(s_{\text{pop},j}^{R+1}, r_j)$ $\Rightarrow \exists i. \text{signal}(s_{\text{pop},i}^{R+1}, \text{False})$ <p><math>\{\text{Postf}_P\}</math></p> $\{ \text{obs}(\{s_{\text{push},v_c}^{\text{max}(n_{\text{push}}, W + 1)}, \dots, s_{\text{push},v_c}^N\}, m) * [\frac{1}{2}] \ell_{v_c}^w \mapsto \text{False} * \widehat{P} \text{Postf}_P \}$ <p><math>b &lt; 0</math></p> $\{\text{Postf}_P\}$ <p>)</p> $\left\{ \begin{array}{l} \exists O'. \text{obs}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}, O') * L_P^{v_c}(\{s_{\text{push},v_c}^1, \dots, s_{\text{push},v_c}^N\}) \\ * [\frac{1}{N}] \text{mutex}(m, P) * O' = \emptyset \end{array} \right\}$ <p><math>\{\text{obs}(\emptyset)\}</math></p>	<p>Continuation of Fig. A.30</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>For definition of this producer's loop invariant <math>L_P^{v_c}(O)</math> cf. Fig. A.28.</p> <p>Justification for application of: PR-AWAIT</p> <p>For definition of lock invariant <math>P \wp \widehat{P}</math> cf. Fig. A.26.</p> <p>PR-EXISTS</p> <p>PR-VIEWSHIFT &amp; VS-SETSIGNAL</p> <p>PR-IF</p> <p>PR-ASSIGNTOHEAP &amp; PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-VIEWSHIFT &amp; VS-SETGCELL</p> <p>PR-VIEWSHIFT &amp; VS-GHOSTLOOP &amp; VS-SETSIGNAL</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>Define <math>P_P^{\text{rest}}</math> such that</p> <p>(<math>\exists i.</math></p> $\text{signal}(s_{\text{pop},i}^{R+1}, \text{False})$ $* P_P^{\text{rest}})$ <p><math>\Rightarrow P</math></p> <p><math>==: \text{Postf}_P</math></p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>Justification of <b>else</b> branch in <math>\text{Postf}_P.</math></p> <p>PR-EXP</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p>
--	--

Figure A.31: Verification of program A.23: Producer loop.

<p>...</p> $\forall \ell_{\text{buf}}, \ell_{\text{mut}}, \ell_c, id_{\text{push},1}^1, id_{\text{pop},1}^1, \dots, id_{\text{push},N}^N, id_{\text{pop},N}^N, \widehat{\ell}_1^w, \widehat{\ell}_1^r, \dots, \widehat{\ell}_N^w, \widehat{\ell}_N^r, v_c.$ $\left\{ \text{obs}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) * [\frac{1}{N}] \text{mutex}(m, P) * [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto \text{False} \right.$ $\left. \left\{ \text{obs}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) * L_c^{v_c}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) \right. \right.$ $\left. \left. * [\frac{1}{N}] \text{mutex}(m, P) \right\} \right\}$ <p><math>m.\text{lev} = 0 &lt; 2 \cdot k - 1 = s_{\text{push},v_c}^k.\text{lev}</math> for <math>1 \leq k \leq N</math></p> <p><b>with mut await</b> (</p> <p style="padding-left: 20px;"><math>\forall O.</math></p> $\left\{ \text{obs}(\boxed{O} \boxplus \{m\}) * L_c^{v_c}(\boxed{O}) * P \right\}$ <p style="padding-left: 20px;"><math>\forall n_{\text{pop}}, v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N.</math></p> <p style="padding-left: 20px;"><math>\widehat{P} := P'(v_{\text{buf}}, W, R, w_1, \dots, w_N, r_1, \dots, r_N)</math></p> $\left\{ \text{obs}(\{s_{\text{pop},v_c}^{n_{\text{pop}}}, \dots, s_{\text{pop},v_c}^N, m\}) * [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto \text{False} * \boxed{P} \right\}$ $\left\{ \text{obs}(\{s_{\text{pop},v_c}^{\max(n_{\text{pop}}, R+1)}, \dots, s_{\text{pop},v_c}^N, m\}) * [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto \text{False} * \widehat{P} \right\}$ $\left\{ [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto \text{False} * [\frac{1}{2}] \widehat{\ell}_{v_c}^r \mapsto r_{v_c} \mid \widehat{\ell}_{v_c}^r \mapsto \text{False} \wedge r_{v_c} = \text{False} \wedge R < N * \dots \right\}$ <p><b>if</b> <math>b \geq 0</math> <b>then</b></p> $\left\{ v_{\text{buf}} \geq 0 * \dots \right\}$ <p style="padding-left: 20px;"><math>[b] := -1</math></p> $\left\{ \ell_{\text{buf}} \mapsto -1 * \dots \right\}$ $\left\{ \widehat{\ell}_{v_c}^r \mapsto \text{True} * \dots \right\}$ $\left\{ \text{obs}(\{s_{\text{pop},v_c}^{\max(n_{\text{pop}}, R+1)}, \dots, s_{\text{pop},v_c}^N, m\}) * \bigstar_{k=\max(n_{\text{pop}}, R+1), \dots, N} \text{signal}(s_{\text{pop},v_c}^k, \text{True}) * \dots \right\}$ $\left\{ \begin{array}{l} \exists O'. \text{obs}(O' \boxplus \{m\}) \\ * \text{if } v_{\text{buf}} \geq 0 \\ \text{then } P * O = \emptyset \\ \text{else } L_c^{v_c}(O') * \exists i. \text{signal}(s_{\text{push},i}^{W+1}, \text{False}) * P_c^{\text{rest}} \\ * L_{\text{push}}^{W+1} \prec_L O' * m.\text{lev} \prec_L O' \end{array} \right\}$ <p><b>else</b></p> $\left\{ v_{\text{buf}} < 0 \wedge R = W < N * \text{obs}(\{s_{\text{pop},v_c}^{\max(n_{\text{pop}}, R+1)}, \dots, s_{\text{pop},v_c}^N, m\}) * \dots \right\}$ $L_{\text{push}}^{W+1} = 2 \cdot (W+1) - 1 = 2 \cdot (R+1) - 1$ $< 2 \cdot \max(n_{\text{pop}}, R+1) \leq L_{\text{pop}}^k \text{ for } \max(n_{\text{pop}}, R+1) \leq k \leq N.$ $m.\text{lev} = 0 < L_{\text{pop}}^k \text{ for any } k.$ $W < N \wedge \widehat{P} \Rightarrow \exists i. w_i = \text{False} \wedge \bigstar_{j=1, \dots, N} \text{signal}(s_{\text{push},j}^{W+1}, w_j)$ $\Rightarrow \exists i. \text{signal}(s_{\text{push},i}^{W+1}, \text{False})$ $\{ \text{Postf}_c \}$ $\left\{ \text{obs}(\{s_{\text{pop},v_c}^{\max(n_{\text{pop}}, R+1)}, \dots, s_{\text{pop},v_c}^N, m\}) * \dots \mid \text{Postf}_c \right\}$ <p><math>b \geq 0</math></p> $\{ \text{Postf}_c \}$ <p>)</p> $\left\{ \begin{array}{l} \exists O'. \text{obs}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) * L_c^{v_c}(\{s_{\text{pop},v_c}^1, \dots, s_{\text{pop},v_c}^N\}) \\ * [\frac{1}{N}] \text{mutex}(m, P) * O' = \emptyset \end{array} \right\}$ <p><math>\{ \text{obs}(\emptyset) \}</math></p>	<p>Continuation of Fig. A.30</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP For definition of this consumer's loop invariant <math>L_c^{v_c}(O)</math> cf. Fig. A.29.</p> <p>Justification for application of: PR-AWAIT For definition of lock invariant <math>P \boxtimes \widehat{P}</math> cf. Fig. A.26.</p> <p>PR-EXISTS</p> <p>PR-VIEWSHIFT &amp; VS-GHOSTLOOP &amp; VS-SET SIGNAL</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-IF</p> <p>PR-ASIGNTOHEAP</p> <p>PR-VIEWSHIFT &amp; VS-SETGCELL</p> <p>PR-VIEWSHIFT &amp; VS-GHOSTLOOP &amp; VS-SET SIGNAL</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP Define <math>P_c^{\text{rest}}</math> such that (<math>\exists i.</math> <math>\text{signal}(s_{\text{push},i}^{W+1}, \text{False})</math> <math>* P_c^{\text{rest}}</math>) <math>\Rightarrow P</math></p> <p><math>\equiv: \text{Postf}_c</math></p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p> <p>Justification of else branch in <math>\text{Postf}_c</math>.</p> <p>PR-EXP</p> <p>PR-VIEWSHIFT &amp; VS-SEMIMP</p>
---	--

Figure A.32: Verification of program A.23: Consumer loop.



# Appendix B

## Ghost Signals: Generalised Logic

This chapter formalizes a generalized version of the logic presented in chapter 3 and presents additional case studies. The contents of this chapter were previously published as stand-alone technical report on Zenodo:

**T. Reinhard** and Bart Jacobs. *Ghost Signals: Verifying Termination of Busy Waiting (Technical Report)*. Zenodo, 2021. <https://doi.org/10.5281/zenodo.4775100>

### B.1 Universe

Throughout this work we assume the existence of the following sets:

- $\mathcal{X}$ : An infinite set of program variables.
- $\mathit{Locs}$ : An infinite set of heap locations.
- $\mathit{Locs}^G$ : An infinite set of ghost locations.
- $\mathit{Levs}, <_{\mathit{L}}$ : An infinite, well-founded partially ordered set of levels.
- $\Delta, <_{\Delta}$ : An infinite, well-founded partially ordered set of degrees.
- $\mathit{ID}$ : An infinite set of IDs.

- $\Theta$ : An infinite, totally ordered and well-founded set of thread IDs.
- *Values*: A set of values which includes:
  - A unit value  $\mathbf{tt} \in \text{Values}$
  - Booleans  $\mathbb{B} = \{\mathbf{True}, \mathbf{False}\} \subset \text{Values}$
  - Heap locations  $\mathcal{Locs} \subset \text{Values}$
- $\text{Values}^G$ : A set of ghost values.
- *Ops*: A set of operations (i.e. partial functions) on values.

We denote program variables by  $x$ , heap locations by  $\ell$ , ghost locations by  $\widehat{\ell}$ , levels by  $L$ , degrees by  $\delta$ , IDs by  $id$ , thread IDs by  $\theta$ , values by  $v$ , ghost values by  $\widehat{v}$ , boolean by  $b$  and operations by  $op$ .

## B.2 General

**Definition B.2.1** (Projections). *For any Cartesian product  $C = \prod_{i \in I} A_i$  and any index  $k \in I$ , we denote the  $k^{\text{th}}$  projection by  $\pi_k^C : \prod_{i \in I} A_i \rightarrow A_k$ . We define*

$$\pi_k^C((a_i)_{i \in I}) := a_k.$$

*In case the domain  $C$  is clear from the context, we write  $\pi_k$  instead of  $\pi_k^C$ .*

In the following we define our notion of *bags*, in the literature also referred to as *multisets*.

**Definition B.2.2** (Bags). *For any set  $X$  we define the set of bags  $\text{Bags}(X)$  and the set of finite bags  $\text{Bags}_{\text{fin}}(X)$  over  $X$  as*

$$\begin{aligned} \text{Bags}(X) &:= X \rightarrow \mathbb{N}, \\ \text{Bags}_{\text{fin}}(X) &:= \{B \in \text{Bags}(X) \mid \{x \in B \mid B(x) > 0\} \text{ finite}\}. \end{aligned}$$

*We define union and subtraction of bags as*

$$\begin{aligned} (B_1 \uplus B_2)(x) &:= B_1(x) + B_2(x), \\ (B_1 \setminus B_2)(x) &:= \max(0, B_1(x) - B_2(x)). \end{aligned}$$

*For finite bags where the domain is clear from the context, we define the following set-like notation:*

$$\begin{aligned} \emptyset &:= x \mapsto 0, \\ \{\{x\}\} &:= \begin{cases} x & \mapsto 1 \\ y & \mapsto 0 \text{ for } y \neq x, \end{cases} \\ \{\{x_1, \dots, x_n\}\} &:= \bigoplus_{i=1}^n \{\{x_i\}\}. \end{aligned}$$

$$v \in \text{Values} \quad x \in \mathcal{X} \quad op \in \text{Ops}$$

$$e \in \text{Exps} ::= x \mid v \mid e = e \mid \neg e \mid op(\bar{e})$$

$$c \in \text{Cmds} ::= \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \mid \mathbf{fork} \ c \mid$$

$$\mathbf{let} \ x := c \ \mathbf{in} \ c \mid \mathbf{if} \ c \ \mathbf{then} \ c \mid$$

$$\mathbf{cons}(e) \mid [e] \mid [e] := e \mid$$

$$\mathbf{new\_mutex} \mid \mathbf{acquire} \ e \mid \mathbf{release} \ e \mid$$

$$e \mid$$

consumeItPerm

intermediate  
representation

Figure B.1: Syntax.

We define the following set-like notations for element and subset relationship:

$$x \in B \quad \Leftrightarrow \quad B(x) > 0,$$

$$B_1 \subseteq B_2 \quad \Leftrightarrow \quad \forall x \in B_1. B_1(x) \leq B_2(x),$$

$$B_1 \subset B_2 \quad \Leftrightarrow \quad \exists C \subseteq B_1. C \neq \emptyset \wedge B_1 = B_2 \setminus C.$$

For any bag  $B \in \text{Bags}(X)$  and predicate  $P \subseteq X$  we define the following refinement:

$$\{\{x \in B \mid P(x)\}\} := \begin{cases} x \mapsto B(x) & \text{if } P(x), \\ x \mapsto 0 & \text{otherwise.} \end{cases}$$

**Definition B.2.3** (Disjoint Union). Let  $A, B$  be sets. We define their disjoint union as

$$A \sqcup B := A \cup B$$

if  $A \cap B = \emptyset$  and leave it undefined otherwise.

## B.3 Syntax

**Definition B.3.1.** We define the sets of commands  $\text{Cmds}$  and expressions  $\text{Exps}$  according to the syntax presented in Figure B.1.

We define  $c ; c'$  as shorthand for  $\mathbf{let} \ x := c \ \mathbf{in} \ c'$  where  $x$  does not occur free in  $c'$  but  $\mathbf{let} \ ; \cdot$  bind stronger. Further, we define  $e \neq e'$  as abbreviation for  $\neg(e = e')$ .

```

let x := cons(0) in
let m := new_mutex in
fork (while (acquire m;
             let y := [x] in
             release m;
             y = 0)
      do skip);
acquire m;
[x] := 1;
release m

```

Figure B.2: Example Program.

## B.4 Example

Figure B.2 presents the example program we plan to verify. For this example we let *Values* include natural numbers.

## B.5 Resources

In this section we define physical resources. We will use the physical resources to define the semantics of our programming language.

**Definition B.5.1** (Physical Resources). *We define the set of physical resources  $\mathcal{R}^{\text{phys}}$  syntactically as follows:*

$$r^{\text{p}} \in \mathcal{R}^{\text{phys}} ::= \ell \mapsto v \mid \text{unlocked}_{\text{pRes}}(\ell) \mid \text{locked}_{\text{pRes}}(\ell)$$

$$\ell \in \mathcal{Locs} \quad v \in \text{Values}$$

**Definition B.5.2** (Physical Heaps). *We define the set of physical heaps as*

$$\text{Heaps}^{\text{phys}} := \mathcal{P}_{\text{fin}}(\mathcal{R}^{\text{phys}})$$

*and the function  $\text{locs}_{\text{pRes}} : \text{Heaps}^{\text{phys}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$  mapping physical heaps to the sets of allocated heap locations as*

$$\text{locs}_{\text{pRes}}(h) ::= \{ \ell \in \mathcal{Locs} \mid \text{unlocked}_{\text{pRes}}(\ell) \in h \vee \text{locked}_{\text{pRes}}(\ell) \in h \vee \exists v \in \text{Values}. \ell \mapsto v \in h \}$$

*We denote physical heaps by  $h$ .*

## B.6 Semantics

**Definition B.6.1** (Evaluation of Closed Expressions). *We define a partial evaluation function  $\llbracket \cdot \rrbracket : Exps \rightarrow Values$  on expressions by recursion on the structure of expressions as follows:*

$$\begin{aligned}
 \llbracket v \rrbracket &::= v && \text{if } v \in Values \\
 \llbracket e = e' \rrbracket &::= \text{True} && \text{if } \llbracket e \rrbracket = \llbracket e' \rrbracket \neq \perp \\
 \llbracket e = e' \rrbracket &::= \text{False} && \text{if } \llbracket e \rrbracket \neq \llbracket e' \rrbracket \wedge \llbracket e \rrbracket \neq \perp \wedge \llbracket e' \rrbracket \neq \perp \\
 \llbracket \neg e \rrbracket &::= \text{False} && \text{if } \llbracket e \rrbracket = \text{True} \\
 \llbracket \neg e \rrbracket &::= \text{True} && \text{if } \llbracket e \rrbracket = \text{False} \\
 \llbracket e \rrbracket &::= \perp && \text{otherwise}
 \end{aligned}$$

We identify closed expressions  $e$  with their ascribed value  $\llbracket e \rrbracket$ .

**Definition B.6.2** (Evaluation Context). *We define the set of evaluation contexts  $EvalCtxts$  as follows:*

$$\begin{aligned}
 E \in EvalCtxts &::= \text{if } \square \text{ then } c \mid \text{let } x := \square \text{ in } c \\
 &c \in Ccmds \quad x \in \mathcal{X}
 \end{aligned}$$

For any  $c \in Ccmds$  and  $E \in EvalCtxts$ , we define  $E[c] := E[\square \mapsto c]$ .

Note that for every  $c \in Ccmds, E \in EvalCtxts$ , we have  $E[c] \in Ccmds$ .

**Definition B.6.3** (Single Thread Reduction Relation). *We define a reduction relation  $\rightsquigarrow_{\text{st}}$  for single threads according to the rules presented in Figure B.3. A reduction step has the form*

$$h, c \rightsquigarrow_{\text{st}} h', c', T$$

for a set of forked threads  $T \subset Ccmds$  with  $|T| \leq 1$ .

For simplicity of notation, we omit  $T$  if it is clear from the context that no thread is forked and  $T = \emptyset$ .

Note that we do not provide a reduction rule for **consumeItPerm**, since we only use it as an intermediate representation for the annotated reduction relation presented in Section B.9.

**Definition B.6.4** (Thread Pools). *We define the set of thread pools  $\mathcal{TP}$  as the set of finite partial functions mapping thread IDs to threads:*

$$\mathcal{TP} := \Theta \rightarrow_{\text{fin}} (Ccmds \cup \{\text{term}\}).$$

$$\frac{\text{ST-RED-EVALCtxt} \quad h, c \rightsquigarrow_{\text{st}} h', c', T}{h, E[c] \rightsquigarrow_{\text{st}} h', E[c'], T}$$

$$\text{ST-RED-FORK} \quad h, \mathbf{fork} \ c \rightsquigarrow_{\text{st}} h, \mathbf{tt}, \{c\}$$

(a) Basic Constructs.

$$\text{ST-RED-WHILE} \quad h, \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rightsquigarrow_{\text{st}} h, \mathbf{if} \ c \ \mathbf{then} \ \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip}$$

$$\text{ST-RED-IFTRUE} \quad h, \mathbf{if} \ \mathbf{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, c$$

$$\text{ST-RED-IFFALSE} \quad h, \mathbf{if} \ \mathbf{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, \mathbf{tt}$$

$$\text{ST-RED-LET} \quad h, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{st}} h, c[x \mapsto v]$$

(b) Control Structures.

$$\text{ST-RED-CONS} \quad \frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{cons}(v) \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \ell}$$

$$\text{ST-RED-READHEAPLOC} \quad \frac{\ell \mapsto v \in h}{h, [\ell] \rightsquigarrow_{\text{st}} h, v}$$

$$\text{ST-RED-ASSIGN} \quad h \sqcup \{\ell \mapsto v'\}, [\ell] := v \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \mathbf{tt}$$

(c) Heap Access.

$$\text{ST-RED-NEWMUTEX} \quad \frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{new\_mutex} \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \ell}$$

$$\text{ST-RED-ACQUIRE} \quad h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \mathbf{acquire} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \mathbf{tt}$$

$$\text{ST-RED-RELEASE} \quad h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \mathbf{release} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \mathbf{tt}$$

(d) Mutexes.

Figure B.3: Single thread reduction rules.

$$\begin{array}{c}
 \text{TP-RED-LIFT} \\
 \frac{P(\theta) = c \quad h, c \rightsquigarrow_{\text{st}} h', c', T}{h, P \rightsquigarrow_{\text{tp}}^{\theta} h', P[\theta := c'] +_{\text{tp}} T}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TP-RED-TERM} \\
 \frac{P(\theta) = v}{h, P \rightsquigarrow_{\text{tp}}^{\theta} h, P[\theta := \text{term}]}
 \end{array}$$

Figure B.4: Thread pool reduction rules.

The symbol `term` represents a terminated thread. We denote thread pools by  $P$ , thread IDs by  $\theta$  and the empty thread pool by  $\emptyset_{\text{tp}}$ , i.e.,

$$\begin{aligned}
 \emptyset_{\text{tp}} &: \Theta \rightarrow_{\text{fin}} (\text{Cmds} \cup \{\text{term}\}), \\
 \text{dom}(\emptyset_{\text{tp}}) &= \emptyset.
 \end{aligned}$$

We define the operation  $+_{\text{tp}} : \mathcal{TP} \times \{C \subset \text{Cmds} \mid |C| \leq 1\} \rightarrow \mathcal{TP}$  as follows:

$$\begin{aligned}
 P +_{\text{tp}} \emptyset &:= P, \\
 P +_{\text{tp}} \{c\} &:= P[\theta_{\text{new}} := c] \quad \text{for } \theta_{\text{new}} := \min(\Theta \setminus \text{dom}(P)).
 \end{aligned}$$

**Definition B.6.5** (Thread Pool Reduction Relation). We define a thread pool reduction relation  $\rightsquigarrow_{\text{tp}}$  according to the rules presented in Figure B.4. A reduction step has the form

$$h, P \rightsquigarrow_{\text{tp}}^{\theta} h', P'.$$

**Definition B.6.6** (Reduction Sequence). Let  $(h_i)_{i \in \mathbb{N}}$  and  $(P_i)_{i \in \mathbb{N}}$  be infinite sequences of physical heaps and thread pools, respectively.

We call  $(h_i, P_i)_{i \in \mathbb{N}}$  a reduction sequence if there exists a sequence of thread IDs  $(\theta_i)_{i \in \mathbb{N}}$  such that

$$h_i, P_i \rightsquigarrow_{\text{tp}}^{\theta_i} h_{i+1}, P_{i+1}$$

holds for every  $i \in \mathbb{N}$ .

**Definition B.6.7** (Fairness). We call a reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  fair iff for all  $i \in \mathbb{N}$  and  $\theta \in \text{dom}(P_i)$  with  $P_i(\theta) \neq \text{term}$  there exists some  $k \geq i$  with

$$h_k, P_k \rightsquigarrow_{\text{tp}}^{\theta} h_{k+1}, P_{k+1}.$$

## B.7 Assertions

**Definition B.7.1** (Fractions). We define the set of fractions as

$$\mathcal{F} := \{f \in \mathbb{Q} \mid 0 < f \leq 1\}.$$

**Definition B.7.2** (Thread Phase IDs). *We define the set of thread phase literals as*

$$\mathcal{T} := \{\text{Forker}, \text{Forkee}\}.$$

*We call a finite sequence of thread phase literals a phase ID and denote it by  $\tau \in \mathcal{T}^*$ . We write  $\tau_1 \sqsubseteq \tau_2$  to express that  $\tau_1$  is a (non-strict) prefix of  $\tau_2$ .*

**Definition B.7.3.** *We define the sets of ghost signals  $\mathcal{S}$ , obligations  $\mathcal{O}$ , wait permission  $\Omega$  and iteration permissions  $\Lambda$  as follows:*

$$\begin{aligned} \mathcal{S} &:= \mathcal{ID} \times \mathcal{Levs}, \\ \mathcal{O} &:= (\mathcal{Locs} \cup \mathcal{ID}) \times \mathcal{Levs}, \\ \Omega &:= \mathcal{T}^* \times \mathcal{ID} \times \Delta, \\ \Lambda &:= \mathcal{T}^* \times \Delta. \end{aligned}$$

*We denote ghost signals by  $s$ , obligations by  $o$ , and bags of obligations by  $O$ . For convenience of notation we define the selector function:*

$$(id, L).id := L.$$

**Definition B.7.4** (Assertions). *We define the set of assertions  $\mathcal{A}$  according to the syntax presented in Figure B.5.<sup>1</sup> Further, we define implication and equivalence as the usual abbreviations:*

$$\begin{aligned} a_1 \rightarrow a_2 &:= \neg a_1 \vee a_2, \\ a_1 \leftrightarrow a_2 &:= (a_1 \rightarrow a_2) \wedge (a_2 \rightarrow a_1). \end{aligned}$$

*Let  $(a(i))_{i \in I}$  be a family of assertions indexed by some set  $I$ . We define quantification over  $I$  as the following abbreviations:*

$$\begin{aligned} \exists i \in I. a(i) &:= \bigvee \{a(i) \mid i \in I\}, \\ \forall i \in I. a(i) &:= \neg \exists i \in I. \neg a(i). \end{aligned}$$

*We omit the index set  $I$  when its choice becomes clear from the context and write  $\exists i. a(i)$  and  $\forall i. a(i)$  instead of  $\exists i \in I. a(i)$  and  $\forall i \in I. a(i)$ , respectively.*

**Definition B.7.5** (Logical Resources). *We define the set of logical resources  $\mathcal{R}^{\text{log}}$  syntactically as follows:*

$$\begin{aligned} r^{\text{log}} \in \mathcal{R}^{\text{log}} &::= \ell \mapsto v \mid \widehat{\ell} \mapsto \widehat{v} \mid \text{signal}_{\text{Res}}((id, L), b) \mid \\ &\text{uninit}_{\text{Res}}(\ell) \mid \text{mutex}_{\text{Res}}((\ell, L), a) \mid \text{locked}_{\text{Res}}((\ell, L), a, f) \mid \\ &\text{phase}_{\text{Res}}(\tau) \mid \text{obs}_{\text{Res}}(O) \mid \text{wperm}_{\text{Res}}(\tau, id, \delta) \mid \\ &\text{itperm}_{\text{Res}}(\tau, \delta) \end{aligned}$$

<sup>1</sup>That is, we define  $\mathcal{A}$  as the least fixpoint of  $F$  where  $F(A) = \{\text{True}, \text{False}\} \cup \{\neg a \mid a \in A\} \cup \{a_1 \wedge a_2 \mid a_1, a_2 \in A\} \cup \dots \cup \{\bigvee A' \mid A' \subseteq A\} \cup \dots$ . Since  $F$  is a monotonic function over a complete lattice, it has a least fixpoint according to the Knaster-Tarski theorem [169].



$$\begin{aligned}
 a \in \mathcal{A} \quad := \quad & \text{True} \mid \text{False} \mid \neg a \mid \\
 & a \wedge a \mid a \vee a \mid a * a \mid [f]\ell \mapsto v \mid [f]\widehat{\ell} \mapsto \widehat{v} \mid \\
 & \bigvee A \mid \\
 & [f]\text{uninit}(\ell) \mid \\
 & [f]\text{mutex}((\ell, L), a) \mid [f]\text{locked}((\ell, L), a, f) \mid \\
 & [f]\text{signal}((id, L), b) \mid \\
 & \text{phase}(\tau) \mid \text{obs}(O) \mid \text{wperm}(\tau, id, \delta) \mid \text{itperm}(\tau, \delta) \\
 \\
 f \in \mathcal{F} \quad v \in \text{Values} \quad \widehat{v} \in \text{Values}^G \quad \ell \in \mathcal{Locs} \quad \widehat{\ell} \in \mathcal{Locs}^G \\
 L \in \mathcal{Levs} \quad id \in \mathcal{ID} \quad b \in \mathbb{B} = \{\text{True}, \text{False}\} \quad \delta \in \Delta \\
 A \subseteq \mathcal{A} \quad O \in \text{Bags}(\mathcal{O}) \quad \tau \in \mathcal{T}^*
 \end{aligned}$$

Figure B.5: Assertion syntax.

Further, we define the functions  $\text{getHLocs}_{\text{IRes}} : \mathcal{R}^{\text{log}} \rightarrow \mathcal{Locs}$  and  $\text{getGLocs}_{\text{IRes}} : \mathcal{R}^{\text{log}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs}^G)$  mapping logical resources to their respective (either empty or singleton) set of involved heap locations and ghost locations, respectively, as

$$\begin{aligned}
 \text{getHLocs}_{\text{IRes}}(\ell \mapsto v) & \quad := \quad \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{uninit}_{\text{IRes}}(\ell)) & \quad := \quad \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{mutex}_{\text{IRes}}((\ell, L), a)) & \quad := \quad \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\text{locked}_{\text{IRes}}((\ell, L), a, f)) & \quad := \quad \{\ell\}, \\
 \text{getHLocs}_{\text{IRes}}(\_) & \quad := \quad \emptyset \quad \text{otherwise}, \\
 \\
 \text{getGLocs}_{\text{IRes}}(\widehat{\ell} \mapsto \widehat{v}) & \quad := \quad \{\widehat{\ell}\}, \\
 \text{getGLocs}_{\text{IRes}}(\_) & \quad := \quad \emptyset \quad \text{otherwise}.
 \end{aligned}$$

**Definition B.7.6 (Mutexes).** We define the set of mutexes as  $\mathcal{M} := \mathcal{Locs} \times \mathcal{Levs}$  and denote mutexes by  $m$ . For convenience of notation we define the selector function

$$(\ell, L).\text{loc} \quad := \quad \ell.$$

**Definition B.7.7 (Logical Heaps).** We define the set of logical heaps as

$$\text{Heaps}^{\text{log}} \quad := \quad \mathcal{R}^{\text{log}} \rightarrow \{q \in \mathbb{Q} \mid q \geq 0\}.$$

We define the empty logical heap  $\emptyset_{\text{log}}$  as the constant zero function

$$\emptyset_{\text{log}} : r^{\text{l}} \mapsto 0.$$

We denote logical heaps by  $H$ , point-wise addition by  $+$  and multiplication with non-negative rationals by  $\cdot$ , i.e.,

$$\begin{aligned}
 (H_1 + H_2)(r^{\text{l}}) & \quad := \quad H_1(r^{\text{l}}) + H_2(r^{\text{l}}), \\
 (q \cdot H)(r^{\text{l}}) & \quad := \quad q \cdot (H(r^{\text{l}}))
 \end{aligned}$$

for  $q \in \mathbb{Q}$  with  $q \geq 0$ . For convenience of notation we represent logical heaps containing finitely many resources by sets of resources and define left-associative functions  $+_{\text{lh}}, -_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \mathcal{R}^{\text{log}} \rightarrow \text{Heaps}^{\text{log}}$  as follows

$$\begin{aligned} \{r_1^l, \dots, r_n^l\} &:= \begin{cases} r_i^l & \mapsto 1 \\ x & \mapsto 0 \text{ if } x \notin \{r_1^l, \dots, r_n^l\}, \end{cases} \\ H +_{\text{lh}} r^l &:= H[r^l := H(r^l) + 1], \\ H -_{\text{lh}} r^l &:= H[r^l := \max(0, H(r^l) - 1)]. \end{aligned}$$

We give  $\cdot$  a higher precedence than  $+$ ,  $+_{\text{lh}}$  and  $-_{\text{lh}}$ .

Further, we define the function  $\text{getGLocs}_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \mathcal{P}(\mathcal{Locs}^{\text{G}})$  mapping logical heaps to their respective set of allocated ghost locations as

$$\text{getGLocs}_{\text{lh}}(H) := \bigcup_{\substack{r^l \in \mathcal{R}^{\text{log}} \\ H(r^l) > 0}} \text{getGLocs}_{\text{Res}}(r^l).$$

We call a logical heap  $H$  complete and write  $\text{complete}_{\text{lh}}(H)$  if it contains exactly one obligations chunk and exactly one phase chunk, i.e., if there exist a bag of obligations  $O$  and a phase ID  $\tau$  with  $H(\text{obs}_{\text{Res}}(O)) = 1$  and  $H(\text{phase}_{\text{Res}}(\tau)) = 1$  and if there do not exist any bag of obligations  $O'$  nor any phase ID  $\tau'$  with (i)  $O \neq O'$  and  $H(\text{obs}_{\text{Res}}(O')) > 0$  or with (ii)  $\tau \neq \tau'$  and  $H(\text{phase}_{\text{Res}}(\tau')) > 0$ .

We call a logical heap  $H$  finite and write  $\text{finite}_{\text{lh}}(H)$  if it contains only finitely many resources, i.e., if the set  $\{r^l \in \mathcal{R}^{\text{log}} \mid H(r^l) > 0\}$  is finite.

We call a logical heap  $H$  consistent and write  $\text{consistent}_{\text{lh}}(H)$  if (i) it contains only full phase, obligations, wait and iteration permission chunks, i.e., if

$$\begin{aligned} H(\text{phase}_{\text{Res}}(\tau)) &\in \mathbb{N}, \\ H(\text{obs}_{\text{Res}}(O)) &\in \mathbb{N}, \\ H(\text{wperm}_{\text{Res}}(\tau, id, \delta)) &\in \mathbb{N}, \\ H(\text{itperm}_{\text{Res}}(\tau, \delta)) &\in \mathbb{N} \end{aligned}$$

holds for all  $\tau \in \mathcal{T}^*$ ,  $O \in \text{Bags}(\mathcal{O})$ ,  $id \in \mathcal{ID}$  and  $\delta \in \Delta$  and if (ii) heap locations and ghost locations are unique in  $H$ , i.e., if there are no  $r_1^l, r_2^l \in \mathcal{R}^{\text{log}}$  with  $r_1^l \neq r_2^l$ ,  $H(r_1^l) > 0$ ,  $H(r_2^l) > 0$  and with  $\text{getHLocs}_{\text{Res}}(r_1^l) \cap \text{getHLocs}_{\text{Res}}(r_2^l) \neq \emptyset$  or  $\text{getGLocs}_{\text{Res}}(r_1^l) \cap \text{getGLocs}_{\text{Res}}(r_2^l) \neq \emptyset$ .

To simplify the specification of logical heaps containing only a single obligations chunk with certain properties, we introduce the abbreviation

$$(H.\text{obs} = O) := (\text{complete}_{\text{lh}}(H) \wedge H(\text{obs}_{\text{Res}}(O)) = 1).$$

**Definition B.7.8** (Assertion Model Relation). We define a model relation  $\models_{\mathcal{A}} \subset \text{Heaps}^{\text{log}} \times \mathcal{A}$  for assertions by recursion on the structure of assertions

$H \models_{\mathcal{A}} \text{True}$	
$H \not\models_{\mathcal{A}} \text{False}$	
$H \models_{\mathcal{A}} \neg a$	if $H \not\models_{\mathcal{A}} a$
$H \models_{\mathcal{A}} a_1 \wedge a_2$	if $H \models_{\mathcal{A}} a_1 \wedge H \models_{\mathcal{A}} a_2$
$H \models_{\mathcal{A}} a_1 \vee a_2$	if $H \models_{\mathcal{A}} a_1 \vee H \models_{\mathcal{A}} a_2$
$H \models_{\mathcal{A}} a_1 * a_2$	if $\exists H_1, H_2 \in \text{Heaps}^{\text{log}}.$ $H = H_1 + H_2 \wedge$ $H_1 \models_{\mathcal{A}} a_1 \wedge H_2 \models_{\mathcal{A}} a_2$
$H \models_{\mathcal{A}} [f]\ell \mapsto v$	if $H(\ell \mapsto v) \geq f$
$H \models_{\mathcal{A}} [f]\widehat{\ell} \mapsto \widehat{v}$	if $H(\widehat{\ell} \mapsto \widehat{v}) \geq f$
$H \models_{\mathcal{A}} \bigvee A$	if $\exists a \in A. H \models_{\mathcal{A}} a$
$H \models_{\mathcal{A}} [f]\text{uninit}(\ell)$	if $H(\text{uninit}_{ \text{Res}}(\ell)) \geq f$
$H \models_{\mathcal{A}} [f]\text{mutex}(m, P)$	if $H(\text{mutex}_{ \text{Res}}(m, P)) \geq f$
$H \models_{\mathcal{A}} [f]\text{locked}(m, P, f_u)$	if $H(\text{locked}_{ \text{Res}}(m, P, f_u)) \geq f$
$H \models_{\mathcal{A}} [f]\text{signal}(s, b)$	if $H(\text{signal}_{ \text{Res}}(s, b)) \geq f$
$H \models_{\mathcal{A}} \text{phase}(\tau)$	if $H(\text{phase}_{ \text{Res}}(\tau)) \geq 1$
$H \models_{\mathcal{A}} \text{obs}(O)$	if $H(\text{obs}_{ \text{Res}}(O)) \geq 1$
$H \models_{\mathcal{A}} \text{wperm}(\tau, id, \delta)$	if $H(\text{wperm}_{ \text{Res}}(\tau, id, \delta)) \geq 1$
$H \models_{\mathcal{A}} \text{itperm}(\tau, \delta)$	if $H(\text{itperm}_{ \text{Res}}(\tau, \delta)) \geq 1$

Figure B.6: Assertion model relation.

according to the rules presented in Figure B.6. We write  $H \models_{\mathcal{A}} a$  to express that logical heap  $H$  models assertion  $a$  and  $H \not\models_{\mathcal{A}} a$  to express that  $H \models_{\mathcal{A}} a$  does not hold.

## B.8 Proof Rules

**Definition B.8.1** (Level Ascriptions). We define a function  $\text{lev} : (\mathcal{ID} \cup \mathcal{Locs}) \times \mathcal{Levs} \rightarrow \mathcal{Levs}$  as

$$\text{lev}(\_, L) := L.$$

**Definition B.8.2** (View Shift). We define a view shift relation  $\Rightarrow \subset \mathcal{A} \times \mathcal{A}$  according to the rules presented in Figure B.7.

**Definition B.8.3** (Proof Relation). We define a proof relation  $\vdash \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$  according to the rules presented in Figures B.8 and B.9.

$$\begin{array}{c}
\text{VS-SEMIMP} \\
\frac{\forall H. \text{consistent}_{\text{lh}}(H) \wedge H \vDash_A A \Rightarrow H \vDash_A B}{A \Rightarrow B} \\
\\
\text{VS-TRANS} \\
\frac{A \Rightarrow C \quad C \Rightarrow B}{A \Rightarrow B} \\
\\
\text{VS-OR} \\
\frac{A_1 \Rightarrow B \quad A_2 \Rightarrow B}{A_1 \vee A_2 \Rightarrow B} \\
\\
\text{VS-NEWSIGNAL} \\
\frac{L \in \mathcal{L}evs}{\text{obs}(O) \Rightarrow \exists id. \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})} \\
\\
\text{VS-SET SIGNAL} \\
\text{obs}(O \uplus \{s\}) * \text{signal}(s, \_) \Rightarrow \text{obs}(O) * \text{signal}(s, \text{True}) \\
\\
\text{VS-WAITPERM} \\
\frac{\delta' <_{\Delta} \delta}{\text{itperm}(\tau', \delta) \Rightarrow \text{wperm}(\tau', id, \delta')} \\
\\
\text{VS-WAIT} \\
\frac{\tau_{\text{anc}} \sqsubseteq \tau \quad \forall o \in O. \text{lev}(s) <_{\text{L}} \text{lev}(o)}{\text{phase}(\tau) * \text{obs}(O) * \text{wperm}(\tau_{\text{anc}}, s, id, \delta) * \text{signal}(s, b) \\
\Rightarrow \text{phase}(\tau) * \text{obs}(O) * \text{wperm}(\tau_{\text{anc}}, s, id, \delta) * \text{signal}(s, b) * (\neg b \leftrightarrow \text{itperm}(\tau, \delta))} \\
\\
\text{VS-SPECITPERM} \\
\frac{\tau_{\text{anc}} \sqsubseteq \tau}{\text{itperm}(\tau_{\text{anc}}, \delta) \Rightarrow \text{itperm}(\tau, \delta)} \\
\\
\text{VS-SPECWAITPERM} \\
\frac{\tau_{\text{anc}} \sqsubseteq \tau}{\text{wperm}(\tau_{\text{anc}}, id, \delta) \Rightarrow \text{wperm}(\tau, id, \delta)} \\
\\
\text{VS-WEAKPERM} \\
\frac{\delta' <_{\Delta} \delta \quad N \in \mathbb{N}}{\text{itperm}(\tau', \delta) \Rightarrow \bigstar_{1, \dots, N} \text{itperm}(\tau', \delta')} \\
\\
\text{VS-MUTINIT} \\
\frac{L \in \mathcal{L}evs}{\text{uninit}(\ell) * P \Rightarrow \text{mutex}((\ell, L), P)} \\
\\
\text{VS-NEWGCELL} \\
\text{True} \Rightarrow \exists \hat{\ell}. \hat{\ell} \mapsto \hat{v} \\
\\
\text{VS-SETGCELL} \\
\hat{\ell} \mapsto \hat{v} \Rightarrow \hat{\ell} \mapsto \hat{v}'
\end{array}$$

Figure B.7: View shift rules.

Note that our proof rules do not allow us to reason about the command **consumeItPerm**, since we only use it as an intermediate representation during reduction.

**Lemma B.8.4.** *We can derive the proof rule presented in Figure B.10.*

*Proof.* Trivial. □

## B.9 Annotated Semantics

**Definition B.9.1** (Annotated Resources). *We define the set of annotated resources  $AnnoRes$  as follows:*

$$r^a \in AnnoRes ::= \ell \mapsto v \mid \text{uninit}_{aRes}(\ell) \mid \\ \text{unlocked}_{aRes}((\ell, L), a, H) \mid \text{locked}_{aRes}((\ell, L), a, f) \mid \\ \text{signal}_{aRes}((id, L), b)$$

**Definition B.9.2** (Annotated Heaps). *We define the set of annotated heaps as*

$$Heaps^{\text{annot}} := \mathcal{P}_{\text{fin}}(AnnoRes),$$

*the function  $\text{locs}_{ah} : Heaps^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$  mapping annotated heaps to the sets of allocated heap locations as*

$$\text{locs}_{ah}(h^a) := \{ \ell \in \mathcal{Locs} \mid \exists v \in \text{Values}. \exists L \in \mathcal{Levs}. \exists a \in \mathcal{A}. \\ \exists H \in Heaps^{\text{log}}. \exists f \in \mathcal{F}. \\ \ell \mapsto v \in h^a \vee \text{uninit}_{aRes}(\ell) \in h^a \vee \\ \text{unlocked}_{aRes}((\ell, L), a, H) \in h^a \vee \\ \text{locked}_{aRes}((\ell, L), a, f) \in h^a \}$$

*and the function  $\text{ids}_{ah} : Heaps^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{ID})$  mapping annotated heaps to sets of allocated signal IDs as*

$$\text{ids}_{ah}(h^a) := \{ id \in \mathcal{ID} \mid \exists L \in \mathcal{Levs}. \exists b \in \mathbb{B}. \text{signal}_{aRes}((id, L), b) \in h^a \}.$$

*We denote annotated heaps by  $h^a$ .*

*We call an annotated heap  $h^a$  finite and write  $\text{finite}_{ah}(h^a)$  if there exists no chunk  $\text{unlocked}_{aRes}((\ell, L), a, H) \in h^a$  for which  $\text{finite}_{eh}(H)$  does not hold.*

**Definition B.9.3** (Annotated Single Thread Reduction Relation). *We define a reduction relation  $\rightsquigarrow_{\text{ast}}$  for annotated threads according to the rules presented in Figures B.11 and B.12. A reduction step has the form*

$$h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T^a$$

$$\begin{array}{c}
\text{PR-FRAME} \\
\frac{\vdash \{A\} c \{B\}}{\vdash \{A * F\} c \{B * F\}} \\
\\
\text{PR-VIEWSHIFT} \\
\frac{A \Rightarrow A' \wedge \text{phase}(\tau) \quad \vdash \{A'\} c \{B'\} \quad \forall \tau'. (B' \wedge \text{phase}(\tau') \wedge \tau \sqsubseteq \tau' \Rightarrow B)}{\vdash \{A\} c \{B\}} \\
\\
\text{PR-VS-SIMP} \qquad \text{PR-EXP} \\
\frac{A \Rightarrow A' \quad \vdash \{A'\} c \{B'\} \quad B' \Rightarrow B}{\vdash \{A\} c \{B\}} \qquad \frac{\llbracket e \rrbracket \in \text{Values}}{\vdash \{\text{True}\} e \{\lambda r. r = \llbracket e \rrbracket\}} \\
\\
\text{PR-EXISTS} \\
\frac{\forall a \in A. \vdash \{a\} c \{B\}}{\vdash \{\bigvee A\} c \{B\}} \\
\\
\text{PR-FORK} \\
\frac{\vdash \{\text{phase}(\tau.\text{Forkee}) * \text{obs}(O_f) * A\} c \{\text{obs}(\emptyset)\}}{\vdash \{\text{phase}(\tau) * \text{obs}(O_m \uplus O_f) * A\} \text{fork} c \{\lambda r. \text{phase}(\tau.\text{Forker}) * \text{obs}(O_m) * r = \text{tt}\}} \\
\\
\text{(a) Basic proof rules.} \\
\\
\text{PR-IF} \\
\frac{\vdash \{A\} c_b \{\lambda b. C(b) \wedge (b = \text{True} \vee b = \text{False})\} \quad \vdash \{C(\text{True})\} c_t \{B\} \quad C(\text{False}) \Rightarrow B}{\vdash \{A\} \text{if } c_b \text{ then } c_t \{B\}} \\
\\
\text{PR-WHILE} \\
\frac{\forall \tau_{\text{it}}. \tau \sqsubseteq \tau_{\text{it}} \Rightarrow \vdash \{\text{phase}(\tau_{\text{it}}) * I(\tau_{\text{it}})\} c_b \left\{ \begin{array}{l} \lambda b. \exists \tau'_{\text{it}}, \tau_{\text{anc}}. \tau_{\text{anc}} \sqsubseteq \tau'_{\text{it}} * \text{phase}(\tau'_{\text{it}}) \\ * (b = \text{True} \vee b = \text{False}) \\ * (b \rightarrow \text{itperm}(\tau_{\text{anc}}, \delta) * I(\tau'_{\text{it}})) \\ * (\neg b \rightarrow B(\tau'_{\text{it}})) \end{array} \right.}{\vdash \{\text{phase}(\tau) * I(\tau)\} \text{while } c_b \text{ do skip } \{\exists \tau'. \tau \sqsubseteq \tau' * \text{phase}(\tau') * B(\tau')\}} \\
\\
\text{PR-LET} \\
\frac{\vdash \{A\} c \{\lambda r. C(r)\} \quad \forall v. \vdash \{C(v)\} c'[x \mapsto v] \{B\}}{\vdash \{A\} \text{let } x := c \text{ in } c' \{B\}}
\end{array}$$

(b) Control structures.

Figure B.8: Proof rules (part 1).

$$\begin{array}{c}
 \text{PR-ACQUIRE} \\
 \frac{m.\text{lev} \prec_{\mathbb{L}} O}{\{ \text{obs}(O) * [f]\text{mutex}(m, P) \}} \\
 \vdash \mathbf{acquire} \ m.\text{loc} \\
 \{ \lambda r. r = \text{tt} * \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * P \} \\
 \\
 \text{PR-RELEASE} \\
 \frac{\text{obs}(O) * A \Rightarrow \text{obs}(O) * P * B}{\{ \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * A \}} \\
 \vdash \mathbf{release} \ m.\text{loc} \\
 \{ \lambda r. r = \text{tt} * \text{obs}(O) * [f]\text{mutex}(m, P) * B \} \\
 \\
 \text{PR-NEWMUTEX} \\
 \vdash \{ \text{True} \} \mathbf{new\_mutex} \ \{ \lambda \ell. \text{uninit}(\ell) \}
 \end{array}$$

(a) Mutexes.

$$\begin{array}{c}
 \text{PR-CONS} \qquad \qquad \qquad \text{PR-READHEAPLOC} \\
 \vdash \{ \text{True} \} \mathbf{cons}(v) \ \{ \lambda \ell. \ell \mapsto v \} \quad \vdash \{ [f]\ell \mapsto v \} \ [\ell] \ \{ \lambda r. r = v * [f]\ell \mapsto v \} \\
 \\
 \text{PR-ASSIGNTOHEAP} \\
 \vdash \{ \ell \mapsto \_ \} \ [\ell] := v \ \{ \lambda r. r = \text{tt} * \ell \mapsto v \}
 \end{array}$$

(b) Heap access.

Figure B.9: Proof rules (part 2).

$$\begin{array}{c}
 \text{PR-WHILE-SIMP} \\
 \frac{\tau_{\text{anc}} \sqsubseteq \tau}{\vdash \{ \text{phase}(\tau) * A \} \ c_b \ \{ \lambda b. \text{phase}(\tau) * (b \rightarrow \text{itperm}(\tau_{\text{anc}}, \delta) * A) * (\neg b \rightarrow B) \}} \\
 \vdash \{ \text{phase}(\tau) * A \} \ \mathbf{while} \ c_b \ \mathbf{do} \ \text{skip} \ \{ \text{phase}(\tau) * B \}
 \end{array}$$

Figure B.10: Derived proof rule.

for a set of annotated forked threads  $T^a \subset \text{Heaps}^{\text{log}} \times \text{Cmds}$  with  $|T^a| \leq 1$ .

It indicates that given annotated heap  $h^a$  and a logical heap  $H$ , command  $c$  can be reduced to annotated heap  $h^{a'}$ , logical heap  $H'$  and command  $c'$ . The either empty or singleton set  $T^a$  represents whether a new thread is forked in this step.

For simplicity of notation we omit  $T^a$  if it is clear from the context that no thread is forked and  $T^a = \emptyset$ .

**Definition B.9.4** (Annotated Thread Pools). We define the set of annotated thread pools  $\mathcal{TP}^a$  as the set of finite partial functions mapping thread IDs to annotated threads:

$$\mathcal{TP}^a := \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds} \cup \{\text{term}\}).$$

We denote annotated thread pools by  $P^a$  and the empty thread pool by  $\emptyset_{\text{atp}}$ , i.e.,

$$\begin{aligned} \emptyset_{\text{atp}} &: \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds} \cup \{\text{term}\}), \\ \text{dom}(\emptyset_{\text{atp}}) &= \emptyset. \end{aligned}$$

We define the modification operations  $+_{\text{atp}}$  and  $-_{\text{atp}}$  analogously to  $+_{\text{tp}}$  and  $-_{\text{tp}}$ , respectively, cf. Definition B.6.4.

For convenience of notation we define selector functions for annotated threads as

$$\begin{aligned} (H, c).\text{heap} &:= H, \\ (H, c).\text{cmd} &:= c. \end{aligned}$$

**Definition B.9.5** (Ghost Reduction Relation). We define a thread pool reduction relation  $\rightsquigarrow_{\text{ghost}}$  according to the rules presented in Figures B.13 and B.14 to express ghost steps. A ghost reduction step has the form

$$h^a, P^a \rightsquigarrow_{\text{ghost}}^{\theta} h^{a'}, P^{a'}.$$

We denote its reflexive transitive closure by  $\rightsquigarrow_{\text{ghost}}^*$ .

**Definition B.9.6** (Non-ghost Thread Pool Reduction Relation). We define a thread pool reduction relation  $\rightsquigarrow_{\text{real}}$  according to the rules presented in Figure B.15 to express real reduction steps. A reduction step has the form

$$h^a, P^a \rightsquigarrow_{\text{real}}^{\theta} h^{a'}, P^{a'}.$$

**Definition B.9.7** (Annotated Thread Pool Reduction Relation). We define the annotated thread pool reduction relation  $\rightsquigarrow_{\text{atp}}$  as

$$\rightsquigarrow_{\text{atp}} := \rightsquigarrow_{\text{ghost}} \cup \rightsquigarrow_{\text{real}}.$$



$$\frac{\text{AST-RED-EVALCTXT} \quad h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T}{h^a, H, E[c] \rightsquigarrow_{\text{ast}} h^{a'}, H', E[c'], T}$$

$$\begin{array}{l} \text{AST-RED-FORK} \\ h^a, H_m + \{\text{phase}_{|\text{Res}}(\tau), \text{obs}_{|\text{Res}}(O_m \uplus O_f)\} + H_f, \mathbf{fork} \ c \\ \rightsquigarrow_{\text{ast}} h^a, H_m + \{\text{phase}_{|\text{Res}}(\tau.\mathbf{Forker}), \text{obs}_{|\text{Res}}(O_m)\}, \mathbf{tt}, \\ \quad \{(\{\text{phase}_{|\text{Res}}(\tau.\mathbf{Forker}), \text{obs}_{|\text{Res}}(O_f)\} + H_f), c\} \end{array}$$

(a) Basic constructs.

$$\begin{array}{l} \text{AST-RED-WHILE} \\ h^a, H, \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \\ \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{if} \ c \ \mathbf{then} \ (\mathbf{consumeItPerm}; \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip}) \end{array}$$

$$\begin{array}{ll} \text{AST-RED-IFTRUE} & \text{AST-RED-IFFALSE} \\ h^a, H, \mathbf{if} \ \mathbf{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, c & h^a, H, \mathbf{if} \ \mathbf{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt} \end{array}$$

$$\begin{array}{l} \text{AST-RED-LET} \\ h^a, H, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{ast}} h^a, H, c[x \mapsto v] \end{array}$$

(b) Control structures.

$$\frac{\text{AST-RED-CONSUMEITPERM} \quad H(\text{phase}_{|\text{Res}}(\tau)) \geq 1 \quad \tau_{\text{anc}} \sqsubseteq \tau}{h^a, H + \{\text{itperm}_{|\text{Res}}(\tau_{\text{anc}}, \delta)\}, \mathbf{consumeItPerm} \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt}}$$

(c) Intermediate representation.

$$\begin{array}{ll} \text{AST-RED-CONS} & \text{AST-RED-READHEAPLOC} \\ \frac{\ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \mathbf{cons}(v)} & \frac{\ell \mapsto v \in h^a}{h^a, H, [\ell] \rightsquigarrow_{\text{ast}} h^a, H, v} \\ \rightsquigarrow_{\text{ast}} h^a \cup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, \ell & \end{array}$$

$$\text{AST-RED-ASSIGN} \quad h \sqcup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, [\ell] := v \rightsquigarrow_{\text{ast}} h \sqcup \{\ell \mapsto v'\}, H + \{\ell \mapsto v'\}, \mathbf{tt}$$

(d) Heap access.

Figure B.11: Annotated single thread reduction rules (part 1).

$$\begin{array}{c}
\text{AST-RED-NEWMUTEX} \\
\frac{\ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \text{new\_mutex} \rightsquigarrow_{\text{ast}} h^a \cup \{\text{uninit}_{\text{aRes}}(\ell)\}, H + \{\text{uninit}_{\text{lRes}}(\ell)\}, \ell} \\
\\
\text{AST-RED-ACQUIRE} \\
\frac{f \in \mathcal{F} \quad m.\text{lev} \prec_{\text{L}} O}{h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, a, H_P)\}, \\ H + \{\text{obs}_{\text{lRes}}(O)\} + f \cdot \{\text{mutex}_{\text{lRes}}(m, P)\}, \\ \text{acquire } m.\text{loc} \\ \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, \\ H + \{\text{obs}_{\text{lRes}}(O \uplus \llbracket m \rrbracket)\}, \text{locked}_{\text{lRes}}(m, P, f)\} + H_P, \\ \text{tt}} \\
\\
\text{AST-RED-RELEASE} \\
\frac{H_P \vDash_{\text{A}} P \quad \text{consistent}_{\text{lh}}(H_P) \\ \exists O. H(\text{obs}_{\text{lRes}}(O)) \geq 1 \quad \exists \tau. H(\text{phase}_{\text{lRes}}(\tau)) \geq 1}{h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, \\ H + \{\text{obs}_{\text{lRes}}(O \uplus \llbracket m \rrbracket)\}, \text{locked}_{\text{lRes}}(m, P, f)\} + H_P, \\ \text{release } m.\text{loc} \\ \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\}, \\ H + \{\text{obs}_{\text{lRes}}(O)\} + f \cdot \{\text{mutex}_{\text{lRes}}(m, P)\}, \\ \text{tt}}
\end{array}$$

(a) Mutexes.

Figure B.12: Annotated single thread reduction rules (part 2).

**Definition B.9.8** (Annotated Reduction Sequence). *Let  $(h^a_i)_{i \in \mathbb{N}}$  and  $(P^a_i)_{i \in \mathbb{N}}$  be infinite sequences of annotated heaps and annotated thread pools, respectively. Let  $\text{sig} : \mathbb{N} \rightarrow \mathcal{S}$  be a partial function mapping indices to signals.*

*We call  $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$  an annotated reduction sequence if there exists a sequence of thread IDs  $(\theta_i)_{i \in \mathbb{N}}$  such that the following holds for every  $i \in \mathbb{N}$ :*

- $h^a_i, P^a_i \rightsquigarrow_{\text{atp}}^{\theta_i} h^a_{i+1}, P^a_{i+1}$
- *If this reduction step results from an application of GTP-RED-WAIT to some signal  $s$ , then  $\text{sig}(i) = s$  holds and otherwise  $\text{sig}(i) = \perp$ .*

*In case the signal annotation  $\text{sig}$  is clear from the context or not relevant, we omit it and write  $(h^a_i, P^a_i)_{i \in \mathbb{N}}$  instead of  $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$ .*

$$\begin{array}{c}
 \text{GTP-RED-NEWSIGNAL} \\
 \frac{P^a(\theta) = (H + \{\text{obs}_{\text{IRes}}(O)\}, c) \quad id \notin \text{ids}_{\text{ah}}(h^a) \quad H' = H + \{\text{signal}_{\text{IRes}}((id, L), \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{id, L\}\})\}}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}((id, L), \text{False})\}, P^a[\theta := (H', c)]} \\
 \\
 \text{GTP-RED-SET SIGNAL} \\
 \frac{P^a(\theta) = (H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{s\}\})\}, c) \quad H' = H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O)\}}{h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{False})\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{True})\}, P^a[\theta := (H', c)]} \\
 \\
 \text{GTP-RED-WAIT PERM} \\
 \frac{\delta' <_{\Delta} \delta \quad P^a(\theta) = (H + \{\text{itperm}_{\text{IRes}}(\tau', \delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{wperm}_{\text{IRes}}(\tau', id, \delta')\}, c)]} \\
 \\
 \text{GTP-RED-WAIT} \\
 \frac{H(\text{phase}_{\text{IRes}}(\tau)) \geq 1 \quad H(\text{wperm}_{\text{IRes}}(\tau_{\text{anc}}, s.\text{id}, \delta)) \geq 1 \quad H(\text{obs}_{\text{IRes}}(O)) \geq 1 \quad \tau_{\text{anc}} \sqsubseteq \tau \quad \forall o \in O. \text{lev}(s) <_{\perp} \text{lev}(o) \quad \text{signal}_{\text{aRes}}(s, \text{False}) \in h^a \quad P^a(\theta) = (H, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{itperm}_{\text{IRes}}(\tau, \delta)\}, c)]} \\
 \\
 \text{GTP-RED-SPECIT PERM} \\
 \frac{\tau_{\text{anc}} \sqsubseteq \tau \quad P^a(\theta) = (H + \{\text{itperm}(\tau_{\text{anc}}, \delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{itperm}(\tau, \delta)\}, c)]} \\
 \\
 \text{GTP-RED-SPECWAIT PERM} \\
 \frac{\tau_{\text{anc}} \sqsubseteq \tau \quad P^a(\theta) = (H + \{\text{wperm}(\tau_{\text{anc}}, id, \delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{wperm}(\tau, id, \delta)\}, c)]} \\
 \\
 \text{GTP-RED-WEAKIT PERM} \\
 \frac{\delta' <_{\Delta} \delta \quad N \in \mathbb{N} \quad P^a(\theta) = (H + \{\text{itperm}_{\text{IRes}}(\tau', \delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + N \cdot \{\text{itperm}_{\text{IRes}}(\tau', \delta')\}, c)]} \\
 \\
 \text{GTP-RED-MUTINIT} \\
 \frac{P^a(\theta) = (H + \{\text{uninit}_{\text{IRes}}(\ell)\} + H_P, c) \quad H_P \vDash_A P \quad \text{consistent}_{\text{Ih}}(H_P) \quad \exists O. H(\text{obs}_{\text{IRes}}(O)) \geq 1 \quad \exists \tau. H(\text{phase}_{\text{IRes}}(\tau)) \geq 1 \quad H' = H + \{\text{mutex}_{\text{IRes}}((\ell, L), H_P)\}}{h^a \sqcup \{\text{uninit}_{\text{aRes}}(\ell)\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}((\ell, L), a, H_P)\}, P^a[\theta := (H', c)]}
 \end{array}$$

Figure B.13: Ghost thread pool reduction rules (part 1).

$$\begin{array}{c}
\text{GTP-RED-NEWGHOSTCELL} \\
\frac{\widehat{\ell} \notin \text{getGLocs}_{\text{lh}}(H) \quad P^a(\theta) = (H, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\widehat{\ell} \mapsto \widehat{v}\}, c)]} \\
\\
\text{GTP-RED-MUTATEGHOSTCELL} \\
\frac{\widehat{\ell} \notin \text{getGLocs}_{\text{lh}}(H) \quad P^a(\theta) = (H + \{\widehat{\ell} \mapsto \widehat{v}\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\widehat{\ell} \mapsto \widehat{v}'\}, c)]}
\end{array}$$

Figure B.14: Ghost thread pool reduction rules (part 2)

$$\begin{array}{c}
\text{RTP-RED-LIFT} \\
\frac{\theta_f = \min(\Theta \setminus \text{dom}(P^a)) \quad P^a(\theta) = (H, c) \quad h^a, H, c \xrightarrow{\ast}_{\text{ast}} h^{a'}, H', c', T^a}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^{a'}, P^a[\theta := (H', c')] +_{\text{atp}} T^a} \\
\\
\text{RTP-RED-TERM} \\
\frac{P^a(\theta) = (H, v) \quad H.\text{obs} = \emptyset}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^a, P^a -_{\text{atp}} \theta}
\end{array}$$

Figure B.15: Non-ghost thread pool reduction rules.

We call  $(h_i^a, P_i^a)$  an annotated machine configuration.

**Lemma B.9.9** (Preservation of Finiteness). *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{finite}_{\text{ah}}(h_0^a)$  and  $\text{finite}_{\text{lh}}(P_0^a(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P_0^a)$ .*

*Then,  $\text{finite}_{\text{lh}}(P_i^a(\theta).\text{heap})$  holds for all  $i \in \mathbb{N}$  and all  $\theta \in \text{dom}(P_i^a)$ .*

*Proof.* Proof by induction on  $i$ . □

**Lemma B.9.10** (Preservation of Completeness). *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{complete}_{\text{lh}}(P_0^a(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P_0^a)$ . Furthermore, let there be no chunk  $\text{unlocked}_{\text{aRes}}(m, P, H_P) \in h_0^a$  such that  $H_P(\text{phase}_{\text{lRes}}(\tau)) > 0$  or  $H_P(\text{obs}_{\text{lRes}}(O)) > 0$  holds for any  $\tau, O$ .*

*Then,  $\text{complete}_{\text{lh}}(P_i^a(\theta).\text{heap})$  holds for every  $i \in \mathbb{N}$  and every  $\theta \in \text{dom}(P_i^a)$ .*

*Proof.* Proof by induction on  $i$ . □

**Definition B.9.11** (Fairness of Annotated Reduction Sequences). *We call an annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  fair iff for all  $i \in \mathbb{N}$  and  $\theta \in \text{dom}(P_i^a)$  with  $P_i^a(\theta).\text{cmd} \neq \text{term}$  there exists some  $k \geq i$  with*

$$h_k^a, P_k^a \xrightarrow{\theta}_{\text{real}} h_{k+1}^a, P_{k+1}^a.$$

Every thread of an annotated thread pool is annotated by a thread-local logical heap that expresses which resources are owned by this thread. In the following we define a function to extract the logical heap expressing which resources are owned by threads of a thread pool (i.e. the sum of all thread-local logical heaps).

**Definition B.9.12.** *We define the function  $\text{ownedResHeap}_{\text{atp}} : \mathcal{TP}^a \rightarrow \text{Heaps}^{\text{log}}$  mapping annotated thread pools to logical heaps as follows:*

$$P^a \mapsto \sum_{\theta \in \text{dom}(P^a)} P^a(\theta).\text{heap}$$

Annotated resources representing unlocked locks, i.e.,  $\text{unlocked}_{\text{aRes}}(m, a, H_a)$ , contain a logical heap  $H_a$  that expresses which resources are protected by this lock. In the following, we define a function that extracts a logical heap from an annotated heap  $h^a$  expressing which resources are protected by unlocked locks in  $h^a$ .

**Definition B.9.13.** *We define the function  $\text{protectedResHeap}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \text{Heaps}^{\text{log}}$  mapping annotated heaps to logical heaps as follows:*

For any annotated heap  $h^a$  let

$$\text{LockInvs}(h^a) := \{ \{ H_P \in \text{Heaps}^{\text{log}} \mid \exists m \in \mathcal{Locs} \times \mathcal{Levs}. \exists P \in \mathcal{A}. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \} \}$$

be the auxiliary set aggregating all logical heaps corresponding to lock invariants of unlocked locks stored in  $h^a$ . We define  $\text{protectedResHeap}_{\text{ah}}$  as

$$h^a \mapsto \sum_{H_P \in \text{LockInvs}(h^a)} H_P.$$

**Definition B.9.14** (Compatibility of Annotated and Logical Heaps). *We inductively define a relation  ${}_{\text{ah}} \sim_{\text{lh}} \subset \text{Heaps}^{\text{annot}} \times \text{Heaps}^{\text{log}}$  between annotated*

and logical heaps such that the following holds

$$\begin{array}{ll}
\emptyset & \text{ah} \sim_{\text{lh}} \emptyset_{\text{log}}, \\
h^a \cup \{\ell \mapsto v\} & \text{ah} \sim_{\text{lh}} H + \{\ell \mapsto v\}, \\
h^a \cup \{\text{uninit}_{\text{aRes}}(\ell)\} & \text{ah} \sim_{\text{lh}} H + \{\text{uninit}_{\text{lRes}}(\ell)\}, \\
h^a \cup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\} & \text{ah} \sim_{\text{lh}} H + \{\text{mutex}_{\text{lRes}}(m, P)\} + H_P, \\
h^a \cup \{\text{locked}_{\text{aRes}}(m, P, f)\} & \text{ah} \sim_{\text{lh}} H + \{\text{locked}_{\text{lRes}}(m, P, f)\} \\
& \quad + (1 - f) \cdot \{\text{mutex}_{\text{lRes}}(m, P)\}, \\
h^a \cup \{\text{signal}_{\text{aRes}}(s, b)\} & \text{ah} \sim_{\text{lh}} H + \{\text{signal}_{\text{lRes}}(s, b)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{phase}_{\text{lRes}}(\tau)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{obs}_{\text{lRes}}(O)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{wperm}_{\text{lRes}}(\tau, \text{id}, \delta)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{itperm}_{\text{lRes}}(\tau, \delta)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\widehat{\ell} \mapsto \widehat{v}\},
\end{array}$$

where  $h^a \in \text{Heaps}^{\text{annot}}$  and  $H \in \text{Heaps}^{\text{log}}$  are annotated and logical heaps with  $\ell, m.\text{loc} \notin \text{locs}_{\text{ah}}(h^a)$ ,  $s.\text{id} \notin \text{ids}_{\text{ah}}(h^a)$  and  $h^a \sim_{\text{ah} \sim_{\text{lh}}} H$ .

We consider a machine configuration  $(h^a, P^a)$  to be *consistent* if it fulfils three criteria: (i) Every thread-local logical heap is consistent, i.e., for all used thread IDs  $\theta$ ,  $P^a(\theta).\text{heap}$  only stores full phase, obligations, wait permission and iteration permission chunks. (ii) Every logical heap protected by an unlocked lock in  $h^a$  is consistent. (iii)  $h^a$  is compatible with the logical heap that represents (a) the resources owned by threads in  $P^a$  and (b) the resources protected by unlocked locks stored in  $h^a$ .

**Definition B.9.15** (Consistency of Annotated Machine Configurations). *We call an annotated machine configuration  $(h^a, P^a)$  consistent and write  $\text{consistent}_{\text{conf}}(h^a, P^a)$  if all of the following hold:*

- $\text{consistent}_{\text{lh}}(P^a(\theta).\text{heap})$  for all  $\theta \in \text{dom}(P^a)$ ,
- $\forall m. \forall P. \forall H_P. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \rightarrow \text{consistent}_{\text{lh}}(H_P)$ ,
- $h^a \sim_{\text{ah} \sim_{\text{lh}}} \text{ownedResHeap}_{\text{atp}}(P^a) + \text{protectedResHeap}_{\text{ah}}(h^a)$ .

**Lemma B.9.16** (Preservation of Consistency). *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $\text{consistent}_{\text{conf}}(h_0^a, P_0^a)$ . Then,  $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$  holds for every  $i \in \mathbb{N}$ .*

*Proof.* Proof by induction on  $i$ . □

## B.10 Hoare Triple Model Relation

**Definition B.10.1** (Command Annotation). *We define the predicate  $\text{annot}_{\text{cmd}} \subseteq \text{Cmds} \times \text{Cmds}$  such that  $\text{annot}_{\text{cmd}}(c', c)$  holds iff  $c'$  results from  $c$  by removing all occurrences of **consumeItPerm**.*

**Definition B.10.2** (Thread Pool Annotation). *We define a predicate  $\text{annot}_{\text{tp}} \subseteq \mathcal{TP}^a \times \mathcal{TP}$  such that:*

$$\begin{aligned} & \text{annot}_{\text{tp}}(P^a, P) \\ & \iff \\ \text{dom}(P^a) = & \text{dom}(P) \wedge \forall \theta \in \text{dom}(P). \text{annot}_{\text{cmd}}(P^a(\theta).\text{cmd}, P(\theta)) \end{aligned}$$

**Definition B.10.3** (Compatibility of Annotated and Physical Heaps). *We inductively define a relation  $\text{ah} \sim_{\text{ph}} \subseteq \text{Heaps}^{\text{annot}} \times \mathcal{R}^{\text{phys}}$  between annotated and physical heaps such that the following holds:*

$$\begin{array}{ll} \emptyset & \text{ah} \sim_{\text{ph}} \emptyset, \\ \ell \mapsto v \cup h^a & \text{ah} \sim_{\text{ph}} \ell \mapsto v \cup h, \\ \text{uninit}_{\text{aRes}}(\ell) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{unlocked}_{\text{aRes}}((\ell, L), P, H_P) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{locked}_{\text{aRes}}((\ell, L), P, f) \cup h^a & \text{ah} \sim_{\text{ph}} \text{locked}_{\text{pRes}}(\ell) \cup h, \\ \text{signal}_{\text{aRes}}(s, b) \cup h^a & \text{ah} \sim_{\text{ph}} h \end{array}$$

where  $h^a \in \text{Heaps}^{\text{annot}}$  and  $h \in \text{Heaps}^{\text{phys}}$  are annotated and physical heaps with  $h^a \text{ ah} \sim_{\text{ph}} h$ .

**Definition B.10.4** (Safety). *We define the safety predicate  $\text{safe} \subseteq \text{Heaps}^{\text{log}} \times \text{Cmds}$  coinductively as the greatest solution (with respect to  $\subseteq$ ) of the following equation:*

$$\begin{aligned} & \text{safe}(H, c) \\ & \iff \\ \text{complete}_{\text{lh}}(H) \rightarrow & \\ \forall P, P'. \forall \theta \in \text{dom}(P). \forall h, h'. \forall P^a. \forall h^a. & \\ \text{consistent}_{\text{conf}}(h^a, P^a) \wedge h^a \text{ ah} \sim_{\text{ph}} h \wedge & \\ P(\theta) = c \wedge P^a(\theta) = (H, c) \wedge \text{annot}_{\text{tp}}(P^a, P) \wedge h, P \overset{\theta}{\rightsquigarrow}_{\text{tp}} h', P' \rightarrow & \\ \exists P^G, P^{a'}. \exists h^G, h^{a'}. & \\ h^a, P^a \overset{\theta}{\rightsquigarrow}_{\text{ghost}}^* h^G, P^G \wedge h^G, P^G \overset{\theta}{\rightsquigarrow}_{\text{real}} h^{a'}, P^{a'} \wedge \text{annot}_{\text{tp}}(P^{a'}, P') \wedge & \\ h^{a'} \text{ ah} \sim_{\text{ph}} h' \wedge & \\ \forall (H_f, c_f) \in \text{range}(P^{a'}) \setminus \text{range}(P^a). \text{safe}(H_f, c_f). & \end{aligned}$$

**Definition B.10.5** (Hoare Triple Model Relation). *We define the model relation for Hoare triples  $\vDash_{\mathbf{H}} \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$  such that:*

$$\begin{aligned} \vDash_{\mathbf{H}} \{A\} c \{\lambda r. B(r)\} \\ \iff \\ \forall H_F. \forall E. (\forall v. \forall H_B. H_B \vDash_{\mathbf{A}} B(v) \rightarrow \text{safe}(H_B + H_F, E[v])) \\ \rightarrow \forall H_A. H_A \vDash_{\mathbf{A}} A \rightarrow \text{safe}(H_A + H_F, E[c]) \end{aligned}$$

We can instantiate context  $E$  in above definition to **let**  $x := \square$  **in** **tt**, which yields the consequent  $\text{safe}(H_A + H_F, \text{let } x := c \text{ in tt})$ . Note that this implies  $\text{safe}(H_A + H_F, c)$ .

**Lemma B.10.6** (Hoare Triple Soundness). *Let  $\vdash \{A\} c \{B\}$  hold, then also  $\vDash_{\mathbf{H}} \{A\} c \{B\}$  holds.*

*Proof.* Proof by induction on the derivation of  $\vdash \{A\} c \{B\}$ . □

**Theorem B.10.7** (Soundness). *Let*

$$\vdash \{\text{phase}(\tau) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$$

*hold. There exists no fair, infinite reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  with  $h_0 = \emptyset$  and  $P_0 = \{(\theta_0, c)\}$  for any choice of  $\theta_0$ .*

## B.11 Soundness

In this section, we prove the soundness theorem B.10.7.

**Lemma B.11.1** (Construction of Annotated Reduction Sequences). *Suppose we can prove  $\vDash_{\mathbf{H}} \{A\} c \{\text{obs}(\emptyset)\}$ . Let  $H_A$  be a logical heap with  $H_A \vDash_{\mathbf{A}} A$  and  $\text{complete}_{\text{lh}}(H_A)$  and  $h_0^a$  an annotated heap with  $h_0^a \text{ah} \sim_{\text{lh}} H_A$ . Let  $(h_i, P_i)_{i \in \mathbb{N}}$  be a fair plain reduction sequence with  $h_0^a \text{ah} \sim_{\text{ph}} h_0$  and  $P_0 = \{(\theta_0, c)\}$  for some thread ID  $\theta_0$  and command  $c$ .*

*Then, there exists a fair annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  with  $P^a = \{(\theta_0, (H_A, c))\}$  and  $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$  for all  $i \in \mathbb{N}$ .*

*Proof.* We can construct the annotated reduction sequence inductively from the plain reduction sequence. □



**Definition B.11.2** (Program Order Graph). *Let  $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$  be an annotated reduction sequence. Let  $N^r$  be the set of names referring to reduction rules defining the relations  $\rightsquigarrow_{\text{real}}$ ,  $\rightsquigarrow_{\text{ghost}}$  and  $\rightsquigarrow_{\text{ast}}$ . We define the set of annotated reduction rule names  $N^a$  where GTP-RED-WAIT is annotated by signals as*

$$N^a := (N^r \setminus \{\text{GTP-RED-WAIT}\}) \cup (\{\text{GTP-RED-WAIT}\} \times \mathcal{S}).$$

We define the program order graph  $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})) = (\mathbb{N}, E)$  with root 0 where  $E \subset \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$ .

A node  $a \in \mathbb{N}$  corresponds to the sequence's  $a^{\text{th}}$  reduction step, i.e.,  $h_a^a, P_a^a \rightsquigarrow_{\text{atp}}^\theta h_{a+1}^a, P_{a+1}^a$  for some  $\theta \in \text{dom}(P_a^a)$ . An edge from node  $a$  to node  $b$  expresses that the  $b^{\text{th}}$  reduction step continues the control flow of step  $a$ . For any  $\ell \in \mathbb{N}$ , let  $\theta_\ell$  denote the ID of the thread reduced in step  $\ell$ . Furthermore, let  $n_\ell^a$  denote the name of the reduction rule applied in the  $\ell^{\text{th}}$  step, in the following sense:

- If  $h_\ell^a, P_\ell^a \rightsquigarrow_{\text{atp}}^\theta h_{\ell+1}^a, P_{\ell+1}^a$  results from an application of RTP-RED-LIFT in combination with single-thread reduction rule  $n^{\text{st}}$ , then  $n_\ell^a = n^{\text{st}}$ .
- If  $h_\ell^a, P_\ell^a \rightsquigarrow_{\text{atp}}^\theta h_{\ell+1}^a, P_{\ell+1}^a$  results from an application of GTP-RED-WAIT, then  $n_\ell^a = (\text{GTP-RED-WAIT}, \text{sig}(\ell))$ .
- Otherwise,  $n^a$  denotes the applied (real or ghost) thread pool reduction rule.

An edge  $(a, \theta, n^a, b) \in \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$  is contained in  $E$  if  $n^a = n_a^a$  and one of the following conditions applies:

- $\theta = \theta_a = \theta_b$  and  $b = \min(\{k > a \mid h_k^a, P_k^a \rightsquigarrow_{\text{atp}}^{\theta_a} h_{k+1}^a, P_{k+1}^a\})$ .  
In this case, the edge expresses that step  $b$  marks the first time that thread  $\theta_a$  is rescheduled for reduction (after step  $a$ ).
- $\text{dom}(P_{a+1}^a) \setminus \text{dom}(P_a^a) = \{\theta\}$  and  
 $b = \min\{k \in \mathbb{N} \mid h_k^a, P_k^a \rightsquigarrow_{\text{atp}}^\theta h_{k+1}^a, P_{k+1}^a\}$ .  
In this case,  $\theta$  identifies the thread forked in step  $a$ . The edge expresses that step  $b$  marks the first reduction of the forked thread.

In case the choice of reduction sequence  $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$  is clear from the context, we write  $G$  instead of  $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig}))$ .

**Observation B.11.3.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence with  $|\text{dom}(P_0^a)| = 1$ . The sequence's program order graph  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  is a binary tree.*

For any reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ , the paths in its program order graph  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  represent the sequence's control flow paths. Hence, we are going to use program order graphs to analyse reduction sequences' control flows.

We refer to a program order graph's edges by the kind of reduction step they represent. For instance, we call edges of the form  $(a, \theta, \text{ST-RED-WHILE}, b)$  *loop edges* because they represent a loop backjump and edges of the form  $(a, \theta, (\text{GTP-RED-WAIT}, s), b)$  *wait edges*. Wait edges of this form represent applications of GTP-RED-WAIT to signal  $s$ .

In the following, we prove that any path in a program order graph that does not involve a loop edge is finite. This follows from the fact that the size of the command reduced along this path decreases with each non-ghost non-loop step.

**Lemma B.11.4.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence. Let  $p = (V, E)$  be a path in  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ . Let  $L = \{e \in E \mid \pi_3(e) = \text{AST-RED-WHILE}\}$  be the set of loop edges contained in  $p$ . Then,  $p$  is infinite if and only if  $L$  is infinite.*

*Proof.* If  $L$  is infinite,  $p$  is obviously infinite as well. So, suppose  $L$  is finite.

For any command, we consider its size to be the number of nodes contained in its abstract syntax tree. By structural induction over the set of commands, it follows that the size of a command  $c = P^a(\theta).\text{cmd}$  decreases in every non-ghost reduction step  $h^a, P^a \xrightarrow{\theta}_{\text{atp}} h^{a'}, P^{a'}$  that is not an application of RTP-RED-LIFT in combination with AST-RED-WHILE.

Since  $L$  is finite, there exists a node  $x$  such that the suffix  $p_{\geq x}$  starting at node  $x$  does not contain any loop edges. By fairness of  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ , every non-empty suffix of  $p_{\geq x}$  contains an edge corresponding to a non-ghost reduction step. For any edge  $e = (i, \theta, n, j)$  consider the command  $c_e = P_i^a(\theta).\text{cmd}$  reduced in this edge. The size of these commands decreases along  $p_{\geq x}$ . So,  $p_{\geq x}$  must be finite and thus  $p$  must be finite as well.  $\square$

**Corollary B.11.5.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence. Let  $p = (V, E)$  be a path in  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ . Let*

$$C = \{e \in E \mid \pi_3(e) = \text{AST-RED-CONSUMEITPERM}\}$$

*be the set of consume edges contained in  $p$ . Then,  $p$  is infinite if and only if  $C$  is infinite.*

*Proof.* Follows from Lemma B.11.4 by the fact that the set  $\{e \in E \mid \pi_3(e) = \text{AST-RED-WHILE}\}$  is infinite if and only if  $C$  is infinite.  $\square$

**Definition B.11.6.** Let  $G = (V, E)$  be a subgraph of some program order graph. We define the function  $\text{waitEdges}_G : \mathcal{S} \rightarrow \mathcal{P}(E)$  mapping any signal  $s$  to the set of wait edges in  $G$  concerning  $s$  as:

$$\text{waitEdges}_G(s) := \{(a, \theta, (\text{GTP-RED-WAIT}, s'), b) \in E \mid s' = s\}.$$

Furthermore, we define the set  $\mathcal{S}_G \subseteq \mathcal{S}$  of signals being waited for in  $G$  and its subset  $\mathcal{S}_G^\infty \subseteq \mathcal{S}_G$  of signals waited-for infinitely often in  $G$  as follows:

$$\begin{aligned} \mathcal{S}_G &:= \{s \in \mathcal{S} \mid \text{waitEdges}_G(s) \neq \emptyset\}, \\ \mathcal{S}_G^\infty &:= \{s^\infty \in \mathcal{S}_G \mid \text{waitEdges}_G(s^\infty) \text{ infinite}\}. \end{aligned}$$

**Definition B.11.7.** Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence and let  $G = (V, E)$  be a subgraph of the sequence's program order graph. We define the function  $\text{itperms}_G : E \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$  mapping any edge  $e$  to the (potentially empty) finite bag of iteration permissions derived in the reduction step corresponding to  $e$  as follows:

Let  $(i, \theta, n, j) \in E$  be an edge.

- If  $n = (\text{GTP-RED-WAIT}, s)$  for some signal  $s \in \mathcal{S}$ , then the  $i^{\text{th}}$  reduction step spawns a single iteration permission  $(\tau, \delta)$ , i.e.,  
 $P_{i+1}^a = P_i^a[\theta := (P_i^a(\theta).\text{heap} + \{\text{itperm}_{\text{Res}}(\tau, \delta)\}, P_i^a(\theta).\text{cmd})]$ .  
 In this case, we define

$$\text{itperms}_G((i, \theta, (\text{GTP-RED-WAIT}, s), j)) := \{(\tau, \delta)\}.$$

- If  $n = \text{GTP-RED-WEAKITPERM}$ , then the  $i^{\text{th}}$  reduction step consumes an iteration permission  $(\tau', \delta)$  and produces  $N$  permissions  $(\tau', \delta')$  of lower degree, i.e.,  $P_i^a(\theta).\text{heap} = H + \{\text{itperm}(\tau', \delta)\}$  for some heap  $H$  and  $P_{i+1}^a = P_i^a[\theta := (H', P_i^a(\theta).\text{cmd})]$  for

$$H' = H + N \cdot \{\text{itperm}_{\text{Res}}(\tau', \delta')\}.$$

In this case, we define

$$\text{itperms}_G((i, \theta, \text{GTP-RED-WEAKITPERM}, j)) := \underbrace{\{(\tau', \delta'), \dots, (\tau', \delta')\}}_{N \text{ times}}.$$

- Otherwise, we define

$$\text{itperms}_G((i, \theta, n, j)) := \emptyset.$$

**Definition B.11.8** (Signal Capacity). Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence and  $G = (V, E)$  be a subgraph of the sequence's program

order graph. We define the function  $\text{sigCap}_G : (\mathcal{S} \setminus S_G^\infty) \times \mathbb{N} \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$  mapping signals and indices to bags of iteration permissions as follows:

$$\text{sigCap}_G(s, i) := \biguplus_{\substack{(a, \theta, n, b) \in \text{waitEdges}_G(s) \\ a \geq i}} \text{itperms}_G((a, \theta, n, b)).$$

We call  $\text{sigCap}_G(s, i)$  the capacity of signal  $s$  at index  $i$ .

Note that the signal capacity above is indeed finite. For every  $G$  and every signal  $s \in \mathcal{S} \setminus S_G^\infty$  the set of wait edges  $\text{waitEdges}_G(s)$  is finite. Hence, the big union above is a finite union over finite iteration permission bags.

**Definition B.11.9** (Partial Order on Permissions). *We define the partial order on iteration permissions  $<_\Lambda \subset \Lambda \times \Lambda$  induced by  $<_\Delta$  such that*

$$(\tau_1, \delta_1) <_\Lambda (\tau_2, \delta_2) \iff \delta_1 <_\Delta \delta_2.$$

**Lemma B.11.10.** *The partial order  $<_\Lambda$  is well-founded.*

*Proof.* Follows directly from well-foundedness of  $<_\Lambda$ .  $\square$

**Definition B.11.11** (Partial Order on Finite Bags). *Let  $X$  be a set and let  $<_X \subset X \times X$  a partial order on  $X$ . We define the partial order  $\prec_X \subset \text{Bags}_{\text{fin}}(X) \times \text{Bags}_{\text{fin}}(X)$  on finite bags over  $X$  as the Dershowitz-Manna ordering [53] induced by  $<_X$ :*

$$\begin{aligned} A \prec_X B \iff \exists C, D \in \text{Bags}_{\text{fin}}(X). \quad & \emptyset \neq C \subseteq B \\ & \wedge A = (B \setminus C) \uplus D \\ & \wedge \forall d \in D. \exists c \in C. d <_X c. \end{aligned}$$

We define  $\preceq_X \subset \text{Bags}_{\text{fin}}(X) \times \text{Bags}_{\text{fin}}(X)$  such that

$$A \preceq_X B \iff A = B \vee A \prec_X B$$

holds.

**Corollary B.11.12.** *The partial order  $\prec_\Lambda \subset \text{Bags}_{\text{fin}}(\Lambda) \times \text{Bags}_{\text{fin}}(\Lambda)$  is well-founded.*

*Proof.* Follows from [53] and Lemma B.11.10.  $\square$

In the following, we view paths in a program order graph as single-branched subgraphs. This allows us to apply above definitions on graphs to paths. In

particular, this allows us to refer to the capacity of a signal  $s$  on a path  $p$  by referring to  $\text{sigCap}_p$ .

For the following definition, remember that a bag  $B \in \text{Bags}(X)$  is a function  $B : X \rightarrow \mathbb{N}$  while a logical heap  $H \in \text{Heaps}^{\text{log}}$  is a function  $H : \mathcal{R}^{\text{log}} \rightarrow \mathbb{Q}_{\geq 0}$ . Also remember the signatures  $\text{ownedResHeap}_{\text{atp}} : \mathcal{TP}^a \rightarrow \text{Heaps}^{\text{log}}$  and  $\text{protectedResHeap}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \text{Heaps}^{\text{log}}$ .

**Definition B.11.13.** *We define the functions  $\text{itperms}_{\text{conf}} : \text{Heaps}^{\text{annot}} \times \mathcal{TP}^a \rightarrow \text{Bags}(\Lambda)$  and  $\text{wperms}_{\text{conf}} : \text{Heaps}^{\text{annot}} \times \mathcal{TP}^a \rightarrow \text{Bags}(\Omega)$  mapping annotated machine configurations to bags of iteration and wait permissions, respectively, as follows:*

$$\begin{aligned} \text{itperms}_{\text{conf}}(h^a, P^a)(\tau, \delta) \\ := \lfloor (\text{ownedResHeap}_{\text{atp}}(P^a) + \text{protectedResHeap}_{\text{ah}}(h^a))(\text{itperm}_{\text{lRes}}(\tau, \delta)) \rfloor, \end{aligned}$$

$$\begin{aligned} \text{wperms}_{\text{conf}}(h^a, P^a)(\tau, \text{id}, \delta) \\ := \lfloor (\text{ownedResHeap}_{\text{atp}}(P^a) + \text{protectedResHeap}_{\text{ah}}(h^a))(\text{wperm}_{\text{lRes}}(\tau, \text{id}, \delta)) \rfloor. \end{aligned}$$

Note that for consistent annotated machine configurations  $(h^a, P^a)$  the above flooring is without any affect.

**Corollary B.11.14.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be an annotated reduction sequence such that  $\text{finite}_{\text{ah}}(h_0^a)$  and  $\text{finite}_{\text{lh}}(P_0^a(\theta).\text{heap})$  hold for every  $\theta \in \text{dom}(P_0^a)$ .*

*Then,  $\text{itperms}_{\text{conf}}(h_i^a, P_i^a)$  and  $\text{wperms}_{\text{conf}}(h_i^a, P_i^a)$  are finite for every choice of  $i \in \mathbb{N}$ .*

*Proof.* Follows by preservation of finiteness, Lemma B.9.9.  $\square$

**Lemma B.11.15.** *Let  $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$  be a program order graph and let  $p = (V, E)$  be a path in  $G$  with  $S_p^\infty = \emptyset$ . For every  $\theta \in \text{dom}(P_0^a)$  let  $P_0^a(\theta).\text{heap}$  be finite and complete. Further, let  $h_0^a$  be finite and contain no chunks  $\text{unlocked}_{\text{aRes}}(m, P, H_P)$  where  $H_P$  contains any phase or obligations chunk.*

*Then,  $p$  is finite.*

*Proof.* Assume  $p$  is infinite. We prove a contradiction by assigning a finite capacity to every node along the path. Let  $\theta_i$  be the ID of the thread reduced in step  $i$ . For every  $\theta \in \text{dom}(P_r^a)$  the logical heap  $P_0^a(\theta).\text{heap}$  is complete and  $h_0^a$  contains no chunks  $\text{unlocked}_{\text{aRes}}(m, P, H_P)$  where  $H_P$  contains any phase or obligations chunk. By preservation of completeness, Lemma B.9.10,  $P_i^a(\theta_i).\text{heap}$

is also complete and hence it contains exactly one phase chunk  $\text{phase}_{\text{Res}}(\tau_i)$ . That is, for every step  $i$ , the phase ID  $\tau_i$  of the thread reduced in step  $i$  is uniquely defined.

Consider the function  $\text{nodeCap} : V \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$  defined as

$$\begin{aligned} \text{nodeCap}(i) := & \{(\tau_{\text{anc}}, \delta) \in \text{itperms}_{\text{conf}}(h_i^a, P_i^a) \mid \tau_{\text{anc}} \sqsubseteq \tau_i\} \\ & \uplus \bigcup_{\substack{id \in \text{waitIDs}(\tau_i) \\ (\tau_{\text{anc}}, id, \delta) \in \text{wperms}_{\text{conf}}(h_i^a, P_i^a) \\ L \in \text{Levs}}} \text{sigCap}_p((id, L), i). \end{aligned}$$

where  $\text{waitIDs}(\tau_i) := \{id \mid \exists \tau_{\text{anc}}. (\tau_{\text{anc}}, id, \_) \in \text{wperms}_{\text{conf}}(h_i^a, P_i^a) \wedge \tau_{\text{anc}} \sqsubseteq \tau_i\}$ .

For every  $i \in V$ , the capacity of node  $i$ , i.e.,  $\text{nodeCap}(i)$ , is the union of two finite iteration permission bags: (i) Above  $\{(\tau_{\text{anc}}, \delta) \in \text{itperms}_{\text{conf}}(h_i^a, P_i^a) \mid \tau_{\text{anc}} \sqsubseteq \tau_i\}$  captures all iteration permissions contained in  $h^a$  and  $P_i^a$  that are qualified by an ancestor  $\tau_{\text{anc}}$  of phase ID  $\tau_i$  and are hence usable by the thread reduced in node  $i$ . This includes the permissions  $(\tau_{\text{anc}}, \delta)$  held by thread  $\theta_i$  as well as such (temporarily) transferred to another thread via a lock invariant. (ii) Below  $\bigcup \text{sigCap}_p((id, L), i)$  captures all iteration permissions that will be created along the suffix of  $p$  that starts at node  $i$  by waiting for signals for which thread  $\theta_i$  already holds a wait permission  $(\tau_{\text{anc}}, id, \delta)$  in step  $i$ .

Note that for every  $i \in V$ , the bag of iteration permissions returned by  $\text{nodeCap}(i)$  is indeed finite. The initial annotated heap and all initial thread-local logical heaps are finite. This allows us to apply Corollary B.11.14, by which we get that  $\text{itperms}_{\text{conf}}(h_i^a, P_i^a)$  and  $\text{wperms}_{\text{conf}}(h_i^a, P_i^a)$  are finite.

Since signal IDs are unique, for every fixed choice of  $i$  and  $id$ , there is at most one level  $L$ , for which  $\text{sigCap}_p((id, L), i) \neq \emptyset$ . By assumption, along  $p$  all signals are waited for only finitely often, i.e.,  $S_p^\infty = \emptyset$ . Hence, also the big union  $\bigcup \text{sigCap}_p((id, L), i)$  is defined and finite.

Consider the sequence  $(\text{nodeCap}(i))_{i \in V}$ . Since every element is a finite bag of permissions, we can order it by  $\prec_\Lambda$ . We are going to prove a contradiction by proving that the sequence is an infinitely descending chain.

Consider any edge  $(i, \theta, n, j) \in E$ . There are only three cases in which  $\text{nodeCap}(i) \neq \text{nodeCap}(j)$  holds.

- $n = \text{GTP-RED-WAITPERM}$ :  
In this case, there are degrees  $\delta, \delta'$  with  $\delta' <_\Delta \delta$ , a signal  $s$  and  $N \in \mathbb{N}$  for

which we get

$$\text{nodeCap}(j) = (\text{nodeCap}(i) \setminus \{\{(\tau', \delta)\}\}) \uplus \underbrace{\{\{(\tau', \delta')\}\}}_{N \text{ times}}.$$

That is,  $\text{nodeCap}(j) \prec_{\Lambda} \text{nodeCap}(i)$ .

- $n = \text{GTP-RED-WEAKITPERM}$ : Same as above.
- $n = \text{AST-RED-CONSUMEITPERM}$ :  
In this case, we know that  $\text{nodeCap}(j) = \text{nodeCap}(i) \setminus \{\{(\tau_{\text{anc}}, \delta)\}\} \prec_{\Lambda} \text{nodeCap}(i)$  holds for some  $\tau_{\text{anc}}$  and  $\delta$ .

(Note that in case of  $n = \text{GTP-RED-WAIT}$ , we have  $\text{nodeCap}(i) = \text{nodeCap}(j)$  since

$$\begin{aligned} \{\{(\tau_{\text{anc}}, \delta) \in \text{itperms}_{\text{conf}}(h_j^a, P_j^a) \mid \tau_{\text{anc}} \sqsubseteq \tau_j\}\} \\ = \\ \{\{(\tau_{\text{anc}}, \delta) \in \text{itperms}_{\text{conf}}(h_i^a, P_i^a) \mid \tau_{\text{anc}} \sqsubseteq \tau_i\}\} \uplus \{\{(\tau, \delta)\}\} \end{aligned}$$

and

$$\bigoplus \text{sigCap}_p((id, L), j) = \left( \bigoplus \text{sigCap}_p((id, L), i) \right) \setminus \{\{(\tau, \delta)\}\}$$

for some  $\delta$ .) So,  $\text{nodeCap}$  is monotonically decreasing.

By assumption  $p$  is infinite. According to Corollary B.11.5 this implies that the path contains infinitely many consume edges, i.e., edges with a labelling  $n = \text{AST-RED-CONSUMEITPERM}$ . Hence, the sequence  $(\text{nodeCap}(i))_{i \in V}$  forms an infinitely descending chain. However, according to Corollary B.11.12,  $\prec_{\Lambda}$  is well-founded. A contradiction.  $\square$

**Lemma B.11.16.** *Let  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  be a fair annotated reduction sequence with  $\text{consistent}_{\text{conf}}(h_0^a, P_0^a)$ ,  $P_0^a = \{(\theta_0, (H_0, c))\}$ ,  $\text{complete}_{\text{lh}}(H_0)$ ,  $\text{finite}_{\text{lh}}(H_0)$  and with  $\text{finite}_{\text{ah}}(h_0^a)$ . Let  $H_0$  contain no signal or wait permission chunks. Further, let  $h_0^a$  contain no chunks  $\text{unlocked}_{\text{aRes}}(m, P, H_P)$  where  $H_P$  contains any obligations, phase or signal chunks. Let  $G$  be the program order graph of  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ . Then,  $S_G^{\infty} = \emptyset$ .*

*Proof.* Suppose  $S_G^{\infty} \neq \emptyset$ . Since  $\mathcal{L}evs$  is well-founded, the same holds for the set  $\{\text{lev}(s) \mid s \in S^{\infty}\}$ . Hence, there is some  $s_{\min} \in S^{\infty}$  for which no  $z \in S^{\infty}$  with  $\text{lev}(z) \prec_{\mathbb{L}} \text{lev}(s_{\min})$  exists.

Since neither the initial logical heap  $H_0$  nor any unlocked lock invariant stored in  $h_0^a$  does contain any signals,  $s_{\min}$  must be created during the

reduction sequence. The reduction step creating signal  $s_{\min}$  is an application of GTP-RED-NEWSIGNAL, which simultaneously creates an obligation to set  $s_{\min}$ . By preservation of completeness, Lemma B.9.10, every thread-local logical heap  $P_i^a(\theta).\text{heap}$  annotating some thread  $\theta$  in some step  $i$  is complete. According to reduction rule GTP-RED-WAIT, every wait edge  $(a, \theta, (\text{GTP-RED-WAIT}, s_{\min}), b)$  implies together with completeness that in step  $a$  (i) thread  $\theta$  does not hold any obligation for  $s_{\min}$  (i.e.  $P_a^a(\theta).\text{heap. obs} = O$  for some bag of obligations  $O$  with  $s_{\min} \notin O$ ) and (ii)  $s_{\min}$  has not been set, yet (i.e.  $\text{signal}_{\text{aRes}}(s_{\min}, \text{False}) \in h_a^a$ ). Hence, in step  $a$  another thread  $\theta_{\text{ob}} \neq \theta$  must hold the obligation for  $s_{\min}$  (i.e.  $P_a^a(\theta_{\text{ob}}).\text{heap. obs} = O$  for some bag of obligations  $O$  with  $s_{\min} \in O$ ). Since there are infinitely many wait edges concerning  $s_{\min}$  in  $G$ , the signal is never set.

By fairness, for every wait edge as above, there must be a non-ghost reduction step  $h_k^a, P_k^a \xrightarrow{\theta_{\text{ob}}}_{\text{atp}} h_{k+1}^a, P_{k+1}^a$  of the thread  $\theta_{\text{ob}}$  holding the obligation for  $s_{\min}$  with  $k \geq a$ . Hence, there exists an infinite path  $p_{\text{ob}}$  in  $G$  where each edge  $(e, \theta_{\text{ob}}, n, f) \in \text{edges}(p_{\text{ob}})$  concerns some thread  $\theta_{\text{ob}}$  holding the obligation for  $s_{\min}$ . (Note that this thread ID does not have to be constant along the path, since the obligation can be passed on during fork steps.)

The path  $p_{\text{ob}}$  does not contain wait edges  $(e, \theta_{\text{ob}}, (\text{GTP-RED-WAIT}, s^\infty), f)$  for any  $s^\infty \in S^\infty$ , since reduction rule GTP-RED-WAIT would (together with completeness of  $P_e^a(\theta_{\text{ob}}).\text{heap}$ ) require  $s^\infty$  to be of a lower level than all held obligations. This restriction implies  $\text{lev}(s^\infty) <_{\text{L}} \text{lev}(s_{\min})$  and would hence contradict the minimality of  $s_{\min}$ . That is,  $S_{p_{\text{ob}}}^\infty = \emptyset$ .

By preservation of finiteness, Lemma B.9.9, we get that every logical heap associated with the root of  $p_{\text{ob}}$  is finite. This allows us to apply Lemma B.11.15, by which we get that  $p_{\text{ob}}$  is finite. A contradiction.  $\square$

**Lemma B.11.17.** *Let*

$$\models_{\text{H}} \{ \text{phase}(\tau_0) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau_0, \delta_i) \} c \{ \text{obs}(\emptyset) \}$$

*hold. There exists no fair, infinite annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  with  $h_0^a = \emptyset$ ,  $P_0^a = \{(\theta_0, (H_0, c))\}$  and*

$$H_0 = \{ \text{phase}_{\text{IRes}}(\tau_0), \text{obs}_{\text{IRes}}(\emptyset), \text{itperm}_{\text{IRes}}(\tau_0, \delta_1), \dots, \text{itperm}_{\text{IRes}}(\tau_0, \delta_N) \}.$$

*Proof.* Suppose a reduction sequence as described above exists. We are going to prove a contradiction by considering its infinite program order graph  $G$ .

According to Observation B.11.3,  $G$  is a binary tree with an infinite set of vertices. By the Weak König's Lemma [163]  $G$  has an infinite branch, i.e. an infinite path  $p$  starting at root 0.



The initial logical heap  $H_0$  is complete and finite and the initial annotated machine configuration  $(h_0^a, P_0^a)$  is consistent. By Lemma B.11.16 we know that  $S_G^\infty = \emptyset$ . Since  $S_p^\infty \subseteq S_G^\infty$ , we get  $S_p^\infty = \emptyset$ . This allows us to apply Lemma B.11.15, by which we get that  $p$  is finite, which is a contradiction.  $\square$

**Theorem B.10.7** (Soundness). *Let*

$$\vdash \{\text{phase}(\tau) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$$

*hold. There exists no fair, infinite reduction sequence  $(h_i, P_i)_{i \in \mathbb{N}}$  with  $h_0 = \emptyset$  and  $P_0 = \{(\theta_0, c)\}$  for any choice of  $\theta_0$ .*

*Proof.* Suppose a reduction sequence as described above exists. Since we can prove  $\vdash \{\text{phase}(\tau) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$ , we can also conclude  $\models_{\text{H}} \{\text{phase}(\tau) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$  by Hoare triple soundness, Lemma B.10.6. Consider the logical heap

$$H_0 = \{\text{phase}_{\text{Res}}(\tau), \text{obs}_{\text{Res}}(\emptyset), \text{itperm}_{\text{Res}}(\tau, \delta_1), \dots, \text{itperm}_{\text{Res}}(\tau, \delta_N)\}$$

and the annotated heap  $h_0^a = \emptyset$ . It holds  $H_0 \models_{\text{A}} \text{phase}(\tau) * \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)$ ,  $h_0^a \text{ah} \sim_{\text{lh}} H_0$  and  $h_0^a \text{ah} \sim_{\text{ph}} h_0$ . This allows us to apply Lemma B.11.1, by which we can construct a corresponding fair annotated reduction sequence  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  that starts with  $h_0^a = \emptyset$  and  $P_0^a = \{(\theta_0, (H_0, c))\}$ . By Lemma B.11.17  $(h_i^a, P_i^a)_{i \in \mathbb{N}}$  does not exist. A contradiction.  $\square$

## B.12 Verification Example

We scaled all proof outlines presented in this section down to fit them into the page layout. You can find the original proof outlines in the official version of the technical report [151].

As in the paper body, we present the proof state in **blue**, applied proof and view shift rules in **purple**, and abbreviations and definition in **red**. Since the verification outlines we present in this section span multiple figures, we include hints and explanations concerning other figures (e.g. a hint pointing to the figure where an invariant was defined). To set these remarks and in general any hints off from the essential parts of the proof, we present them in a *grey and italic font*. We occasionally remind the reader of earlier proof steps performed in a previous figure by repeating them in *the same grey font* at the beginning of the current figure.

Further, we highlight how our proof steps effect the proof state as follows: Consider a proof state of the form  $\{A * \ell \mapsto v\}$ . (i) When a proof step adds a new chunk  $C$ , we highlight it in **green**, i.e.  $\{A * \ell \mapsto v * C\}$ . (ii) When a proof step removes the chunk  $A$ , we highlight this change by underlying the removed part of the assertion with a **dark grey** background, i.e.,  $\{\overline{A} * \ell \mapsto v\}$ . Note that in this case, the greyed out  $\overline{A} *$  is not a part of the proof state anymore. (iii) When a proof step changes only part of a chunk, we highlight this change in **yellow**. For instance, if the step changes the value of heap location  $\ell$  from  $v$  to  $v'$ , we highlight it in the resulting state as  $\{A * \ell \mapsto v'\}$ .

### B.12.1 Minimal Example

Figures B.16 and B.17 sketch the verification of the example program presented in Figure B.2. For this verification we let the set of values *Values* include natural numbers and choose  $\mathcal{L}evs = \Delta = \mathbb{N}$ .

### B.12.2 Bounded FIFO

For this section, we let the set of values *Values* include natural numbers and finite sequences, aka lists, of natural numbers. Further, the set of operations *Ops* includes the canonical operations on natural numbers and lists, i.e., (i)  $<$ ,  $\leq$ ,  $-$  and (ii) list concatenation  $l_1 \cdot l_2$ , prepending an element  $e :: l$ , getting the head and tail of a list **head**( $l$ ) (defined for non-empty  $l$ ), **tail**( $l$ ) and getting the size of a list **size**( $l$ ). We denote the empty list by **nil**. We use the abbreviation  $a R_1 b R_2 c$  for  $R_1, R_2 \in \{<, \leq\}$  to denote  $a R_1 b * b R_2 c$ . Furthermore, we choose  $\mathcal{L}evs = \Delta = \mathbb{N}$ . Figure B.19 presents an example program involving a bounded FIFO.

To simplify its verification, we refine the process of creating a new ghost signal, i.e., we split it in two steps: allocating a new signal ID and initializing a signal. To implement this, we replace view shift rule VS-NEWSIGNAL by the rules VS-ALLOCSIGID and VS-SIGINIT presented in Figure B.18. This way we can fix the IDs of all the signals we need throughout the proof at its beginning. This refinement does not affect the soundness of our verification approach. Figures B.20 – B.30 sketch the program’s verification using fine-grained signals.



<pre> ... ∀ℓ<sub>X</sub>, id<sub>S</sub>. ...  {phase(τ.Forkee) * obs(∅) * itperm((), 1) * [½]mutex((ℓ<sub>m</sub>, 0), P)}  {phase(τ.Forkee) * obs(∅) * itperm((), 1) * wperm((), id<sub>S</sub>, 0) * [½]mutex((ℓ<sub>m</sub>, 0), P)} while {phase(τ.Forkee) * obs(∅) * wperm((), id<sub>S</sub>, 0) * [½]mutex((ℓ<sub>m</sub>, 0), P)} acquire m; {phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0)}) * wperm((), id<sub>S</sub>, 0) * locked((ℓ<sub>m</sub>, 0), P, ½) * P } let y := {   phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0)}) * wperm((), id<sub>S</sub>, 0) * locked((ℓ<sub>m</sub>, 0), P, ½)   * ∃v<sub>X</sub>. ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) }   ∀v<sub>X</sub>.   {ℓ<sub>X</sub> ↦ v<sub>X</sub>}   [x]   {λr. r = v<sub>X</sub> * ℓ<sub>X</sub> ↦ v<sub>X</sub>}   {λr. ∃v<sub>X</sub>. r = v<sub>X</sub> * ℓ<sub>X</sub> ↦ v<sub>X</sub>}   {     λ r. phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0)}) * wperm((), id<sub>S</sub>, 0) * locked((ℓ<sub>m</sub>, 0), P, ½)     * ∃v<sub>X</sub>. ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) * r = v<sub>X</sub>   } in ∀v<sub>y</sub>. {   phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0)}) * wperm((), id<sub>S</sub>, 0) * locked((ℓ<sub>m</sub>, 0), P, ½)   * ∃v<sub>X</sub>. ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) * v<sub>X</sub> = v<sub>y</sub> } release m; {   phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0) ∅}) * wperm((), id<sub>S</sub>, 0)   * locked((ℓ<sub>m</sub>, 0), P, ½) * ∃v<sub>X</sub>. ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) * v<sub>X</sub> = v<sub>y</sub> } ∀v<sub>X</sub>. {   phase(τ.Forkee) * obs({∅}) * wperm((), id<sub>S</sub>, 0)   * ∃v<sub>X</sub>. ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) * v<sub>X</sub> = v<sub>y</sub> } {   phase(τ.Forkee) * obs({∅}) * wperm((), id<sub>S</sub>, 0)   * ℓ<sub>X</sub> ↦ v<sub>X</sub> * signal((id<sub>S</sub>, 1), v<sub>X</sub> ≠ 0) * v<sub>X</sub> = v<sub>y</sub> * (v<sub>X</sub> = 0 ↔ itperm(τ.Forkee, 0)) } {   phase(τ.Forkee) * obs({∅}) * wperm((), id<sub>S</sub>, 0) * P   * (v<sub>y</sub> = 0 → itperm(τ.Forkee, 0)) } {   phase(τ.Forkee) * obs({(ℓ<sub>m</sub>, 0) ∅}) * wperm((), id<sub>S</sub>, 0) * [½]mutex((ℓ<sub>m</sub>, 0), P)   * P * (v<sub>y</sub> = 0 → itperm(τ.Forkee, 0)) } y = 0 {   λb. phase(τ.Forkee) * obs(∅) * wperm((), id<sub>S</sub>, 0) * [½]mutex((ℓ<sub>m</sub>, 0), P)   * (v<sub>y</sub> = 0 → itperm(τ.Forkee, 0)) * b = [v<sub>y</sub> = 0] } {   λb. phase(τ.Forkee) * (¬b → obs(∅))   * (b → obs(∅) * wperm((), id<sub>S</sub>, 0) * [½]mutex((ℓ<sub>m</sub>, 0), P) * itperm(τ.Forkee, 0)) } do skip {obs(∅)} </pre>	<p><i>Continuation of Figure B.16.</i>  <math>\ell_X, id_S</math> universally quantified below  <math>P :=</math>  <math>\exists v_X. \ell_X \mapsto v_X * \text{signal}((id_S, 1), v_X \neq 0)</math>      PR-VS-SIMP      &amp; VS-WAITPERM      &amp; PR-FRAME</p> <p>PR-WHILE-SIMP</p> <p>PR-ACQUIRE</p> <p>PR-LET      PR-EXISTS      &amp; PR-FRAME      P  <math>v_X</math> quantified in local scope.</p> <p>PR-VS-SIMP      &amp; VS-SEMIMP</p> <p>P</p> <p><math>v_y</math> represents value bound to <math>y</math>.</p> <p>PR-RELEASE</p> <p><i>Release view shift</i>      PR-EXISTS</p> <p><math>v_X</math> quantified in local scope.      PR-VS-SIMP      &amp; VS-WAIT      &amp; PR-FRAME</p> <p>PR-VS-SIMP      &amp; VS-SEMIMP</p> <p>PR-EXP      &amp; PR-FRAME</p> <p>PR-VS-SIMP      &amp; VS-SEMIMP</p>
---	---

Figure B.17: Verification sketch of busy-waiting thread of example program presented in Figure B.2. For readability we omit information about a command's return value if it is not relevant to the proof.

$$\begin{array}{ll}
\text{VS-ALLOCSIGID} & \text{VS-SIGINIT} \\
\text{True} \Rightarrow \exists id. \text{uninitSig}(id) & \text{obs}(O) * \text{uninitSig}(id) \\
& \Rightarrow \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})
\end{array}$$

Figure B.18: Fine-grained view shift rules for signal creation.

```

let fifo10 := cons(nil) in
let m := new_mutex in
let cp := cons(100) in
let cc := cons(100) in
fork (
  while (
    acquire m;
    let f := [fifo10] in
    if size(f) < 10 then (
      let c := [cp] in
      [fifo10] := f · (c :: nil);
      [cp] := c - 1
    );
    release m;
    let c := [cp] in
    c ≠ 0
  ) do skip;
);
while (
  acquire m;
  let f := [fifo10] in
  if size(f) > 0 then (
    let c := [cc] in
    [fifo10] := tail(f);
    [cc] := c - 1
  );
  release m;
  let c := [cc] in
  c ≠ 0
) do skip

```

Figure B.19: Example program with two threads communicating via a shared bounded FIFO with maximal size 10. Producer thread writes numbers 100, ..., 1 to shared FIFO and busy-waits until FIFO is not full and next element can be pushed. Consumer thread pops 100 numbers from FIFO and busy-waits for next number to arrive.

$\{ \text{phase}() * \text{obs}(\emptyset) * \text{itperm}(), 2 \}$   
**let**  $\text{fifo}_{10} := \text{cons}(\text{nil})$  **in** **let**  $m := \text{new\_mutex}$  **in**  
 $\forall \ell_{\text{fifo}_{10}}, \ell_m.$

$\{ \text{phase}() * \text{obs}(\emptyset) * \text{itperm}(), 2 * \ell_{\text{fifo}_{10}} \mapsto \text{nil} * \text{uninit}(\ell_m) \}$   
 $\forall id_{\text{pop}}^1, \dots, id_{\text{pop}}^{100}, id_{\text{push}}^1, \dots, id_{\text{push}}^{100}.$

$\{ *_{i=1, \dots, 100} \text{uninitSig}(id_{\text{pop}}^i) * *_{i=1, \dots, 100} \text{uninitSig}(id_{\text{push}}^i) * \dots \}$   
 $L_{\text{pop}}^i := 102 - i, \quad L_{\text{push}}^i := 101 - i \quad \text{for } 1 \leq i \leq 100$   
*(Later  $L_{\text{pop}}^{i+10} < L_{\text{push}}^i$  and  $L_{\text{push}}^i < L_{\text{pop}}^i$  must hold, cf. Figures B.27 and B.28.)*

$s_{\text{push}}^i := (id_{\text{push}}^i, L_{\text{push}}^i), \quad s_{\text{pop}}^i := (id_{\text{pop}}^i, L_{\text{pop}}^i) \quad \text{for } 1 \leq i \leq 100$

$\left\{ \begin{array}{l} \text{uninitSig}(id_{\text{pop}}^{100}) * \text{uninitSig}(id_{\text{push}}^{100}) \quad \text{signal}(s_{\text{pop}}^{100}, \text{False}) * \text{signal}(s_{\text{push}}^{100}, \text{False}) \\ * \text{obs}(\{ s_{\text{pop}}^{100}, s_{\text{push}}^{100} \}) * \dots \end{array} \right\}$

$\{ \text{itperm}(), 2 * *_{1, \dots, 400} \text{itperm}(), 1 * \dots \}$

$\left\{ \begin{array}{l} *_{1, \dots, 200} \text{itperm}(), 1 \\ * *_{1, \dots, 200} \text{itperm}(), 1 * *_{i=1, \dots, 100} (\text{wperm}(), id_{\text{pop}}^i, 0) * \text{wperm}(), id_{\text{push}}^i, 0) \\ * \dots \end{array} \right\}$

*Later each thread uses  $*_{1, \dots, 100} \text{itperm}(), 1$  to justify productive iterations.*

**let**  $c_p := \text{cons}(100)$  **in** **let**  $c_c := \text{cons}(100)$  **in**  
 $\forall \ell_{c_p}, \ell_{c_c}.$

$\{ \ell_{c_p} \mapsto 100 * \ell_{c_c} \mapsto 100 * \dots \}$

$\left\{ \begin{array}{l} \text{phase}() * \text{obs}(\{ s_{\text{push}}^{100}, s_{\text{pop}}^{100} \}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * [\frac{1}{2}] \ell_{c_c} \mapsto 100 * \text{uninit}(\ell_m) \\ * *_{i=1, \dots, 100} (\text{wperm}(), id_{\text{pop}}^i, 0) * \text{wperm}(), id_{\text{push}}^i, 0) \\ * *_{1, \dots, 200} \text{itperm}(), 1 * P_m \\ * *_{i=1, \dots, 99} (\text{uninitSig}(id_{\text{push}}^i) * \text{uninitSig}(id_{\text{pop}}^i)) \end{array} \right\}$

$\text{mut} := (\ell_m, 0)$  *(Later  $\text{lev}(\text{mut}) < L_{\text{push}}^i$  and  $\text{lev}(\text{mut}) < L_{\text{pop}}^i$  must hold for all  $1 \leq i \leq 100$ , cf. Figures B.25 and B.28.)*

$\{ \text{uninit}(\ell_m) * P_m * \text{mutex}(\text{mut}, P_m) * \dots \}$

$\left\{ \begin{array}{l} \text{phase}() * \text{obs}(\{ s_{\text{push}}^{100}, s_{\text{pop}}^{100} \}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * [\frac{1}{2}] \ell_{c_c} \mapsto 100 \\ * *_{1, \dots, 200} \text{itperm}(), 1 \\ * *_{i=1, \dots, 100} (\text{wperm}(), id_{\text{pop}}^i, 0) * \text{wperm}(), id_{\text{push}}^i, 0) \\ * *_{i=1, \dots, 99} (\text{uninitSig}(id_{\text{push}}^i) * \text{uninitSig}(id_{\text{pop}}^i)) \\ * [\frac{1}{2}] \text{mutex}(\text{mut}, P_m) * [\frac{1}{2}] \text{mutex}(\text{mut}, P_m) \end{array} \right\}$

$\dots$

PR-LET (2x)  
 & PR-CONS  
 & PR-NEWMUTEX  
  
 PR-VS-SIMP  
 & VS-ALLOCSIGID  
 & PR-EXISTS (200x)  
  
 PR-VS-SIMP  
 & VS-SIGINT  
  
 PR-VS-SIMP  
 & VS-WEAKPERM  
  
 PR-VS-SIMP  
 & VS-WAIT  
  
 PR-LET  
 & PR-CONS (2x)  
  
 PR-VS-SIMP  
 & VS-SEMIIMP  
  
 For definition of  
 lock invariant  $P_m$   
 cf. Figure B.22.  
 PR-VS-SIMP  
 & VS-MUTINT  
  
 PR-VS-SIMP  
 & VS-SEMIIMP  
  
 Continued in  
 Figure B.21.

Figure B.20: Verification example bounded FIFO, initialisation. To lighten the notation, we do not show applications of the frame rule.

... B.20 ...

$$\left\{ \begin{array}{l} \text{phase}(\text{()}) * \text{obs}(\{s_{\text{push}}^{100}, s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * [\frac{1}{2}] \ell_{c_c} \mapsto 100 * \text{itperm}(\text{()}, 1)_{1, \dots, 200} \\ * \text{wperm}(\text{()}, id_{\text{pop}}^i, 0) * \text{wperm}(\text{()}, id_{\text{push}}^i, 0)_{i=1, \dots, 100} \\ * \text{uninitSig}(id_{\text{push}}^i) * \text{uninitSig}(id_{\text{pop}}^i) * [\frac{1}{2}] \text{mutex}(mut, P_m) * [\frac{1}{2}] \text{mutex}(mut, P_m)_{i=1, \dots, 99} \end{array} \right\}$$

fork (

$$\left\{ \begin{array}{l} \text{phase}(\text{(Forker)}) * \text{obs}(\{s_{\text{push}}^{100}\}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * \text{itperm}(\text{()}, 1)_{i=1, \dots, 100} \\ * \text{wperm}(\text{()}, id_{\text{pop}}^i, 0) * \text{uninitSig}(s_{\text{push}}^i) * [\frac{1}{2}] \text{mutex}(mut, P_m)_{i=1, \dots, 99} \\ \text{phase}(\text{(Forker)}) * \text{obs}(\{s_{\text{push}}^{100}\}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * \text{itperm}(\text{()}, 1)_{i=1, \dots, 100} \\ * \text{wperm}(\text{()}, id_{\text{pop}}^i, 0) * \text{uninitSig}(s_{\text{push}}^i) * [\frac{1}{2}] \text{mutex}(mut, P_m)_{i=1, \dots, 99} \\ \exists v_{c_p}, O_p. L_p(v_{c_p}, O_p) * v_{c_p} \neq 0 \end{array} \right\}$$

... B.25 ...

$$\{ \text{phase}(\text{(Forker)}) * \exists v_{c_p}, O_p. L_p(v_{c_p}, O_p) * v_{c_p} \neq 0 \text{ obs}(\emptyset) \}$$

$$\{ \text{phase}(\text{(Forker)}) * \text{obs}(\emptyset) \}$$

);

$$\left\{ \begin{array}{l} \text{phase}(\text{()}) * \text{obs}(\{s_{\text{push}}^{100}, s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \ell_{c_p} \mapsto 100 * [\frac{1}{2}] \ell_{c_c} \mapsto 100 \\ * \text{itperm}(\text{()}, 1)_{1, \dots, 100} * \text{itperm}(\text{()}, 1)_{1, \dots, 100} \\ * \text{wperm}(\text{()}, id_{\text{push}}^i, 0) * \text{wperm}(\text{()}, id_{\text{pop}}^i, 0)_{i=1, \dots, 100} \\ * \text{uninitSig}(id_{\text{pop}}^i) * \text{uninitSig}(id_{\text{push}}^i)_{i=1, \dots, 99} \\ * [\frac{1}{2}] \text{mutex}(mut, P_m) * [\frac{1}{2}] \text{mutex}(mut, P_m) \\ \text{phase}(\text{()}) \\ * \text{obs}(\{s_{\text{pop}}^{100}\}) * [\frac{1}{2}] \ell_{c_c} \mapsto 100 * \text{itperm}(\text{()}, 1)_{1, \dots, 100} \\ * \text{wperm}(\text{()}, id_{\text{push}}^i, 0)_{i=1, \dots, 100} \\ * \text{uninitSig}(id_{\text{pop}}^i) * [\frac{1}{2}] \text{mutex}(mut, P_m)_{i=1, \dots, 99} \\ \exists v_{c_c}, O_c. L_c(v_{c_c}, O_c) * v_{c_c} \neq 0 \end{array} \right\}$$

... B.28 ...

$$\{ \text{phase}(\text{(Forker)}) * \exists v_{c_c}, O_c. L_c(v_{c_c}, O_c) * v_{c_c} \neq 0 \text{ obs}(\emptyset) \}$$

$$\{ \text{phase}(\text{(Forker)}) * \text{obs}(\emptyset) \}$$

Continuation of Figure B.20.

PR-FORK

Resources transferred to producer thread. PR-VS-SIMP & VS-SEMIMP

For definition of producer loop invariant  $L_p(n, O)$  cf. Figure B.23.

Producer loop on Figure B.25. PR-VS-SIMP & VS-SEMIMP

Resources remaining with consumer thread. PR-VS-SIMP & VS-SEMIMP

For definition of consumer loop invariant  $L_c(n, O)$  cf. Figure B.24.

Consumer loop on Figure B.28. PR-VS-SIMP & VS-SEMIMP

Figure B.21: Verification example bounded FIFO, forking.





$$\begin{aligned}
 L_c(n, O_c) &:= \\
 & [\frac{1}{2}] \ell_{c_c} \mapsto n * 0 \leq n \leq 100 * [\frac{1}{2}] \text{mutex}(mut, P_m) \\
 & * \text{obs}(O_c) * (n > 0 \leftrightarrow O_c = \{s_{pop}^n\}) * (n = 0 \leftrightarrow O_c = \emptyset) \\
 & * \bigstar_{i=1, \dots, n} \text{itperm}(\cdot, 1) \\
 & * \bigstar_{i=1, \dots, 100} \text{wperm}(\cdot, id_{push}^i, 0) \\
 & * \bigstar_{i=1, \dots, n-1} \text{unitSig}(id_{pop}^i) \\
 \\
 L_c^{\text{locked}}(n, O_p) &:= \\
 & [\frac{1}{2}] \ell_{c_c} \mapsto n * 0 \leq n \leq 100 \\
 & * [\frac{1}{2}] \text{mutex}(mut, P_m) \text{ locked}(mut, P_m, \frac{1}{2}) * \text{obs}(O_c \ \psi \{mut\}) \\
 & * (n > 0 \leftrightarrow O_c = \{s_{pop}^n\}) * (n = 0 \leftrightarrow O_c = \emptyset) \\
 & * \bigstar_{i=1, \dots, n} \text{itperm}(\cdot, 1) * \bigstar_{i=1, \dots, 100} \text{wperm}(\cdot, id_{push}^i, 0) \\
 & * \bigstar_{i=1, \dots, n-1} \text{unitSig}(id_{pop}^i) \\
 \\
 \text{no:mutex} \\
 L_c^{\text{no:obs}}(n, O_p) &:= \exists id_{pop}^n. \\
 & [\frac{1}{2}] \ell_{c_c} \mapsto n * 0 \leq n \leq 100 * [\frac{1}{2}] \text{mutex}(mut, P_m) * \text{obs}(O_c) \\
 & * (n > 0 \leftrightarrow O_c = \{s_{pop}^n\}) * (n = 0 \leftrightarrow O_c = \emptyset) \\
 & * \bigstar_{i=1, \dots, n} \text{itperm}(\cdot, 1) * \bigstar_{i=1, \dots, 100} \text{wperm}(\cdot, id_{push}^i, 0) \\
 & * \bigstar_{i=1, \dots, n-1} \text{unitSig}(id_{pop}^i)
 \end{aligned}$$

Loop invariant of consumer.

Iteration permissions consumed by productive loop iterations, i.e., by iterations which decrease the consumer counter  $c_c$ .

Used to generate iteration permissions to justify unproductive loop iterations.

Remaining allocated signal IDs used to initialize new signal after next pop.

Shorthand for invariant with acquired mutex.

Shorthand for invariant without mutex chunk and without obligations chunk.

Figure B.24: Consumer's loop invariant.

$\forall \ell_{\text{fifo10}}, \ell_m, \ell_{c_p}, \ell_{c_c}.$ $\dots$ $\{\text{phase}(\text{Forkee}) * \exists v_{c_p}, O_p. L_p(v_{c_p}, O_p) * v_{c_p} \neq 0\}$ <p><b>while</b> (<math>\forall v_{c_p}, O_p.</math> <span style="float: right; font-size: small;"><i>For definition of producer loop invariant <math>L_p(n, O)</math>, lock invariant <math>F_m</math> and variations cf. Figures B.23 and B.22.</i></span>)</p> $\left\{ \text{phase}(\text{Forkee}) * \exists v_{c_p}, O_p. L_p(v_{c_p}, O_p) * v_{c_p} \neq 0 * \forall o \in O_p. \text{lev}(o) = L_{\text{push}}^{v_{c_p}} \right\}$ <p><math>\text{lev}(mut) = 0 &lt; 101 - v_{c_p} = L_{\text{push}}^{v_{c_p}}</math></p> <p><b>acquire</b> <math>m;</math></p> $\left\{ \begin{array}{l} \text{phase}(\text{Forkee}) * L_p(v_{c_p}, O_p) * L_p^{\text{locked}}(v_{c_p}, O_p) * F_m * v_{c_p} \neq 0 \\ * \forall o \in O_p. \text{lev}(o) = \dots \end{array} \right\}$ <p><math>\forall v_{\text{fifo10}}^m.</math></p> $\left\{ F_m * F'_m(v_{\text{fifo10}}^m) * \dots \right\}$ <p><b>let</b> <math>f := [\text{fifo10}]</math> <b>in</b></p> $\{\text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * v_{c_p} \neq 0\}$ <p><b>if</b> <math>\text{size}(f) &lt; 10</math> <b>then</b> (<math>\dots</math> <span style="float: right; font-size: small;"><i>Production step presented on Figure B.26.</i></span>)</p> $\left\{ \begin{array}{l} \text{size}(v_{\text{fifo10}}^m) < 10 * \text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * v_{c_p} \neq 0 \\ \exists v'_{c_p}, O'_p. \text{obs}(O'_p \uplus \{mut\}) * \text{locked}(mut, F_m, \frac{1}{2}) * v_{c_p} \neq 0 \\ * \left( \text{size}(v_{\text{fifo10}}^m) < 10 \rightarrow L_p^{\text{no:mutext}}(v'_{c_p}, O'_p) * F_m * \text{itperm}(\cdot, 1) \right) \\ * \left( \text{size}(v_{\text{fifo10}}^m) = 10 \rightarrow L_p^{\text{no:obs}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * O'_p = O_p \right) \end{array} \right\}$ <p>) <b>else</b> (<math>\dots</math>)</p> $\left\{ \begin{array}{l} \text{size}(v_{\text{fifo10}}^m) = 10 * \text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * v_{c_p} \neq 0 \\ \text{size}(v_{\text{fifo10}}^m) = 10 * \text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * v_{c_p} \neq 0 \\ v_{c_p} \neq 0 \exists v'_{c_p}, O'_p. \text{obs}(O'_p \uplus \{mut\}) * \text{locked}(mut, F_m, \frac{1}{2}) * \text{Postf}_p \end{array} \right\}$ <p>);</p> $\left\{ \begin{array}{l} \text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{c_p}, O_p) * F'_m(v_{\text{fifo10}}^m) * v_{c_p} \neq 0 \\ \exists v'_{c_p}, O'_p. \text{obs}(O'_p \uplus \{mut\}) * \text{locked}(mut, F_m, \frac{1}{2}) * \text{Postf}_p \end{array} \right\}$ <p><math>\forall v'_{c_p}, O'_p.</math></p> <p><b>release</b> <math>m;</math> <span style="float: right; font-size: small;"><i>Wait step presented on Figure B.27, i.e., view shift performed after releasing <math>mut</math> but before consuming <math>F_m</math>.</i></span></p> <p><math>\forall \tau_p^{\text{anc}}, \delta_p.</math></p> <p><b>let</b> <math>c := [c_p]</math> <b>in</b> <math>c \neq 0</math></p> $\left\{ \begin{array}{l} \tau_p^{\text{anc}} \sqsubseteq (\text{Forkee}) * \text{phase}(\text{Forkee}) \\ * (v'_{c_p} \neq 0 \rightarrow \exists v_{c_p}, O_p. L_p(v_{c_p}, O_p) * v_{c_p} \neq 0 * \text{itperm}(\tau_p^{\text{anc}}, \delta_p)) \\ * (v'_{c_p} = 0 \rightarrow \text{obs}(\emptyset)) \end{array} \right\}$ <p>) <b>do skip</b></p> $\{\text{phase}(\text{Forkee}) * \exists v_{c_p}. L_p(v_{c_p}) * v_{c_p} \neq 0 * \text{obs}(\emptyset)\}$ <p><math>\dots</math></p>	<p style="font-size: small;"><i>Continuation of Figure B.21.</i></p> <p style="font-size: small;">PR-WHILE-SIMP &amp; PR-EXISTS (2x) &amp; PR-VS-SIMP &amp; VS-SEMIMP</p> $O_p = \{s_{\text{push}}^{v_{c_p}}\} \vee O_p = \emptyset$ <p style="font-size: small;"><i>Justification for application of:</i> PR-ACQUIRE</p> <p style="font-size: small;">PR-EXISTS</p> <p style="font-size: small;">PR-LET &amp; PR-READHEAPLOC</p> <p style="font-size: small;">PR-IF</p> <p style="font-size: small;">Define <math>\text{Postf}_p</math> such that:</p> $= \left\{ \begin{array}{l} \text{phase}(\text{Forkee}) \\ * \exists v'_{c_p}, O'_p. \\ \text{obs}(O'_p \uplus \{mut\}) \\ * \text{locked}(mut, F_m, \frac{1}{2}) \\ * \text{Postf}_p \end{array} \right\}$ <p style="font-size: small;">PR-VS-SIMP &amp; VS-SEMIMP</p> <p style="font-size: small;">PR-EXISTS (2x)</p> <p style="font-size: small;">PR-RELEASE &amp; PR-EXISTS (2x)</p> <p style="font-size: small;">PR-LET &amp; PR-READHEAPLOC &amp; PR-EXP</p> <p style="font-size: small;"><i>Reestablished loop invariant.</i></p> <p style="font-size: small;"><i>Continued in Figure B.21.</i></p>
--	--

Figure B.25: Verification example bounded FIFO, producer loop.

$\forall \ell_{\text{fifo10}}, \ell_m, \ell_{\text{cp}}, \ell_{\text{cc}}, v_{\text{cp}}, O_p, v_{\text{fifo10}}^m$ .  
 $\dots$

For definition of  $F_m$ ,  $F_p^f(v)$ ,  $L_{\text{locked}}(n, O)$  and  $L_p^{\text{no:obs}}(n, O)$  cf. Figures B.22 and B.23.

$\{\text{size}(v_{\text{fifo10}}^m) < 10 * \text{phase}(\text{Forkee}) * L_p^{\text{locked}}(v_{\text{cp}}, O_p) * P'_m(v_{\text{fifo10}}^m) * v_{\text{cp}} \neq 0\}$

$$\left\{ \begin{array}{l} [\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}} * [\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}}^m \quad \ell_{\text{cp}} \mapsto v_{\text{cp}} * v_{\text{cp}} = v_{\text{cp}}^m \\ * (v_{\text{cp}} > 0 \rightarrow \text{signal}(s_{\text{push}}^{v_{\text{cp}}}, \text{False})) \quad \text{signal}(s_{\text{push}}^{v_{\text{cp}}}, \text{False}) \\ * (v_{\text{cp}} > 0 \leftrightarrow O_p = \{s_{\text{push}}^{v_{\text{cp}}}\}) * (v_{\text{cp}} = 0 \leftrightarrow O_p = \emptyset) \quad O_p = \{s_{\text{push}}^{v_{\text{cp}}}\} \\ \dots \end{array} \right\}$$

let  $c := [\text{cp}]$  in  
 $[\text{fifo10}] := f \cdot (c :: \text{nil}); [\text{cp}] := c - 1$   
 $\{\ell_{\text{fifo10}} \mapsto v_{\text{fifo10}}^m \cdot (v_{\text{cp}} :: \text{nil}) * \ell_{\text{cp}} \mapsto v_{\text{cp}} - 1 \dots\}$   
 $\{\text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}, \text{mut}\}) * \text{signal}(s_{\text{push}}^{v_{\text{cp}}}, \text{True}) * \dots\}$   
 $\{(v_{\text{cp}} - 1 = 0 \vee v_{\text{cp}} > 0) * \dots\}$

case:  $v_{\text{cp}} - 1 = 0$  *Last iteration, nothing left to do.*

$$\left\{ \begin{array}{l} \text{phase}(\text{Forkee}) * \exists v'_{\text{cp}}, O'_p. \text{obs}(O'_p \uplus \{\text{mut}\}) * \text{locked}(\text{mut}, F_m, \frac{1}{2}) \\ * v_{\text{cp}} \neq 0 \\ * (\text{size}(v_{\text{fifo10}}^m) < 10 \rightarrow L_p^{\text{no:obs}}(v'_{\text{cp}}, O'_p) * F_m * \text{itperm}(\cdot, 1)) \\ * (\text{size}(v_{\text{fifo10}}^m) = 10 \rightarrow L_p^{\text{no:obs}}(v_{\text{cp}}, O_p) * P'_m(v_{\text{fifo10}}^m) * O'_p = O_p) \end{array} \right\}$$

case:  $v_{\text{cp}} - 1 > 0$  *Must create signal for next iteration.*

$$\left\{ \begin{array}{l} \text{obs}(\{id_{\text{push}}^{v_{\text{cp}}-1}, \text{mut}\}) * \bigstar_{i=1, \dots, v_{\text{cp}}-2} \text{uninitSig}(id_{\text{push}}^i) \\ * \text{uninitSig}(id_{\text{push}}^{v_{\text{cp}}-1}) \quad \text{signal}(s_{\text{push}}^{v_{\text{cp}}-1}, \text{False}) * \dots \end{array} \right\}$$

$\{\text{phase}(\text{Forkee}) * \exists v'_{\text{cp}}, O'_p. \text{obs}(O'_p \uplus \{\text{mut}\}) * \text{locked}(\text{mut}, F_m, \frac{1}{2}) * \text{Postf}_p\}$   
 $\dots$

*Continuation of Figure B.25.*

PR-VS-SIMP & VS-SEMIMP

PR-LET & PR-READHEAPLOC  
PR-ASSIGNTOHEAP (2x)

PR-VS-SIMP & VS-SETIGNAL

PR-VS-SIMP & VS-SEMIMP

PR-VS-SIMP & VS-OR

PR-VS-SIMP & VS-SEMIMP

$$= \left\{ \begin{array}{l} \text{phase}(\text{Forkee}) \\ * \exists v'_{\text{cp}}, O'_p. \\ \text{obs}(O'_p \uplus \{\text{mut}\}) \\ * \text{locked}(\text{mut}, F_m, \frac{1}{2}) \\ * \text{Postf}_p \end{array} \right\}$$

*For definition of Postf<sub>p</sub>*  
*cf. Figure B.25.*

PR-VS-SIMP & VS-SIGINIT

PR-VS-SIMP & VS-SEMIMP

*Continued in Figure B.25.*

Figure B.26: Verification example bounded FIFO, producer thread's production step.

$\forall \ell_{\text{fifo10}}, \ell_m, \ell_{\text{cp}}, \ell_{\text{cc}}, v_{\text{cp}}, O_p, v_{\text{fifo10}}^m, v'_{\text{cp}}, O'_p.$ <p>...</p> <p>For definition of <math>F_m, F'_m(v), L_p^{\text{locked}(n, O)}, L_p^{\text{no:obs}(n, O)}, \text{Postf}_p</math> cf. Figures B.22, B.23 and B.25.</p> $\{ \text{phase}(\text{Forkee}) * \text{obs}(O'_p \uplus \{mut\}) * \text{locked}(mut, F_m, \frac{1}{2}) * \text{Postf}_p \}$ <p><b>release m</b> <i>PR-RELEASE allows view shift to happen after mutex mut was released but before lock invariant <math>F_m</math> is consumed.</i></p> $\{ \text{phase}(\text{Forkee}) * \text{obs}(O'_p \uplus \{mut\}) * \text{locked}(mut, F_m, \frac{1}{2}) * \text{Postf}_p \}$ $\{ (\text{size}(v_{\text{fifo10}}^m) < 10 \vee \text{size}(v_{\text{fifo10}}^m) = 10) * \dots \}$ <p>case: <math>\text{size}(v_{\text{fifo10}}^m) &lt; 10</math> <i>Production step already performed, nothing left to do.</i></p> $\{ \text{size}(v_{\text{fifo10}}^m) < 10 * \text{phase}(\text{Forkee}) * \text{obs}(O'_p) * \text{Postf}_p \}$ $\left\{ \begin{array}{l} \text{obs}(O'_p) * F_m * \exists \tau_p^{\text{anc}}, \delta_p. \tau_p^{\text{anc}} \sqsubseteq (\text{Forkee}) * \text{phase}(\text{Forkee}) \\ \text{no:mutex} \\ * L_p^{\text{no:obs}}(v'_{\text{cp}}, O'_p) * \text{itperm}(\tau_p^{\text{anc}}, \delta_p) \end{array} \right\}$ <p>case: <math>\text{size}(v_{\text{fifo10}}^m) = 10</math> <i>No production step performed. Must wait to generate permission.</i></p> $\{ \text{size}(v_{\text{fifo10}}^m) = 10 * \text{phase}(\text{Forkee}) * \text{obs}(O'_p) * \text{Postf}_p \}$ $\left\{ \begin{array}{l} \text{no:mutex} \\ \text{Postf}_p \\ L_p^{\text{no:obs}}(v_{\text{cp}}, O_p) * F'_m(v_{\text{fifo10}}^m) * O'_p = O_p = \{s_{\text{push}}^{v_{\text{cp}}}\} \\ * \text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}\}) * v_{\text{cp}} \neq 0 * \text{phase}(\text{Forkee}) \end{array} \right\}$ $\forall v_{\text{cp}}^m, v_{\text{cc}}^m.$ $\left\{ \begin{array}{l} \exists v_{\text{cp}}^m. [\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}} * [\frac{1}{2}] \ell_{\text{cp}} \mapsto v_{\text{cp}}^m * v_{\text{cp}} = v_{\text{cp}}^m \\ * v_{\text{cc}}^m = v_{\text{cp}} + 10 \\ * \exists v_{\text{cc}}^m. (v_{\text{cc}}^m > 0 \rightarrow \text{signal}(s_{\text{pop}}^m, \text{False})) \text{signal}(s_{\text{pop}}^m, \text{False}) \\ * \bigstar_{i=1, \dots, 100} \text{wperm}(\ell, id_{\text{pop}}^i, 0) * \text{obs}(\{s_{\text{push}}^{v_{\text{cp}}}\}) * \text{phase}(\text{Forkee}) \\ * \dots \end{array} \right\}$ $\text{lev}(s_{\text{pop}}^m) = L_{\text{pop}}^m = L_{\text{pop}}^m \frac{v_{\text{cp}} + 10}{v_{\text{cp}}} = 102 - (v_{\text{cp}} + 10)$ $< 101 - v_{\text{cp}} = L_{\text{push}}^m \frac{v_{\text{cp}}}{v_{\text{cp}}} = \text{lev}(s_{\text{push}}^{v_{\text{cp}}})$ $\{ \text{itperm}(\text{Forkee}, 0) * \dots \}$ $\{ \text{obs}(O'_p) * F_m * \text{PostReleaseVS}_p \}$ $\left\{ \begin{array}{l} \text{phase}(\text{Forkee}) * \text{obs}(O'_p) \\ * \text{Postf}_p * (\text{size}(v_{\text{fifo10}}^m) < 10 \vee \text{size}(v_{\text{fifo10}}^m) = 10) \\ * F_m * \text{PostReleaseVS}_p \end{array} \right\}$ <p><math>\{ \text{obs}(O'_p) * F_m * \text{PostReleaseVS}_p \}</math> <i>Lock invariant <math>F_m</math> consumed by PR-RELEASE.</i></p> $\left\{ \begin{array}{l} \exists \tau_p^{\text{anc}}, \delta_p. \tau_p^{\text{anc}} \sqsubseteq (\text{Forkee}) * \text{phase}(\text{Forkee}) \\ * (v'_{\text{cp}} \neq 0 \rightarrow \exists v_{\text{cp}}, O_p. L_p(v_{\text{cp}}, O_p) * v_{\text{cp}} \neq 0 * \text{itperm}(\tau_p^{\text{anc}}, \delta_p)) \\ * (v'_{\text{cp}} = 0 \rightarrow \text{obs}(\emptyset)) \end{array} \right\}$ <p>...</p>	<p><i>Continuation of Figure B.25.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-VS-SIMP &amp; VS-OR</p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>Define <math>\text{PostReleaseVS}_p</math> such that: = <math>\{ \text{obs}(O'_p) * F_m * \text{PostReleaseVS}_p \}</math></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-EXISTS (2x) &amp; PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-VS-SIMP &amp; VS-WAIT</p> <p><i>Justification for application of VS-WAIT.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p><i>Conclusion of VS-OR application.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p><i>Reestablished loop invariant.</i></p> <p><i>Continued in Figure B.25.</i></p>
--	--

Figure B.27: Verification example bounded FIFO, producer's wait step.

<pre> <math>\forall \ell_{\text{fifo10}}, \ell_m, \ell_{\text{cp}}, \ell_{\text{cc}}.</math> ... {phase((Forker)) * <math>\exists v_{\text{cc}}, O_c. L_c(v_{\text{cc}}, O_c) * v_{\text{cc}} \neq 0</math>} while (   <math>\forall v_{\text{cc}}, O_c.</math>   {phase((Forker)) * <math>\exists v_{\text{cp}}, O_c. L_c(v_{\text{cc}}, O_c) * v_{\text{cc}} \neq 0 * \forall o \in O_c. \text{lev}(o) = L_{\text{pop}}^{v_{\text{cc}}}</math>}   lev(mut) = 0 &lt; 102 - v_{\text{cc}} = L_{\text{pop}}^{v_{\text{cc}}}   acquire m   {     phase((Forker)) * <math>L_c(v_{\text{cc}}, O_c)</math> * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m * v_{\text{cc}} \neq 0</math>     * <math>\forall o \in O_c. \text{lev}(o) = \dots</math>   }   <math>\forall v_{\text{fifo10}}^m.</math>   { <math>P_m</math> * <math>P_m'(v_{\text{fifo10}}^m)</math> * ... }   let f := [fifo10] in   {phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math>}   if size(f) &gt; 0 then (     { size(v_{\text{fifo10}}^m) &gt; 0 * phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math> }     ... B.29 ...     {       size(v_{\text{fifo10}}^m) &gt; 0 * phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math>       * <math>\exists v'_{\text{cc}}, O'_c. \text{obs}(O'_c \uplus \{mut\}) * \text{locked}(mut, P_m, \frac{1}{2}) * v_{\text{cc}} \neq 0</math>       * (size(v_{\text{fifo10}}^m) &gt; 0 <math>\rightarrow L_c^{\text{no:mutex}}(v'_{\text{cc}}, O'_c) * P_m * \text{itperm}(\cdot, 1)</math>)       * (size(v_{\text{fifo10}}^m) = 0 <math>\rightarrow L_c^{\text{no:mutex, no:obs}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * O'_c = O_c</math>)     }   ) else (     { size(v_{\text{fifo10}}^m) = 0 * phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math> }     {       size(v_{\text{fifo10}}^m) = 0 * phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math>       * <math>\exists v'_{\text{cc}}, O'_c. \text{obs}(O'_c \uplus \{mut\}) * \text{locked}(mut, P_m, \frac{1}{2}) * \text{Postf}_c</math>     }   );   {     phase((Forker)) * <math>L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m'(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0</math>     * <math>\exists v'_{\text{cc}}, O'_c. \text{obs}(O'_c \uplus \{mut\}) * \text{locked}(mut, P_m, \frac{1}{2}) * \text{Postf}_c</math>   }   <math>\forall v'_{\text{cc}}, O'_c.</math>   release m;   Wait step presented on Figure B.30, i.e., view shift performed   after releasing mut but before consuming P_m.   <math>\forall \tau_c^{\text{anc}}, \delta_c.</math>   let c := [c] in c <math>\neq 0</math>   {     <math>\tau_c^{\text{anc}} \sqsubseteq (\text{Forker}) * \text{phase}(\text{Forker})</math>     * (v'_{\text{cc}} <math>\neq 0 \rightarrow \exists v_{\text{cp}}, O_c. L_c(v_{\text{cc}}, O_c) * v_{\text{cc}} \neq 0 * \text{itperm}(\tau_c^{\text{anc}}, \delta_c)</math>)     * (v'_{\text{cc}} = 0 <math>\rightarrow \text{obs}(\emptyset)</math>)   } ) do skip {phase((Forker)) * <math>\exists v_{\text{cp}}, L_c(v_{\text{cp}}) * v_{\text{cc}} \neq 0</math> * <math>\text{obs}(\emptyset)</math> } ... </pre>	<p>Continuation of Figure B.21.</p> <p>PR-WHILE-SIMP &amp; PR-EXISTS (2x) &amp; PR-VS-SIMP &amp; VS-SEMIMP</p> <p><math>O_c = \{s_{\text{pop}}^{v_{\text{cc}}}\} \vee O_c = \emptyset</math></p> <p>Justification for application of: PR-ACQUIRE</p> <p>PR-EXISTS</p> <p>PR-LET &amp; PR-READHEAPLOC</p> <p>PR-IF</p> <p>Define Postf<sub>c</sub> such that:</p> $= \left\{ \begin{array}{l} \text{phase}(\text{Forker}) \\ * \exists v'_{\text{cc}}, O'_c. \\ \text{obs}(O'_c \uplus \{mut\}) \\ * \text{locked}(mut, P_m, \frac{1}{2}) \\ * \text{Postf}_c \end{array} \right\}$ <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-EXISTS (2x)</p> <p>PR-RELEASE &amp; PR-EXISTS (2x)</p> <p>PR-LET &amp; PR-READHEAPLOC &amp; PR-EXP</p> <p>Reestablished loop invariant.</p> <p>Continued in Figure B.21.</p>
---	---

Figure B.28: Verification example bounded FIFO, consumer loop.

$\forall \ell_{\text{fifo10}}, \ell_m, \ell_{\text{cp}}, \ell_{\text{cc}}, v_{\text{cc}}, O_c, v_{\text{fifo10}}^m$   
 $\dots$

For definition of  $P_m, P_m^f(v), L_c^{\text{locked}}(n, O)$  and  $L_c^{\text{no:obs}}(n, O)$  cf. Figures B.22 and B.24.

$\{\text{size}(v_{\text{fifo10}}^m) > 0 * \text{phase}(\text{Forker}) * L_c^{\text{locked}}(v_{\text{cc}}, O_c) * P_m^f(v_{\text{fifo10}}^m) * v_{\text{cc}} \neq 0\}$   
 $\forall v_{\text{cc}}^m$

$$\left\{ \begin{array}{l} [\frac{1}{2}] \ell_{\text{cc}} \mapsto v_{\text{cc}} * [\frac{1}{2}] \ell_{\text{cc}} \mapsto v_{\text{cc}}^m \quad \ell_{\text{cc}} \mapsto v_{\text{cc}} * v_{\text{cc}} = v_{\text{cc}}^m \\ * (v_{\text{cc}}^m > 0 \rightarrow \text{signal}(s_{\text{pop}}^{v_{\text{cc}}^m}, \text{False})) \quad \text{signal}(s_{\text{pop}}^{v_{\text{cc}}}, \text{False}) \\ * (v_{\text{cc}} > 0 \leftrightarrow O_c = \{s_{\text{pop}}^{v_{\text{cc}}}\}) * (v_{\text{cc}} = 0 \leftrightarrow O_c = \emptyset) \quad O_c = \{s_{\text{pop}}^{v_{\text{cc}}}\} \end{array} \right\}$$

$\text{let } c := [c_c] \text{ in}$   
 $[\text{fifo10}] := \text{tail}(f); [c_c] := c - 1$   
 $\{\ell_{\text{fifo10}} \mapsto \text{tail}(v_{\text{fifo10}}^m) * \ell_{\text{cc}} \mapsto v_{\text{cc}} - 1 * \dots\}$   
 $\{\text{obs}(\{s_{\text{pop}}^{v_{\text{cc}}}, \text{mut}\}) * \text{signal}(s_{\text{pop}}^{v_{\text{cc}}}, \text{True}) * \dots\}$   
 $\{(v_{\text{cc}} - 1 = 0 \vee v_{\text{cc}} > 0) * \dots\}$

case:  $v_{\text{cc}} - 1 = 0$  *Last iteration, nothing left to do.*

$$\left\{ \begin{array}{l} \text{phase}(\text{Forker}) * \exists v_{\text{cp}}', O_c'. \text{obs}(O_c' \uplus \{\text{mut}\}) * \text{locked}(\text{mut}, P_m, \frac{1}{2}) \\ * v_{\text{cc}} \neq 0 \\ * (\text{size}(v_{\text{fifo10}}^m) > 0 \rightarrow L_c^{\text{no:obs}}(v_{\text{cp}}', O_c') * P_m * \text{itperm}(\cdot, 1)) \\ * (\text{size}(v_{\text{fifo10}}^m) = 0 \rightarrow L_c^{\text{no:obs}}(v_{\text{cc}}, O_c) * P_m^f(v_{\text{fifo10}}^m) * O_c' = O_c) \end{array} \right\}$$

case:  $v_{\text{cc}} - 1 > 0$  *Must create signal for next iteration.*

$$\left\{ \begin{array}{l} \text{obs}(\{id_{\text{pop}}^{v_{\text{cc}}-1}, \text{mut}\}) * \text{uninitSig}(id_{\text{pop}}^i) \\ * \text{uninitSig}(id_{\text{pop}}^{v_{\text{cc}}-1}) \quad \text{signal}(s_{\text{pop}}^{v_{\text{cc}}-1}, \text{False}) * \dots \\ i=1, \dots, v_{\text{cc}}-1 \end{array} \right\}$$

$\{\text{phase}(\text{Forker}) * \exists v_{\text{cc}}', O_c'. \text{obs}(O_c' \uplus \{\text{mut}\}) * \text{locked}(\text{mut}, P_m, \frac{1}{2}) * \text{Postf}_c\}$   
 $\{\text{phase}(\text{Forker}) * \exists v_{\text{cc}}', O_c'. \text{obs}(O_c' \uplus \{\text{mut}\}) * \text{locked}(\text{mut}, P_m, \frac{1}{2}) * \text{Postf}_c\}$   
 $\dots$

*Continuation of Figure B.28.*

PR-EXISTS  
 PR-VS-SIMP & VS-SEMIMP  
  
 PR-LET & PR-READHEAPLOC  
 PR-ASSIGNTOHEAP (2x)  
 PR-VS-SIMP & VS-SET SIGNAL  
  
 PR-VS-SIMP & VS-SEMIMP  
  
 PR-VS-SIMP & VS-OR  
  
 PR-VS-SIMP & VS-SEMIMP  

$$= \left\{ \begin{array}{l} \text{phase}(\text{Forker}) \\ * \exists v_{\text{cc}}', O_c'. \\ \text{obs}(O_c' \uplus \{\text{mut}\}) \\ * \text{locked}(\text{mut}, P_m, \frac{1}{2}) \\ * \text{Postf}_c \end{array} \right\}$$

*For definition of Postf<sub>c</sub>*  
*cf. Figure B.28.*

PR-VS-SIMP & VS-SIGINT  
  
 PR-VS-SIMP & VS-SEMIMP

*Continued in Figure B.28.*

Figure B.29: Verification example bounded FIFO, consumer thread's consumption step.

$\forall \ell_{\text{fifo10}}, \ell_m, \ell_{cc}, \ell_{cc}, v_{cc}, O_c, v_{\text{fifo10}}^m, v'_{cc}, O'_c.$ <p>...</p> <p>For definition of <math>P_m, P'_m(v), L^{\text{locked}}(n, O), L_c^{\text{no:obs}}(n, O), \text{Postf}_c</math> cf. Figures B.22, B.24 and B.28.  <math>\{ \text{phase}(\text{Forker}) * \text{obs}(O'_c \uplus \{mut\}) * \text{locked}(mut, P_m, \frac{1}{2}) * \text{Postf}_c \}</math>  <b>release m</b> <math>\quad</math> PR-RELEASE allows view shift to happen after mutex <math>mut</math> was released but before lock invariant <math>P_m</math> is consumed.</p> $\{ \text{phase}(\text{Forker}) * \text{obs}(O'_c \uplus \{mut\}) * \text{locked}(mut, P_m, \frac{1}{2}) * \text{Postf}_c \}$ $\{ (\text{size}(v_{\text{fifo10}}^m) > 0 \vee \text{size}(v_{\text{fifo10}}^m) = 0) * \dots \}$ <p>case: <math>\text{size}(v_{\text{fifo10}}^m) &gt; 0</math> <span style="float: right;"><i>Consumption step already performed, nothing left to do.</i></span></p> $\{ \text{size}(v_{\text{fifo10}}^m) > 0 * \text{phase}(\text{Forker}) * \text{obs}(O'_c) * \text{Postf}_c \}$ $\left\{ \begin{array}{l} \text{obs}(O'_c) * P_m * \exists \tau_c^{\text{anc}}, \delta_c. \tau_c^{\text{anc}} \sqsubseteq (\text{Forker}) * \text{phase}(\text{Forker}) \\ * L_c^{\text{no:obs}}(v'_{cc}, O'_c) * \text{itperm}(\cdot, 1) \end{array} \right\}$ <p>case: <math>\text{size}(v_{\text{fifo10}}^m) = 0</math> <span style="float: right;"><i>No production step performed. Must wait to generate permission.</i></span></p> $\{ \text{size}(v_{\text{fifo10}}^m) = 0 * \text{phase}(\text{Forker}) * \text{obs}(O'_c) * \text{Postf}_c \}$ $\left\{ \begin{array}{l} \text{Postf}_c \quad L_c^{\text{no:obs}}(v_{cc}, O_c) * P'_m(v_{\text{fifo10}}^m) * O'_c = O_c = \{s_{\text{pop}}^{v_{cp}}\} \\ * \text{obs}(\{s_{\text{pop}}^{v_{cp}}\}) * v_{cc} \neq 0 * \text{phase}(\text{Forker}) \end{array} \right\}$ <p><math>\forall v_{cc}^m, v_{cp}^m.</math></p> $\left\{ \begin{array}{l} \exists v_{cc}^m, v_{cp}^m. [\frac{1}{2}] \ell_{cc} \mapsto v_{cc} * [\frac{1}{2}] \ell_{cc} \mapsto v_{cc}^m * v_{cc} = v_{cc}^m \\ * v_{cc} = v_{cp}^m + 0 \\ (v_{cp}^m > 0 \rightarrow \text{signal}(s_{\text{push}}^{v_{cp}}, \text{False})) \quad \text{signal}(s_{\text{push}}^{v_{cp}}, \text{False}) \\ * \bigstar_{i=1, \dots, 100} \text{wperm}(\cdot, id_{\text{push}}^i, 0) * \text{obs}(\{s_{\text{pop}}^{v_{cp}}\}) * \text{phase}(\text{Forker}) \\ * \dots \end{array} \right\}$ <p><math>\text{lev}(s_{\text{push}}^{v_{cp}}) = L_{\text{push}}^{v_{cp}} = L_{\text{push}}^{v_{cc}} = 101 - v_{cc} &lt; 102 - v_{cc}</math></p> $= L_{\text{pop}}^{v_{cc}} = \text{lev}(s_{\text{pop}}^{v_{cc}})$ $\{ \text{itperm}(\text{Forker}, 0) * \dots \}$ $\{ \text{obs}(O'_c) * P_m * \text{PostReleaseVS}_c \}$ $\left\{ \begin{array}{l} \text{phase}(\text{Forker}) * \text{obs}(O'_c) \\ * \text{Postf}_c * (\text{size}(v_{\text{fifo10}}^m) < 10 \vee \text{size}(v_{\text{fifo10}}^m) = 10) \\ * P_m * \text{PostReleaseVS}_c \end{array} \right\}$ <p><math>\{ \text{obs}(O'_c) * P_m * \text{PostReleaseVS}_c \}</math> <span style="float: right;"><i>Lock invariant <math>P_m</math> consumed by PR-RELEASE.</i></span></p> $\left\{ \begin{array}{l} \exists \tau_c^{\text{anc}}, \delta_c. \tau_c^{\text{anc}} \sqsubseteq (\text{Forker}) * \text{phase}(\text{Forker}) \\ * (v_{cc}^m \neq 0 \rightarrow \exists v_{cc}, O_c. L_c(v_{cc}, O_c) * v_{cc} \neq 0 * \text{itperm}(\tau_c^{\text{anc}}, \delta_c)) \\ * (v_{cc}^m = 0 \rightarrow \text{obs}(\emptyset)) \end{array} \right\}$ <p>...</p>	<p><i>Continuation of Figure B.28.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-VS-SIMP &amp; VS-OR</p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>Define <math>\text{PostReleaseVS}_c</math> such that:  <math>= \{ \text{obs}(O'_c) * P_m * \text{PostReleaseVS}_c \}</math></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-EXISTS (2x)          &amp; PR-VS-SIMP &amp; VS-SEMIMP</p> <p>PR-VS-SIMP &amp; VS-WAIT</p> <p><i>Justification for application of VS-WAIT.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p><i>Conclusion of VS-OR application.</i></p> <p>PR-VS-SIMP &amp; VS-SEMIMP</p> <p><i>Reestablished loop invariant.</i></p> <p><i>Continued in Figure B.28.</i></p>
--	--

Figure B.30: Verification example bounded FIFO, consumer's wait step.





# Appendix C

## Completeness Thresholds: Formalization

This chapter formalizes the work described in chapter 4. The contents of this chapter were published as technical report on arXiv:

**T. Reinhard.** *Completeness Thresholds for Memory Safety of Array Traversing Programs: Early Technical Report.* arXiv, CoRR, abs/2211.11885, 2022. <https://doi.org/10.48550/arXiv.2211.11885>

### C.1 Introduction

In this early technical report on an ongoing project, we present – to the best of our knowledge – the first study of completeness thresholds for memory safety proofs. Specifically we consider heap-manipulating programs that iterate over arrays without allocating or freeing memory. We present the first notion of completeness thresholds for program verification which reduce unbounded memory safety proofs to bounded ones. Moreover, we present some preliminary ideas on how completeness thresholds can be computed for concrete programs.

**Unbounded vs Bounded Proofs** Memory safety is a very basic property we want to hold for every critical program, regardless of its nature or purpose. Yet, it remains hard to prove and in general requires us to write tedious, inductive

proofs. One way to automate the verification process is to settle on bounded proofs and accept bounded guarantees.

Consider a program  $c$  that searches through an array of size  $s$ . An unbounded memory safety proof for  $c$  would yield that the program is safe for any possible input, in particular for any array size, i.e.,  $\forall s. \text{safe}(c)$ . A bounded proof that only considers input sizes smaller than 10 would only guarantee that the program is safe for any such bounded array, i.e.,  $\forall s < 10. \text{safe}(c)$ .

**Completeness Thresholds** Approximating unbounded proofs by bounded ones is a technique often used in model checking. Hence, the relationship between bounded and unbounded proofs about finite state transition systems has been studied extensively [15, 40, 109, 28, 2, 87, 10, 126]. For a finite transition system  $T$  and a property of interest  $\phi$ , a *completeness threshold* is any number  $k$  such that we can prove  $\phi$  by only examining path prefixes of length  $k$  in  $T$ , i.e.,  $T \models_k \phi \Rightarrow T \models \phi$  [40]<sup>1</sup>. Over the years, various works characterised over-approximations of least completeness thresholds for different types of properties  $\phi$ . These over-approximations are typically described in terms of key attributes of the transition system  $T$ , such as the recurrence diameter [109].

Heap-manipulating programs are essentially infinite state transition systems. Hence, in general, these key attributes are infinite. This vast structural difference between the programs we are interested in and the transition systems for which completeness thresholds have been studied prevents us from reusing any of the existing definitions or results.

In § C.2 we start by presenting basic definitions and notations that are used throughout this work. In § C.3 and § C.4 we present the syntax and semantics of the programming language we consider. In § C.5 and § C.6 we formalise our assertion language as well as the notion of memory safety we consider. Our study of completeness thresholds relies on verification conditions, which we define in § C.7. We introduce completeness thresholds and study their properties in § C.8.

## C.2 General Notation and Basic Definitions

The following definitions and notations will be used throughout this work.

---

<sup>1</sup>Note that the term completeness threshold is used inconsistently in literature. Some papers such as [40] use the definition above, according to which completeness thresholds are not unique. Others such as [109] define them as the minimal number  $k$  such that  $T \models_k \phi \Rightarrow T \models \phi$ , which makes them unique.

**Definition C.2.1** (Tuples). For any set  $X$  we denote the set of tuples over  $X$  as

$$X^* := \bigcup_{n \in \mathbb{N}} \{(x_1, \dots, x_n) \mid x_1, \dots, x_n \in X\}.$$

We denote tuples by overlining the variable name, i.e.,  $\bar{x} \in X^*$ , except if it is clear from the context.

**Definition C.2.2** (Non-Empty Tuples). For any set  $X$  we denote the set of non-empty tuples over  $X$  as

$$X^+ := \{t \in X^* \mid \text{len}(t) > 0\}.$$

**Definition C.2.3** (Disjoint Union). Let  $A, B$  be sets. We define their disjoint union as

$$A \sqcup B := A \cup B$$

if  $A \cap B = \emptyset$  and leave it undefined otherwise.

**Notation C.2.4** (Homomorphic & Isomorphic). Let  $A, B$  be algebraic structures. We define the following notations:

- $A \sim B$  expresses that  $A$  and  $B$  are homomorphic.
- $A \simeq B$  expresses that  $A$  and  $B$  are isomorphic.

**Remark C.2.5** (Canonical Homomorphism from Tuples to Sets). Let  $X$  be a set. Then,  $X^*$  with concatenation and  $\mathcal{P}(X)$  with union are both monoids and the canonical homomorphism from  $X^*$  to  $\mathcal{P}(X)$  is  $(x_1, \dots, x_n) \mapsto \{x_1, \dots, x_n\}$ .

**Definition C.2.6** (Congruence Relation between Tuples and Sets). Let  $X$  be a set and let  $h_X : X^* \rightarrow \mathcal{P}(X)$  be the canonical homomorphism. We define the congruence relation  $\cong_X \subseteq (X^* \times \mathcal{P}(X)) \cup (\mathcal{P}(X) \times X^*)$  such that the following holds for all  $\bar{t} \in X^*$  and all  $S \in \mathcal{P}(X)$ :

$$S \cong_X \bar{t} \iff \bar{t} \cong_X S \iff h_X(\bar{t}) = S$$

Whenever the base set  $X$  is clear from the context, we write  $\cong$  instead of  $\cong_X$ .

### C.3 Syntax

In this section we define the syntax of the programming language that we use in the rest of this work.

**Definition C.3.1** (Variables). *We define  $\mathcal{X}$  to be an infinite set of variable symbols.*

Our language allows for simple pointer arithmetic of the form  $\ell +_{\mathcal{L}ocs} \mathbb{Z} Z$  where  $\ell$  is a heap location and  $Z$  is an offset.

**Definition C.3.2** (Structured Heap Locations). *We define the sets of heap objects  $HObjs$  and heap indices  $HIdxs$  to be infinite sets with  $\mathbb{N} \subset HIdxs$ . Further, we define the set of heap locations as*

$$\mathcal{L}ocs := HObjs \times HIdxs.$$

We denote heap locations by  $\ell$ , heap objects by  $o$  and heap indices by  $idx$ .

**Definition C.3.3** (Unit Type). *We define the unit type as  $\text{Unit} := \{()\}$ .*

**Definition C.3.4** (Types). *We define the set of types  $Types$  syntactically as follows:*

$$T \in Types := \mathcal{L}ocs \mid HObjs \mid \mathbb{Z} \mid \mathbb{B} \mid \text{Unit}$$

Operations are pure functions that map inputs to output and cannot access the heap.

**Definition C.3.5** (Operations). *We define  $Ops$  to be a set of operations with  $Ops \subseteq \{f : D \rightarrow C \mid D \in Types^*, C \in Types\}$  and with  $<, + \in Ops$ .*

*For each  $f : D \rightarrow C \in Ops$ , we use the following notation  $\text{dom}(f) = D$  and  $\text{codom}(f) = C$ .*

While this work currently only deals with arrays, our plan is to investigate completeness thresholds for programs that deal with arbitrary tree-like inductive data structures. In order to keep the semantics of our language modular and to allow for easy extensions, we introduce an unspecified set of *heap commands* that captures the APIs of the data structures we are interested in.

**Definition C.3.6** (Heap Commands). *We define  $HCmds$  to be a set of symbols. Every  $hc \in HCmds$  represents a command that accesses the heap.*

**Definition C.3.7** (Program syntax). *We define the set of commands  $Cmds$ , expressions  $Exps$  and values  $Values$  syntactically by the grammar presented in Fig. C.1.*

$\ell \in \mathcal{Locs}$		Heap locations
$o \in \mathcal{HObjs}$		Heap Objects
$Z \in \mathbb{Z}$		
$b \in \mathbb{B}$		
$x \in \mathcal{X}$		Variables
$op \in \mathcal{Ops}$		Primitive Operations
$hc \in \mathcal{HCmds}$		Heap Commands
$v \in \mathcal{Values}$	$::= \ell \mid o \mid Z \mid b \mid ()$	Values
$e \in \mathcal{Exps}$	$::= v \mid x \mid op(\bar{e})$	Expressions
$c \in \mathcal{Cmds}$	$::= e \mid \mathbf{let} \ x := c \ \mathbf{in} \ c \mid$ $\mathbf{if} \ e \ \mathbf{then} \ c \ \mathbf{else} \ c \mid$ $\mathbf{while} \ !e \ \mathbf{do} \ c \mid$ $\mathbf{for} \ x \ \mathbf{in} \ [e \ \mathbf{to} \ e] \ \mathbf{do} \ c \mid$ $!e \mid !e := e \mid hc(\bar{e})$	Heap access
	(a) Values, expressions and commands.	
$c; c'$	$:= \mathbf{let} \ x := c \ \mathbf{in} \ c'$	where $x$ is not free in $c'$
$\mathbf{if} \ e \ \mathbf{then} \ c$	$:= \mathbf{if} \ e \ \mathbf{then} \ c \ \mathbf{else} \ ()$	
	(b) Syntactic sugar.	

Figure C.1: Program syntax.

## C.4 Dynamic Semantics

In order to keep things simple, our language uses a heap but no store. Hence, variables are actually constants that can be bound to values via **let** commands. As a consequence, the entire evaluation state of a program is represented by the heap and the program itself. Further, expressions are pure, hence their evaluation does not depend on the heap.

**Definition C.4.1** (Evaluation of Closed Expressions). *We define a partial evaluation function  $\llbracket \cdot \rrbracket : \mathcal{Exps} \rightarrow \mathcal{Values}$  on expressions by recursion on the structure of expressions as follows:*

$\llbracket v \rrbracket$	$:= v$	<i>if <math>v \in \mathcal{Values}</math>,</i>
$\llbracket op(e_1, \dots, e_n) \rrbracket$	$:= op(\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket)$	<i>if <math>\perp \notin \{\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket\}</math></i>
		<i>and <math>(\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket) \in \text{dom}(op)</math>,</i>
$\llbracket e \rrbracket$	$:= \perp$	<i>otherwise.</i>

We identify closed expressions  $e$  with their ascribed value  $\llbracket e \rrbracket$ .

**Definition C.4.2** (Evaluation Context). *We define the set of evaluation contexts  $EvalCtxts$  syntactically as follows:*

$$E \in EvalCtxts \quad ::= \text{let } x := \square \text{ in } c$$

For any  $c \in Ccmds$  and  $E \in EvalCtxts$ , we define  $E[c] := E[\square \mapsto c]$ .

**Definition C.4.3** (Free Variables (Commands)). *We define free variables in the usual way. For any command  $c$  we denote the set of variables that occur freely in  $c$  by  $freeVars(c)$ .*

**Definition C.4.4** (Substitution (Commands)). *We define substitution in the usual way. For any command  $c$ , variable  $x$  and expression  $e$ , we denote the result of substituting every free occurrence of  $x$  in  $c$  with  $e$  by  $c[x \mapsto e]$ . Further, we extend substitutions to tuples of variables and expressions in the canonical way.*

We explicitly model memory errors in our operational semantics. This way we know that (i) any execution which ends in a value does not involve memory errors and (ii) any execution that does involve memory errors ends in the dedicated error state `error`.

**Definition C.4.5** (Memory Errors). *We denote the memory error state by `error` and the set of potentially erroneous commands by  $Ccmds^+ := Ccmds \cup \{\text{error}\}$ . We denote potentially erroneous commands by  $c^+$ .*

Heaps are finite collections of resources that can be manipulated by commands.

**Definition C.4.6** (Physical Resources). *We define the set of physical resources  $\mathcal{R}^{\text{phys}}$  syntactically as follows:*

$$r^{\text{p}} \in \mathcal{R}^{\text{phys}} \quad ::= \ell \mapsto v$$

$$\ell \in \mathcal{Locs} \quad v \in \text{Values}$$

**Definition C.4.7** (Physical Heaps). *We define the set of physical heaps as*

$$Heaps^{\text{phys}} := \mathcal{P}_{\text{fin}}(\mathcal{R}^{\text{phys}})$$

and the function  $\text{locs}_{\text{pRes}} : Heaps^{\text{phys}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$  mapping physical heaps to the sets of allocated heap locations as

$$\text{locs}_{\text{pRes}}(h) := \{\ell \in \mathcal{Locs} \mid \exists v \in \text{Values}. \ell \mapsto v \in h\}.$$

We denote physical heaps by  $h$ .

**Definition C.4.8** (Basic Commands). *We define the set of basic commands  $BCmds$  syntactically as follows:*

$$bc \in BCmds ::= e \mid \mathbf{let} \ x := bc \ \mathbf{in} \ bc \mid \\ \mathbf{if} \ e \ \mathbf{then} \ bc \ \mathbf{else} \ bc \mid \\ \mathbf{while} \ !e \ \mathbf{do} \ bc \mid \\ \mathbf{for} \ x \ \mathbf{in} \ [e \ \mathbf{to} \ e] \ \mathbf{do} \ bc \mid \\ !e \mid !e := e$$

The set of basic commands is the subset of  $Cmds$  that consists exactly of those commands that do not involve any heap command call  $hc(\bar{e})$ . Remember that the set of heap commands captures the APIs of data structures. We do not want to change our operational semantics each time we want to consider a new data structure. Hence, we assume that there exists an interpretation for each heap command that describes its behaviour in terms of basic commands.

**Assumption C.4.9** (Heap Command Interpretation). *We assume that there exists a function  $\mathcal{I}_{hcmds} : HCmds \rightarrow (BCmds \times \mathcal{X}^*)$  that maps each heap command to a basic command and a vector of variables. Further, for every mapping of the form  $\mathcal{I}_{hcmds}(hc) = (bc, (x_1, \dots, x_n))$  the following two properties hold:*

- $\bigwedge_{i \neq j} x_i \neq x_j$
- $\text{freeVars}(bc) \subseteq \{x_1, \dots, x_n\}$ .

**Definition C.4.10** (Command Reduction Relation). *We define a command reduction relation  $\rightsquigarrow_{\text{cmd}}$  according to the rules presented in Fig. C.2. A reduction step has the form*

$$h, c \rightsquigarrow_{\text{cmd}} h', c'.$$

*We define  $\rightsquigarrow_{\text{cmd}}^*$  as the reflexive transitive closure of  $\rightsquigarrow_{\text{cmd}}$ .*

## C.5 Assertion Language

In the previous sections we introduced heap commands that capture the APIs of data structures. In a similar way, we introduce heap predicates that describe their memory layout.

**Assumption C.5.1** (Heap Predicates). *We assume that there is a set of symbols  $P$ . Every  $p \in P$  represents a predicate characterising the heap. Further,*

$$\begin{array}{c}
\text{CMDRED-EVALCTXT} \\
\frac{h, c \rightsquigarrow_{\text{cmd}} h', c'}{h, E[c] \rightsquigarrow_{\text{cmd}} h', E[c']} \\
\\
\text{CMDRED-EVALCTXT-FAIL} \\
\frac{h, c \rightsquigarrow_{\text{cmd}} h', \text{error}}{h, E[c] \rightsquigarrow_{\text{cmd}} h', \text{error}} \\
\\
\text{CMDRED-IFTRUE} \qquad \text{CMDRED-IFFALSE} \\
h, \text{if True then } c_t \text{ else } c_f \rightsquigarrow_{\text{cmd}} h, c_t \qquad h, \text{if False then } c_t \text{ else } c_f \rightsquigarrow_{\text{cmd}} h, c_f \\
\\
\text{CMDRED-WHILE} \\
\frac{x \notin \text{freeVars}(c)}{h, \text{while } !\ell \text{ do } c \rightsquigarrow_{\text{cmd}} h, \text{let } x := !\ell \text{ in if } x \text{ then } (c; \text{while } !\ell \text{ do } c)} \\
\\
\text{CMDRED-FOR} \\
\frac{h, \text{for } x \text{ in } [n \text{ to } n'] \text{ do } c}{\rightsquigarrow_{\text{cmd}} h, \text{if } n \leq n' \text{ then } (c; \text{for } x \text{ in } [n + 1 \text{ to } n'] \text{ do } c)} \\
\\
\text{CMDRED-LET} \qquad \text{CMDRED-READHEAPLOC} \\
h, \text{let } x := v \text{ in } c \rightsquigarrow_{\text{cmd}} h, c[x \mapsto v] \qquad \frac{\ell \mapsto v \in h}{h, !\ell \rightsquigarrow_{\text{cmd}} h, v} \\
\\
\text{CMDRED-READHEAPLOC-FAIL} \qquad \text{CMDRED-ASSIGNHEAPLOC} \\
\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, !\ell \rightsquigarrow_{\text{cmd}} h, \text{error}} \qquad h \sqcup \{\ell \mapsto \_ \}, !\ell := v \rightsquigarrow_{\text{cmd}} h \sqcup \{\ell \mapsto v\}, () \\
\\
\text{CMDRED-ASSIGNHEAPLOC-FAIL} \\
\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, !\ell := v \rightsquigarrow_{\text{cmd}} h, \text{error}} \\
\\
\text{CMDRED-DESUGAR-HEAPCMDCALL} \\
\frac{\mathcal{I}_{\text{hcnds}}(hc) = (bc, (x_1, \dots, x_n))}{h, hc(v_1, \dots, v_n) \rightsquigarrow_{\text{cmd}} h, bc[(x_1, \dots, x_n) \mapsto (v_1, \dots, v_n)]}
\end{array}$$

Figure C.2: Command reduction rules.



$$\begin{aligned}
 & e \in \mathit{Exps} \\
 & p \in P \\
 & A \subseteq \mathcal{A} \\
 & \text{Index set } I \subseteq \mathbb{Z} \\
 \\
 a \in \mathcal{A} \quad ::= & \text{ True } \mid \text{ False } \mid e \mid \neg a \mid a \wedge a \mid a \vee a \mid a * a \mid a \rightarrow a \mid \\
 & e \mapsto e \mid p(\bar{e}) \mid \bigvee A \mid \square a
 \end{aligned}$$

(a) Assertion syntax.

$$\begin{aligned}
 a_1 \rightarrow a_2 & \quad := \quad \neg a_1 \vee a_2 \\
 a_1 \leftrightarrow a_2 & \quad := \quad (a_1 \rightarrow a_2) \wedge (a_2 \rightarrow a_1) \\
 \exists i \in I. a(i) & \quad := \quad \bigvee \{a(i) \mid i \in I\} \\
 \forall i \in I. a(i) & \quad := \quad \neg \exists i \in I. \neg a(i)
 \end{aligned}$$

(b) Syntactic sugar.

Figure C.3: Assertions.

we assume that there is a function  $\mathcal{I}_{\text{hpreds}} : P \rightarrow \mathcal{P}(\text{Heaps}^{\text{phys}} \times \text{Values}^*)$  that maps each heap predicate symbol to a predicate over heaps and value tuples.

**Definition C.5.2** (Assertions). We define the set of assertions  $\mathcal{A}$  according to the syntax presented in Figure C.3.

We omit the index set  $I$  in quantifications when its choice becomes clear from the context and write  $\exists i. a(i)$  and  $\forall i. a(i)$  instead of  $\exists i \in I. a(i)$  and  $\forall i \in I. a(i)$ ,

**Definition C.5.3.** *Assertion Model Relation* We define the assertion model relation  $\models_{\mathcal{A}} \subseteq \text{Heaps}^{\text{phys}} \times \mathcal{A}$  by recursion over the structure of assertions according to the rules presented in Fig. C.4.

**Definition C.5.4** (Free Variables (Assertions)). We define the notion of free variables for assertions analogously to that of commands (cf. C.4.3).

**Definition C.5.5** (Substitution (Assertions)). We define substitution for assertions analogously to substitution for commands (cf. C.4.4).

**Notation C.5.6** (Free Variables (Tuples)). For convenience we define the following notation for any tuple of commands and assertions  $(y_1, \dots, y_n) \in (\text{Cmds} \cup \mathcal{A})^*$ :

$$\text{freeVars}(y_1, \dots, y_n) := \text{freeVars}(y_1) \cup \dots \cup \text{freeVars}(y_n).$$

$h \models_{\mathcal{A}} \text{True}$	
$h \not\models_{\mathcal{A}} \text{False}$	
$h \models_{\mathcal{A}} e$	iff $\emptyset \models_{\mathcal{A}} \llbracket e \rrbracket$
$h \models_{\mathcal{A}} \neg a$	iff $h \not\models_{\mathcal{A}} a$
$h \models_{\mathcal{A}} a_1 \wedge a_2$	iff $h \models_{\mathcal{A}} a_1 \wedge h \models_{\mathcal{A}} a_2$
$h \models_{\mathcal{A}} a_1 \vee a_2$	iff $h \models_{\mathcal{A}} a_1 \vee h \models_{\mathcal{A}} a_2$
$h \models_{\mathcal{A}} a_1 * a_2$	iff $\exists h_1, h_2. h = h_1 \sqcup h_2$ $\quad \wedge h_1 \models_{\mathcal{A}} a_1 \wedge h_2 \models_{\mathcal{A}} a_2$
$h \models_{\mathcal{A}} a_1 \multimap a_2$	iff $\forall h_1. h_1 \cap h = \emptyset \wedge h_1 \models_{\mathcal{A}} a_1$ $\quad \rightarrow h_1 \sqcup h \models_{\mathcal{A}} a_2$
$h \models_{\mathcal{A}} \ell \mapsto v$	iff $\ell \mapsto v \in h$
$h \models_{\mathcal{A}} p(e_1, \dots, e_n)$	iff $\mathcal{I}_{\text{hpreds}}(p)(h, \llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket$
$h \models_{\mathcal{A}} \bigvee A$	iff $\exists a \in A. h \models_{\mathcal{A}} a$
$h \models_{\mathcal{A}} \Box a$	iff $\emptyset \models_{\mathcal{A}} a$

Figure C.4: Assertion model relation. We write  $h \not\models_{\mathcal{A}} a$  if  $h \models_{\mathcal{A}} a$  does not hold.

**Definition C.5.7** (Validity). *Let  $a \in \mathcal{A}$  be an assertion with  $\text{freeVars}(a) \cong \bar{x} = (x_i)_i$ . For each  $i = 1, \dots, n$ , let  $T_i$  be the type of variables  $x_i$  and let  $\bar{T} = (T_i)_i$ . We call assertion  $a$  valid if the following holds:*

$$\forall h. \forall \bar{v} \in \bar{T}. h \models_{\mathcal{A}} a[\bar{x} \mapsto \bar{v}]$$

We denote validity of  $a$  by writing  $\models_{\mathcal{A}} a$ .

## C.6 Memory Safety

**Definition C.6.1** (Memory Safety of Commands and Heaps). *We define the safety relation for commands  $\text{safe} \subseteq \text{Heaps}^{\text{phys}} \times \text{Cmds}$  as follows:*

*Let  $\bar{x} = (x_i)_i \cong \text{freeVars}(c)$  be the variables occurring freely in  $c$ . For each  $i$ , let  $T_i$  be the type of variable  $x_i$  and let  $\bar{T} = (T_i)_i$ . Then,*

$$\text{safe}(h, c) \Leftrightarrow \forall \bar{v} \in \bar{T}. \neg \exists h'. h, c[\bar{x} \mapsto \bar{v}], \rightsquigarrow_{\text{cmd}}^* h', \text{error}$$

*We say that a command  $c$  is safe under a physical heap  $h$  if  $\text{safe}(h, c)$  holds.*

We consider a command safe under a heap if its execution does not lead to a memory error. Note that a command's execution can get stuck without any memory error occurring. Such cases arise for not-well-typed commands such as

**if 13 then** ... For this work, we only consider well-typed programs. Hence, we do not care about cases in which a program gets stuck as long as no memory error occurs.

**Definition C.6.2** (Memory Safety of Commands and Assertions). *We define the safety relation for commands  $\text{safe}_A \subseteq \mathcal{A} \times \text{Cmds}$  as follows:*

*Let  $\bar{x} = (x_i)_i \cong \text{freeVars}(a, c)$  be the variables occurring freely in  $a$  and  $c$ . For each  $i$ , let  $T_i$  be the type of variable  $x_i$  and let  $\bar{T} = (T_i)_i$ . Then,*

$$\text{safe}_A(a, c) \iff \forall \bar{v} \in \bar{T}. \forall h. (h \vDash_A a[\bar{x} \mapsto \bar{v}] \Rightarrow \text{safe}(h, c[\bar{x} \mapsto \bar{v}]))$$

*We say that a command  $c$  is safe under a assertion  $a$  if  $\text{safe}_A(a, c)$  holds.*

**Notation C.6.3.** *We denote preconditions by  $M$ . Further, we aggregate preconditions  $M$  and programs  $c$  into tuples  $\{M\} c$ .*

## C.7 Verification Conditions

A common approach in program verification is to derive a verification condition  $vc$  from the program  $c$  and correctness property  $\phi$  in question [71, 143]. Instead of verifying the program directly, we prove the verification condition. In general,  $vc$  describes an over-approximation of all possible program behaviours. The process is sound iff truth of the verification condition indeed implies that our program is correct, i.e.,  $\vDash vc \Rightarrow c \vDash \phi$ . We proceed analogously during our study of completeness thresholds.

**Definition C.7.1** (Verification Condition). *We call an assertion  $vc \in \mathcal{A}$  a verification condition for  $\{M\} c$  if the following holds:*

$$\vDash_A vc \Rightarrow \text{safe}_A(M, c)$$

*We denote verification conditions by  $vc$ .*

**Definition C.7.2** (Precise Verification Conditions). *Let  $vc$  be a verification condition for  $\{M\} c$  and let  $x \in \text{freeVars}(M, c)$  be a free variable of type  $T$ . We call  $vc$  precise in  $x$  for  $\{M\} c$  if the following holds for every value  $v \in T$ :*

$$\text{safe}_A(M[x \mapsto v], c[x \mapsto v]) \Rightarrow \vDash_A vc[x \mapsto v]$$

**Assumption C.7.3** (Array Predicate). *We assume that there exists a predicate symbol  $\text{array} \in P$ . We further assume that  $\mathcal{I}_{\text{hpreds}}(\text{array}) \subseteq \text{Heaps}^{\text{phys}} \times \text{HObj}s \times \mathbb{N}$  is the minimal relation for which the following holds:*

$$\mathcal{I}_{\text{hpreds}}(\text{array})(h, a, s) \iff h \vDash_A \bigstar_{0 \leq i < s} (a, i) \mapsto \_$$

We identify the assertion  $\text{array}(a, s)$  with the assertion  $\bigstar_{0 \leq i < s} (a, i) \mapsto \_.$

## C.7.1 VC Generation

We use weakest liberal preconditions [71, 61] as verification conditions.

**Definition C.7.4.** *Assertion Lambdas* We define the set of assertion lambda terms as  $\mathcal{A}_\lambda := \{\lambda x. a \mid x \in \mathcal{X}, a \in \mathcal{A}\}$ . For convenience we define the notation  $\lambda\_. a := \lambda x_{fresh}. a$  for  $x_{fresh} \notin \text{freeVars}(a)$ .

We use assertion lambda terms  $\lambda r. Q$  to denote postconditions referring to a result value  $r$ .

**Definition C.7.5.** *Weakest Liberal Precondition* We define the weakest liberal precondition function  $wlp : \text{Cmds} \times \mathcal{A}_\lambda \rightarrow \mathcal{A}$  by recursion over the structure of commands as follows:

$$\begin{aligned}
wlp(e, \lambda r. Q) &:= Q[r \mapsto e], \\
wlp(\text{let } x := c_1 \text{ in } c_2, \lambda r. Q) &:= wlp(c_1, \lambda x. wlp(c_2, \lambda r. Q)), \\
wlp(\text{if } e \text{ then } c_t \text{ else } c_f, \lambda r. Q) &:= (e \rightarrow wlp(c_t, \lambda r. Q)) \\
&\quad \wedge (\neg e \rightarrow wlp(c_f, \lambda r. Q)), \\
wlp(\text{while } !e \text{ inv } I \text{ do } c, \lambda r. Q) &:= I \\
&\quad * \square(I \text{ -* } e \mapsto \_) \\
&\quad * \square(I \wedge e \mapsto \text{True} \\
&\quad \quad \text{ -* } wlp(c, \lambda\_. I)) \\
&\quad * (I \wedge e \mapsto \text{False} \text{ -* } Q), \\
wlp(\text{for } x \text{ in } [e_1 \text{ to } e_2] \text{ inv } I \text{ do } c, &:= I \\
\lambda r. Q) &\quad * \square(\forall x. e_1 \leq x < e_2 \wedge I \\
&\quad \quad \text{ -* } wlp(c, \lambda\_. I)) \\
&\quad * (I \text{ -* } Q), \\
wlp(!e, \lambda r. Q) &:= \exists y. e \mapsto y \wedge Q[r \mapsto y] \\
&\quad \text{where } y \notin \text{freeVars}(e, Q), \\
wlp(!e_1 := e_2, \lambda r. Q) &:= e_1 \mapsto \_ * (e_1 \mapsto e_2 \text{ -* } Q), \\
wlp(\text{hc}(\bar{e}), \lambda r. Q) &:= wlp(\text{bc}[\bar{x} \mapsto \bar{e}], \lambda r. Q) \\
&\quad \text{where } \mathcal{I}_{\text{hcmds}}(\text{hc}) = (\text{bc}, \bar{x})
\end{aligned}$$

Note that our weakest precondition for **for** loops does not allow the postcondition to depend on the knowledge that any work was done. This simplified precondition is sufficient to reason about the memory safety of array-traversing programs that do not allocate nor free memory.

## C.8 Completeness Thresholds

Now that we have an intuition for what a completeness threshold should be and for how we want to use it, let's formalise this intuition.

**Definition C.8.1** (Completeness Thresholds for Quantified Assertions). *Let  $Q \in \{\forall, \exists\}$  be a quantifier and let  $Qx \in T$ .  $a$  be a quantified assertion. Further, let  $Q \subseteq T$  be a restriction of the domain of  $x$ . We call  $Q$  a completeness threshold for  $a$  if the following holds:*

$$\models_A Qx \in Q. a \quad \Rightarrow \quad \models_A Qx \in T. a.$$

**Definition C.8.2** (Completeness Thresholds for Programs). *Let  $T$  be the type of  $x$  in  $\{M\} c$  and let  $Q \subseteq T$  be a restriction of this type. We call  $Q$  a completeness threshold for  $x$  in  $\{M\} c$  if the following holds:*

$$\forall v \in Q. \text{safe}_A(M[x \mapsto v], c[x \mapsto v]) \quad \Rightarrow \quad \forall v \in T. \text{safe}_A(M[x \mapsto v], c[x \mapsto v])$$

**Theorem C.8.3** (Soundness). *Let  $\forall x \in T. vc$  be a verification condition for  $\{M\} c$  and let  $Q \subseteq T$  be a completeness threshold for  $vc$ . Further, let  $vc$  be precise in  $x$  for  $\{M\} c$ . Then  $Q$  is a completeness threshold for  $\{M\} c$ .*

*Proof.* Assume that  $\forall v \in Q. \text{safe}_A(M[x \mapsto v], c[x \mapsto v])$  holds. We have to prove that the program is safe for the unrestricted domain, i.e.,  $\text{safe}_A(M, c)$ .

Together with the precision of  $vc$  in  $x$ , this assumption implies  $\forall v \in Q. \models_A vc[x \mapsto v]$  (cf. Def. C.7.2), which is equivalent to  $\models_A \forall x \in Q. vc$ .

$Q$  is a completeness threshold for  $vc$ . According to Def. C.8.1, this means  $\models_A \forall x \in T. vc$ . Since the latter is a verification condition for  $\{M\} c$ , proposition  $\text{safe}_A(M, c)$  holds by Def. C.7.1. Hence,  $Q$  is a completeness threshold for  $\{M\} c$ .  $\square$

Finding completeness threshold for a non-precise verification condition  $\forall x. vc$  does not allow us to conclude that we found a completeness threshold for the actual program  $\{M\} c$ . However, as long as our ultimate goal is to verify the program by proving an assertion  $A$  that is at least as strong as our verification condition, i.e.,  $A \Rightarrow \forall x. vc$ , we can leverage the completeness threshold. Hence, it makes sense to first concentrate on completeness thresholds for verification conditions. Later, we can try to relate our results to completeness thresholds for programs.

**Lemma C.8.4.** *Let  $\forall x \in T. a$  be an assertion. Let  $Q \subseteq T$  be such that*

$$\forall v_1, v_2 \in Q. (\models_A a[x \mapsto v_1] \Leftrightarrow \models_A a[x \mapsto v_2]).$$

Then, for every  $q \in Q$  it holds that

$$\vDash_A \forall x \in (T \setminus Q) \cup \{q\}. a \quad \Leftrightarrow \quad \vDash_A \forall x \in T. a.$$

*Proof.* Let  $q \in Q$ . With the assumption from the lemma, we get

$$\forall v \in Q. (\vDash_A a[x \mapsto q] \Leftrightarrow \vDash_A a[x \mapsto v])$$

and hence

$$\vDash_A a[x \mapsto q] \Leftrightarrow \vDash_A \forall x \in Q. a$$

Further,  $\forall x \in T. a$  is valid iff both  $\forall x \in (T \setminus Q). a$  and  $\forall x \in Q. a$  are valid. As we can reduce validity of the latter to validity of  $a[x \mapsto q]$ , we get

$$\begin{aligned} & (\vDash_A \forall x \in (T \setminus Q). a) \wedge (\vDash_A a[x \mapsto q]) \\ \Leftrightarrow & (\vDash_A \forall x \in (T \setminus Q). a) \wedge (\vDash_A \forall x \in Q. a) \\ \Leftrightarrow & \vDash_A \forall x \in T. a. \end{aligned}$$

□

**Corollary C.8.5.** *Let  $\forall x \in T. a$  be an assertion. Let  $Q \subseteq T$  be such that*

$$\forall v_1, v_2 \in Q. (\vDash_A a[x \mapsto v_1] \Leftrightarrow \vDash_A a[x \mapsto v_2]).$$

*Then, for every  $q \in Q$ , the set  $(T \setminus Q) \cup \{q\}$  is a completeness threshold for assertion  $a$ .*

*Proof.* Follows from Lem. C.8.4 and Def. C.8.1. □

Consider a verification condition  $\forall x \in T. vc$  and suppose we are interested in a completeness threshold for  $x$ . By definition, the threshold is a restriction of  $x$ 's domain, i.e.,  $Q \subseteq T$ . The lemmas and corollary above show us that one way forward is to identify a validity-preserving subset of  $T$ . That is, we need to look for a subset  $Q \subseteq T$  of the domain within which the concrete choice for  $x$  does not affect the validity of the verification condition. Once we got this, we can collapse  $Q$  to any representative  $q \in Q$  and we found our completeness threshold  $(T \setminus Q) \cup \{q\}$ .

Notice that validity preservation of domain restrictions is a transitive property. Hence, we can easily turn the search for completeness thresholds for a fixed variable into an iterative approach.

**Lemma C.8.6** (Transitivity of Completeness Thresholds for Fixed Variable). *Let  $Q_0, Q_1, Q_2$  be sets with  $Q_2 \subseteq Q_1 \subseteq Q_0$ . Let  $a_i = \forall x \in Q_i. a$  be assertions. Let  $Q_1$  and  $Q_2$  be completeness thresholds for  $x$  in  $a_0$  and  $a_1$ , respectively. Then,  $Q_2$  is also a completeness threshold for  $x$  in  $a_0$ .*

*Proof.* Since  $Q_1$  is a completeness threshold for  $x$  in  $a_0 = \forall x \in Q_0. a$ , we get

$$\models_{\mathbf{A}} \forall x \in Q_1. a \quad \Rightarrow \quad \models_{\mathbf{A}} \forall x \in Q_0. a.$$

Since  $Q_2$  is a completeness threshold for  $x$  in  $a_1 = \forall x \in Q_1. a$ , we get

$$\models_{\mathbf{A}} \forall x \in Q_2. a \quad \Rightarrow \quad \models_{\mathbf{A}} \forall x \in Q_1. a.$$

That is,

$$\models_{\mathbf{A}} \forall x \in Q_2. a \quad \Rightarrow \quad \models_{\mathbf{A}} \forall x \in Q_1. a \quad \Rightarrow \quad \models_{\mathbf{A}} \forall x \in Q_0. a$$

and hence  $Q_2 \subset Q_0$  is a completeness threshold for  $x$  in  $a_0 = \forall x \in Q_0. a$   $\square$

**Corollary C.8.7.** *Let  $T$  be a type and let  $(Q_i)_i$  be a family of sets with  $Q_0 = T$  and  $Q_{i+1} \subseteq Q_i$ . Let  $(a_i)_i = (\forall x \in Q_i. a)_i$  be a family of assertions such that each  $Q_{i+1}$  is a completeness threshold for  $x$  in  $a_i$ . Then, each  $Q_i$  is a completeness threshold for  $x$  in  $a_0 = \forall x \in T. a$ .*

*Proof.* Follows from Lem. C.8.6 by induction.  $\square$

Consider a program that traverses an array  $a$  of size  $s_a$  and an array  $b$  of size  $s_b$ . When we analyse the verification condition of this program we find that it contains distinct parts that describe memory safety of the accesses to array  $a$  and distinct parts for the accesses to  $b$ . Since both arrays describe separate parts of the heap, we can bring the verification condition into a form that reflects this. Thereby, we get a formula of the form  $vc \equiv vc_a * vc_b$  where  $vc_a$  and  $vc_b$  describe memory safety in respect to  $a$  and  $b$ , respectively.

Suppose, we want to find a completeness threshold for  $s_a$ . In some cases, the manipulation of both arrays is entangled which means that  $s_a$  potentially affects the validity of  $vc_b$ . In such a case, we have no choice but to analyse the entire formula to find our completeness threshold. However, often that's not the case and  $s_a$  only shows up in the subformula  $vc_a$  that actually concerns array  $a$ . In this case, it is sufficient to analyse  $vc_a$  in order to find a completeness threshold for  $s_a$ .

**Lemma C.8.8** (Elimination). *Let  $a, a_x, a'$  be assertions with  $\forall x \in T. a \equiv \forall x \in T. a_x * a'$ . Suppose the choice of  $x$  does not affect the validity of  $a'$ , i.e.,*

$$\forall v \in T. (\models_{\mathbf{A}} a' \Leftrightarrow \models_{\mathbf{A}} a'[x \mapsto v]).$$

*Let  $Q \subseteq T$  be a completeness threshold for  $x$  in  $\forall x \in T. a_x$ . Then,  $Q$  is also a completeness threshold for  $x$  in  $\forall x \in T. a$ .*

*Proof.* Since  $Q$  is a completeness threshold for  $x$  in  $\forall x \in Q. a_x$ , we get

$$\begin{aligned} & \models_{\mathbf{A}} \forall x \in Q. a_x * a' \\ \Rightarrow & \models_{\mathbf{A}} (\forall x \in Q. a_x) * (\forall x \in Q. a') \\ \Rightarrow & \models_{\mathbf{A}} (\forall x \in T. a_x) * (\forall x \in Q. a'). \end{aligned}$$

By using the assumption that the choice of  $x$  does not affect the validity of  $a'$  we can conclude

$$\begin{aligned} & \models_{\mathbf{A}} (\forall x \in T. a_x) * (\forall x \in Q. a') \\ \Rightarrow & \models_{\mathbf{A}} (\forall x \in T. a_x) * (\forall x \in T. a') \\ \Rightarrow & \models_{\mathbf{A}} \forall x \in T. a_x * a'. \end{aligned}$$

□

**Corollary C.8.9** (VC Slicing). *Let  $a, a_x, a'$  be assertions with  $\forall x \in T. a \equiv \forall x \in T. a_x * a'$ . Suppose  $x$  is not free in  $a'$ . Let  $Q \subseteq T$  be a completeness threshold for  $x$  in  $\forall x \in T. a_x$ . Then,  $Q$  is also a completeness threshold for  $x$  in  $\forall x \in T. a$ .*

*Proof.* Follows from Lem. C.8.8. □

## C.8.1 Iteratively Extracting Completeness Thresholds

What we saw so far, gives us the tools to define an iterative process to extract completeness thresholds.

**Workflow C.8.10.** *Let  $\{M\} c$  be a program with variables  $x_1, \dots, x_n$  for which we would like to extract completeness thresholds.*

1. *Compute a verification condition for  $\{M\} c$ , e.g., by using weakest preconditions. The result has the form  $\forall \bar{x} \in \bar{T}. vc$ .*
2. *Iteratively extract completeness thresholds for each  $x_i$ .*  
For all  $i \in \{1, \dots, n\}$ :

(a) *Let  $\bar{Q} \subseteq \bar{T}$  be the completeness thresholds extracted so far. (Initially  $\bar{Q} = \bar{T}$ .)*

(b) *Bring the verification condition into the form*  

$$\forall \bar{x} \in \bar{Q}. vc \equiv \forall \bar{x} \in \bar{Q}. \bigstar_{0 \leq j \leq m} vc_j.$$

(c) *Identify a subformula  $vc'$  whose validity is not affected by the choice of  $x_i$ . Bring the verification condition into the form*  

$$\forall \bar{x} \in \bar{Q}. vc \equiv \forall \bar{x} \in \bar{Q}. vc_i * vc'.$$
*In the remaining steps it suffices to analyse  $\forall \bar{x} \in \bar{Q}. vc_i$ .*



- (d) *Examine  $vc_i$  and extract a completeness threshold. This can either be done purely manually or by identifying patterns for which we previously proved that we can extract completeness thresholds. This step yields a completeness threshold  $Q'_i$  which, in the worst case, does not yield an improvement, i.e.,  $Q'_i = Q_i$ .*
- (e) *Repeat this process iteratively until the extracted completeness threshold does not improve in respect to the last iteration.*

## C.8.2 Iterating over Arrays

In the following, we study patterns encountered in verification conditions of programs that iterate over arrays. The goal of this section is to formulate reusable lemmas that allow us to automate the extraction of completeness thresholds.

**Lemma C.8.11** (VC-CT for Bounded, Unconditional Array Access). *Let  $Z, a, b \in \mathbb{Z}$  be constants and  $vc(s) = \forall i \in \mathbb{Z}. a \leq i < s + b \rightarrow 0 \leq i + Z < s$ . Then, for every  $q \in \mathbb{Z}$  with  $q > a - b$  it holds that*

$$\models_{\mathbf{A}} \forall s \in \mathbb{Z}. vc(s) \quad \Leftrightarrow \quad \models_{\mathbf{A}} vc(q)$$

*That is,  $\{q\}$  is a completeness threshold for  $vc(s)$ .*

*Proof.* For  $s \leq a - b$  it holds  $a \leq i < s + b \equiv \text{False}$  and hence  $vc(s) \equiv \text{True}$ .

For  $s > a - b$  we get:

$$\begin{aligned} vc(s) &\equiv \forall i. a \leq i < s + b \rightarrow 0 \leq i + Z < s \\ &\equiv \forall i. (a \leq i < s + b \rightarrow 0 \leq i + Z) \wedge (a \leq i < s + b \rightarrow i + Z < s) \\ &\equiv \forall i. (a \leq i \rightarrow 0 \leq i + Z) \wedge (i < s + b \rightarrow i + Z < s) \\ &\equiv \forall i. (a \leq i \rightarrow 0 \leq i + Z) \wedge \forall i. (i < s + b \rightarrow i + Z < s) \\ &\equiv \forall i. (a \leq i \rightarrow 0 \leq i + Z) \wedge \forall i. (i < b \rightarrow i + Z < 0) \\ &=: vc^+ \end{aligned}$$

$s$  does not occur in  $vc^+$  and, hence, the concrete value of  $s$  does not impact the validity of  $vc^+$ . By Lem. C.8.5, for any choice of  $q \in \mathbb{Z}$  with  $q > l - r$ , the set  $\{q\}$  is a completeness threshold for  $vc(s)$ .  $\square$



# Appendix D

## Generalising Completeness Thresholds

In the previous chapters § 4 and § C we studied completeness thresholds for memory safety proofs. In this chapter we drop this limitation and show that we can straightforwardly generalise the theoretic foundations of our work to arbitrary correctness properties.

### D.1 Basic Definitions

Since we are generalising our work formalised in appendix § C, we reuse most of that chapter’s definitions throughout this chapter. This includes but it not limited to the definitions presented in § C.2. Deviations from previous definitions are explicitly redefined.

**Definition D.1.1** (Equivalence Classes). *Let  $R \subseteq X \times X$  be an equivalence relation. For any  $x \in X$ , we define the notation*

$$\langle x \rangle_R := \{x' \in X \mid R(x, x')\} \in X/R$$

*We omit  $R$  whenever it is clear from the context and write  $\langle x \rangle$  instead of  $\langle x \rangle_R$ .*

**Definition D.1.2** (Equivalence Relation Refinement). *Let  $R, S \subseteq X \times X$  be equivalence relations. We call  $S$  a refinement of  $R$  if the following holds:*

$$\forall x \in X. \langle x \rangle_S \subseteq \langle x \rangle_R.$$

**Definition D.1.3** (Cover). *Let  $X$  be a set. We call a family of subsets  $(Y_i)_{i \in I}$  a cover of  $X$  if*

$$\bigcup_{i \in I} Y_i = X.$$

**Definition D.1.4** (Equivalence Cover). *Let  $R \subseteq X \times X$  be an equivalence relation over a set  $X$ . We call a subset  $Y \subseteq X$  an  $R$ -cover of  $X$  if the family  $(\langle y \rangle_R)_{y \in Y}$  is a cover of  $X$ . When  $R$  is clear from the context or irrelevant, we also call  $Y$  an equivalence cover of  $X$ .*

**Notation D.1.5** (Logical Term Operator Variables). *We denote variables ranging over quantifiers  $\{\forall, \exists\}$  by  $\nabla$ . Likewise, we denote variables ranging over the logical connectors  $\{\wedge, \vee\}$  by  $\oplus$ .*

## D.2 Generalised Completeness Thresholds

In this chapter we treat correctness as an abstract property of programs. Hence, the following definition of correctness encompasses both safety properties like memory safety and liveness properties such as termination. Moreover, it also covers more complex functional correctness properties.

Without loss of generality, we can assume that any correctness property  $\phi$  is given in prenex normal form  $\nabla_1 \dots \nabla_n. \psi$  for quantifiers  $\nabla \in \{\forall, \exists\}$ . To simplify the notation in subsequent definitions and lemmas, we call the quantifier-free core a *correctness predicate* and denote it by  $\psi$ . Meanwhile, we call the full correctness specification, where every variable except for the target program is explicitly quantified, a *correctness property* and denote it by  $\phi$ .

**Definition D.2.1** (Correctness Predicate). *Let  $X_1, \dots, X_n$  be variable domains. Let  $\psi(c, x_1, \dots, x_n) \subseteq \text{Cmds} \times X_1 \times \dots \times X_n$  be a predicate over commands and typed values. We call  $\psi$  a correctness predicate.*

$\models \psi(c, v_1, \dots, v_n)$  means that every run of  $c[x_1, \dots, x_n \mapsto v, \dots, v_n]$  that takes place under the conditions described by  $\psi$  is correct in some sense described by  $\psi$ .

**Definition D.2.2** (Correctness Property). *Let  $\psi(c, x_1, \dots, x_n) \subseteq \text{Cmds} \times X_1 \times \dots \times X_n$  be a correctness predicate and let  $\nabla_1, \dots, \nabla_n \in \{\forall, \exists\}$  be quantifiers. We call a predicate  $\phi(c) \subseteq \text{Cmds}$  of the form*

$$\phi(c) \quad := \quad \nabla_1 x_1 \in X_1. \dots \nabla_n x_n \in X_n. \psi(c, x_1, \dots, x_n)$$

a correctness property. We write  $c \models \phi$  to express that  $\phi(c)$  holds.

The following definitions are straightforward generalisations of Def. C.7.1, Def C.7.2, Def C.8.1, Def C.8.2.

**Definition D.2.3** (Verification Condition). *We call an assertion  $vc \in \mathcal{A}$  a verification condition for a correctness property  $\phi$  if the following holds:*

$$\models_{\mathbf{A}} vc \quad \Rightarrow \quad c \models \phi$$

We call a VC precise in  $x$  for some correctness property  $\phi$  if it does not over-approximate the influence of  $x$  on  $\phi$ . Note that precision is a significantly weaker property than stating that the VC as a whole does not over-approximate.

**Definition D.2.4** (Precise Verification Conditions). *Let  $\nabla x \in X$ .  $vc$  be a verification condition for a correctness property  $\nabla x \in X$ .  $\psi$  and a command  $c$ . We call  $vc$  precise in  $x$  for  $\phi$  and  $c$  if the following holds for every value  $v \in X$ :*

$$c \models \psi[x \mapsto v] \quad \Rightarrow \quad \models_{\mathbf{A}} vc[x \mapsto v]$$

*We also say that  $vc$  does not over-approximate  $\psi$  and  $c$  with respect to  $x$ .*

**Definition D.2.5** (Completeness Thresholds for Quantified Assertions). *Let  $\nabla x \in X$ .  $a$  be a quantified assertion. Further, let  $Q \subseteq X$  be a subdomain. We call  $Q$  a completeness threshold for  $x$  in  $a$  if the following holds:*

$$\models_{\mathbf{A}} \nabla x \in Q. a \quad \Rightarrow \quad \models_{\mathbf{A}} \nabla x \in X. a.$$

**Observation D.2.6** (Trivial CTs for  $\exists$ -Assertions). *For existentially quantified assertions of the form  $\exists x \in X$ .  $a$  any subdomain  $Y \subseteq X$  is a completeness threshold for  $x$ .*

**Definition D.2.7** (Completeness Thresholds for Correctness Properties and Programs). *Let  $\nabla x \in X$ .  $\psi$  be a correctness property and  $c$  a comand. We call a subdomain  $Q \subseteq X$  a completeness threshold for  $x$  in  $\psi$  if the following holds:*

$$c \models \nabla x \in Q. \psi \quad \Rightarrow \quad c \models \nabla x \in X. \psi$$

**Observation D.2.8** (Trivial CTs for  $\exists$ -Correctness-Properties). *For existentially quantified correctness properties of the form  $\exists x \in X$ .  $\psi$  any subdomain  $Y \subseteq X$  is a completeness threshold for  $x$ .*

**Notation D.2.9.** *We denote completeness thresholds by  $Q$  and singleton ones by  $\{q\}$ .*

**Terminology D.2.10.** *We call a completeness threshold minimal if it is minimal wrt.  $\subseteq$ .*

The following theorem states that precision in some variable  $x$  is a sufficiently strong requirement to ensure the soundness of any CT for  $x$  extracted via the precise VC.

**Theorem D.2.11** (Soundness). *Let  $\nabla x \in X$ .  $vc$  be a verification condition for a correctness property  $\nabla x \in X$ .  $\psi$  and a command  $c$ . Let  $Q \subseteq X$  be a completeness threshold for  $x$  in  $vc$  and let  $vc$  be precise in  $x$ . Then,  $Q$  is also a completeness threshold for  $x$  in  $\psi$ .*

*Proof.* We can summarize the proof with this commuting diagram:

$$\begin{array}{ccc}
 c \models \nabla x \in Q. \psi & \xrightarrow{Q \text{ is CT for } x \text{ in } \psi, \text{ cf. Def. D.2.7}} & c \models \nabla x \in X. \psi \\
 \Downarrow \text{vc is precise in } x, \text{ cf. Def. D.2.4} & & \Uparrow \\
 \nabla v \in Q. (\models_{\mathbf{A}} vc[x \mapsto v]) & & \text{vc is VC, Def. D.2.3} \\
 \Updownarrow & & \Uparrow \\
 \models_{\mathbf{A}} \nabla x \in Q. vc & \xrightarrow{Q \text{ is CT for } x \text{ in } vc, \text{ cf. Def. D.2.5}} & \models_{\mathbf{A}} \nabla x \in X. vc
 \end{array}$$

Assume that  $c \models \nabla x \in Q. \psi$  holds. We have to show that the correctness predicate  $\psi$  holds for the unrestricted domain  $X$ , i.e.,  $c \models \nabla x \in X. \psi$ .

By our assumption and the precision of  $vc$  in  $x$  (cf. Def. D.2.4), we get  $\nabla v \in Q. (\models_{\mathbf{A}} vc[x \mapsto v])$ , which is equivalent to  $\models_{\mathbf{A}} \nabla x \in Q. vc$ . Since  $Q$  is a completeness threshold for  $x$  in  $vc$  (cf. Def. D.2.5), the latter implies  $\models_{\mathbf{A}} \nabla x \in X. vc$ . Further,  $\nabla x \in X. vc$  is a verification condition for  $\nabla x \in X. \psi$ . Per Def. D.2.3, this means that  $c \models \nabla x \in X. \psi$ .  $\square$

**Lemma D.2.12** (Subdomain Reduction). *Let  $\nabla x \in X$ .  $a$  be an assertion. Let  $Y \subseteq X$  be such that*

$$\nabla v_1, v_2 \in Y. (\models_{\mathbf{A}} a[x \mapsto v_1] \Leftrightarrow \models_{\mathbf{A}} a[x \mapsto v_2]).$$

*Then, for every  $y \in Y$  it holds that*

$$\models_{\mathbf{A}} \nabla x \in (X \setminus Y) \cup \{y\}. a \Leftrightarrow \models_{\mathbf{A}} \nabla x \in X. a.$$

*Proof.* Let  $y \in Y$ . With the assumption from the lemma, we get

$$\nabla v \in Y. (\models_{\mathbf{A}} a[x \mapsto y] \Leftrightarrow \models_{\mathbf{A}} a[x \mapsto v])$$

and hence

$$\models_{\mathbf{A}} a[x \mapsto y] \Leftrightarrow \models_{\mathbf{A}} \nabla x \in Y. a$$

If  $\nabla = \forall$ , we get that  $\nabla x \in X. a$  is valid iff both  $\forall x \in (X \setminus Y). a$  and  $\forall x \in Y. a$  are valid. In this case we define  $\oplus := \wedge$ . If  $\nabla = \exists$ , we get that  $\nabla x \in X. a$  is valid iff  $\exists x \in (X \setminus Y). a$  or  $\exists x \in Y. a$  is valid. In this case we define  $\oplus := \vee$ .

As we can reduce validity of the  $\nabla x \in Y. a$  to validity of  $a[x \mapsto y]$ , we get

$$\begin{aligned} & (\models_{\mathbf{A}} \nabla x \in (X \setminus Y). a) \oplus (\models_{\mathbf{A}} a[x \mapsto y]) \\ \iff & (\models_{\mathbf{A}} \nabla x \in (X \setminus Y). a) \oplus (\models_{\mathbf{A}} \nabla x \in Y. a) \\ \iff & \models_{\mathbf{A}} \nabla x \in X. a. \end{aligned}$$

□

**Corollary D.2.13** (CT Extraction). *Let  $\nabla x \in X. a$  be an assertion. Let  $Y \subseteq X$  be such that*

$$\forall v_1, v_2 \in Y. (\models_{\mathbf{A}} a[x \mapsto v_1] \iff \models_{\mathbf{A}} a[x \mapsto v_2]).$$

*Then, for every  $y \in Y$ , the set  $(X \setminus Y) \cup \{y\}$  is a completeness threshold for assertion  $a$ .*

*Proof.* Follows from the Subdomain Reduction Lem. D.2.12 and Def. D.2.5. □

This corollary shows that we can extract CTs by factoring out validity preserving subdomains. Later on we will see that it is indeed sufficient to factor out equivalence classes that capture all potential errors. For now, it is sufficient to note that we can reduce CTs by dropping elements that provably do not cause any errors.

**Corollary D.2.14.** *Let  $\nabla x \in X. a$  be an assertion. Let  $Y \subseteq X$  be such that*

$$\forall y \in Y. \models_{\mathbf{A}} a[x \mapsto y].$$

*Then,  $(X \setminus Y)$  is a completeness threshold for assertion  $a$ .*

*Proof.* Let  $y \in Y$ . By Cor. D.2.13, the set  $(X \setminus Y) \cup \{y\}$  is a CT. By assumption,  $\models_{\mathbf{A}} a[x \mapsto y]$ . Hence we get

$$\nabla x \in (X \setminus Y). a \quad \Rightarrow \quad \nabla x \in (X \setminus Y \cup \{y\}). a \quad \Rightarrow \quad \nabla x \in X. a.$$

□

Intuitively, CTs are subdomains that cover all possible types of potential errors. In the following, we formalise this intuition by describing them as covers of (refinements of) a specific validity equivalence relation  $\approx_a^x$ .

**Definition D.2.15** (Equivalence Relation for Assertion Arguments). *Let  $a$  be an assertion,  $x$  a variable with domain  $X$ . We define the relation  $\approx_a^x \subseteq X \times X$  as follows:*

$$v_1 \approx_a^x v_2 \iff (\models_A a[x \mapsto v_1] \Leftrightarrow \models_A a[x \mapsto v_2])$$

We omit the variable  $x$  whenever it is clear from the context and write  $\approx_a$ .

**Corollary D.2.16.** *For any choice of  $a$  and  $x$ , the relation  $\approx_a^x$  is an equivalence relation.*

**Theorem D.2.17** (Equivalence Cover CT). *Let  $\nabla x \in X. a$  be an assertion and let  $Y \subseteq X$  be a  $\approx_a^x$ -cover of  $X$ . Then,  $Y$  is a completeness threshold for  $x$  in  $a$ .*

*Proof.* Suppose the claim does not hold. By Obs. D.2.6 we know that  $\nabla = \forall$ . Further, it must hold that  $\models_A \forall x \in Y. a$  and there must exist a  $v_X \in X$  for which  $\not\models_A a[x \mapsto v_X]$ .

Since  $Y$  is a  $\approx_a^x$ -cover of  $X$ , there exists a  $y \in Y$  with  $v_X \in \langle y \rangle_{\approx_a^x}$ . By the definition of  $\approx_a^x$  we get  $\not\models_A a[x \mapsto y]$ , which contradicts  $\models_A \forall x \in Y. a$ .  $\square$

**Corollary D.2.18** (Refinement Cover CT). *Let  $\nabla x \in X. a$  be an assertion, let  $\sim$  be a  $\approx_a^x$ -refinement and let  $Y \subseteq X$  be a  $\sim$ -cover of  $X$ . Then,  $Y$  is a completeness threshold for  $x$  in  $a$ .*

*Proof.* Since  $\sim$  is a  $\approx_a^x$ -refinement, every  $\sim$ -cover of  $X$  is also a  $\approx_a^x$ -cover of  $X$ . The rest follows by Theo. D.2.17.  $\square$

**Definition D.2.19** (VC Error Equivalence Class). *Let  $\nabla x \in X. a$  be an assertion, let  $\sim$  be a  $\approx_a^x$ -refinement. For any  $v \in X$ , we call  $\langle v \rangle_{\sim}$  a VC error equivalence class if  $\not\models_A a[x \mapsto v]$ .*

**Theorem D.2.20** (Error Cover CT). *Let  $\nabla x \in X. a$  be an assertion, let  $\sim$  be a  $\approx_a^x$ -refinement and let  $Y \subseteq X$  cover all  $\sim$ -error equivalence classes. Then,  $Y$  is a completeness threshold for  $x$  in  $a$ .*

*Proof.* Let  $\langle e_1 \rangle_{\sim}, \dots, \langle e_n \rangle_{\sim}$  be all the  $\sim$ -error equivalence classes. Without loss of generality, we can choose the representatives such that  $e_i \in Y$  for all  $i$ . Since  $\sim$  is a  $\approx_a^x$ -refinement, their union is exactly the  $\approx_a^x$ -error equivalence class  $E$ . (Note that  $\approx_a^x$  is an equivalence relation with only two equivalence classes.) Let  $s \in X \setminus E$  be an element for which  $a$  holds. Then  $\{e_1, \dots, e_n, s\}$  is a  $\approx_a^x$ -cover. By Theo. D.2.17 the set  $\{e_1, \dots, e_n, s\}$  is a CT. By Cor. D.2.14 we know that we can drop elements for which the VC holds. It follows that  $\{e_1, \dots, e_n\}$  is a CT. As  $\{e_1, \dots, e_n\} \subseteq Y$ , we also know that  $Y$  is a CT.  $\square$



**Lemma D.2.21** (Transitivity of Completeness Thresholds for Fixed Variable). *Let  $X_0, X_1, X_2$  be sets with  $X_0 \supseteq X_1 \supseteq X_2$ . Let  $a_i = \nabla x \in X_i. a$  be assertions. Let  $X_1$  and  $X_2$  be completeness thresholds for  $x$  in  $a_0$  and  $a_1$ , respectively. Then,  $X_2$  is also a completeness threshold for  $x$  in  $a_0$ .*

*Proof.* Since  $X_1$  is a completeness threshold for  $x$  in  $a_0 = \nabla x \in X_0. a$ , we get

$$\vDash_{\mathbf{A}} \nabla x \in X_1. a \quad \Rightarrow \quad \vDash_{\mathbf{A}} \nabla x \in X_0. a.$$

Since  $X_2$  is a completeness threshold for  $x$  in  $a_1 = \nabla x \in X_1. a$ , we get

$$\vDash_{\mathbf{A}} \nabla x \in X_2. a \quad \Rightarrow \quad \vDash_{\mathbf{A}} \nabla x \in X_1. a.$$

That is,

$$\vDash_{\mathbf{A}} \nabla x \in X_2. a \quad \Rightarrow \quad \vDash_{\mathbf{A}} \nabla x \in X_1. a \quad \Rightarrow \quad \vDash_{\mathbf{A}} \nabla x \in X_0. a$$

and hence  $X_2 \subseteq X_0$  is a completeness threshold for  $x$  in  $a_0 = \nabla x \in X_0. a$   $\square$

**Corollary D.2.22** (CT Refinement). *Let  $X$  be a set and let  $(Y_i)_i$  be a family of sets with  $Q_0 = X$  and  $Q_i \supseteq Q_{i+1}$ . Let  $(a_i)_i = (\nabla x \in Q_i. a)_i$  be a family of assertions such that each  $Q_{i+1}$  is a completeness threshold for  $x$  in  $a_i$ . Then, each  $Q_i$  is a completeness threshold for  $x$  in  $a_0 = \forall x \in T. a$ .*

*Proof.* Follows from the Transitivity Lem. D.2.21 by induction.  $\square$

**Lemma D.2.23** (Elimination). *Let  $a, a_x, a'$  be assertions with  $\nabla x \in X. a \equiv \nabla x \in X. a_x * a'$ . Suppose the choice of  $x$  does not affect the validity of  $a'$ , i.e.,*

$$\forall v \in X. (\vDash_{\mathbf{A}} a' \Leftrightarrow \vDash_{\mathbf{A}} a'[x \mapsto v]).$$

*Let  $Q \subseteq X$  be a completeness threshold for  $x$  in  $\forall x \in X. a_x$ . Then,  $Q$  is also a completeness threshold for  $x$  in  $\forall x \in X. a$ .*

*Proof.* Since  $Q$  is a completeness threshold for  $x$  in  $\nabla x \in X. a_x$ , we get

$$\begin{aligned} & \vDash_{\mathbf{A}} \nabla x \in Q. a_x * a' \\ \Rightarrow & \vDash_{\mathbf{A}} (\nabla x \in Q. a_x) * (\nabla x \in Q. a') \\ \Rightarrow & \vDash_{\mathbf{A}} (\nabla x \in X. a_x) * (\nabla x \in Q. a'). \end{aligned}$$

By using the assumption that the choice of  $x$  does not affect the validity of  $a'$  we can conclude

$$\begin{aligned} & \vDash_{\mathbf{A}} (\nabla x \in X. a_x) * (\nabla x \in Q. a') \\ \Rightarrow & \vDash_{\mathbf{A}} (\nabla x \in X. a_x) * (\nabla x \in X. a') \\ \Rightarrow & \vDash_{\mathbf{A}} \nabla x \in X. a_x * a'. \end{aligned}$$

$\square$

**Corollary D.2.24** (VC Slicing). *Let  $a, a_x, a'$  be assertions with  $\nabla x \in X$ .  $a \equiv \nabla x \in X. a_x * a'$ . Suppose  $x$  is not free in  $a'$ . Let  $Q \subseteq X$  be a completeness threshold for  $x$  in  $\nabla x \in X. a_x$ . Then,  $Q$  is also a completeness threshold for  $x$  in  $\nabla x \in X. a$ .*

*Proof.* Follows from the Elimination Lem. [D.2.23](#). □

# Bibliography

- [1] AWS encryption sdk for Dafny. URL <https://github.com/aws/aws-encryption-sdk-dafny#aws-encryption-sdk-for-dafny>. Accessed: 2023-10-18.
- [2] Mohammad Abdulaziz, Michael Norrish, and Charles Gretton. Formally verified algorithms for upper-bounding state space diameters. *Journal of Automated Reasoning*, 61:485–520, 2018.
- [3] Omar M. Alhawi, Herbert Rocha, Mikhail R. Gadelha, Lucas C. Cordeiro, and Eddie Batista de Lima Filho. Verification and refutation of C programs based on k-induction and invariant inference. *Int. J. Softw. Tools Technol. Transf.*, 23(2):115–135, 2021. doi: 10.1007/S10009-020-00564-1. URL <https://doi.org/10.1007/s10009-020-00564-1>.
- [4] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, 1985. ISSN 0020-0190. doi: [https://doi.org/10.1016/0020-0190\(85\)90056-0](https://doi.org/10.1016/0020-0190(85)90056-0). URL <https://www.sciencedirect.com/science/article/pii/0020019085900560>.
- [5] Benjamin Aminof, Swen Jacobs, Ayrat Khalimov, and Sasha Rubin. Parameterized model checking of token-passing systems. In Kenneth L. McMillan and Xavier Rival, editors, *Verification, Model Checking, and Abstract Interpretation - 15th International Conference, VMCAI 2014, San Diego, CA, USA, January 19-21, 2014, Proceedings*, volume 8318 of *Lecture Notes in Computer Science*, pages 262–281. Springer, 2014. doi: 10.1007/978-3-642-54013-4\_15. URL [https://doi.org/10.1007/978-3-642-54013-4\\_15](https://doi.org/10.1007/978-3-642-54013-4_15).
- [6] Paul Ammann and Jeff Offutt. *Introduction to Software Testing*. Cambridge University Press, 2008. ISBN 978-0-521-88038-1. doi: 10.1017/CBO9780511809163. URL <https://doi.org/10.1017/CBO9780511809163>.

- [7] Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.*, 23(5):657–683, September 2001. doi: 10.1145/504709.504712.
- [8] Andrew W. Appel, Robert Dockins, Aquinas Hobor, Lennart Beringer, Josiah Dodds, Gordon Stewart, Sandrine Blazy, and Xavier Leroy. *Semantic model and soundness of Verifiable C*, page 362–362. Cambridge University Press, 2014. doi: 10.1017/CBO9781107256552.046.
- [9] Irina Mariuca Asavaoae, Mihail Asavaoae, and Adrián Riesco. Slicing from formal semantics: Chisel - a tool for generic program slicing. *Int. J. Softw. Tools Technol. Transf.*, 20(6):739–769, 2018. doi: 10.1007/s10009-018-0500-y. URL <https://doi.org/10.1007/s10009-018-0500-y>.
- [10] Mohammad Awedh and F. Somenzi. Proving more properties with bounded model checking. In *CAV*, 2004.
- [11] Vasanth Balasundaram and Ken Kennedy. Compile-time detection of race conditions in a parallel program. In *Proceedings of the 3rd International Conference on Supercomputing*, pages 175–185, 1989.
- [12] Thomas Ball and Sriram K. Rajamani. The slam project: Debugging system software via static analysis. In *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '02*, page 1–3, New York, NY, USA, 2002. Association for Computing Machinery. ISBN 1581134509. doi: 10.1145/503272.503274. URL <https://doi.org/10.1145/503272.503274>.
- [13] Dirk Beyer and Matthias Dangl. Software verification with PDR: an implementation of the state of the art. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2020. doi: 10.1007/978-3-030-45190-5\_1. URL [https://doi.org/10.1007/978-3-030-45190-5\\_1](https://doi.org/10.1007/978-3-030-45190-5_1).
- [14] Dirk Beyer, Nian-Ze Lee, and Philipp Wendler. Interpolation and SAT-based model checking revisited: Adoption to software verification. *CoRR*, abs/2208.05046, 2022. doi: 10.48550/arXiv.2208.05046. URL <https://doi.org/10.48550/arXiv.2208.05046>.
- [15] Armin Biere, Alessandro Cimatti, Edmund M. Clarke, and Yunshan Zhu. Symbolic model checking without bdds. In *International Conference on Tools and Algorithms for Construction and Analysis of Systems*, 1999.

- [16] Per Bjesse and Koen Claessen. Sat-based verification without state space traversal. In Warren A. Hunt Jr. and Steven D. Johnson, editors, *Formal Methods in Computer-Aided Design, Third International Conference, FMCAD 2000, Austin, Texas, USA, November 1-3, 2000, Proceedings*, volume 1954 of *Lecture Notes in Computer Science*, pages 372–389. Springer, 2000. doi: 10.1007/3-540-40922-X\_23. URL [https://doi.org/10.1007/3-540-40922-X\\_23](https://doi.org/10.1007/3-540-40922-X_23).
- [17] Andreas Björklund, Thore Husfeldt, and Sanjeev Khanna. Approximating longest directed paths and cycles. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 222–233. Springer, 2004. doi: 10.1007/978-3-540-27836-8\_21. URL [https://doi.org/10.1007/978-3-540-27836-8\\_21](https://doi.org/10.1007/978-3-540-27836-8_21).
- [18] Stefan Blom, Saeed Darabi, Marieke Huisman, and Wytse Oortwijn. The VerCors tool set: Verification of parallel and concurrent software. In Nadia Polikarpova and Steve A. Schneider, editors, *Integrated Formal Methods - 13th International Conference, IFM 2017, Turin, Italy, September 20-22, 2017, Proceedings*, volume 10510 of *Lecture Notes in Computer Science*, pages 102–110. Springer, 2017. doi: 10.1007/978-3-319-66845-1\_7. URL [https://doi.org/10.1007/978-3-319-66845-1\\_7](https://doi.org/10.1007/978-3-319-66845-1_7).
- [19] 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 für Informatik, 2015. doi: 10.4230/LIPICs.ECOOP.2015.639.
- [20] Maha Boussabbeh, Mohamed Tounsi, Mohamed Mosbah, and Ahmed Hadj Kacem. Formal proofs of termination detection for local computations by refinement-based compositions. In Michael Butler, Klaus-Dieter Schewe, Atif Mashkoor, and Miklos Biro, editors, *Abstract State Machines, Alloy, B, TLA, VDM, and Z*, pages 198–212, Cham, 2016. Springer International Publishing.
- [21] Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership types for safe programming: Preventing data races and deadlocks. In *OOPSLA*, 2002. doi: 10.1145/582419.582440.
- [22] J. Boyland. Checking interference with fractional permissions. In *SAS*, 2003. doi: 10.1007/3-540-44898-5\_4.

- [23] Marius Bozga, Peter Habermehl, Radu Iosif, Filip Konečný, and Tomás Vojnar. Automatic verification of integer array programs. In Ahmed Bouajjani and Oded Maler, editors, *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings*, volume 5643 of *Lecture Notes in Computer Science*, pages 157–172. Springer, 2009. doi: 10.1007/978-3-642-02658-4\_15. URL [https://doi.org/10.1007/978-3-642-02658-4\\_15](https://doi.org/10.1007/978-3-642-02658-4_15).
- [24] Aaron R. Bradley. SAT-based model checking without unrolling. In Ranjit Jhala and David A. Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings*, volume 6538 of *Lecture Notes in Computer Science*, pages 70–87. Springer, 2011. doi: 10.1007/978-3-642-18275-4\_7. URL [https://doi.org/10.1007/978-3-642-18275-4\\_7](https://doi.org/10.1007/978-3-642-18275-4_7).
- [25] Aaron R. Bradley, Zohar Manna, and Henny B. Sipma. What’s decidable about arrays? In E. Allen Emerson and Kedar S. Namjoshi, editors, *Verification, Model Checking, and Abstract Interpretation, 7th International Conference, VMCAI 2006, Charleston, SC, USA, January 8-10, 2006. Proceedings*, volume 3855 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2006. doi: 10.1007/11609773\_28. URL [https://doi.org/10.1007/11609773\\_28](https://doi.org/10.1007/11609773_28).
- [26] S. Brookes and P. O’Hearn. Concurrent separation logic. *ACM SIGLOG News*, 3:47–65, 2016. doi: 10.1145/2984450.2984457.
- [27] Michael C. Browne, Edmund M. Clarke, and Orna Grumberg. Reasoning about networks with many identical finite state processes. *Inf. Comput.*, 81(1):13–31, 1989. doi: 10.1016/0890-5401(89)90026-6. URL [https://doi.org/10.1016/0890-5401\(89\)90026-6](https://doi.org/10.1016/0890-5401(89)90026-6).
- [28] Daniel Bundala, Joël Ouaknine, and James Worrell. On the magnitude of completeness thresholds in bounded model checking. *27th Annual IEEE Symposium on Logic in Computer Science*, pages 155–164, 2012.
- [29] Qinxiang Cao, Lennart Beringer, Samuel Gruetter, Josiah Dodds, and Andrew W. Appel. Vst-floyd: A separation logic tool to verify correctness of C programs. *J. Autom. Reason.*, 61(1-4):367–422, 2018. doi: 10.1007/S10817-018-9457-5. URL <https://doi.org/10.1007/s10817-018-9457-5>.
- [30] Supratik Chakraborty, Ashutosh Gupta, and Divyesh Unadkat. Verifying array manipulating programs with full-program induction. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and*

- Analysis of Systems - 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 22–39. Springer, 2020. doi: 10.1007/978-3-030-45190-5\_2. URL [https://doi.org/10.1007/978-3-030-45190-5\\_2](https://doi.org/10.1007/978-3-030-45190-5_2).
- [31] Zhe Chen, Yi Gu, Zhiqiu Huang, Jun Zheng, Chang Liu, and Ziyi Liu. Model checking aircraft controller software: a case study. *Software: Practice and Experience*, 45(7):989–1017, 2015.
- [32] Nathan Chong, Byron Cook, Konstantinos Kallas, Kareem Khazem, Felipe R. Monteiro, Daniel Schwartz-Narbonne, Serdar Tasiran, Michael Tautschnig, and Mark R. Tuttle. Code-level model checking in the software development workflow. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Practice, ICSE-SEIP '20*, page 11–20, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450371230. doi: 10.1145/3377813.3381347. URL <https://doi.org/10.1145/3377813.3381347>.
- [33] Koen Claessen and John Hughes. Quickcheck: a lightweight tool for random testing of haskell programs. In Martin Odersky and Philip Wadler, editors, *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18-21, 2000*, pages 268–279. ACM, 2000. doi: 10.1145/351240.351266. URL <https://doi.org/10.1145/351240.351266>.
- [34] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In E. Allen Emerson and Aravinda Prasad Sistla, editors, *Computer Aided Verification*, pages 154–169, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- [35] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In E. Allen Emerson and Aravinda Prasad Sistla, editors, *Computer Aided Verification*, pages 154–169, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- [36] Edmund M. Clarke. Model checking – my 27-year quest to overcome the state explosion problem. In Iliano Cervesato, Helmut Veith, and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, pages 182–182, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. ISBN 978-3-540-89439-1.
- [37] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In Dexter

- Kozen, editor, *Logics of Programs*, pages 52–71, Berlin, Heidelberg, 1982. Springer Berlin Heidelberg.
- [38] Edmund M. Clarke, Daniel Kroening, and Karen Yorav. Behavioral consistency of C and verilog programs using bounded model checking. In *Proceedings of the 40th Design Automation Conference, DAC 2003, Anaheim, CA, USA, June 2-6, 2003*, pages 368–371. ACM, 2003. doi: 10.1145/775832.775928. URL <https://doi.org/10.1145/775832.775928>.
- [39] Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. A tool for checking ANSI-C programs. In Kurt Jensen and Andreas Podelski, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 10th International Conference, TACAS 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*, volume 2988 of *Lecture Notes in Computer Science*, pages 168–176. Springer, 2004. doi: 10.1007/978-3-540-24730-2\_15. URL [https://doi.org/10.1007/978-3-540-24730-2\\_15](https://doi.org/10.1007/978-3-540-24730-2_15).
- [40] Edmund M. Clarke, Daniel Kroening, Joël Ouaknine, and Ofer Strichman. Completeness and complexity of bounded model checking. In *International Conference on Verification, Model Checking and Abstract Interpretation*, 2004.
- [41] Edmund M. Clarke, Muralidhar Talupur, Tayssir Touili, and Helmut Veith. Verification by network decomposition. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 276–291. Springer, 2004. doi: 10.1007/978-3-540-28644-8\_18. URL [https://doi.org/10.1007/978-3-540-28644-8\\_18](https://doi.org/10.1007/978-3-540-28644-8_18).
- [42] Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. *Model checking, 2nd Edition*. MIT Press, 2018. ISBN 978-0-262-03883-6. URL <https://mitpress.mit.edu/books/model-checking-second-edition>.
- [43] Ernie Cohen, Markus Dahlweid, Mark A. Hillebrand, Dirk Leinenbach, Michal Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies. VCC: A practical system for verifying concurrent C. 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 23–42. Springer, 2009. doi:



- 10.1007/978-3-642-03359-9\_2. URL [https://doi.org/10.1007/978-3-642-03359-9\\_2](https://doi.org/10.1007/978-3-642-03359-9_2).
- [44] Byron Cook. Formal reasoning about the security of amazon web services. In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification*, pages 38–47, Cham, 2018. Springer International Publishing. ISBN 978-3-319-96145-3.
- [45] Thierry Coquand and Gérard Huet. The calculus of constructions. *Information and Computation*, 76(2):95–120, 1988. ISSN 0890-5401. doi: [https://doi.org/10.1016/0890-5401\(88\)90005-3](https://doi.org/10.1016/0890-5401(88)90005-3). URL <https://www.sciencedirect.com/science/article/pii/0890540188900053>.
- [46] Lucas C. Cordeiro and Bernd Fischer. Verifying multi-threaded software using smt-based context-bounded model checking. In Richard N. Taylor, Harald C. Gall, and Nenad Medvidovic, editors, *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu , HI, USA, May 21-28, 2011*, pages 331–340. ACM, 2011. doi: 10.1145/1985793.1985839. URL <https://doi.org/10.1145/1985793.1985839>.
- [47] The MITRE Corporation. Cwe-193: Off-by-one error, 2006. URL <https://cwe.mitre.org/data/definitions/193.html>.
- [48] C. Cowan, F. Wagle, Calton Pu, S. Beattie, and J. Walpole. Buffer overflows: attacks and defenses for the vulnerability of the decade. In *Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00*, volume 2, pages 119–129 vol.2, 2000. doi: 10.1109/DISCEX.2000.821514.
- [49] Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. Tada: A logic for time and data abstraction. In *ECOOP 2014–Object-Oriented Programming: 28th European Conference, Uppsala, Sweden, July 28–August 1, 2014. Proceedings 28*, pages 207–231. Springer, 2014.
- [50] Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. Modular termination verification for non-blocking concurrency. In *ESOP*, volume 9632 of *Lecture Notes in Computer Science*, pages 176–201. Springer, 2016. doi: 10.1007/978-3-662-49498-1\_8.
- [51] Pedro de Carvalho Gomes, Dilian Gurov, Marieke Huisman, and Cyrille Artho. Specification and verification of synchronization with condition variables. *Sci. Comput. Program.*, 163:174–189, 2018. doi: 10.1016/J.SCICO.2018.05.001. URL <https://doi.org/10.1016/j.scico.2018.05.001>.

- [52] Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The lean theorem prover (system description). In Amy P. Felty and Aart Middeldorp, editors, *Automated Deduction - CADE-25*, pages 378–388, Cham, 2015. Springer International Publishing.
- [53] N. Dershowitz and Z. Manna. Proving termination with multiset orderings. In *ICALP*, 1979. doi: 10.1007/3-540-09510-1\_15.
- [54] Coq Developers. Coq website, 1989. URL <https://coq.inria.fr/>. Accessed: 2023-11-07.
- [55] FreeRTOS Developers. FreeRTOS website, 2003. URL <https://www.freertos.org>. Accessed: 2023-11-01.
- [56] FreeRTOS Developers. Repository of FreeRTOS’s TCP/IP stack, 2020. URL <https://github.com/FreeRTOS/FreeRTOS-Plus-TCP>. Accessed: 2023-11-01.
- [57] Lean Developers. Lean website, 1989. URL <https://lean-lang.org/>. Accessed: 2023-11-07.
- [58] David Dice, Virendra J. Marathe, and Nir Shavit. Lock cohorting: a general technique for designing NUMA locks. In J. Ramanujam and P. Sadayappan, editors, *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2012, New Orleans, LA, USA, February 25-29, 2012*, pages 247–256. ACM, 2012. doi: 10.1145/2145816.2145848. URL <https://doi.org/10.1145/2145816.2145848>.
- [59] David Dice, Virendra J. Marathe, and Nir Shavit. Lock cohorting: A general technique for designing NUMA locks. *ACM Trans. Parallel Comput.*, 1(2):13:1–13:42, 2015. doi: 10.1145/2686884. URL <https://doi.org/10.1145/2686884>.
- [60] Edsger W. Dijkstra. Notes on structured programming. 1970. URL <https://api.semanticscholar.org/CorpusID:8242220>.
- [61] Edsger W. Dijkstra. A discipline of programming. Pentice Hall, 1976.
- [62] Mike Dodds, Xinyu Feng, Matthew J. Parkinson, and Viktor Vafeiadis. Deny-guarantee reasoning. 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 363–377. Springer, 2009. doi: 10.1007/978-3-642-00590-9\_26. URL [https://doi.org/10.1007/978-3-642-00590-9\\_26](https://doi.org/10.1007/978-3-642-00590-9_26).

- [63] Alastair F. Donaldson, Leopold Haller, Daniel Kroening, and Philipp Rümmer. Software verification using k-induction. In Eran Yahav, editor, *Static Analysis - 18th International Symposium, SAS 2011, Venice, Italy, September 14-16, 2011. Proceedings*, volume 6887 of *Lecture Notes in Computer Science*, pages 351–368. Springer, 2011. doi: 10.1007/978-3-642-23702-7\_26. URL [https://doi.org/10.1007/978-3-642-23702-7\\_26](https://doi.org/10.1007/978-3-642-23702-7_26).
- [64] Emanuele D’Osualdo, Julian Sutherland, Azadeh Farzan, and Philippa Gardner. Tada live: Compositional reasoning for termination of fine-grained concurrent programs. *ACM Trans. Program. Lang. Syst.*, 43(4): 16:1–16:134, 2021. doi: 10.1145/3477082. URL <https://doi.org/10.1145/3477082>.
- [65] Vijay Victor D’Silva, Daniel Kroening, and Georg Weissenbacher. A survey of automated techniques for formal software verification. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 27(7):1165–1178, 2008. doi: 10.1109/TCAD.2008.923410. URL <https://doi.org/10.1109/TCAD.2008.923410>.
- [66] Niklas Eén, Alan Mishchenko, and Robert K. Brayton. Efficient implementation of property directed reachability. In Per Bjesse and Anna Slobodová, editors, *International Conference on Formal Methods in Computer-Aided Design, FMCAD ’11, Austin, TX, USA, October 30 - November 02, 2011*, pages 125–134. FMCAD Inc., 2011. URL <http://dl.acm.org/citation.cfm?id=2157675>.
- [67] E. Allen Emerson and Kedar S. Namjoshi. Reasoning about rings. In Ron K. Cytron and Peter Lee, editors, *Conference Record of POPL’95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, pages 85–94. ACM Press, 1995. doi: 10.1145/199448.199468. URL <https://doi.org/10.1145/199448.199468>.
- [68] Barry S. Fagin and Martin C. Carlisle. Provably secure DNS: A case study in reliable software. In Hubert B. Keller, Erhard Plödereder, Peter Dencker, and Herbert Klenk, editors, *Reliable Software Technologies - Ada-Europe 2013, 18th Ada-Europe International Conference on Reliable Software Technologies, Berlin, Germany, June 10-14, 2013. Proceedings*, volume 7896 of *Lecture Notes in Computer Science*, pages 81–93. Springer, 2013. doi: 10.1007/978-3-642-38601-5\_6. URL [https://doi.org/10.1007/978-3-642-38601-5\\_6](https://doi.org/10.1007/978-3-642-38601-5_6).
- [69] Justus Fasse and Bart Jacobs. Expressive modular verification of termination for busy-waiting programs (talk). Iris Workshop 2023. URL <https://iris-project.org/workshop-2023/slides/fasse.pdf>.

- [70] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 - where programs meet provers. 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 125–128. Springer, 2013. doi: 10.1007/978-3-642-37036-6\_8. URL [https://doi.org/10.1007/978-3-642-37036-6\\_8](https://doi.org/10.1007/978-3-642-37036-6_8).
- [71] Cormac Flanagan and James B. Saxe. Avoiding exponential explosion: generating compact verification conditions. In *ACM-SIGACT Symposium on Principles of Programming Languages*, 2001.
- [72] Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for java. In *PLDI '02*, 2002. doi: 10.1145/512529.512558.
- [73] Jason Franklin, Sagar Chaki, Anupam Datta, and Arvind Seshadri. Scalable parametric verification of secure systems: How to verify reference monitors without worrying about data structure size. In *31st IEEE Symposium on Security and Privacy, SP 2010, 16-19 May 2010, Berkeley/Oakland, California, USA*, pages 365–379. IEEE Computer Society, 2010. doi: 10.1109/SP.2010.29. URL <https://doi.org/10.1109/SP.2010.29>.
- [74] Jason Franklin, Sagar Chaki, Anupam Datta, Jonathan M. McCune, and Amit Vasudevan. Parametric verification of address space separation. In Pierpaolo Degano and Joshua D. Guttman, editors, *Principles of Security and Trust*, pages 51–68, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [75] Desheng Fu and Feiyue Shi. Buffer overflow exploit and defensive techniques. *2012 Fourth International Conference on Multimedia Information Networking and Security*, pages 87–90, 2012. URL <https://api.semanticscholar.org/CorpusID:41074172>.
- [76] Mikhail Y. R. Gadelha, Hussama Ibrahim Ismail, and Lucas C. Cordeiro. Handling loops in bounded model checking of C programs via k-induction. *Int. J. Softw. Tools Technol. Transf.*, 19(1):97–114, 2017. doi: 10.1007/S10009-015-0407-9. URL <https://doi.org/10.1007/s10009-015-0407-9>.
- [77] Mikhail Y. R. Gadelha, Felipe R. Monteiro, Lucas C. Cordeiro, and Denis A. Nicole. ESBMC v6.0: Verifying C programs using k-induction and invariant inference - (competition contribution). In Dirk Beyer,

- Marieke Huisman, Fabrice Kordon, and Bernhard Steffen, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 25 Years of TACAS: TOOLympics, Held as Part of ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part III*, volume 11429 of *Lecture Notes in Computer Science*, pages 209–213. Springer, 2019. doi: 10.1007/978-3-030-17502-3\_15. URL [https://doi.org/10.1007/978-3-030-17502-3\\_15](https://doi.org/10.1007/978-3-030-17502-3_15).
- [78] Adam Geras, Michael R. Smith, and James Miller. A prototype empirical evaluation of test driven development. *10th International Symposium on Software Metrics, 2004. Proceedings.*, pages 405–416, 2004. URL <https://api.semanticscholar.org/CorpusID:18768790>.
- [79] Alexey Gotsman, Byron Cook, Matthew Parkinson, and Viktor Vafeiadis. Proving that non-blocking algorithms don't block. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '09, page 16–28, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605583792. doi: 10.1145/1480881.1480886. URL <https://doi.org/10.1145/1480881.1480886>.
- [80] Mohamed G. Gouda, Eric G. Manning, and Yao-Tin Yu. On the progress of communications between two finite state machines. *Inf. Control.*, 63(3):200–216, 1984. doi: 10.1016/S0019-9958(84)80014-5. URL [https://doi.org/10.1016/S0019-9958\(84\)80014-5](https://doi.org/10.1016/S0019-9958(84)80014-5).
- [81] Susanne Graf and Hassen Saidi. Construction of abstract state graphs with pvs. In Orna Grumberg, editor, *Computer Aided Verification*, pages 72–83, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg. ISBN 978-3-540-69195-2.
- [82] Ronghui Gu, Jérémie Koenig, Tahina Ramananandro, Zhong Shao, Xiongnan (Newman) Wu, Shu-Chun Weng, Haozhong Zhang, and Yu Gu. Deep specifications and certified abstraction layers. 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 595–608. ACM, 2015. doi: 10.1145/2676726.2676975. URL <https://doi.org/10.1145/2676726.2676975>.
- [83] Hui Guo, Chunyan Hou, Jinsong Wang, and Chen Chen. A k-induction method extended with value analysis for C program safety verification. In *IEEE International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom 2022, Wuhan, China, December 9-11, 2022*, pages 1072–1077. IEEE, 2022. doi: 10.

- 1109/TRUSTCOM56396.2022.00147. URL <https://doi.org/10.1109/TrustCom56396.2022.00147>.
- [84] Jafar Hamin and Bart Jacobs. Deadlock-free monitors. In *ESOP*, 2018. doi: 10.1007/978-3-319-89884-1\_15.
- [85] Jafar Hamin and Bart Jacobs. Transferring Obligations Through Synchronizations. In *33rd European Conference on Object-Oriented Programming (ECOOP 2019)*, volume 134 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:58, 2019. ISBN 978-3-95977-111-5. doi: 10.4230/LIPIcs.ECOOP.2019.19.
- [86] Inzemamul Haque, Deepak D’Souza, Habeeb P, Arnab Kundu, and Ganesh Babu. Verification of a generative separation kernel. In Dang Van Hung and Oleg Sokolsky, editors, *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*, volume 12302 of *Lecture Notes in Computer Science*, pages 305–322. Springer, 2020. doi: 10.1007/978-3-030-59152-6\_17. URL [https://doi.org/10.1007/978-3-030-59152-6\\_17](https://doi.org/10.1007/978-3-030-59152-6_17).
- [87] Keijo Heljanko, Tommi A. Junttila, and Timo Latvala. Incremental and complete bounded model checking for full pltl. In *International Conference on Computer Aided Verification*, 2005.
- [88] Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Grégoire Sutre. Lazy abstraction. 37(1):58–70, jan 2002. ISSN 0362-1340. doi: 10.1145/565816.503279. URL <https://doi.org/10.1145/565816.503279>.
- [89] Joseph Herkert, Jason Borenstein, and Keith W. Miller. The boeing 737 MAX: lessons for engineering ethics. *Science and Engineering Ethics*, 26(6):2957–2974, 2020. doi: 10.1007/s11948-020-00252-y. URL <https://doi.org/10.1007/s11948-020-00252-y>.
- [90] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming, Revised Reprint*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 2012. ISBN 9780123973375.
- [91] Michael Hilton, Jonathan Bell, and Darko Marinov. A large-scale study of test coverage evolution. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE ’18*, page 53–63, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450359375. doi: 10.1145/3238147.3238183. URL <https://doi.org/10.1145/3238147.3238183>.
- [92] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12:576–580, 1968. doi: 10.1145/363235.363259.

- [93] Krystof Hoder and Nikolaj S. Bjørner. Generalized property directed reachability. In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference, Trento, Italy, June 17-20, 2012. Proceedings*, volume 7317 of *Lecture Notes in Computer Science*, pages 157–171. Springer, 2012. doi: 10.1007/978-3-642-31612-8\_13. URL [https://doi.org/10.1007/978-3-642-31612-8\\_13](https://doi.org/10.1007/978-3-642-31612-8_13).
- [94] Aurélie Hurault and Philippe Quéinnec. Proving a non-blocking algorithm for process renaming with TLA+. In Dirk Beyer and Chantal Keller, editors, *Tests and Proofs - 13th International Conference, TAP@FM 2019, Porto, Portugal, October 9-11, 2019, Proceedings*, volume 11823 of *Lecture Notes in Computer Science*, pages 147–166. Springer, 2019. doi: 10.1007/978-3-030-31157-5\_10. URL [https://doi.org/10.1007/978-3-030-31157-5\\_10](https://doi.org/10.1007/978-3-030-31157-5_10).
- [95] Vinay M. Ijure, Sean A. Laughter, and Ronald D. Williams. Security issues in scada networks. *Computers and Security*, 25(7):498–506, 2006. ISSN 0167-4048. doi: <https://doi.org/10.1016/j.cose.2006.03.001>. URL <https://www.sciencedirect.com/science/article/pii/S0167404806000514>.
- [96] Erdal Irmak and İsmail Erkek. An overview of cyber-attack vectors on scada systems. In *2018 6th international symposium on digital forensic and security (ISDFS)*, pages 1–5. IEEE, 2018.
- [97] Bart Jacobs. Modular verification of liveness properties of the I/O behavior of imperative programs. In *ISoLA*, 2020. doi: 10.1007/978-3-030-61362-4\_29.
- [98] Bart Jacobs, editor. *VeriFast 21.04*. Zenodo, 2021. doi: 10.5281/zenodo.4705416.
- [99] Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. Verifast: A powerful, sound, predictable, fast verifier for C and Java. In M. Bobaru, K. Havelund, G.J. Holzmann, and R. Joshi, editors, *NASA Formal Methods (NFM 2011)*, volume 6617, pages 41–55. Springer, 2011. ISBN 978-3-642-20397-8. doi: 10.1007/978-3-642-20398-5\_4. URL <https://lirias.kuleuven.be/95720>.
- [100] Bart Jacobs, Dragan Bosnacki, and Ruurd Kuiper. Modular termination verification of single-threaded and multithreaded programs. *ACM Trans. Program. Lang. Syst.*, 40:12:1–12:59, 2018. doi: 10.1145/3210258.
- [101] Ranjit Jhala and Kenneth L. McMillan. Array abstractions from proofs. In Werner Damm and Holger Hermanns, editors, *Computer Aided*

- Verification, 19th International Conference, CAV 2007, Berlin, Germany, July 3-7, 2007, Proceedings*, volume 4590 of *Lecture Notes in Computer Science*, pages 193–206. Springer, 2007. doi: 10.1007/978-3-540-73368-3\_23. URL [https://doi.org/10.1007/978-3-540-73368-3\\_23](https://doi.org/10.1007/978-3-540-73368-3_23).
- [102] Taylor T. Johnson and Sayan Mitra. A small model theorem for rectangular hybrid automata networks. In Holger Giese and Grigore Rosu, editors, *Formal Techniques for Distributed Systems - Joint 14th IFIP WG 6.1 International Conference, FMOODS 2012 and 32nd IFIP WG 6.1 International Conference, FORTE 2012, Stockholm, Sweden, June 13-16, 2012. Proceedings*, volume 7273 of *Lecture Notes in Computer Science*, pages 18–34. Springer, 2012. doi: 10.1007/978-3-642-30793-5\_2. URL [https://doi.org/10.1007/978-3-642-30793-5\\_2](https://doi.org/10.1007/978-3-642-30793-5_2).
- [103] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. Higher-order ghost state. *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, 2016. doi: 10.1145/2951913.2951943.
- [104] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.*, 28:e20, 2018. doi: 10.1017/S0956796818000151.
- [105] Ioannis T. Kassios. Dynamic frames: Support for framing, dependencies and sharing without restrictions. In Jayadev Misra, Tobias Nipkow, and Emil Sekerinski, editors, *FM 2006: Formal Methods*, pages 268–283, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-37216-5.
- [106] Jieung Kim, Vilhelm Sjöberg, Ronghui Gu, and Zhong Shao. Safety and liveness of mcs lock—layer by layer. In *Asian Symposium on Programming Languages and Systems*, 2017.
- [107] Naoki Kobayashi. A new type system for deadlock-free processes. In *CONCUR*, 2006. doi: 10.1007/11817949\_16.
- [108] Daniel Kroening. Computing over-approximations with bounded model checking. In Armin Biere and Ofer Strichman, editors, *Proceedings of the Third International Workshop on Bounded Model Checking, BMC@CAV 2005, Edinburgh, UK, July 11, 2005*, volume 144 of *Electronic Notes in Theoretical Computer Science*, pages 79–92. Elsevier, 2005. doi: 10.1016/J.ENTCS.2005.07.021. URL <https://doi.org/10.1016/j.entcs.2005.07.021>.
- [109] Daniel Kroening and Ofer Strichman. Efficient computation of recurrence diameters. In Lenore D. Zuck, Paul C. Attie, Agostino Cortesi, and



- Supratik Mukhopadhyay, editors, *Verification, Model Checking, and Abstract Interpretation, 4th International Conference, VMCAI 2003, New York, NY, USA, January 9-11, 2002, Proceedings*, volume 2575 of *Lecture Notes in Computer Science*, pages 298–309. Springer, 2003. doi: 10.1007/3-540-36384-X\_24. URL [https://doi.org/10.1007/3-540-36384-X\\_24](https://doi.org/10.1007/3-540-36384-X_24).
- [110] Daniel Kroening, Joël Ouaknine, Ofer Strichman, Thomas Wahl, and James Worrell. Linear completeness thresholds for bounded model checking. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 557–572. Springer, 2011. doi: 10.1007/978-3-642-22110-1\_44. URL [https://doi.org/10.1007/978-3-642-22110-1\\_44](https://doi.org/10.1007/978-3-642-22110-1_44).
- [111] Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2):125–143, 1977. doi: 10.1109/TSE.1977.229904. URL <https://doi.org/10.1109/TSE.1977.229904>.
- [112] Jacob Thamsborg Lars Birkedal, Kristian Støvring. The category-theoretic solution of recursive metric-space equations. *Theoretical Computer Science*, 411(47):4102 – 4122, 2010. doi: 10.1016/j.tcs.2010.07.010.
- [113] Mark Lawford and Alan Wassynq. Formal verification of nuclear systems: Past, present, and future. In *1st International Workshop on Critical Infrastructure Safety and Security (CrISS-DESSERT'11)*, volume 1, pages 43–51, 2012.
- [114] Nico Lehmann, Adam T. Geller, Niki Vazou, and Ranjit Jhala. Flux: Liquid types for rust. *Proc. ACM Program. Lang.*, 7(PLDI), jun 2023. doi: 10.1145/3591283. URL <https://doi.org/10.1145/3591283>.
- [115] K. Leino and Peter Müller. A basis for verifying multi-threaded programs. In *ESOP '09 Proceedings of the 18th European Symposium on Programming Languages and Systems: Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009*, page 378, March 2009. doi: 10.1007/978-3-642-00590-9\_27.
- [116] K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16, Dakar, Senegal, April 25-May 1, 2010, Revised Selected Papers*, volume 6355 of *Lecture Notes in Computer Science*, pages 348–370. Springer, 2010. doi: 10.1007/978-3-642-17511-4\_20. URL [https://doi.org/10.1007/978-3-642-17511-4\\_20](https://doi.org/10.1007/978-3-642-17511-4_20).

- [117] K. Rustan M. Leino and Francesco Logozzo. Loop invariants on demand. In Kwangkeun Yi, editor, *Programming Languages and Systems, Third Asian Symposium, APLAS 2005, Tsukuba, Japan, November 2-5, 2005, Proceedings*, volume 3780 of *Lecture Notes in Computer Science*, pages 119–134. Springer, 2005. doi: 10.1007/11575467\_9. URL [https://doi.org/10.1007/11575467\\_9](https://doi.org/10.1007/11575467_9).
- [118] K. Rustan M. Leino, Peter Müller, and Jan Smans. Deadlock-free channels and locks. In *Proceedings of the 19th European Conference on Programming Languages and Systems*, ESOP’10, page 407–426, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 3642119565. doi: 10.1007/978-3-642-11957-6\_22.
- [119] Ruy Ley-Wild and Aleksandar Nanevski. Subjective auxiliary state for coarse-grained concurrency. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’13, page 561–574, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450318327. doi: 10.1145/2429069.2429134. URL <https://doi.org/10.1145/2429069.2429134>.
- [120] Hongjin Liang and Xinyu Feng. A program logic for concurrent objects under fair scheduling. In *POPL*, 2016. doi: 10.1145/2837614.2837635.
- [121] Hongjin Liang and Xinyu Feng. Progress of concurrent objects with partial methods. *Proc. ACM Program. Lang.*, 2:20:1–20:31, 2017. doi: 10.1145/3158108.
- [122] Didrik Lundberg, Roberto Guanciale, Andreas Lindner, and Mads Dam. Hoare-style logic for unstructured programs. In Frank de Boer and Antonio Cerone, editors, *Software Engineering and Formal Methods*, pages 193–213, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58768-0.
- [123] Isabella Mastroeni and Damiano Zanardini. Data dependencies and program slicing: from syntax to abstract semantics. In Robert Glück and Oege de Moor, editors, *Proceedings of the 2008 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation, PEPM 2008, San Francisco, California, USA, January 7-8, 2008*, pages 125–134. ACM, 2008. doi: 10.1145/1328408.1328428. URL <https://doi.org/10.1145/1328408.1328428>.
- [124] Umang Mathur, Adithya Murali, Paul Krogmeier, P. Madhusudan, and Mahesh Viswanathan. Deciding memory safety for single-pass heap-manipulating programs. *Proc. ACM Program. Lang.*, 4(POPL):35:1–35:29, 2020. doi: 10.1145/3371103. URL <https://doi.org/10.1145/3371103>.
- [125] John W McCormick and Peter C Chapin. *Building high integrity applications with SPARK*. Cambridge University Press, 2015.

- [126] Kenneth L. McMillan. Interpolation and sat-based model checking. In *International Conference on Computer Aided Verification*, 2003.
- [127] 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 <https://doi.org/10.1145/103727.103729>.
- [128] John M. Mellor-Crummey and Michael L. Scott. Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Trans. Comput. Syst.*, 9:21–65, 1991. doi: 10.1145/103727.103729.
- [129] Barton P. Miller, Lars Fredriksen, and Bryan So. An empirical study of the reliability of unix utilities. *Commun. ACM*, 33(12):32–44, dec 1990. ISSN 0001-0782. doi: 10.1145/96267.96279. URL <https://doi.org/10.1145/96267.96279>.
- [130] Mark Moir and James H. Anderson. Wait-free algorithms for fast, long-lived renaming. *Science of Computer Programming*, 25(1):1–39, 1995. ISSN 0167-6423. doi: [https://doi.org/10.1016/0167-6423\(95\)00009-H](https://doi.org/10.1016/0167-6423(95)00009-H). URL <https://www.sciencedirect.com/science/article/pii/016764239500009H>.
- [131] Louise E. Moser. Data dependency graphs for ada programs. *IEEE Trans. Software Eng.*, 16(5):498–509, 1990. doi: 10.1109/32.52773. URL <https://doi.org/10.1109/32.52773>.
- [132] Peter Müller, Malte Schwerhoff, and Alexander J. Summers. Viper: A verification infrastructure for permission-based reasoning. In Barbara Jobstmann and K. Rustan M. Leino, editors, *Verification, Model Checking, and Abstract Interpretation*, pages 41–62, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. ISBN 978-3-662-49122-5.
- [133] K. Mühlemann. Method for reducing memory conflicts caused by busy waiting in multiple processor synchronisation. *IEE Proceedings E - Computers and Digital Techniques*, 127(3):85–87, 1980. doi: 10.1049/ip-e.1980.0017.
- [134] Kedar S. Namjoshi. Symmetry and completeness in the analysis of parameterized systems. In Byron Cook and Andreas Podelski, editors, *Verification, Model Checking, and Abstract Interpretation*, pages 299–313, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- [135] P. O’Hearn. A primer on separation logic (and automatic program verification and analysis). In *Software Safety and Security*, 2012.

- [136] Peter W. O’Hearn. Separation logic. *Commun. ACM*, 62(2):86–95, 2019. doi: 10.1145/3211968. URL <https://doi.org/10.1145/3211968>.
- [137] Peter W. O’Hearn. Incorrectness logic. *Proc. ACM Program. Lang.*, 4 (POPL), dec 2019. doi: 10.1145/3371078. URL <https://doi.org/10.1145/3371078>.
- [138] Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *CSL*, 2001. doi: 10.1007/3-540-44802-0\_1.
- [139] S. Owicki and D. Gries. Verifying properties of parallel programs: an axiomatic approach. *Commun. ACM*, 19:279–285, 1976.
- [140] Paolo Panaroni, Giovanni Sartori, Fabrizio Fabbrini, Mario Fusani, and Giuseppe Lami. Safety in automotive software: An overview of current practices. In *2008 32nd Annual IEEE International Computer Software and Applications Conference*, pages 1053–1058, 2008. doi: 10.1109/COMPSAC.2008.139.
- [141] Panos M. Pardalos and Athanasios Migdalas. A note on the complexity of longest path problems related to graph coloring. *Appl. Math. Lett.*, 17(1):13–15, 2004. doi: 10.1016/S0893-9659(04)90003-1. URL [https://doi.org/10.1016/S0893-9659\(04\)90003-1](https://doi.org/10.1016/S0893-9659(04)90003-1).
- [142] D.Y.W. Park, U. Stern, J.U. Skakkebaek, and D.L. Dill. Java model checking. In *Proceedings ASE 2000. Fifteenth IEEE International Conference on Automated Software Engineering*, pages 253–256, 2000. doi: 10.1109/ASE.2000.873671.
- [143] Gaurav Parthasarathy, Peter Müller, and Alexander J. Summers. Formally validating a practical verification condition generator. In *International Conference on Computer Aided Verification*, 2021.
- [144] Christopher Pulte, Dhruv C. Makwana, Thomas Sewell, Kayvan Memarian, Peter Sewell, and Neel Krishnaswami. CN: verifying systems C code with separation-logic refinement types. *Proc. ACM Program. Lang.*, 7(POPL):1–32, 2023. doi: 10.1145/3571194. URL <https://doi.org/10.1145/3571194>.
- [145] Shengchao Qin, Guanhua He, Chenguang Luo, and Wei-Ngan Chin. Loop invariant synthesis in a combined domain. In Jin Song Dong and Huibiao Zhu, editors, *Formal Methods and Software Engineering*, pages 468–484, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

- [146] J. P. Queille and J. Sifakis. Specification and verification of concurrent systems in cesar. In Mariangiola Dezani-Ciancaglini and Ugo Montanari, editors, *International Symposium on Programming*, pages 337–351, Berlin, Heidelberg, 1982. Springer Berlin Heidelberg.
- [147] Tobias Reinhard. Completeness thresholds for memory safety of array traversing programs: Early technical report. *CoRR*, abs/2211.11885, 2022. doi: 10.48550/arXiv.2211.11885. URL <https://doi.org/10.48550/arXiv.2211.11885>.
- [148] Tobias Reinhard and Bart Jacobs. VeriFast proof of safety for CLH lock. 2020. URL <https://github.com/verifast/verifast/blob/master/examples/busywaiting/clhlock/clhlock.c>.
- [149] Tobias Reinhard and Bart Jacobs. VeriFast proof of termination for consumer-producer problem with bounded FIFO. 2020. URL [https://github.com/verifast/verifast/blob/master/examples/busywaiting/bounded\\_fifo.c](https://github.com/verifast/verifast/blob/master/examples/busywaiting/bounded_fifo.c).
- [150] Tobias Reinhard and Bart Jacobs. VeriFast proof of I/O liveness for a simple server with a receiver and a responder thread communicating via a shared buffer. 2020. URL [https://github.com/verifast/verifast/blob/master/examples/busywaiting/ioliveness/echo\\_live\\_mt.c](https://github.com/verifast/verifast/blob/master/examples/busywaiting/ioliveness/echo_live_mt.c).
- [151] Tobias Reinhard and Bart Jacobs. Ghost signals: Verifying termination of busy waiting (technical report). Zenodo, 2021. doi: 10.5281/zenodo.4775181.
- [152] Tobias Reinhard and Bart Jacobs. Ghost signals: Verifying termination of busy waiting (extended version). *CoRR*, abs/2010.11762, 2021. URL <https://arxiv.org/abs/2010.11762>.
- [153] Tobias Reinhard, Mark Tuttle, and FreeRTOS team. VeriFast proof of memory and thread safety of the FreeRTOS multicore scheduler. URL [https://github.com/Tobias-internship-AWS-2022/FreeRTOS-Kernel/tree/verifast\\_switch\\_context/Test/VeriFast/tasks/vTaskSwitchContext](https://github.com/Tobias-internship-AWS-2022/FreeRTOS-Kernel/tree/verifast_switch_context/Test/VeriFast/tasks/vTaskSwitchContext). Accessed: 2023-11-03.
- [154] Tobias Reinhard, Amin Timany, and Bart Jacobs. A separation logic to verify termination of busy-waiting for abrupt program exit. In Wytse Oortwijn, editor, *FTfJP 2020: Proceedings of the 22nd ACM SIGPLAN International Workshop on Formal Techniques for Java-Like Programs, Virtual Event, USA, July 23, 2020*, pages 26–32. ACM, 2020. doi: 10.1145/3427761.3428345. URL <https://doi.org/10.1145/3427761.3428345>.

- [155] John C. Reynolds. Separation logic: a logic for shared mutable data structures. *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, 2002. doi: 10.1109/LICS.2002.1029817.
- [156] Henry Gordon Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical society*, 74(2):358–366, 1953.
- [157] Grigore Roşu, Wolfram Schulte, and Traian-Florin Serbanuta. Runtime verification of c memory safety. In *Runtime Verification*, 2009. URL <https://api.semanticscholar.org/CorpusID:1002782>.
- [158] Neha Rungta. A billion smt queries a day (invited paper). In Sharon Shoham and Yakir Vizel, editors, *Computer Aided Verification*, pages 3–18, Cham, 2022. Springer International Publishing. ISBN 978-3-031-13185-1.
- [159] Reiner Sailer, Trent Jaeger, Enriquillo Valdez, Ramón Cáceres, Ronald Perez, Stefan Berger, John Linwood Griffin, and Leendert van Doorn. Building a mac-based security architecture for the xen open-source hypervisor. In *21st Annual Computer Security Applications Conference (ACSAC 2005), 5-9 December 2005, Tucson, AZ, USA*, pages 276–285. IEEE Computer Society, 2005. doi: 10.1109/CSAC.2005.13. URL <https://doi.org/10.1109/CSAC.2005.13>.
- [160] Harshad Sathaye, Domien Schepers, Aanjhan Ranganathan, and Guevara Noubir. Wireless attacks on aircraft landing systems: Demo. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '19, page 295–297, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367264. doi: 10.1145/3317549.3326298. URL <https://doi.org/10.1145/3317549.3326298>.
- [161] Arvind Seshadri, Mark Luk, Ning Qu, and Adrian Perrig. Secvisor: A tiny hypervisor to provide lifetime kernel code integrity for commodity oses. *SIGOPS Oper. Syst. Rev.*, 41(6):335–350, oct 2007. ISSN 0163-5980. doi: 10.1145/1323293.1294294. URL <https://doi.org/10.1145/1323293.1294294>.
- [162] Mary Sheeran, Satnam Singh, and Gunnar Stålmarck. Checking safety properties using induction and a sat-solver. In Warren A. Hunt Jr. and Steven D. Johnson, editors, *Formal Methods in Computer-Aided Design, Third International Conference, FMCAD 2000, Austin, Texas, USA, November 1-3, 2000, Proceedings*, volume 1954 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2000. doi: 10.1007/3-540-40922-X\_8. URL [https://doi.org/10.1007/3-540-40922-X\\_8](https://doi.org/10.1007/3-540-40922-X_8).

- [163] S. Simpson. Subsystems of second order arithmetic. In *Perspectives in mathematical logic*, 1999. doi: 10.1017/CBO9780511581007.
- [164] Rohit Sinha, Cynthia Sturton, Petros Maniatis, Sanjit A. Seshia, and David A. Wagner. Verification with small and short worlds. In Gianpiero Cabodi and Satnam Singh, editors, *Formal Methods in Computer-Aided Design, FMCAD 2012, Cambridge, UK, October 22-25, 2012*, pages 68–77. IEEE, 2012. URL <https://ieeexplore.ieee.org/document/6462557/>.
- [165] Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames: Combining dynamic frames and separation logic. In Sophia Drossopoulou, editor, *ECOOP 2009 – Object-Oriented Programming*, pages 148–172, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-03013-0.
- [166] Youngju Song, Minki Cho, Dongjae Lee, Chung-Kil Hur, Michael Sammler, and Derek Dreyer. Conditional contextual refinement. *Proc. ACM Program. Lang.*, 7(POPL):1121–1151, 2023. doi: 10.1145/3571232. URL <https://doi.org/10.1145/3571232>.
- [167] Luca Spalazzi and Francesco Spegni. Parameterized model-checking of timed systems with conjunctive guards. In Dimitra Giannakopoulou and Daniel Kroening, editors, *Verified Software: Theories, Tools and Experiments - 6th International Conference, VSTTE 2014, Vienna, Austria, July 17-18, 2014, Revised Selected Papers*, volume 8471 of *Lecture Notes in Computer Science*, pages 235–251. Springer, 2014. doi: 10.1007/978-3-319-12154-3\_15. URL [https://doi.org/10.1007/978-3-319-12154-3\\_15](https://doi.org/10.1007/978-3-319-12154-3_15).
- [168] Simon Spies, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. Transfinite iris: Resolving an existential dilemma of step-indexed separation logic. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2021*, page 80–95, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383912. doi: 10.1145/3453483.3454031. URL <https://doi.org/10.1145/3453483.3454031>.
- [169] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955. doi: 10.2307/2963937.
- [170] Joseph Tassarotti, Ralf Jung, and Robert Harper. A higher-order logic for concurrent termination-preserving refinement. In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences*

- on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10201 of *Lecture Notes in Computer Science*, pages 909–936. Springer, 2017. doi: 10.1007/978-3-662-54434-1\_34. URL [https://doi.org/10.1007/978-3-662-54434-1\\_34](https://doi.org/10.1007/978-3-662-54434-1_34).
- [171] André Teixeira, György Dán, Henrik Sandberg, and Karl H. Johansson. A cyber security study of a scada energy management system: Stealthy deception attacks on the state estimator\*. *IFAC Proceedings Volumes*, 44 (1):11271–11277, 2011. ISSN 1474-6670. doi: <https://doi.org/10.3182/20110828-6-IT-1002.02210>. URL <https://www.sciencedirect.com/science/article/pii/S147466701645425X>. 18th IFAC World Congress.
- [172] Amin Timany, Simon Oddershede Gregersen, Léo Stefanescu, Jonas Kastberg Hinrichsen, Léon Gondelman, Abel Nieto, and Lars Birkedal. Trillium: Higher-order concurrent and distributed separation logic for intensional refinement. *Proc. ACM Program. Lang.*, (POPL), 2024.
- [173] Alan Turing. Checking a large routine. In *Report of a Conference on High Speed Automatic Calculating Machines*, pages 67–69, June 1949.
- [174] 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 <https://doi.org/10.1145/2500365.2500600>.
- [175] 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 <https://doi.org/10.1145/2429069.2429111>.
- [176] Viktor Vafeiadis. Concurrent separation logic and operational semantics. *Electronic Notes in Theoretical Computer Science*, 276:335–351, 2011. ISSN 1571-0661. doi: 10.1016/j.entcs.2011.09.029. Twenty-seventh Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVII).
- [177] Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In Luís Caires and Vasco T. Vasconcelos, editors, *CONCUR 2007 – Concurrency Theory*, pages 256–271, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.



- [178] Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. Refinement types for haskell. *SIGPLAN Not.*, 49(9): 269–282, aug 2014. ISSN 0362-1340. doi: 10.1145/2692915.2628161. URL <https://doi.org/10.1145/2692915.2628161>.
- [179] Alan Wassyng, Mark S. Lawford, and Thomas S.E. Maibaum. Software certification experience in the canadian nuclear industry: Lessons for the future. In *Proceedings of the Ninth ACM International Conference on Embedded Software*, EMSOFT '11, page 219–226, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450307147. doi: 10.1145/2038642.2038676. URL <https://doi.org/10.1145/2038642.2038676>.
- [180] Mark D. Weiser. Program slicing. *IEEE Trans. Software Eng.*, 10(4): 352–357, 1984. doi: 10.1109/TSE.1984.5010248. URL <https://doi.org/10.1109/TSE.1984.5010248>.
- [181] Felix A. Wolf, Malte Schwerhoff, and Peter Müller. Concise outlines for a complex logic: A proof outline checker for tada. In Marieke Huisman, Corina Păsăreanu, and Naijun Zhan, editors, *Formal Methods*, pages 407–426, Cham, 2021. Springer International Publishing.
- [182] Felix A Wolf, Malte Schwerhoff, and Peter Müller. Concise outlines for a complex logic: a proof outline checker for tada. *Formal Methods in System Design*, pages 1–27, 2023.
- [183] Haoze Wu, Clark W. Barrett, and Nina Narodytska. Lemur: Integrating large language models in automated program verification. *CoRR*, abs/2310.04870, 2023. doi: 10.48550/ARXIV.2310.04870. URL <https://doi.org/10.48550/arXiv.2310.04870>.
- [184] Qiwen Xu, Willem Paul de Roever, and Jifeng He. The rely-guarantee method for verifying shared variable concurrent programs. *Form. Asp. Comput.*, 9(2):149–174, mar 1997. ISSN 0934-5043. doi: 10.1007/BF01211617. URL <https://doi.org/10.1007/BF01211617>.
- [185] Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. LeanDojo: Theorem proving with retrieval-augmented language models. *CoRR*, abs/2306.15626, 2023. doi: 10.48550/arXiv.2306.15626. URL <https://doi.org/10.48550/arXiv.2306.15626>.
- [186] Zhe Zhou, Ashish Mishra, Benjamin Delaware, and Suresh Jagannathan. Covering all the bases: Type-based verification of test input generators. *Proc. ACM Program. Lang.*, 7(PLDI), jun 2023. doi: 10.1145/3591271. URL <https://doi.org/10.1145/3591271>.





FACULTY OF ENGINEERING SCIENCE  
DEPARTMENT OF COMPUTER SCIENCE  
DISTRINET  
Celestijnenlaan 200A box 2402  
B-3001 Leuven  
reito.eu

