

TaDA Live: Compositional Reasoning for Termination of Fine-grained Concurrent Programs

EMANUELE D’OSUALDO, Imperial College London and MPI-SWS Saarbrücken

JULIAN SUTHERLAND, Imperial College London

AZADEH FARZAN, University of Toronto

PHILIPPA GARDNER, Imperial College London

We present TaDA Live, a concurrent separation logic for reasoning compositionally about the termination of blocking fine-grained concurrent programs. The crucial challenge is how to deal with *abstract atomic blocking*: that is, abstract atomic operations that have blocking behaviour arising from busy-waiting patterns as found in, for example, fine-grained spin locks. Our fundamental innovation is with the design of abstract specifications that capture this blocking behaviour as liveness assumptions on the environment. We design a logic that can reason about the termination of clients which use such operations without breaking their abstraction boundaries, and the correctness of the implementations of the operations with respect to their abstract specifications. We introduce a novel semantic model using layered subjective obligations to express liveness invariants, and a proof system that is sound with respect to the model. The subtlety of our specifications and reasoning is illustrated using several case studies.

Additional Key Words and Phrases: fine-grained concurrency, linearizability, busy-waiting, termination, liveness, concurrent separation logics

1 INTRODUCTION

Compositional reasoning for fine-grained concurrent programs interacting with shared memory is a fundamental, open research problem. We are beginning to obtain a good understanding of how to reason about *safety properties* of concurrent programs: i.e. if the program terminates and the input satisfies the precondition, then the program does not fault and the result satisfies the postcondition. O’Hearn and Brookes [4, 36] introduced concurrent separation logic for reasoning compositionally about course-grained concurrent programs. Since then, there has been a flowering of work on modern concurrent separation logics for reasoning compositionally about safety properties of fine-grained concurrent programs: e.g. CAP [10], TaDA [7], Iris [24] and FCSL [34]. With these modern logics, it is possible to provide abstract specifications that match the intuitive software interface understood by the developer, and to verify both implementations and client programs.

We have comparatively little understanding of how to reason compositionally about *progress (liveness) properties* for fine-grained concurrent algorithms: i.e. something good eventually happens. Examples of progress properties include termination, livelock-freedom, or that every user request is eventually served. The intricacies of the design of concurrent programs often arise precisely from the need to make the program correct with respect to progress properties. The goal of this paper is to design a program logic to reason compositionally about the safety and termination of fine-grained concurrent programs: i.e. to be able to prove that if the input satisfies the precondition, then the program terminates without faulting and the result satisfies the postcondition. As with safety, the aim is to provide abstract specifications, and to verify implementations and clients.

A truly compositional approach would achieve *proof scalability* through the reduction of large complex proofs into a composition of smaller, more tractable proofs, and *proof reuse* through the ability to define abstract interfaces between independent sub-proofs. Proof scalability for concurrent

systems is achieved through *thread-local* reasoning: i.e. the proof of the parallel composition of threads should be the composition of smaller, separate proofs of each thread. Proof reuse is achieved when the right *abstract interface* for a module is identified, so that the proof of correctness of the implementation of the module and the proof of its clients is decoupled: a proof of a client can be reused when swapping the implementation of the module for one satisfying the same specification; a proof of an implementation can be reused when the specification is general enough to support arbitrary correct clients.

For safety, thread-local reasoning can be obtained through rely/guarantee proofs: a protocol on shared state is specified in terms of the set of *allowed* updates, and each thread is verified to respect the protocol under the assumption that the environment respects the protocol. There have been successful attempts at using rely/guarantee reasoning to prove progress properties, such as termination, of *non-blocking* concurrent programs [5, 8, 13, 14, 21, 32], which are the programs where the progress of a thread does not depend on the progress of other threads. For example, the Total TaDA concurrent separation logic [8] was introduced to provide compositional reasoning about the safety and termination of non-blocking programs. It provided thread-local reasoning and abstract specification of module interfaces, without the need to extend the rely/guarantee reasoning.

Standard rely/guarantee reasoning is not enough to prove progress properties for *blocking* programs. In a blocking program, termination of a thread may depend on other threads performing some updates to the shared state. For example, if a thread t is requesting a lock that has been acquired by another thread, then the lack of progress of the thread currently owning the lock would hinder the progress of t . Thread t is blocked, waiting for the lock owner to release the lock. In such situation, a safety abstraction of the environment is insufficient to support a termination argument for t : knowing that the release of the lock is *always allowed to happen* does not imply that it is *eventually happening*.

There has been some work [3, 22, 27] on proving progress properties for programs where blocking is caused solely by *blocking primitives* such as built-in locks or channels. However, it is very common, especially for fine-grained programs, to use ad hoc busy-waiting patterns. For example, consider a thread running `while (v ≠ 1) {v := [x]}`. The termination of this thread is entirely dependent on the environment eventually storing 1 in x . This form of blocking is completely different from a call to a blocking primitive that cannot take a step in the current state. It instead corresponds to code executing steps without making real progress. We call this pattern of behaviour *abstract blocking*.

We have identified two ways to reason about progress in the presence of abstract blocking in the literature: the history-based approach and the refinement-based approach. The history-based approach [15, 25, 37] is very general but results in complex and indirect specifications with complicated reasoning involving explicit trace manipulations. We discuss this approach further in Section 6. In the refinement-based approach, the LiLi logic [30, 31] is the work most closely related to our goals. LiLi extends rely/guarantee with liveness information, to prove a *progress-preserving* contextual refinement between the implementation of a module’s operations and simpler code representing their specifications. LiLi’s extension of rely/guarantee requires, however, heavy use of global auxiliary shared state manipulated through ghost code, which makes the proofs less local. Moreover, the specification code associated with abstractly atomic operations that are blocking is not atomic and exposes implementation details, which hinders scalability and reuse. We give a detailed comparison with our work and LiLi in Sections 2, 5.4 and 6.

The refinement approach does not prove termination directly, but instead relates termination of implementation code with termination of specification code. By contrast, our goal is to develop a program logic with which we are able to verify specifications that describe termination directly, without the manipulation of histories, with proofs that keep auxiliary state as local as possible

without requiring the addition of ghost code, and with specifications that allow the abstraction of implementation details while representing precisely the abstract termination guarantees.

Contributions. Our starting observation is that just as safety rely/guarantee arguments are centred around *invariants*, i.e. facts of the form *always P*, so liveness rely/guarantee arguments for proving progress in the presence of blocking should be centred around *liveness invariants*, i.e. facts of the form *always eventually P*. TaDA Live’s design is based on the idea that this is not a fluke: the dependence on liveness invariants might be considered a *definition* of abstract blocking. To capture this observation within a program logic, we introduce a number of key innovations:

- **subjective obligations**, a new form of logical ghost state to express liveness invariants in a thread-local way without the need for ghost code;
- **obligation layers**, to express dependencies between liveness invariants and avoid unsound circular reasoning;
- **abstract specifications for atomic blocking operations**, to express termination guarantees conditionally on an *environment liveness assumption* of the form “always eventually *P*”.

We obtain TaDA Live, a concurrent separation logic which uses liveness invariants to provide compositional reasoning for establishing safety and termination for blocking programs. The logic makes extensive use of abstract specifications for atomic blocking operations to achieve proof scalability and reuse. This paper presents the following contributions:

- the TaDA Live logic and its specification format;
- a novel semantic model and soundness proof for the logic: the new model is a substantial re-definition of the TaDA model to allow for the non-trivial extensions needed to incorporate the liveness content of the TaDA Live specifications;
- TaDA Live proofs for several paradigmatic case studies: two fine-grained implementations of locks showcase abstraction in the specifications and the obligation mechanism; a program mixing locks and busy-waiting illustrates common proof patterns for clients; two counter modules illustrate TaDA Live’s ability to hide internal blocking and proof reuse; and a set module using a lock-coupling pattern illustrates the generality of the layer system.

Outline. Section 2 provides an example-driven overview of the main innovations of TaDA Live. Section 3 introduces the assertion language and the semantics of the TaDA Live specifications. Section 4 presents the crucial proof rules of TaDA Live, with a running example to illustrate their use. Section 5 presents TaDA Live proofs of several key case studies and a discussion on the limitations of the TaDA Live reasoning. Section 6 contains related work and Section 7 ends with conclusions and future work.

2 AN OVERVIEW OF TADA LIVE

We introduce the main ideas of TaDA Live in this section, leaving the complex technical details for the following sections. Consider a simple example program with non-primitive blocking behaviour:

$$C_1 \left\{ \begin{array}{l} \text{var } v = 0 \text{ in} \\ \text{while}(v \neq 1) \{ \\ \quad v := [x] \\ \} \end{array} \right\} \parallel [x] := 1 \left. \right\} C_2$$

We use a first-order, fine-grained concurrent while language for manipulating shared state. The shared state comprises heap cells which have addresses and store values (addresses, integers, booleans). The $[x]$ notation denotes the value stored at the heap cell with address x . The thread on

the left (\mathbb{C}_1) is busy-waiting on the value stored at the shared heap cell at x . Under fair scheduling, the program is guaranteed to terminate: eventually, the right-hand thread (\mathbb{C}_2) will be scheduled, and will set the heap cell to 1; after that, eventually the left-hand thread will read the value 1 into the local variable v and the while loop will terminate. Since we are aiming at a thread-local proof method, we should be able to break the proof of termination of the program into two separate proofs for the two threads.

We first explore how to provide a thread-local proof of *safety* for this example program using the TaDA logic [7]. We then extend the reasoning with the ingredients needed to prove termination. TaDA is a concurrent separation logic so it uses the standard separation logic assertions. Let us assume the precondition $P = \exists v. x \mapsto v$ and, for simplicity, aim at the postcondition True . TaDA uses the standard parallel rule for concurrent separation logics, where the precondition is separated into two preconditions $P = P_1 * P_2$, one for each thread. Since both threads dereference x , we need a means to share the heap cell in the assertions, turning $x \mapsto v$ into a duplicable assertion, called a *shared region* in TaDA. For our example, we define a shared region $\text{ex}_r(x, v)$ with an associated *interpretation* $\mathcal{I}(\text{ex}_r(x, v)) \triangleq x \mapsto v$, which specifies which resource is being shared. The *region type* ex (for “example”) is the name associated with this interpretation, and the *region identifier* r is an abstract identifier associated with this specific instance of the region type ex . The arguments (x, v) of the region are called the *abstract state* of the region. The definition of a region is completed by an *interference protocol* \mathcal{T}_{ex} which restricts, in rely/guarantee style, the allowed updates to the abstract state. Here, we encode the facts that (a) only \mathbb{C}_2 can update x and (b) v can only be updated to 1. Although such strong invariants are not required to just prove safety, they will be useful for the termination proof later. To encode fact (a), we introduce a form of ghost state called a *guard*, \mathbf{E} , which gives **EX**clusive permission to update x . Formally, guards (probably first introduced in deny-guarantee reasoning [11]) form a partial commutative monoid (PCM), where in this case $\mathbf{E} \bullet \mathbf{E}$ is undefined to capture exclusive permission: if a thread owns \mathbf{E} then no other thread can own it at the same time. To link \mathbf{E} with the ability to change x , the protocol \mathcal{T}_{ex} allows the guarded update

$$\mathbf{E} : (x, v) \rightsquigarrow (x, 1) \quad (1)$$

Fact (b) is encoded by this being the only allowed update.

In TaDA and other modern separation logics such as Iris, implication is generalised to the *viewshift* construct (\Rightarrow) from [9], which can be used to consistently update ghost information, purely within the logic (as opposed to through ghost code). Here, it can be used to turn the owned resource $x \mapsto v$ into a shared resource $P = \exists v. x \mapsto v \Rightarrow \exists r. (\exists v. \text{ex}_r(x, v) * [\mathbf{E}]_r) \equiv \exists r. (P_1 * P_2)$ where $P_1 = \exists v. \text{ex}_r(x, v)$ and $P_2 = \exists v. \text{ex}_r(x, v) * [\mathbf{E}]_r$. The *guard assertion* $[\mathbf{E}]_r$ indicates ownership of the guard \mathbf{E} for the region with identity r . Using standard reasoning, one can then prove $\vdash \{P_1\} \mathbb{C}_1 \{\text{True}\}$ and $\vdash \{P_2\} \mathbb{C}_2 \{\text{True}\}$, which entails, by the parallel rule $\vdash \{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{\text{True} * \text{True}\}$. By consequence and existential elimination on r , we obtain our goal $\vdash \{P\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{\text{True}\}$.

Let us now turn to termination. A thread-local approach would proceed by first proving that \mathbb{C}_1 and \mathbb{C}_2 terminate separately, and then concluding that their parallel composition terminates. In the case of *non-blocking* code, it is possible to obtain a proof of this form: by definition, a non-blocking thread does not need the progress of another thread in order to terminate. For non-blocking code, a rely/guarantee protocol that only asserts safety facts about the extent of the interference of the threads is all that is needed to prove termination. This is exploited by virtually all the program logics which prove total specifications for non-blocking programs [5, 8, 21, 32]. The non-blocking

case allows the use of a while rule which is essentially the one of total Hoare logics:

$$\frac{\forall \beta \leq \beta_0. \vdash \{P(\beta) \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma < \beta\}}{\vdash \{P(\beta_0)\} \mathbf{while}(\mathbb{B})\{\mathbb{C}\} \{\exists \gamma. P(\gamma) \wedge \neg \mathbb{B} \wedge \gamma \leq \beta_0\}} \text{WHILENB}$$

Here β is an *ordinal-valued variant* which is shown to strictly decrease after each iteration. By well-foundedness of ordinals, there can only be finitely many iterations, and hence the loop terminates. However, this rule is completely inadequate for blocking code: in our example, the loop of \mathbb{C}_1 admits no variant, since the iterations do not achieve any sort of progress. Indeed, none of the cited works can handle this simple example. Reasoning about progress for blocking programs requires a whole set of new reasoning principles, and a genuine extension of rely/guarantee with liveness information.

In TaDA Live, the while rule has a more general form:¹

$$\frac{\begin{array}{c} \square L \Rightarrow \diamond \square T \\ \forall \beta \leq \beta_0. \vdash \{P(\beta) \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma \leq \beta\} \\ \forall \beta \leq \beta_0. \vdash \{P(\beta) * T \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma < \beta\} \end{array}}{\vdash \{P(\beta_0) * L\} \mathbf{while}(\mathbb{B})\{\mathbb{C}\} \{\exists \gamma. P(\gamma) * L \wedge \neg \mathbb{B} \wedge \gamma \leq \beta_0\}} \text{WHILEB}$$

The crucial difference is that the rule uses a set of *target* states T : when an iteration starts in a target state, the variant must be shown to strictly decrease, $\gamma < \beta$ (i.e. the iteration needs to produce measurable progress); when an iteration starts from a non-target state, the variant is only required not to increase, $\gamma \leq \beta$ (i.e. no progress is undone). These two conditions alone do not prove the termination of the loop: the execution may be constantly in a non-target state. In our example, the T is $\mathbf{ex}_r(x, 1)$. To conclude that the loop terminates, the first premise requires $\square L \Rightarrow \diamond \square T$: that is, in traces where L holds constantly, with the help of the environment, we will be *eventually always* in a target state. The assertion L captures facts that hold at any point of the iterations of the loop, as it is in the triple of the conclusion but framed off the triples in the premises. When T finally happens, by fairness of the scheduler the loop will execute, and will do so from a state where, by the third premise, the iterations will make progress towards termination.

To make this reasoning work, the first problem we encounter is that none of the information in a standard rely/guarantee specification supports proving $\diamond \square T$. Indeed, nothing in the protocol defined by \mathbf{ex} expresses the idea that at some point the environment will help \mathbb{C}_1 by setting x to 1. A safety rely merely expresses that an update is *allowed*, not that it will be eventually executed. In other words, a safety rely alone is too imprecise an abstraction: it cannot distinguish between environments that make the local thread terminate from the ones that do not. The first question we have to answer is: *how can “help” from the environment be represented in a rely/guarantee proof?*

Innovation 1: Subjective Obligations For Liveness Invariants

Safety arguments are centred around *invariants*: that is, facts of the form *always* P , encoded using regions in TaDA. TaDA Live’s basic observation is that to represent help from the environment, all that is needed is *liveness invariants*: that is, facts of the form *always eventually* P . By combining liveness invariants and safety invariants one can encode more complex progress conditions such as $\diamond \square T$. To represent liveness invariants in a thread-local way, TaDA Live introduces a new kind of ghost state called *obligations*. Similarly to guards, they form a PCM. The interference protocol is augmented by a component that explains how an update affects the obligations. In our example, we want to represent the liveness invariant always eventually $\mathbf{ex}_r(x, 1)$ which, together with the

¹We simplify the rule for this introductory section, informally using the standard LTL notation $\square P$ for *always* P (i.e. P holds at every point of a trace) and $\diamond P$ for *eventually* P (i.e. P holds at some point of a trace). The full rule is given in Section 4.6.

invariant that x can only be set to 1, implies $\diamond\Box(\text{ex}_r(x, 1))$. We therefore introduce an obligation \mathbf{u} (for **u**ppdate-to-1), where again $\mathbf{u} \bullet \mathbf{u}$ undefined captures exclusivity, and extend the protocol to link \mathbf{u} to the update:

$$\mathbf{E} : ((x, v), \mathbf{u}) \rightsquigarrow ((x, 1), \mathbf{0}) \quad (2)$$

This transition to update the region can be executed by a thread with both the \mathbf{E} guard and the \mathbf{u} obligation; the effect of the update is to “consume” the \mathbf{u} resource, as the obligation resulting from the update is the unit $\mathbf{0}$. We say the update *fulfils* the obligation \mathbf{u} .

A safety rely, as expressed by specification (1), says: verify a thread under the assumption that the environment steps will obey the protocol. As a first approximation, our liveness rely, as expressed by (2), additionally says: verify a thread under the assumption that the environment will always eventually fulfil the obligations it owns. (We will refine this idea in the next section to avoid unsound circular reasoning). We say an obligation O is *assumed live* if the environment always eventually fulfils O . In other words: if, at any time, the environment owns O , it eventually fulfils O .

This idea introduces a complication; we need to locally keep track of which (relevant) obligations are owned by the environment, in order to make use of the liveness rely assumption. We solve this problem by taking inspiration from the concept of subjective separation of [29]. We introduce subjective obligation assertions: local obligations, $\lfloor \mathbf{u} \rfloor_r^L$, asserting local ownership of the obligation \mathbf{u} associated with region r , and environmental obligations, $\lfloor \mathbf{u} \rfloor_r^E$, asserting environment ownership of the obligation \mathbf{u} . What makes these assertions interesting is the way they compose: that is, $\lfloor \mathbf{u} \rfloor_r^L \Leftrightarrow \lfloor \mathbf{u} \rfloor_r^L * \lfloor \mathbf{u} \rfloor_r^E$. If we start with local obligation \mathbf{u} and we want to fork into two threads, we use $*$ to give responsibility of \mathbf{u} to one thread and knowledge that the environment has this responsibility to the other.

To complete the proof sketch for our example, we first need to extend the region interpretation by adding the obligation protocol:²

$$\mathcal{I}(\text{ex}_r(x, v)) \triangleq x \mapsto v * (v = 1 \Rightarrow \lfloor \mathbf{u} \rfloor_r^L)$$

When the value at x is 1, the obligation \mathbf{u} is owned by the interpretation, and hence owned by no thread. A thread owning \mathbf{u} and setting x to 1 fulfils the obligation precisely by leaving it inside the interpretation. There is no other way of losing ownership of an obligation because we adopt a classical interpretation of separation: that is, $P * \lfloor \mathbf{u} \rfloor_r^L \not\Rightarrow P$. For soundness, the interpretation of a region with id r is only allowed to own obligations of r .

The TaDA Live proof starts by using viewshift to transform the resource in the precondition into this new region that is shared between the two threads:

$$\exists v. x \mapsto v \quad \Rightarrow \quad \exists r. (\exists v. \text{ex}_r(x, v) * \lceil \mathbf{E} \rceil_r * (v \neq 1 \Rightarrow \lfloor \mathbf{u} \rfloor_r^L)) \quad \equiv \quad \exists r. (P_1 * P_2)$$

where $P_1 = \exists v. \text{ex}_r(x, v) * v \neq 1 \Rightarrow \lfloor \mathbf{u} \rfloor_r^E$ and $P_2 = \exists v. \text{ex}_r(x, v) * \lceil \mathbf{E} \rceil_r * v \neq 1 \Rightarrow \lfloor \mathbf{u} \rfloor_r^L$ are the preconditions of the proofs of \mathbb{C}_1 and \mathbb{C}_2 respectively. To discharge $\diamond\Box(\text{ex}_r(x, 1))$ in the proof of the while loop of \mathbb{C}_1 , we can use $L = \exists v. \text{ex}_r(x, v) * v \neq 1 \Rightarrow \lfloor \mathbf{u} \rfloor_r^E$, which holds throughout the loop: if we are in a state $\text{ex}_r(x, v)$ either $v = 1$, in which case we are in a target state and the value of v will remain 1 forever; or $v \neq 1$ in which case we know $\lfloor \mathbf{u} \rfloor_r^E$. By the liveness rely, when the environment owns \mathbf{u} , it will eventually fulfil it, which by (2) can only be done by setting $v = 1$. Section 4 explains in detail how this argument is carried out formally in TaDA Live.

At this point, we are able to prove the *total* triples $\vdash \{P_1\} \mathbb{C}_1 \{\text{True}\}$ and $\vdash \{P_2\} \mathbb{C}_2 \{\text{True}\}$. However, the standard parallel rule is unsound in the sense that the two triples can be proven even with $\mathbb{C}_2 = \text{skip}$ but, in this case, the parallel composition would not terminate! TaDA Live’s parallel rule can recover soundness by checking that the postconditions of the two threads do not own

²The assertion $\mathbb{B} \Rightarrow Q$ stands for $(\mathbb{B} \wedge Q) \vee (\neg \mathbb{B} \wedge \text{emp})$.

pending obligations, which we can show by proving the stronger triples $\vdash \{P_1\} \mathbb{C}_1 \{ex_r(x, 1)\}$ and $\vdash \{P_2\} \mathbb{C}_2 \{ex_r(x, 1) * [E]_r\}$. This condition is too restrictive in general, and we will relax it appropriately in the next section.

Innovation 2: Obligation Layers To Avoid Circular Arguments

Structuring liveness invariants through obligations, as sketched, presents a significant problem for soundness due to the possibility of making unsound circular liveness assumptions. Consider the following variant of our busy-waiting example:

$$\mathbb{C}'_1 \left\{ \begin{array}{l} \text{var } v_1 = 0 \text{ in} \\ \text{while}(v_1 \neq 1) \{ \\ \quad v_1 := [x_1] \\ \} \\ [x_2] := 1 \end{array} \right\} \parallel \left\{ \begin{array}{l} \text{var } v_2 = 0 \text{ in} \\ \text{while}(v_2 \neq 1) \{ \\ \quad v_2 := [x_2] \\ \} \\ [x_1] := 1 \end{array} \right\} \mathbb{C}'_2$$

There are two shared heap cells at x_1 and x_2 respectively. The thread on the left (\mathbb{C}'_1) is busy-waiting on x_1 which is supposed to be set by the thread on the right (\mathbb{C}'_2), and vice versa, causing a classic high-level deadlock³ situation: the program does not terminate.

Let us try to replicate the argument we used for the busy-waiting example. We require a region sharing both cells, $\text{dex}_r(x_1, x_2, v_1, v_2)$, where v_i is the value stored at x_i . We use two guards E_1 and E_2 , and two obligations, U_1 and U_2 linked to the update of x_1 and x_2 respectively:

$$E_1 : ((x_1, x_2, v_1, v_2), U_1) \rightsquigarrow ((x_1, x_2, 1, v_2), \mathbf{0}) \quad (3)$$

$$E_2 : ((x_1, x_2, v_1, v_2), U_2) \rightsquigarrow ((x_1, x_2, v_1, 1), \mathbf{0}) \quad (4)$$

Without additional precautions, we would be able to derive the triples (for $i = 1, 2$)

$$\vdash \{P_i\} \mathbb{C}'_i \{ \text{dex}_r(x_1, x_2, 1, 1) * [E_i]_r \} \quad (5)$$

where $P_i = \exists v_1, v_2. \text{dex}_r(x_1, x_2, v_1, v_2) * [E_i]_r * (v_i \neq 1 \Rightarrow [U_i]_r^E) * (v_{3-i} \neq 1 \Rightarrow [U_{3-i}]_r^L)$. Given the interpretation we sketched earlier, these triples mean: thread i terminates provided its environment (i.e. thread $3 - i$) always eventually fulfils obligation U_{3-i} . This leads, in the application of the parallel rule, to an unsound circular argument: to show thread i fulfils obligation U_i , thread i is relying on the assumption about the eventual fulfilment of U_{3-i} by the environment, which in turn relies on the eventual fulfilment of U_i by thread i itself. The question is then: *how can we rule out circular arguments, while keeping the proof thread-local?* In particular, we want a solution that allows us to keep the abstraction of the environment as local and abstract as possible, without revealing unnecessary structure of the other threads.

Our solution is to specify dependencies between liveness invariants. We do this by imposing a partial order on obligations: each obligation O is associated with a *layer*, denoted $\text{lay}(O)$, which is an element of a user-defined well-founded partial order, \mathcal{L} . Using layers, we can refine our reasoning principle and gain soundness: to be allowed to assume O is live, one has to show all the locally owned obligations have layers greater than $\text{lay}(O)$. The intuition is that local fulfilment of O_2 can depend on the environment's fulfilment of O_1 *only if* $\text{lay}(O_1) < \text{lay}(O_2)$.

In our deadlocking example, layers expose the circularity issue and prevent the triples (5) from being derivable. Specifically, the proof of the loop of \mathbb{C}'_1 requires us to prove $\diamond \square (\text{dex}_r(x_1, x_2, 1, _))$. At this point we are continuously holding the obligation U_2 so, to be able to assume U_1 live, we

³This liveness form of deadlock is also known as “livelock” since every thread is always taking steps, although no global progress is made by any of those steps. This is not to be confused with the safety property of “global” deadlock as found in languages with blocking primitives.

require $\text{lay}(u_1) < \text{lay}(u_2)$. However, the proof of the loop of C'_2 would require the symmetric constraint, $\text{lay}(u_2) < \text{lay}(u_1)$, leading to a contradiction.

If we replace C'_2 with $C''_2 \triangleq ([x_1] := 1; \text{var } v_2 = 0 \text{ in while}(v_2 \neq 1)\{v_2 := [x_2]\})$, the program $C'_1 \parallel C''_2$ terminates and indeed the proof goes through with $\text{lay}(u_1) < \text{lay}(u_2)$. This is because the first instruction of C''_2 fulfils u_1 so the loop no longer constantly owns it whilst assuming u_2 live. The structure of C''_2 does not impose any dependency on the two liveness invariants.

The generalisation of the liveness rely to use obligations with layers, enables us to give a general parallel rule: instead of just forbidding pending obligations in the postconditions, we require that the postcondition of each thread only owns obligations with layers greater than the layers of obligations assumed live in the other thread’s proof.

Let us contrast our layered obligations with other solutions found in the literature. The LiLi logic cannot verify the above examples as it lacks support for parallel composition.⁴ LiLi’s while rule does share the same high-level structure as **WHILEB**, a structure that can be traced as far back as [37]. The main crucial difference is in how $\diamond\Box T$ is proven. LiLi proposes the idea of *definite actions*, a reincarnation of “leads-to” assertions of [37], to build a liveness rely. Definite actions require the identification of a logical global “queue” of threads where the thread at the front is always able to execute its action and that action implies global progress. In LiLi, the target states are the ones where the local thread is at the head of this queue, and the $\diamond\Box T$ condition is proven by showing that when the head of the queue executes an action, there is some local well-founded progress measure that decreases. Definite actions have a number of drawbacks:

- they require heavy introduction of ghost code for manipulating globally shared ghost state in order to construct the queue of threads; and
- the progress reasoning on the queue requires analysing all possible ways the other threads may finally produce the target states.

Layered obligations are key to resolving these problems:

- they remove the need for ghost code altogether, and make liveness invariants local using the local/environmental obligation assertions; and
- by only relying on the eventual fulfilment of layered obligations, the proof of $\diamond\Box T$ can ignore *how* the environment is going to implement such fulfilment; the only important fact to retain about the *how* is which liveness invariants are assumed to guarantee the fulfilment.

There has been work on proving various safety (e.g. global deadlock-freedom) [16, 27] and progress (e.g. deadlock-freedom, termination) [3, 22, 26, 28] properties of concurrent programs, which assume the only source of blocking behaviour comes from the use of blocking primitives (e.g. built-in locks or channels). Although none of them can handle busy-waiting patterns like our previous examples, they typically detect deadlocks using “tokens” (often also called *obligations*) which represent the responsibility to call a blocking primitive. These tokens are arranged in an acyclic graph of dependencies. Superficially, these tokens are related to our layered obligations in that they both are devices to rule out cyclical dependencies. There are, however, deep differences between the two. Tokens are linked (ad hoc in the operational semantics and through ghost code) to blocking primitive operations calls, and dependencies between the tokens represent causal dependencies between these primitive *events*. By contrast, our layers represent dependencies between *liveness assumptions*, and reflect a purely logical structure. This makes our layered obligations particularly general and flexible: they are able to express arbitrary high-level blocking patterns and not just primitive blocking operations, enabling truly abstract specifications.

⁴Indeed, LiLi’s goal is limited to proving that a module’s implementation refines its specification. The code of the module cannot fork threads but any multi-threaded client of the module is guaranteed not to be able to distinguish the implementation from the specification.

Spin Lock	CLH Lock
<pre> 1 def lock(x){ 2 var d=0 in 3 while(d=0){ 4 d:=CAS(x,0,1) 5 } 6 } 7 8 def unlock(x) { [x] := 0 } 9 10 def makeLock() { var x in 11 x := alloc(1); 12 [x] := 0; 13 ret := x 14 } </pre>	<pre> 1 def lock(x) { var c,p,v in 2 c := alloc(1); [c] := 1; 3 p := FAS(x+1, c); 4 v := [p]; 5 while(v≠0) { v := [p] } 6 [x] := c; 7 dealloc(p) 8 } 9 def unlock(x) { var h in h := [x]; [h] := 0 } 10 def makeLock() { var x,h in 11 h := alloc(1); [h] := 0; 12 x := alloc(2); [x] := h; [x+1] := h; 13 ret := x 14 } </pre>

Fig. 1. Two fine-grained lock implementations.

Innovation 3: Abstract Atomic Specifications for Blocking Operations

Understanding blocking behaviour as the need for an abstraction of the environment that includes liveness invariants unlocks a novel approach in giving abstract, precise and reusable total specifications for abstractly atomic operations. Building on Total TaDA, we propose a new specification format that expresses the atomic effect of a linearizable operation, and succinctly states the liveness invariant required for ensuring termination, at the right level of abstraction. To see the problem and our solution, let us consider the paradigmatic example of two fine-grained implementations of a lock module.

Two Lock Implementations. Consider the *spin lock* and the *CLH lock* given in Fig. 1. The implementations enable threads to compete for the acquisition of a lock at address x by running concurrent invocations of the $\text{lock}(x)$ operation. Only one thread will succeed, leaving the others to wait until the $\text{unlock}(x)$ operation is called by the winning thread.

The primitive commands, such as assignment, lookup and mutation, are primitive atomic and non-blocking: every primitive command, if given a CPU cycle, will terminate in one step. Since reads and writes may race, the language is equipped with a *compare-and-swap* primitive command, $\text{CAS}(x, v_1, v_2)$, which checks if the value stored at x is v_1 : if so, it atomically stores v_2 at x and returns 1; otherwise it just returns 0. Similarly, the *fetch-and-set* primitive command, $\text{FAS}(x, v)$, stores v at x returning the value that was stored at x just before overwriting it.

The spin lock in Fig. 1 is standard. Its state comprises a heap cell at x which stores either 0 (unlocked) or 1 (locked). The Craig-Landin-Hagersten (CLH) lock [18] in Fig. 1 serves threads competing for the lock in a FIFO order. It queues requests, keeping a head and a tail pointer (at x and $x+1$ respectively). The predecessor pointers are stored in each thread's local state (in p). The lock can be acquired by a thread once its predecessor signals release of the lock by setting its queue node to 0. Unlocking the lock corresponds to setting the queue's head node value to 0.

Let us focus on the lock operation of the CLH lock. The interesting aspect is that lock displays blocking behaviour that is observable by the client of the module (it is indeed the quintessence of blocking). We cannot just provide a total triple for it: the operation does not always terminate. The challenge is to design a specification format that accurately captures the abstract functionality of the operation and its subtle termination properties.

First off, one would like a specification that hides the implementation details and only exposes the abstract state of the lock to the client: a lock instance is represented by an abstract resource $L(x, l)$ ⁵ where $l = 1$ indicates the lock is locked, and $l = 0$ means it is unlocked. It is worth noting that traditional Hoare triples are not able to represent the useful behaviour of $\text{lock}(x)$. The triple $\vdash \{L(x, 0)\} \text{lock}(x) \{L(x, 1)\}$ requires the client to establish that the lock is unlocked *before* calling the operation, defying the very purpose of the operation’s functionality. The triple $\vdash \{L(x, 0) \vee L(x, 1)\} \text{lock}(x) \{L(x, 1)\}$ allows the operation to be called in the locked state, but is not precise enough since the same triple holds for a simple assignment $[x] := 1$. It does not express the property that, upon termination of the operation, we can claim that we have acquired the lock. A partial specification of a lock is already a challenge; a total specification more difficult still.

Proposed solutions in the literature can be divided into history-based, refinement-based and abstract atomicity-based approaches. The history-based approach (e.g. [38] for safety, [15, 25] for progress) is expressive but at the price of complex and indirect specifications; the verification requires explicit manipulation of the histories, complicating client reasoning. The only progress-aware refinement-based approach that can modularly verify the CLH lock is the LiLi logic [31]. LiLi’s refinement \sqsubseteq is progress-preserving and contextual, allowing the result to be reused in arbitrary client contexts. For example, the LiLi proof for CLH lock (under weak fair scheduling) shows that

$$\text{lock}(x) \sqsubseteq \text{spec_lock}(x)$$

where $\text{spec_lock}(x)$ is defined (in pseudocode) as⁶

```
spec_lock(x) {
  enqueue(x.queue, self);
  await (head(x.queue) = self ∧ x.state = 0) {
    ⟨x.state := self; x.queue := tail(x.queue)⟩
  }
}
```

The abstract state of the lock is represented by $x.state$, but to represent the fact that threads will not be starved, an abstract FIFO queue at $x.queue$ keeps track of the threads to be served; **self** is the thread id of the caller. The command $\text{await}(\mathbb{B})\{\mathbb{C}\}$ is a blocking *primitive* introduced to express the non-primitive blocking of the implementation. The potential absence of progress of the implementation’s busy-waiting steps is represented by potential absence of a step ahead in the specification.

LiLi’s specification style has three major drawbacks:

- (1) the specification code is not much simpler than the original implementation, and is not able to hide the implementation detail of the thread queue;
- (2) the specification code is *not atomic*: it produces one step for entering the queue, and one step for acquiring the lock;
- (3) since the termination properties are represented through the behaviour of code, a client proof that wants to make use of these properties must reprove them on the specification code before being able to use them in the argument.

These problems limit the abstraction capabilities, proof reuse and scalability of the approach.

The abstract atomicity approach has been pioneered by the TaDA logic. It directly influenced logical atomicity in Iris [24], and was extended to provide total specification for non-blocking programs in Total TaDA [8]. The aim of the TaDA approach is to keep the Hoare-triple style of

⁵We omit the region identifier to simplify the discussion.

⁶In [31], this is the result of applying the appropriate wrapper to the lock specification: $\text{wr}_{\text{PSF}}^{\text{fair}}(\text{await}(1=0)\{l := \text{cid}\})$.

specification, whilst being able to give precise and abstract specifications to fine-grained code like CLH lock. The TaDA solution is to provide a Hoare triple for `lock` which embraces the fact that, between the invocation of the operation and the execution of the atomic update of the lock, there is a phase of *interference* where the environment can change the value of the lock. It is important to be able to distinguish the imprecise precondition that holds during the interference phase, $L(x, 0) \vee L(x, 1)$, and the precise precondition, $L(x, 0)$, that holds *just before* the atomic update performed by the lock operation at its *linearization point* [20].

The TaDA safety specification for `lock` is the partial *atomic triple*:

$$\vdash \forall l \in \{0, 1\}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle \quad (6)$$

The *interference precondition* $\forall l \in \{0, 1\}. \langle L(x, l) \rangle$ describes the interference phase. It states that the environment must preserve the existence of the lock at x but may change the value of l , and the implementation of the lock must tolerate these environmental changes. The *pseudo-quantifier* $\forall l \in \{0, 1\}$ is unusual, behaving like an evolving universal quantifier in that the environment is able to keep changing l over time and behaving like an existential quantifier in that the implementation can assume that the lock always exists with $l \in \{0, 1\}$. The triple (6) states that, if the environment satisfies the interference precondition and the operation terminates, then the implementation guarantees that, just before the linearization point, the lock must have been available for locking ($l = 0$) and, just afterwards, the lock has been locked by the operation ($L(x, 1)$). Exclusive ownership of the lock after the operation terminates can be derived from the $l = 0$ assertion in the postcondition: just before we locked it, nobody else could claim that they owned the lock. The TaDA safety specification for `unlock` is the partial atomic triple

$$\vdash \forall l \in \{1\}. \langle L(x, l) \rangle \text{unlock}(x) \langle L(x, 0) \rangle$$

This triple⁷ states that, to be used correctly, the `unlock` operation requires the lock to be locked and not changed by the environment during the interference phase; in return, the operation promises to atomically set the lock to be unlocked.

TaDA Live builds on the TaDA specification format. To turn the TaDA triple for `lock` into a total specification, the termination guarantee must depend on the environment: if the environment decides to hold the lock indefinitely, no lock implementation should allow the `lock` operation to terminate. Hence, we express blocking as a *liveness condition* on the *environment* during the interference phase of an abstractly atomic operation. The CLH lock operation will terminate under weak fairness, provided that, if the lock is locked by the environment during the interference phase, the environment will *eventually* unlock it. In general, a blocking operation will require an environment that is *live*: it will always eventually bring the abstract state to a *good* (e.g. unlocked) state.

The TaDA Live total specification of the CLH lock operation is:

$$\vdash \forall l \in \{0, 1\} \rightarrow \{0\}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle \quad (7)$$

The interference precondition is $\forall l \in \{0, 1\} \rightarrow \{0\}. \langle L(x, l) \rangle$ with the pseudo-quantifier now incorporating the environment liveness condition. As well as stating that the environment can keep changing the lock, the interference precondition also states that if the lock is in a bad state ($l \in \{0, 1\} \setminus \{0\}$) then the environment must always eventually change it to a good state ($l \in \{0\}$). The implementation needs to ensure termination under the assumption that the lock always eventually returns to the unlocked state. Note that the environment is allowed to change l back to 1 arbitrarily many times, provided it always eventually sets it back to 0. To see why this is enough to ensure termination, consider Fig. 2(a) where we chart the evolution of the abstract state induced by a live environment in the interference phase of `lock`. Progress towards termination of `lock`

⁷We typically omit the pseudo-quantifier in the case where the set has just one element, e.g. $\vdash \langle L(x, 1) \rangle \text{unlock}(x) \langle L(x, 0) \rangle$.




Fig. 2. Live environment (a); measure of progress for CLH lock where q is the number of threads ahead in the queue (b); live environment with bounded impedance (c); measure of progress for spin lock (d).

is guaranteed by the progress measure charted in Fig. 2(b): every time the environment unlocks, the value of l decreases from 1 to 0; when the environment locks, although l increases to 1, the number q of threads in front of us in the queue decreases. One crucial aspect of our specification design is that we do not want to expose the progress argument to the client *unless part of the argument needs to be made by the client*. With CLH, the part of the argument appealing to the queue of threads is completely internal to the implementation of the operation, while the argument for the environment’s liveness must be provided by the client (the implementation has no power over this). We prove this formally in Section 5.

Now let us consider the spin lock implementation. The spin lock operation cannot promise to terminate just by relying on a live environment. The problem is that when the environment locks the lock, there is no measure of progress that decreases: we are genuinely delayed by this action. We call this effect *impedance*. We conceptualise impedance as a greater *leaking* of the progress argument to the client. In the spin lock example, the whole of the progress argument needs to be provided by the client: the client needs to ensure that the environment will always eventually unlock the lock, and that it will only impede the operation a bounded number of times. To represent this extra *bounded impedance* requirement (depicted in Fig. 2(c)), we extend the abstract state of the lock with an ordinal α , an *impedance budget* that strictly decreases when the lock state is set to 1. We arrive at the following TaDA Live specification for spin lock:

$$\forall \phi. \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}, \alpha. \langle L(x, l, \alpha) \wedge \phi(\alpha) < \alpha \rangle \text{lock}(x) \langle L(x, 1, \phi(\alpha)) \wedge l = 0 \rangle \quad (8)$$

The lock is now represented by the predicate assertion $L(x, l, \alpha)$ with ordinal α , which can also be changed by the environment during the interference phase. As well as expressing the dependency on a live environment on l , this triple states that every lock operation consumes the budget α by a non-trivial amount, thus providing a logical measure of progress from good to bad states. The initial value of the budget and the function ϕ from ordinals to ordinals is determined by the client, which must demonstrate that the budget is enough to make all its calls.

The TaDA Live total specification of `unlock` for the CLH lock is the same as the TaDA partial specification. By contrast, the TaDA Live specification of `unlock` for the spin lock needs to incorporate the ordinals: $\vdash \langle L(x, 1, \alpha) \rangle \text{unlock}(x) \langle L(x, 0, \alpha) \rangle$. The impedance budget α is preserved by `unlock`. This encodes the fact that `unlock` does not impede the other operations, but also that by unlocking we cannot increase the budget. By combining these assumptions about the budget (it decreases when locking, stays constant when unlocking), it is possible to conclude that the implementation of the spin lock terminates using the progress measure in Fig. 2(d). Crucially, for spin lock, the *whole* of the progress argument is provided by (and thus visible to) the client.

The impedance budget technique was first introduced to concurrent separation logics for non-blocking operations in Total TaDA [8]. Here, we smoothly integrate ordinals into TaDA Live that fully supports blocking.

2.1 Abstraction and proof reuse

The TaDA Live program logic works with *hybrid* triples of the form:

$$\vdash \forall x \in X \rightarrow X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \langle Q_h(x) \mid Q_a(x) \rangle$$

which generalises both Hoare triples and abstract atomic triples. This triple comprises: a pseudo-quantifier with its environment liveness condition; atomic pre/post-conditions $P_a(x)$ and $Q_a(x)$; and Hoare pre/post-conditions, P_h and $Q_h(x)$. The Hoare pre/post-conditions describe stable resources that are owned locally by \mathbb{C} and can be updated non-atomically. Hoare triples correspond to the case where $X = X' = \{1\}$ and $P_a = Q_a = \text{emp}$. Abstract atomic triples correspond to the case when P_h and $Q_h(x)$ are empty. We have omitted some details from the hybrid triples, such as layers and levels, since they are not important for the ideas of this section; the full details are given in Section 3.8.

The integration of the liveness annotations in triples achieves the goal of keeping the specification abstract and atomic. To obtain the goal of reuse of proofs, there are two missing ingredients: a mechanism to make use of the $X \rightarrow X'$ assumption in a proof of an implementation of the specification; and a way to reuse the specification in an arbitrary client context.

Imagine proving the CLH lock implementation correct with respect to specification (7). The while loop needs to discharge that “finally, the current thread is at the head of the queue, and the lock is unlocked”. This can only be proven with the help of the $l \in \{0, 1\} \rightarrow \{0\}$ liveness assumption coming from the lock specification. To this end, in addition to liveness assumptions given by obligation assertions, TaDA Live extends judgements to allow contexts with $X \rightarrow X'$ liveness assumptions, used to discharge the $\diamond \square T$ condition in the while rule. The full details are given in Section 3.8.

Now consider proving a client of a lock using the specifications of the lock operations for the calls to these operations. This requires the LIVENESS Check rule:

$$\frac{\begin{array}{l} \square L \Rightarrow \square \diamond T \quad \forall x \in X. \vdash P_a(x) * T \Rightarrow x \in X' \\ \vdash \forall x \in X \rightarrow X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \langle Q_h(x) \mid Q_a(x) \rangle \end{array}}{\vdash \forall x \in X. \langle P_h * L \mid P_a(x) \rangle \mathbb{C} \langle Q_h(x) * L \mid Q_a(x) \rangle} \text{LIVEC'}$$

The rule’s crucial effect is to remove the liveness annotation $X \rightarrow X'$, which can only be done in a situation where the corresponding liveness assumption $\square(x \in X) \Rightarrow \square \diamond(x \in X')$ is satisfied. Just like the **WHILEB** rule, we frame an assertion L which is information that holds for the duration of the call. Typically, L asserts the existence of some shared region and that the environment holds some obligations depending on the state of the region. We also need to provide a set of target states T capturing when $x \in X'$ (second premise). The crucial check of the rule is the first premise, which examines the traces where L holds everywhere, and asks us to prove that in those traces we see T satisfied infinitely often (and thus $x \in X'$ infinitely often). If that is true, we can conclude that the command terminates in the current context without the extra assumption in the pseudo-quantification. The resulting triple can be then manipulated using standard TaDA reasoning.

Take the typical use of (CLH) locks $\mathbb{C} = \text{lock}(x); \dots; \text{unlock}(x)$ in a client $\mathbb{C} \parallel \dots \parallel \mathbb{C}$. To share the lock resource $L(x, l)$, the client proof would specify some region $\mathbf{client}_r(x, l)$ where l is the abstract state of the lock. A typical client would include the abstract state of other shared resources too, but for simplicity we focus here on the lock. The client needs to specify in its protocol

that the lock will be always eventually unlocked by the threads sharing it. We therefore introduce an exclusive obligation κ (the **key** of the lock) which is obtained when locking the lock and fulfilled when unlocking it:

$$((x, 0), \mathbf{0}) \rightsquigarrow ((x, 1), \kappa) \qquad ((x, 1), \kappa) \rightsquigarrow ((x, 0), \mathbf{0})$$

The protocol is mirrored in the region’s interpretation $\mathcal{I}(\mathbf{client}_r(x, l)) \triangleq L(x, l) * (l = 0 \Rightarrow \lfloor \kappa \rfloor_r^l)$. With the application of standard TaDA reasoning, it is possible to derive

$$\frac{\vdash \forall l \in \{0, 1\} \rightarrow \{0\}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle}{\vdash \forall l \in \{0, 1\} \rightarrow \{0\}. \langle \text{emp} \mid \mathbf{client}_r(x, l) \rangle \text{lock}(x) \langle \lfloor \kappa \rfloor_r^l \mid \mathbf{client}_r(x, 1) \wedge l = 0 \rangle}$$

which amounts to saying that if $\text{lock}(x)$ atomically locks the lock region, then it also atomically updates the client region containing the lock. Notice that the $\{0, 1\} \rightarrow \{0\}$ annotation is propagated as is. In other words, the update on the lock is put in the context of the current client. In such context, we can set the frame L to be $\exists l \in \{0, 1\}. \mathbf{client}_r(x, l) * l = 1 \Rightarrow \lfloor \kappa \rfloor_r^E$: according to the protocol of the current client, the environment holds an obligation κ when $l = 1$. Because of the liveness invariant encoded by κ , it is true that the environment will always eventually unlock the lock, allowing us to discharge the side condition of **LIVEC’**:

$$\square(\exists l \in \{0, 1\}. \mathbf{client}_r(x, l) * l = 1 \Rightarrow \lfloor \kappa \rfloor_r^E) \Rightarrow \square \diamond (\mathbf{client}_r(x, 0))$$

Indeed, if $l = 1$ the precondition gives us $\lfloor \kappa \rfloor_r^E$ which means that the environment owns κ and will therefore eventually fulfil it, which can only be done by setting $l = 0$. The environment is allowed to then lock x again, but that is fine: as we discussed, a CLH lock can promise termination under this milder condition.

Thanks to the smooth integration of liveness annotations in the specifications and liveness invariants expressed as obligations, TaDA Live proofs can properly abstract and encapsulate behaviour. Consider a module implementing a counter which can be safely used concurrently thanks to the internal use of locks to protect access to the shared cell holding the value of the counter. For example, the increment operation can be implemented as

```
def incr(x){var v in lock(x); v := [x+1]; [x+1] := v+1; unlock(x)}
```

While the use of locks involves blocking behaviour, the blocking is handled completely internally and a client of the counter cannot observe it. The TaDA Live specification of the increment operation thus does not leak this implementation detail:

$$\vdash \forall n \in \mathbb{N}. \langle C(x, n) \rangle \text{incr}(x) \langle C(x, n+1) \rangle$$

A client of the counter does not need to worry about the internal blocking since the specification does not entail any liveness proof obligation. The proof of `incr` discharges the liveness assumption of the specification of `lock` by using obligations analogous to κ above, specified in an internal protocol that is not exposed to the client proof. In Section 4.9, we discuss the encapsulation properties of TaDA Live’s specifications in more detail.

Our approach contrasts significantly with previous work [3, 22] where blocking is represented in specifications by the acquisition of tokens acting as obligations. In this work, the specification style fixes an expected protocol to be followed by the client. For example, the axiom for a built-in lock acquisition operation returns a built-in token representing the need for calling a lock release primitive.

In contrast, our lock specification does not impose on the client any particular way in which its environment liveness assumption should be enforced. It is the job of the client to devise a protocol that ensures the environment liveness assumptions of the lock specifications will be provable. For locks, this is indeed often achieved by making sure every thread that locks a lock eventually

unlocks it. Such a protocol is encoded by the liveness invariants of the client’s region (e.g. **client** in the example above) and the κ -obligation pattern. The specification of the lock, however, does not transfer obligations to the client, leaving open the possibility for clients to use completely different protocols. The following example client illustrates the added flexibility of our approach:

$$\begin{array}{l} \text{lock}(x); \\ [y] := 1 \end{array} \parallel \begin{array}{l} \mathbf{var} \ b = 0 \ \mathbf{in} \\ \mathbf{while}(b \neq 1) \{ \ b := [y] \ }; \\ \text{unlock}(x) \end{array}$$

The code assumes a lock has been allocated at x , and y initially stores 0. In the specification style where the expected (liveness) protocol is built-in, the lock call in the left-hand thread would return a built-in token which can only be consumed by calling `unlock`. This in turn requires an extension of the logic –as done e.g. in [17]– providing some mechanism for the sound delegation of tokens from one thread to the other. In TaDA Live, there is no need for such an extension. The protocol of this client does not need to associate obligations with the lock; one can simply define an obligation (owned initially by the left-hand thread) that is fulfilled when y is set to 1, and use it to prove the appropriate environment liveness conditions for the proof.

In this informal overview, we used temporal logic formulas to represent the key liveness conditions in the **WHILEB** and **LIVEC**’ rules. The formal versions of these TaDA Live rules, however, implement those checks with what we call the *environment liveness condition*, which reduces these liveness properties to safety checks via a dedicated set of rules (explained in Section 4). Remarkably, the liveness checks of both rules can be phrased in terms of the environment liveness condition, which therefore provides a uniform proof principle for blocking termination.

2.2 A Guide For the Reader

The rest of the paper proceeds by introducing the assertion language and the semantic model of TaDA Live in full detail (Section 3), then presenting the proof rules through the proof of an example (Section 4), then walking through the proofs of our case studies and commenting on limitations (Section 5), and finally discussing related work (Section 6). A reader interested in the proof rules can skim through Sections 3.3 to 3.5 and 3.8 and the beginning of Section 3.9 to familiarise with the basic definitions, and then move to Section 4 to understand how the rules themselves work, and the typical proof patterns.

3 THE TADA LIVE SEMANTIC MODEL

We introduce the semantic model which justifies TaDA Live, defining:

- the operational semantics of commands and their fair traces;
- the assertion language, regions, guards, obligations and protocols;
- the semantics of assertions and viewshifts;
- the specification format; and
- the trace semantics of specifications.

In Section 2, we introduced hybrid triples which generalise Hoare and atomic triples. For our formal semantics, we separate triples into two components: the command \mathbb{C} , and the specification \mathbb{S} comprising the pseudo-quantifier, the precondition and the postcondition. We introduce the *semantic judgement*, $\vDash \mathbb{C} : \mathbb{S}$ with $\mathbb{S} = \forall x \in X \rightarrow X'. \langle P_h \mid P_a(x) \rangle \cdot \langle Q_h(x) \mid Q_a(x) \rangle$, which captures the semantic properties of a command that satisfies a specification: i.e. safety and termination of its fair traces. This required a complete reformulation of the model of TaDA. First, we give a trace semantics to specifications independently of commands. This enables us to define the semantic judgement to hold when $\llbracket \mathbb{C} \rrbracket \subseteq \llbracket \mathbb{S} \rrbracket$: that is, when the concrete traces of a command are allowed by

$\mathbb{C} ::= \mathbf{skip}$	(skip)	Domains:
$x := \mathbb{E}$	(assignment)	$v \in \text{Val} \triangleq \mathbb{Z} \cup \text{Bool} \supseteq \text{Addr} \triangleq \mathbb{N}$
$x := [\mathbb{E}]$	(read)	$\sigma \in \text{Store} \triangleq \text{PVar} \rightarrow \text{Val}$
$[\mathbb{E}] := \mathbb{E}$	(write)	$h \in \text{Heap} \triangleq \text{Addr} \rightarrow_f \text{Val}$
$x := \mathbf{CAS}(\mathbb{E}, \mathbb{E}, \mathbb{E})$	(compare-and-swap)	$\varphi \in \text{FImpl} \triangleq \text{FName} \rightarrow (\text{PVar}^*, \text{Cmd})$
$x := \mathbf{FAS}(\mathbb{E}, \mathbb{E}, \mathbb{E})$	(fetch-and-set)	$\text{PState} \ni \mathbb{C} ::= \checkmark \mid \mathbb{C} \mid \mathbb{C} \parallel \mathbb{C} \mid \mathbb{C}; \mathbb{C}$
$x := \mathbf{alloc}(\mathbb{E})$	(allocate)	$(\sigma, \mathbb{C}) \mid \mathbf{let} \ f(\vec{x}) = \mathbb{C} \ \mathbf{in} \ \mathbb{C}$
$\mathbf{dealloc}(\mathbb{E})$	(deallocate)	$\mathbf{c} \in \text{PConf} \triangleq (\text{Store} \times \text{Heap} \times \text{PState}) \uplus \{\perp\}$
$\mathbb{C}; \mathbb{C}$	(sequential composition)	
$\mathbb{C} \parallel \mathbb{C}$	(parallel composition)	
$\mathbf{let} \ f(\vec{x}) = \mathbb{C} \ \mathbf{in} \ \mathbb{C}$	(function definition)	
$\mathbf{var} \ x = \mathbb{E} \ \mathbf{in} \ \mathbb{C}$	(local variable binding)	
$\mathbf{if}(\mathbb{B})\{\mathbb{C}\}\mathbf{else}\{\mathbb{C}\}$	(if)	
$\mathbf{while}(\mathbb{B})\{\mathbb{C}\}$	(while loop)	
$x := f(\vec{\mathbb{E}})$	(function call)	
$\langle \mathbb{C} \rangle$	(primitive atomic block)	
$\mathbb{E} ::= v \mid x \mid \mathbb{E} + \mathbb{E} \mid \mathbb{E} * \mathbb{E} \mid \dots$	(numeric expressions)	
$\mathbb{B} ::= b \mid x \mid \neg \mathbb{B} \mid \mathbb{E} \leq \mathbb{E} \mid \dots$	(Boolean expressions)	

Fig. 3. Syntax of commands $\mathbb{C} \in \text{Cmd}$ and basic semantic domains.

the specification traces. This approach is unusual for separation logics based on Hoare-style triples, and brings the semantics nearer to approaches based on refinement. Second, the trace model is an “open-world” semantics where traces include both individual local steps made by the command and individual arbitrary environment steps. Other models typically model the environment interference indirectly, representing a sequence of environment steps as a single big jump. Our “open-world” approach is crucial to capture the assumptions on the liveness of the environment stipulated by the specifications. Third, the trace semantics of the specification is given in a style that is closely related to alternating automata [40]. The specification is seen as an automaton which traverses a concrete trace and only accepts those traces that satisfy the specification. This enables us to cleanly separate the (alternating) safety constraints from the (linear time) liveness constraints, imposed by a specification.

3.1 Notation

We write $X \rightarrow Y$ for the set of partial functions from X to Y , and $X \rightarrow_f Y$ for the set of finite partial functions. Given $f: X \rightarrow Y$, we write $f(x) = \perp$ if f is undefined on x , and $\text{dom}(f) \triangleq \{x \mid f(x) \neq \perp\}$. We write $[x_1 \mapsto y_1, \dots, x_n \mapsto y_n]$ for the finite function that maps each of the x_i to y_i and is undefined on any other input. We write $f[x \mapsto y]$ for the partial function which coincides with f except on x where it returns y , and write $f[x \mapsto \perp]$ analogously. The disjoint union between partial functions $f \uplus g$ is defined if their domains are disjoint. In contexts where the expected type is a function, we write \emptyset for the empty function.

3.2 Fair Trace Semantics of Commands

We present a standard first-order imperative language, called `WHILE`, with shared-memory concurrency and fine-grained non-blocking primitives, and define the fair concrete trace semantics of its commands. Our `WHILE` language is parametrised by the following sets: the *Booleans*, $\text{Bool} \triangleq \{\text{true}, \text{false}\} \ni b$; the *values*, $\text{Val} \triangleq \mathbb{Z} \cup \text{Bool} \ni v$; the *program variables*, $\text{PVar} \ni x, y, \dots$; and the *function names*, $\text{FName} \ni f$. The set PVar contains a special element, `ret`, the name of a local variable that holds a function’s return value.

Definition 3.1 (Commands). The set of commands, $\text{Cmd} \ni \mathbb{C}$, is defined by the grammar in Fig. 3 where $x \in \text{PVar}$, $\vec{x} \in \text{PVar}^*$ is a list of pairwise distinct variables, and $f \in \text{FName}$. The notation $\text{var } x_1, x_2, \dots, x_n \text{ in } \mathbb{C}$ denotes $\text{var } x_1 = \emptyset \text{ in } \text{var } x_2 = \emptyset \text{ in } \dots \text{var } x_n = \emptyset \text{ in } \mathbb{C}$.

We place some restrictions on these commands to simplify exposition. We write $\text{pv}(\mathbb{C})$ for the free program variables of a command. The set $\text{mods}(\mathbb{C})$ is the set of free variables that are potentially modified by a command, i.e. any free x of \mathbb{C} appearing in instructions of the form $x := \dots$; in particular, $\text{mods}(\text{var } x = \mathbb{E} \text{ in } \mathbb{C}) = \text{mods}(\mathbb{C}) \setminus \{x\}$. In a command $\mathbb{C}_1 \parallel \mathbb{C}_2$, we apply the mild syntactic restriction that $\text{mods}(\mathbb{C}_1) = \text{mods}(\mathbb{C}_2) = \emptyset$. Each individual thread is still able to modify variables that are created locally and to modify shared heap cells, but are not allowed to modify the free variables.⁸ In a function definition $\text{let } f(x_1, \dots, x_n) = \mathbb{C}_1 \text{ in } \mathbb{C}_2$, we use the natural restriction $\text{pv}(\mathbb{C}_1) \subseteq \{x_1, \dots, x_n, \text{ret}\}$. Also for simplicity, we assume each function name is given at most one definition. The function $\text{fn}: \text{Cmd} \rightarrow \wp(\text{FName})$ returns the function names occurring in Cmd that are not bound by a **let**. Although function definitions may be recursive, we will disallow recursion in our logical rules to simplify the development. In the programs we consider, all potentially divergent behaviour stems from **while**. It is straightforward to reformulate the **WHILE** rule into a **LET** rule that supports terminating recursion.

Commands manipulate heaps $h \in \text{Heap} \triangleq \text{Addr} \rightarrow_f \text{Val}$ (where $\text{Addr} \triangleq \mathbb{N}$ and \emptyset is the empty heap) and *local variable stores*, $\sigma \in \text{Store} \triangleq \text{PVar} \rightarrow \text{Val}$. A command can contain free function names, so we use a *function implementation context* $\varphi \in \text{FImp} \triangleq \text{FName} \rightarrow (\text{PVar}^*, \text{Cmd})$, to map function names to pairs comprising a finite list of distinct variables (the formal arguments) and a command (the body of the function).

A command induces transitions over *program configurations* $c \in \text{PConf} \triangleq (\text{Store} \times \text{Heap} \times \text{PState}) \uplus \{\frac{!}{!}\}$ which keep track of the current variable store and global heap, and the program states $C \in \text{PState}$ (see Fig. 3) which represent the set of the active threads and their execution state. The $\frac{!}{!}$ program configuration represents a faulty configuration, e.g. the one reached after dereferencing an unallocated address. For the details of program states we refer to Appendix D; what is relevant is that \checkmark is the program state of a terminated thread, and we can define a function $\text{threads}: \text{PConf} \rightarrow \wp(\text{Tid})$ that computes the set of thread identifiers ($t \in \text{Tid}$) of the active threads of a program configuration (details in Appendix).

To model fair traces of commands, we use a small-step operational semantics, parametrised by a function implementation context φ and defined by a relation $\longrightarrow_\varphi \subseteq \text{PConf} \times \text{Sched} \times \text{PConf}$. In a transition (c_1, π, c_2) , the scheduling annotation, $\pi \in \text{Sched}$, keeps track of who executed the step:

$$\text{Sched} \triangleq \{\text{loc}_t \mid t \in \text{Tid}\} \uplus \{\text{env}\}$$

that is, either a local active thread t or the environment. Environment steps can have arbitrary effects on the heap and can generate faults at any time:

$$\frac{h' \in \text{Heap}}{\sigma, h, C \xrightarrow{\text{env}}_\varphi \sigma, h', C} \quad \frac{c \in \text{PConf}}{c \xrightarrow{\text{env}}_\varphi \frac{!}{!}}$$

The full definition of the transition semantics is defined in Fig. 36 and Fig. 37 of the Appendix.

We call *program traces* the infinite sequences of the form $c_0 \pi_0 c_1 \pi_1 \dots$ where, for all $i \in \mathbb{N}$, $c_i \in \text{PConf}$ and $\pi_i \in \text{Sched}$. We use τ to range over infinite suffixes of program traces and PTrace for the set of all program traces. We define the *set of φ -program traces*

$$\text{PTrace}_\varphi \triangleq \{c_0 \pi_0 c_1 \pi_1 \dots \mid \forall i \in \mathbb{N}. c_i \xrightarrow{\pi_i}_\varphi c_{i+1}\}.$$

⁸To lift this restriction, one can use the “variables as resources” technique [2]. Our restriction simplifies the handling of the local state without sacrificing expressivity: any local variable in the scope common to both threads that needs to be modified can be instead implemented by using a shared memory cell.

Definition 3.2 (Fairness). A φ -program trace $(c_0 \pi_0 c_1 \pi_1 \dots) \in \text{PTrace}_\varphi$ is *fair* if:

$$\forall i \in \mathbb{N}. \forall t \in \text{threads}(c_i). \exists j \geq i. (\pi_j = \text{loc}_t \vee c_j = \downarrow) \quad (9)$$

$$\forall i \in \mathbb{N}. \exists j \geq i. \pi_j = \text{env} \quad (10)$$

That is: a trace is fair if, at any point in time, every thread that can take a step (and the environment) will eventually be scheduled.

The open-world program semantics defines the behaviour of a command when run concurrently with an arbitrary environment. It corresponds to the fair program traces of a command, with the information about program states and thread identifiers removed.

Definition 3.3 (Open World Semantics). We call *traces* the infinite sequences $c_0 \pi_0 c_1 \pi_1 \dots$ where, for all $i \in \mathbb{N}$, $c_i \in \text{Conf} \triangleq (\text{Store} \times \text{Heap}) \cup \{\downarrow\}$ and $\pi_i \in \{\text{loc}, \text{env}\}$. We use τ for ranging over infinite suffixes of traces and Trace for the set of all traces. For a trace $\tau = c_0 \pi_0 c_1 \pi_1 \dots$, we define $\tau(i) \triangleq (c_i, \pi_i)$, and $\tau_{/i} \triangleq c_i \pi_i c_{i+1} \pi_{i+1} \dots$. The function $[\cdot]: \text{PTrace} \rightarrow \text{Trace}$ is defined by $[\mathbf{c}_0 \pi_0 \mathbf{c}_1 \pi_1 \dots] \triangleq c_0 \pi_0 c_1 \pi_1 \dots$ where

$$c_i \triangleq \begin{cases} (\sigma, h) & \text{if } c_i = (\sigma, h, _) \\ \downarrow & \text{if } c_i = \downarrow \end{cases} \quad \pi_i \triangleq \begin{cases} \text{loc} & \text{if } \pi_i \in \text{Sched} \setminus \{\text{env}\} \\ \text{env} & \text{if } \pi_i = \text{env} \end{cases}$$

The *open-world program semantics function*, $[\cdot]_\varphi: \text{Cmd} \rightarrow \wp(\text{Trace})$ is defined by

$$[\mathbf{C}]_\varphi \triangleq \{ [\mathbf{c}_0 \tau] \mid (c_0 \tau) \in \text{PTrace}_\varphi, \text{fv}(\mathbf{C}) \subseteq \text{dom}(\sigma_0), c_0 = (\sigma_0, _, \mathbf{C}), c_0 \tau \text{ is fair} \}$$

The notation $[\mathbf{C}]$ is syntactic sugar for $[\mathbf{C}]_\emptyset$.

The goal of TaDA Live is to prove termination of the local command.

Definition 3.4 (Local termination). A trace $\tau \in \text{Trace}$ is *locally terminating*, written $\text{lterm}(\tau)$, if it contains finitely many occurrences of loc .

It might seem odd that our program semantics only contains infinite traces, since our goal is proving termination. Traces that locally terminate simply have an infinite tail of environment steps. To simulate a closed system one can select for the traces where the environment steps are all identity steps.

Remark 1 (On primitive blocking). It is important to remember that the primitives of our programming language are non-blocking, in the sense that they can always take a step if scheduled: for all $h \in \text{Heap}$, $C \in \text{PState}$, for all σ with $\text{dom}(\sigma) \supseteq \text{pv}(C)$, and every $t \in \text{threads}(C)$, there is a $c \in \text{PConf}$ such that $(\sigma, h, C) \xrightarrow{\text{loc}_t}_\varphi c$. Hence, a trace is locally terminating only if all the threads terminated.

For languages which have blocking primitives (e.g. built-in locks/channels), traces may be locally terminating because a non-terminated thread may not have a local successor (i.e. it is not enabled) at any point in the future (e.g. if a built-in lock remains locked forever, an acquire operation would not have local successors). With blocking primitives, fairness also comes in two variants: strong and weak. Strong fairness requires that if an operation is infinitely often enabled it is infinitely often executed. Strong and weak fairness coincide for languages like ours where every primitive is enabled at all times.

Notice that our lack of blocking primitives does not make our setting less general: blocking primitives can be implemented on top of non-blocking ones, both with weak and strong fairness assumptions for termination, as illustrated by our spin and ticket lock examples. In other words, blocking primitives can be given TaDA Live specifications and be treated uniformly by the logic. The addition of built-in blocking primitives to the language does not pose new challenges.

3.3 TaDA Live Assertions and Worlds

We formally introduce the TaDA Live assertion language, and its semantics in terms of its models, called *worlds*. The TaDA Live assertions are built from the standard *classical* connectives and quantifiers of separation logic,⁹ TaDA region and guard assertions, and new TaDA Live obligation and layer assertions. To formalise the assertions, we assume a number of basic domains:

- a set of *logical variables*, denoted LVar , disjoint from PVar ;
- an enumerable set of *region types*, $\text{RType} \ni \mathbf{t}$;
- an enumerable set of *region identifiers*, $\text{RId} \ni r$;
- the set of *levels*, $\text{Lvl} \triangleq \mathbb{N} \ni \lambda$, to stratify regions to avoid the problem of re-entrancy¹⁰ (explained in Remark 2);
- a set of *abstract states*, $\text{AState} \ni a$, including sets and lists of values;
- a set of *guards*, $\text{Guard} \ni G$, which will offer the support for the *guard algebras* defined later;
- a well-founded partial order $(\mathcal{L}, \leq, \top, \perp)$ of *layers*, which will be associated to special guards called *obligations*; and
- a set of ordinals, \mathbb{O} .

For layers, we use the abbreviations $k_1 < k_2 \triangleq (k_1 \leq k_2 \wedge k_2 \not\leq k_1)$ and $k \geq n \triangleq (\forall k' > k. k' \geq n)$. The *set of abstract values* is $\text{AVal} \triangleq \text{Val} \cup \text{AState} \cup \text{Guard} \cup \text{RId} \cup \mathcal{L}$.

As is standard, when used in assertions, we extend numeric and Boolean expressions to use logical variables and abstract values too. A *logical variable store*, $l \in \text{LStore} \triangleq \text{LVar} \rightarrow \text{AVal}$, assigns values to logical variables. Given a logical and a program variable store l, σ , the evaluation of expressions $\mathcal{E}[\![\mathbb{E}]\!]_{l, \sigma} \in \text{AVal}$ and of Boolean expressions $\mathcal{B}[\![\mathbb{B}]\!]_{l, \sigma} \in \text{Bool}$ are standard.

Assertions and worlds are built using partial commutative monoids.

Definition 3.5 (PCM). A (multi-unit) *partial commutative monoid* (PCM) is a tuple (X, \bullet, E) comprising a set X , a binary *partial* composition operation $\bullet: X \times X \rightarrow X$ and a set of unit elements E , such that the following axioms are satisfied (where either both sides are defined and equal, or both sides are undefined):

$$\begin{aligned} \forall x, y, z \in X. \quad (x \bullet y) \bullet z &= x \bullet (y \bullet z) && \text{(associativity)} \\ \forall x, y \in X. \quad x \bullet y &= y \bullet x && \text{(commutativity)} \\ \forall x \in X. \exists e \in E. \quad x \bullet e &= x && \text{(identity)} \end{aligned}$$

For $x, y \in X$, we write $x \# y$ if $x \bullet y \neq \perp$, and $x \sqsubseteq y$ if $\exists x_1. y = x \bullet x_1$. A PCM is *cancellative* when, for any $x, y_1, y_2 \in X$, if $x \bullet y_1 = x \bullet y_2$ then $y_1 = y_2$.

The partial heaps form a PCM $(\text{Heap}, \uplus, \{\emptyset\})$, as standard in separation logics. We also use *guard algebras* and *obligation algebras* which are PCMs for describing auxiliary ghost state, specified by the user of the logic.

Definition 3.6 (Guard Algebras). A *guard algebra* is a PCM $(\text{Grd}, \bullet, \{\mathbf{0}\})$ with $\text{Grd} \subseteq \text{Guard}$. TaDA Live is parametrised by a function $\mathcal{G}(\cdot)$ mapping a region type \mathbf{t} to a guard algebra $\mathcal{G}(\mathbf{t}) = (\mathcal{G}_{\mathbf{t}}, \bullet_{\mathbf{t}}, \{\mathbf{0}_{\mathbf{t}}\})$. The \mathbf{t} subscript is omitted from $\bullet_{\mathbf{t}}$ and $\mathbf{0}_{\mathbf{t}}$ when it is clear from the context.

As discussed, the obligations represent ghost state for describing liveness invariants. They form an *obligation algebra* which is little more complicated to define due to the association of obligations with layers.

⁹TaDA interprets the separating conjunction intuitionistically. With TaDA Live, we interpret it classically in order to not lose information about the obligations.

¹⁰In Iris, levels roughly correspond to masks.

$$\begin{aligned}
P ::= & \mathbb{B} \mid \exists x. P \mid \mathbb{E} \in X \mid P \Rightarrow Q \mid P \wedge Q \mid \text{emp} \mid P * Q \mid \mathbb{E} \mapsto \mathbb{E} \mid \mathfrak{t}_r^\lambda(\mathbb{E}) \mid r \Rightarrow d \\
& \mid \lceil \mathbb{E} \rceil_r \mid \lfloor \mathbb{E} \rfloor_r^L \mid \lfloor \mathbb{E} \rfloor_r^E \mid \text{emp}_{\text{Ob}}^R \mid \text{emp}_{\text{Ob}}^\lambda \mid r \triangleright m \\
d ::= & \blacklozenge \mid \diamond \mid (\mathbb{E}, \mathbb{E}) \qquad \mathfrak{t} \in \text{RType}, \lambda \in \text{Lvl}, r \in \text{RId} \cup \text{LVar}, R \subseteq \text{RId}, m \in \mathcal{L}.
\end{aligned}$$

Fig. 4. Syntax of Assertions. Logical expressions, \mathbb{E} , and logical Boolean expressions, \mathbb{B} , are standard.

Definition 3.7 (Obligation Algebras). TaDA Live is parametrised by a set of *atoms* AOb and a *layered obligation structure*: that is, a pair $(\text{Oblig}, \text{lay})$ where $\text{Oblig} = \wp(\text{AOb}) \subseteq \text{Guard}$ and $\text{lay}: \text{Oblig} \rightarrow \mathcal{L}$ such that $\forall O \in \text{Oblig}. \perp < \text{lay}(O) \leq \top$. We will implicitly coerce atoms $a \in \text{AOb}$ into obligations $\{a\} \in \text{Oblig}$. An *obligation algebra* is a guard algebra $(\text{Obl}, \bullet, \{\mathbf{0}\})$ where $\text{Obl} \subseteq \text{Oblig}$, $\mathbf{0} = \emptyset$, \bullet is union of disjoint sets and $\forall O_1, O_2 \in \text{Obl}. O_1 \sqsubseteq O_2 \Rightarrow \text{lay}(O_1) \geq \text{lay}(O_2)$.

TaDA Live is parametrised by a function $O(\cdot)$ mapping a region type \mathfrak{t} to an obligation algebra $O(\mathfrak{t}) = (O_{\mathfrak{t}}, \bullet_{\mathfrak{t}}, \{\mathbf{0}_{\mathfrak{t}}\})$. The \mathfrak{t} subscript is omitted from $\bullet_{\mathfrak{t}}$ and $\mathbf{0}_{\mathfrak{t}}$ when its clear from the context.

In Section 2, we have seen examples of obligation algebras. For instance, the $\mathbb{C}'_1 \parallel \mathbb{C}''_2$ example used two atoms \mathbf{u}_1 and \mathbf{u}_2 , giving rise to the obligation algebra with elements $\{\mathbf{u}_1\}$, $\{\mathbf{u}_2\}$, and $\{\mathbf{u}_1, \mathbf{u}_2\}$. As mentioned, we make no difference between an atom \mathbf{u}_1 and the obligation $\{\mathbf{u}_1\}$ using the symbol of the former for both. For our examples, it is enough to assign layers to atoms, e.g. $\text{lay}(\mathbf{u}_1) < \text{lay}(\mathbf{u}_2)$, and extend the layers to obligations by taking the minimum layer of the composed atoms, for example $\text{lay}(\mathbf{u}_1 \bullet \mathbf{u}_2) = \text{lay}(\mathbf{u}_1)$. Note that, by construction, each obligation is incompatible with itself: $O \bullet O = \perp$.

Definition 3.8 (TaDA Live Assertions). The set of TaDA Live assertions, $\text{Assrt} \ni P, Q, \dots$, is defined by the grammar in Fig. 4. The only binder is \exists . The function $\text{fv}: \text{Assrt} \rightarrow (\text{PVar} \uplus \text{LVar})$ returns the free variables of an assertion and its definition is standard. We also define $\text{pv}(P) \triangleq \text{fv}(P) \cap \text{PVar}$ and $\text{lv}(P) \triangleq \text{fv}(P) \cap \text{LVar}$. We write $P(x_1, \dots, x_n)$ to indicate that $\text{lv}(P) \subseteq \{x_1, \dots, x_n\}$ and, for $v_1, \dots, v_n \in \text{AVal}$, write $P(v_1, \dots, v_n)$ for $P[v_1/x_1, \dots, v_n/x_n]$.

We summarise the intuitive meaning of our assertions before giving their formal semantics.

- TaDA *region assertion* $\mathfrak{t}_r^\lambda(a)$ asserts the existence of a shared region with type \mathfrak{t} , identity r , level λ and abstract state a . Region assertions represent shared resources and, hence, are duplicable. We have $\vdash \mathfrak{t}_r^\lambda(a) \Leftrightarrow \mathfrak{t}_r^\lambda(a) * \mathfrak{t}_r^\lambda(a)$.
- TaDA *atomicity tracking assertion* $r \Rightarrow \blacklozenge$ gives permission to perform a single atomic change of the state of region r . Once the change is performed, the assertion becomes $r \Rightarrow (a_1, a_2)$ recording the abstract states just before and after the change (the linearization point). The assertion $r \Rightarrow \diamond$ asserts that the environment has the permission to do the atomic update. We have $\vdash r \Rightarrow \blacklozenge * r \Rightarrow \blacklozenge \Rightarrow \text{False}$, and $\vdash r \Rightarrow \blacklozenge \Leftrightarrow (r \Rightarrow \blacklozenge * r \Rightarrow \diamond)$.
- TaDA *guard assertion* $\lceil G \rceil_r$ asserts that the guard G is held locally. Guard composition is reflected by separation: $\vdash \lceil G_1 \bullet G_2 \rceil_r \Leftrightarrow \lceil G_1 \rceil_r * \lceil G_2 \rceil_r$.
- TaDA Live *local obligation assertion* $\lfloor O \rfloor_r^L$ asserts that obligation O is held locally. We have $\vdash \lfloor O_1 \bullet O_2 \rfloor_r^L \Leftrightarrow \lfloor O_1 \rfloor_r^L * \lfloor O_2 \rfloor_r^L$. Separating conjunction is interpreted classically precisely so that we do not lose local obligation information: that is, $\vdash \lfloor O \rfloor_r^L \neq \text{emp}$. It is often useful to use the same guard algebra for guards and obligations. We write $\lfloor O \rfloor_r^L \triangleq \lceil O \rceil_r * \lfloor O \rfloor_r^L$.
- TaDA Live *environment obligation assertion* $\lfloor O \rfloor_r^E$ asserts that O is held by the environment: $\vdash \lfloor O_1 \bullet O_2 \rfloor_r^E \Leftrightarrow \lfloor O_1 \rfloor_r^E * \lfloor O_2 \rfloor_r^E$. Unlike for local obligations, it is possible to lose this information, $\vdash \lfloor O \rfloor_r^E \Rightarrow \text{emp}$, because we do not need to keep track of the full obligations held by the environment, just a lower bound. The composition of environment and local obligation assertions is subtle, inspired by the subjective separation of [29]. The existence of

the local obligation can be recorded in a frame: $\vdash [O]_r^L \Leftrightarrow [O]_r^L * [O]_r^E$. We also have the derived law $\vdash [O_1 \bullet O_2]_r^L \Leftrightarrow ([O_1]_r^L * [O_2]_r^E) * ([O_1]_r^E * [O_2]_r^L)$, giving knowledge to each thread of the obligations delegated to the other.

- TaDA Live *empty obligation assertion* emp_{Ob}^R (resp. $\text{emp}_{\text{Ob}}^\lambda$) asserts that no obligation is locally held for regions with identifiers in R (resp. regions of level $< \lambda$).
- TaDA Live *layer assertion* $r \triangleright m$ asserts that the layer of the obligations held locally for region with identifier r is greater or equal than m . We often use notation such as $r \triangleright m \leq m'$ to denote $r \triangleright m \wedge m \leq m'$.

We introduce the *worlds* of TaDA Live, which are instrumented heaps providing the models of the assertions of TaDA Live. A world is a *local* model in the sense that it reflects the state as seen from the perspective of a single thread. It is built from a local heap, and a set of shared regions with associated guards and obligations. Worlds are parametrised by a set of region identifiers \mathcal{R} which, intuitively, are the regions which the current operation is supposed to update abstractly exactly once. We say the regions in \mathcal{R} are *tracked* for proving atomicity, using special ghost state given by the *atomicity tracking algebra* that supports the semantics of the atomicity tracking assertions.

Definition 3.9 (Atomicity Tracking Algebra). The *atomicity tracking algebra* is a PCM defined by $\text{ATrack} \triangleq ((\text{AState} \times \text{AState}) \uplus \{\diamond, \diamond\}, \cdot, \text{Emp}_\diamond)$, where the composition is $\diamond \cdot \diamond = \diamond = \diamond \cdot \diamond$, $\diamond \cdot \diamond = \diamond$ and $\forall a, b \in \text{AState}. (a, b) \cdot (a, b) = (a, b)$ (undefined otherwise), and the set of unit elements is $\text{Emp}_\diamond \triangleq (\text{AState} \times \text{AState}) \uplus \{\diamond\}$. The expression evaluation function is extended to map expressions d in the atomicity tracking assertions to the corresponding elements of ATrack : $\mathcal{E}[\diamond]_{\mathcal{C}} = \diamond$, $\mathcal{E}[\diamond]_{\mathcal{C}} = \diamond$, $\mathcal{E}[(\mathbb{E}_1, \mathbb{E}_2)]_{\mathcal{C}} = (\mathcal{E}[\mathbb{E}_1]_{\mathcal{C}}, \mathcal{E}[\mathbb{E}_2]_{\mathcal{C}})$.

Definition 3.10 (Worlds). Let $\mathcal{R} \subseteq \text{RId}$. A *world*, $w \in \text{World}_{\mathcal{R}}$, is a tuple $w = (h, \rho, \gamma, \chi, \theta, \xi)$ where

- $h \in \text{Heap}$ is the local heap, i.e. the cells owned locally;
- $\rho \in \text{RMap} \triangleq \text{RId} \rightarrow_f (\text{RType} \times \text{Lvl} \times \text{AState})$ describes the shared regions;
- $\gamma \in \text{GMap} \triangleq \text{RId} \rightarrow_f \text{Guard}$ describes the local guards;
- $\chi \in \text{AMap}_{\mathcal{R}} \triangleq \mathcal{R} \rightarrow \text{ATrack}$ describes the local atomicity tracking components;
- $\theta \in \text{OMap} \triangleq \text{RId} \rightarrow_f \text{Oblig}$ describes the local obligations;
- $\xi \in \text{OMap} \triangleq \text{RId} \rightarrow_f \text{Oblig}$ describes the environment obligations, known to be held locally by the environment;

satisfying the following well-formedness constraints:

- $\text{dom}(\rho) = \text{dom}(\gamma) = \text{dom}(\theta) = \text{dom}(\xi) \supseteq \mathcal{R}$,
- $\forall r \in \text{RId}$. if $\rho(r) = (\mathbf{t}, _ , _)$ then $\gamma(r) \in \mathcal{G}_{\mathbf{t}}$, $\theta(r) \in \mathcal{O}_{\mathbf{t}}$, $\xi(r) \in \mathcal{O}_{\mathbf{t}}$,
- $\forall r \in \text{dom}(\theta)$. $\theta(r) \# \xi(r)$.

A shared region with identifier r , given by $\rho(r) = (\mathbf{t}, \lambda, a)$, has type \mathbf{t} and abstract state a . For a world w , we write h_w and ρ_w and so on, for the corresponding components of w . We also define $\text{rty}_w(r) \triangleq \mathbf{t}$, $\text{lvl}_w(r) \triangleq \lambda$, and $\text{ast}_w(r) \triangleq a$, if $\rho_w(r) = (\mathbf{t}, \lambda, a)$.

We define a PCM on worlds (called *world algebra*). We define how worlds compose by first defining composition on each component of a world. Heap composition is disjoint union. Region maps only compose if they are equal. Given $\rho \in \text{RMap}$, the compositions $\bullet_\rho : \text{GMap} \times \text{GMap} \rightarrow \text{GMap}$ and $\circ_{\mathcal{R}} : \text{AMap}_{\mathcal{R}} \times \text{AMap}_{\mathcal{R}} \rightarrow \text{AMap}_{\mathcal{R}}$ are:

$$\begin{aligned} \gamma_1 \bullet_\rho \gamma_2 &\triangleq \lambda r \in \text{dom}(\rho). \gamma_1(r) \bullet_{\mathbf{t}} \gamma_2(r) && \text{if } \forall r \in \text{dom}(\rho). \rho(r) = (\mathbf{t}, _ , _) \wedge \gamma_1(r) \bullet_{\mathbf{t}} \gamma_2(r) \neq \perp \\ \chi_1 \circ_{\mathcal{R}} \chi_2 &\triangleq \lambda r \in \mathcal{R}. \chi_1(r) \cdot \chi_2(r) && \text{if } \forall r \in \mathcal{R}. \chi_1(r) \cdot \chi_2(r) \neq \perp \end{aligned}$$

and undefined otherwise. The composition \bullet_ρ on OMap is defined analogously to \bullet_ρ on GMap .

The local and environment obligation maps compose in a subtle way inspired by the subjective separation of [29]. To express this interaction, we define a composition on *pairs* of local/environment

obligation maps. Given $\theta_1, \theta_2, \xi_1, \xi_2 \in \text{OMap}$, we define

$$(\theta_1, \xi_1) \odot_\rho (\theta_2, \xi_2) \triangleq \begin{cases} (\theta_1 \bullet_\rho \theta_2, \xi) & \text{if } \xi = \min_{\sqsubseteq} \{\xi \mid \xi_1 \sqsubseteq (\theta_2 \bullet_\rho \xi) \wedge \xi_2 \sqsubseteq (\theta_1 \bullet_\rho \xi)\} \\ & \text{and } (\theta_1 \bullet_\rho \theta_2) \neq \perp \\ \perp & \text{otherwise} \end{cases}$$

Note that, for obligation algebras, the minimum taken by the definition is always unique if it exists. Indeed, in general one can set $\xi(r) = \xi_1(r) \setminus \theta_2(r) \cup \xi_2(r) \setminus \theta_1(r)$. For example, assuming **A**, **B**, **C**, **D**, **E**, and **F** are distinct atoms, we have

$$([r \mapsto \mathbf{A} \bullet \mathbf{B}], [r \mapsto \mathbf{C} \bullet \mathbf{E}]) \odot_\rho ([r \mapsto \mathbf{C} \bullet \mathbf{D}], [r \mapsto \mathbf{A} \bullet \mathbf{F}]) = ([r \mapsto \mathbf{A} \bullet \mathbf{B} \bullet \mathbf{C} \bullet \mathbf{D}], [r \mapsto \mathbf{E} \bullet \mathbf{F}])$$

provided the composition $\mathbf{A} \bullet \mathbf{B} \bullet \mathbf{C} \bullet \mathbf{D}$ is defined. Furthermore, this definition supports the implication $[O]_r^L \Rightarrow [O]_r^L * [O]_r^E$ since $([r \mapsto O], [r \mapsto \mathbf{0}]) \odot_\rho ([r \mapsto \mathbf{0}], [r \mapsto O]) = ([r \mapsto O], [r \mapsto \mathbf{0}])$.

Definition 3.11 (World Algebras). The PCM of *world algebras*, $(\text{World}_{\mathcal{R}}, \odot, \text{Emp}_{\mathcal{R}})$, is defined by the set of worlds $\text{World}_{\mathcal{R}}$,

- the *subjective world composition*, \odot , given by:

$$(h_1, \rho_1, \gamma_1, \chi_1, \theta_1, \xi_1) \odot (h_2, \rho_2, \gamma_2, \chi_2, \theta_2, \xi_2) = (h_1 \uplus h_2, \rho, \gamma_1 \bullet_\rho \gamma_2, \chi_1 \circ_{\mathcal{R}} \chi_2, \theta, \xi)$$

if $h_1 \# h_2$, $\rho = \rho_1 = \rho_2$, $\gamma_1 \bullet_\rho \gamma_2 \neq \perp$, $\chi_1 \circ_{\mathcal{R}} \chi_2 \neq \perp$, and $(\theta_1, \xi_1) \odot_\rho (\theta_2, \xi_2) = (\theta, \xi)$, undefined otherwise; and

- the set of unit elements given by:

$$\text{Emp}_{\mathcal{R}} \triangleq \left\{ (\emptyset, \rho, \gamma, \chi, \theta, \xi) \in \text{World}_{\mathcal{R}} \mid \begin{array}{l} \forall r. \rho(r) = (\mathbf{t}, _ , _) \Rightarrow \gamma(r) = \mathbf{0}_t \wedge \theta(r) = \mathbf{0}_t, \\ \forall r \in \mathcal{R}. \chi(r) \in \text{Emp}_{\blacklozenge} \end{array} \right\}$$

Notice that the units are worlds with arbitrary shared regions, atomicity components from $\text{Emp}_{\blacklozenge}$, and arbitrary environment obligations.

Subjective composition of worlds (\odot) is lifted to composition of sets of worlds ($*$), defined as $p_1 * p_2 \triangleq \{w_1 \odot w_2 \mid w_1 \in p_1, w_2 \in p_2, w_1 \# w_2\}$.

We want the region and environment obligations assertions to enjoy the elimination rule, e.g. $\mathbf{t}_r^\lambda(a) * Q \Rightarrow Q$. Assertions therefore denote sets of worlds that are upward-closed with respect to adding regions and adding environment obligations. Formally, we define the *world ordering* $\leq_{\mathcal{R}}$ as the smallest reflexive and transitive relation such that:

$$\begin{aligned} (h, \rho, \gamma, \chi, \theta, \xi) &\leq_{\mathcal{R}} (h, \rho[r \mapsto (\mathbf{t}, \lambda, a)], \gamma[r \mapsto \mathbf{0}_t], \chi, \theta[r \mapsto \mathbf{0}_t], \xi[r \mapsto \mathbf{0}_t]) && r \notin \text{dom}(\rho) \\ (h, \rho, \gamma, \chi, \theta, \xi) &\leq_{\mathcal{R}} (h, \rho, \gamma, \chi, \theta, \xi[r \mapsto \xi(r) \bullet O]) && \theta(r) \# O \# \xi(r) \end{aligned}$$

The upward-closed sets of worlds $\text{World}_{\mathcal{R}}^\dagger \triangleq \{p \subseteq \text{World}_{\mathcal{R}} \mid \forall w, w'. w \leq_{\mathcal{R}} w' \wedge w \in p \Rightarrow w' \in p\}$ are the semantic domain of our assertions.

Definition 3.12 (Satisfaction Relation). Let $\zeta: (\text{PVar} \uplus \text{LVar}) \rightarrow \text{AVal}$ be the union of a program and logic variable store. For a world $w \in \text{World}_{\mathcal{R}}$ and an assertion P , the *assertion satisfaction* relation, $\zeta, w \vDash_{\mathcal{R}} P$, is defined in Fig. 5.

We write $\vdash_{\mathcal{R}} P$ if, for $\forall \zeta: (\text{PVar} \uplus \text{LVar}) \rightarrow \text{AVal}, w \in \text{World}_{\mathcal{R}}$ we have $\zeta, w \vDash_{\mathcal{R}} P$, and write $\mathcal{W}[[P]]_{\mathcal{R}}^\zeta \triangleq \{w \mid \zeta, w \vDash_{\mathcal{R}} P\}$ for any assertion P . It is easy to check that $\mathcal{W}[[P]]_{\mathcal{R}}^\zeta \in \text{World}_{\mathcal{R}}^\dagger$ for every P and ζ .

$$\begin{aligned}
\varsigma, w \vDash_{\mathcal{R}} \text{emp} & \text{ iff } w \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} \mathbb{B} & \text{ iff } \mathcal{B}[\mathbb{B}]_{\varsigma} \\
\varsigma, w \vDash_{\mathcal{R}} \mathbb{E}_1 \mapsto \mathbb{E}_2 & \text{ iff } h_w = [\mathcal{E}[\mathbb{E}_1]_{\varsigma} \mapsto \mathcal{E}[\mathbb{E}_2]_{\varsigma}] \wedge (\emptyset, \rho_w, \gamma_w, \chi_w, \theta_w, \xi_w) \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} P \Rightarrow Q & \text{ iff } \forall w'. w \leq_{\mathcal{R}} w' \wedge \varsigma, w' \vDash_{\mathcal{R}} P \Rightarrow \varsigma, w' \vDash_{\mathcal{R}} Q \\
\varsigma, w \vDash_{\mathcal{R}} \exists x. P & \text{ iff } \exists v \in \text{AVal}. \varsigma[x \mapsto v], w \vDash_{\mathcal{R}} P \\
\varsigma, w \vDash_{\mathcal{R}} \mathbb{E} \in X & \text{ iff } \mathcal{E}[\mathbb{E}]_{\varsigma} \in X \\
\varsigma, w \vDash_{\mathcal{R}} P_1 \wedge P_2 & \text{ iff } (\varsigma, w \vDash_{\mathcal{R}} P_1) \wedge (\varsigma, w \vDash_{\mathcal{R}} P_2) \\
\varsigma, w \vDash_{\mathcal{R}} P_1 * P_2 & \text{ iff } \exists w_1, w_2. w = w_1 \odot w_2 \wedge (\varsigma, w_1 \vDash_{\mathcal{R}} P_1) \wedge (\varsigma, w_2 \vDash_{\mathcal{R}} P_2) \\
\varsigma, w \vDash_{\mathcal{R}} \mathfrak{t}_r^{\lambda}(\mathbb{E}) & \text{ iff } \rho_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma}) = (\mathfrak{t}, \lambda, \mathcal{E}[\mathbb{E}]_{\varsigma}) \wedge w \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} r \Rightarrow d & \text{ iff } \chi_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma}) = \mathcal{E}[\mathfrak{t}_d]_{\varsigma} \wedge (h_w, \rho_w, \gamma_w, \chi_w[\mathcal{E}[\mathfrak{t}_r]_{\varsigma} \mapsto \diamond], \theta_w, \xi_w) \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} [\mathbb{E}]_r & \text{ iff } \gamma_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma}) = \mathcal{E}[\mathbb{E}]_{\varsigma} \wedge (h_w, \rho_w, \gamma_w[\mathcal{E}[\mathfrak{t}_r]_{\varsigma} \mapsto \mathbf{0}], \chi_w, \theta_w, \xi_w) \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} [\mathbb{E}]_r^{\perp} & \text{ iff } \theta_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma}) = \mathcal{E}[\mathbb{E}]_{\varsigma} \wedge (h_w, \rho_w, \gamma_w, \chi_w, \theta_w[\mathcal{E}[\mathfrak{t}_r]_{\varsigma} \mapsto \mathbf{0}], \xi_w) \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} [\mathbb{E}]_r^{\mathbb{E}} & \text{ iff } \mathcal{E}[\mathbb{E}]_{\varsigma} \sqsubseteq \xi_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma}) \wedge w \in \text{Emp}_{\mathcal{R}} \\
\varsigma, w \vDash_{\mathcal{R}} \text{emp}_{\text{Ob}}^R & \text{ iff } \forall r \in R. \rho_w(r) = (\mathfrak{t}, _, _) \Rightarrow \theta_w(r) = \mathbf{0} \\
\varsigma, w \vDash_{\mathcal{R}} \text{emp}_{\text{Ob}}^{\lambda} & \text{ iff } \forall r, \lambda' < \lambda. \rho_w(r) = (\mathfrak{t}, \lambda', _) \Rightarrow \theta_w(r) = \mathbf{0} \\
\varsigma, w \vDash_{\mathcal{R}} r \triangleright m & \text{ iff } \text{lay}(\theta_w(\mathcal{E}[\mathfrak{t}_r]_{\varsigma})) \geq m
\end{aligned}$$

Fig. 5. Definition of assertion satisfaction.

3.4 Protocols: Interference and World Rely

A world describes the state of the current thread, both the local state owned by the thread (the heap, guards, local obligations and atomity tracking components), the shared state (the regions) and the environment obligations describing obligations owned locally by the environment. We define the *world rely relation* which describes how the world may change as a result of the “well-behaved” interference of the environment characterised by the region interference relations, the atomicity tracking components and the environment obligations. To define the world rely, we need to introduce two other components of TaDA Live: the region protocols, expressed by the *region interference function*, and atomicity contexts.

The type of each region is associated with a *region interference function* which establishes which updates to a shared region are allowed by the owner of which guards.

Definition 3.13 (Region Interference). TaDA Live is parametrised by the *region interference function*, \mathcal{T} , which takes a region type $\mathfrak{t} \in \text{RType}$ and returns a function $\mathcal{T}_{\mathfrak{t}}: \mathcal{G}_{\mathfrak{t}} \rightarrow \wp^2((\text{AState} \times O_{\mathfrak{t}}) \times (\text{AState} \times O_{\mathfrak{t}}))$. Every function $\mathcal{T}_{\mathfrak{t}}$ is required to satisfy three properties:

- reflexivity: $((a, \mathbf{0}_{\mathfrak{t}}), (a, \mathbf{0}_{\mathfrak{t}})) \in \mathcal{T}_{\mathfrak{t}}(\mathbf{0}_{\mathfrak{t}})$, for all $a \in \text{AState}$;
- monotonicity in the guards: $\forall G_1, G_2 \in \mathcal{G}_{\mathfrak{t}}. G_1 \sqsubseteq G_2 \Rightarrow \mathcal{T}_{\mathfrak{t}}(G_1) \subseteq \mathcal{T}_{\mathfrak{t}}(G_2)$;
- closure under obligation frames: for all $O_1, O_2, O \in O_{\mathfrak{t}}$, if $((a_1, O_1), (a_2, O_2)) \in \mathcal{T}_{\mathfrak{t}}(G)$ and $O_1 \# O$ and $O_2 \# O$, then $((a_1, O_1 \bullet_{\mathfrak{t}} O), (a_2, O_2 \bullet_{\mathfrak{t}} O)) \in \mathcal{T}_{\mathfrak{t}}(G)$.

We write $\mathcal{T}_{\mathfrak{t}}(_)$ for $\bigcup_{G \in \mathcal{G}_{\mathfrak{t}}} \mathcal{T}_{\mathfrak{t}}(G)$. For any $T \subseteq (\text{AState} \times \text{Oblig}) \times (\text{AState} \times \text{Oblig})$, we write $\text{io}(T) \triangleq \{(a, b) \mid ((a, _), (b, _)) \in T\}$.

The final concept we need before introducing the world rely relation is the *atomicity context*, \mathcal{A} . In TaDA Live proofs, we keep in the context of the judgment information about which updates we are currently proving are abstractly atomic. The rule driving this bookkeeping is the MkATOM rule.

Although we will properly explain the rule in Section 4.7, we sketch the main idea as a motivation for the atomicity context now. The relevant “skeleton” of the rule is as follows:

$$\frac{r \notin \text{dom}(\mathcal{A}) \quad T \subseteq \mathcal{T}_t(G) \quad R = \text{io}(T) \quad \dots}{m; \lambda'; \mathcal{A}[r \mapsto (X, k, X', T)] \vdash \{\exists x \in X. t_r^\lambda(x) * r \Rightarrow \diamond\} \mathbb{C} \{\exists x, y. R(x, y) \wedge r \Rightarrow (x, y)\}}{m; \lambda'; \mathcal{A} \vdash \forall x \in X \rightarrow_k X'. \langle t_r^\lambda(x) * [G]_r \rangle \mathbb{C} \langle \exists y. t_r^\lambda(y) * [G]_r \wedge R(x, y) \rangle}$$

The judgments include the context information such as the layer m , the level λ' and the atomicity context \mathcal{A} , and the pseudo-quantifier includes a layer k . We formally introduce these details in Section 3.8. Here, we focus on motivating the use of the atomicity context \mathcal{A} . This rule describes how an update to the state of a region r can be declared atomic even if it was realised through a series of steps. It does this by converting a Hoare triple to an atomic triple, provided the Hoare triple bears evidence (through the atomicity tracking assertions of the premise) that, although many steps might have been taken, the abstract state was changed by the command exactly once. The atomic triple may constrain the environment interference with a non-trivial pseudo-quantifier. The proof of the premise in general needs to make use of these assumptions about the environment, but the conversion to a Hoare triple means we cannot use pseudo-quantification to represent them. These assumptions are instead made available to the proof of the Hoare triple using the atomicity context, that records the (X, k, X') information from the pseudo-quantifier and the relation T which stores the update that we are proving happens atomically.

Definition 3.14 (Atomicity Context). An *atomicity context* \mathcal{A} is a finite partial function from RId to tuples of the form (X, k, X', T) where $X, X' \subseteq \text{AState}$, $k \in \mathcal{L}$, and $T \subseteq (\text{AState} \times \text{Oblig}) \times (\text{AState} \times \text{Oblig})$ is closed under obligation frames (as in Definition 3.13).

Assuming $\mathcal{A}(r) = (X, k, X', T)$, we write $\text{safe}(\mathcal{A}, r) \triangleq X$, $\text{good}(\mathcal{A}, r) \triangleq X'$, $\text{live}(\mathcal{A}, r) \triangleq (X, k, X')$ which we write $X \rightarrow_k X'$, and $\text{tr}(\mathcal{A}, r) \triangleq T$. For every $r \in \text{dom}(\mathcal{A})$, we require $\{x \mid (x, _) \in \text{io}(T)\} \subseteq \text{safe}(\mathcal{A}, r)$. The set $\text{dom}(\mathcal{A})$ declares the regions for which we are tracking atomicity: for $r \in \text{dom}(\mathcal{A})$, the environment will only change the abstract state within $\text{safe}(\mathcal{A}, r)$ and will obey the liveness condition given by $\text{live}(\mathcal{A}, r)$ that the environment will always eventually return to a good state in $\text{good}(\mathcal{A}, r) \triangleq X'$; and the local thread will only change the abstract state at most once according to the relation $\text{io}(\text{tr}(\mathcal{A}, r))$. We write $\vDash_{\mathcal{A}}$ for $\vDash_{\text{dom}(\mathcal{A})}$, and similarly for $\vdash_{\mathcal{A}}$, $\mathcal{W}[[P]]_{\mathcal{A}}^{\mathcal{C}}$, $\text{World}_{\mathcal{A}}$ and $\text{Emp}_{\mathcal{A}}$.

Definition 3.15 (World Rely). The *world rely relation*, $\mathbf{R}_{\mathcal{A}} \subseteq \text{World}_{\mathcal{A}} \times \text{World}_{\mathcal{A}}$, is the smallest reflexive and transitive relation satisfying the rules in Fig. 6.

Rule WR_1 describes the case where the environment can update the abstract state of a region according to the interference relation \mathcal{T}_t . Notice that, for this rule, when $\chi(r) \in \{\diamond, \diamond\}$, the environment can only change the abstract state to something in $\text{safe}(\mathcal{A}, r)$. When $\chi(r)$ is undefined or a pair of abstract states, then the environment does not have this restriction and can do any update consistent with \mathcal{T}_t . Also, notice how the environment obligations map ξ is affected by the transition. Rule WR_2 describes the case where the atomic update given by \mathcal{A} has been delegated to the environment ($\chi[r \mapsto \diamond]$) in which case the current thread can observe the abstract state change corresponding to the update.

So far, we have introduced assertions, and worlds as their models. These structures express information mostly over *ghost state*, that is, state that is purely logical and has no representation in concrete executions. For example, the notion that there is some shared region is purely fictional, as in the concrete machine there is no special way to mark a portion of the heap as shared. We introduced interference protocols and the world rely, as a way to specify the expected well-behaved transformations shared resources may be subjected to. Since well-behaved interference from the

$$\frac{\frac{\gamma(r) \# G \quad ((a_1, O_1), (a_2, O_2)) \in \mathcal{T}_t(G) \quad \chi(r) \in \{\diamond, \diamond\} \Rightarrow a_2 \in \text{safe}(\mathcal{A}, r) \quad O_2 \# \theta(r)}{(h, \rho[r \mapsto (t, \lambda, a_1)], \gamma, \chi, \theta, \xi[r \mapsto O_1]) \mathbf{R}_{\mathcal{A}} (h, \rho[r \mapsto (t, \lambda, a_2)], \gamma, \chi, \theta, \xi[r \mapsto O_2])} \text{WR}_1}{((a_1, O_1), (a_2, O_2)) \in \text{tr}(\mathcal{A}, r) \quad O_2 \# \theta(r)} \text{WR}_2}$$

Fig. 6. World Rely rules

environment can change the state of shared regions, a single world (describing a single state for each region) cannot capture the logical state we may be in, when interleaved with environment actions. Views are the sets of worlds that can explain the logical state we may be in after being suspended for an arbitrary number of environment steps. Views represent information about the logical state, that cannot be invalidated by a well-behaved environment.

Definition 3.16 (Views, Stability). An upward-closed set of worlds, $p \in \text{World}_{\mathcal{A}}^{\uparrow}$, is an \mathcal{A} -view if it is closed under $\mathbf{R}_{\mathcal{A}}$: that is, $\forall w \in p, w' \in \text{World}_{\mathcal{A}}. w \mathbf{R}_{\mathcal{A}} w' \Rightarrow w' \in p$. An assertion P is \mathcal{A} -stable, written $\mathcal{A} \vDash P$ stable, if and only if, for all $\varsigma: (\text{PVar} \uplus \text{LVar}) \rightarrow \text{AVal}$, $\mathcal{W}[\![P]\!]_{\mathcal{A}}^{\varsigma}$ is an \mathcal{A} -view.

We write $\text{View}_{\mathcal{A}}$ for the set of all \mathcal{A} -views and $\text{Stable}_{\mathcal{A}}$ for the set of all \mathcal{A} -stable assertions.

Definition 3.17 (View Algebra). The PCM of *view algebras*, $(\text{View}_{\mathcal{A}}, *, \{\text{Emp}_{\mathcal{A}}\})$, is formed from the set $\text{View}_{\mathcal{A}}$, and the composition $p_1 * p_2 \triangleq \{w_1 \odot w_2 \mid w_1 \in p_1, w_2 \in p_2, w_1 \# w_2\}$.

Notice that the composition of views always gives rise to a view: in the case where there are no compatible pairs of worlds in the views, the result is the empty view (the denotation of False).

On checking stability. TaDA Live proofs require checking stability of assertions in some crucial steps. The notion of stability of Definition 3.16 is given in terms of the semantics of assertions, but it is possible, in principle, to provide a set of lemmas to prove stability of common cases without reasoning at the level of the model. For example, any traditional separation logic assertion (such as $\text{emp}, x \mapsto v$, pure formulas) is always stable; guard and local obligation assertions are also automatically stable; stability is preserved by $*$, \wedge , \vee , and existential quantification. The crucial sources of instability are region assertions, environment obligation assertions, and $r \mapsto \diamond$. Stability of the first two can be established by inspecting the protocol of regions. A rule that would be expressive enough to prove most stability checks for our examples is:

$$\frac{\forall x \in X, x', G', O'. (G' \# G(x)) \wedge ((x, O(x)), (x', O')) \in \mathcal{T}_t(G') \Rightarrow x' \in X \wedge O' = O(x')}{\mathcal{A} \vDash \exists x \in X. \mathfrak{t}_r^{\lambda}(x) * [G(x)]_r * [O(x)]_r^E \text{ stable}}$$

It is similarly easy to extract from Fig. 6 rules involving the atomicity context information:

$$\frac{\text{safe}(\mathcal{A}, r) = X}{\mathcal{A} \vDash \exists x \in X. \mathfrak{t}_r^{\lambda}(x) * r \mapsto \diamond \text{ stable}} \qquad \frac{r \in \text{dom}(\mathcal{A})}{\mathcal{A} \vDash r \mapsto \diamond \vee r \mapsto (_, _) \text{ stable}}$$

3.5 Linking Levels of Abstraction: Interpretations and Reification

As we mentioned, worlds and views represent ghost information about state. Ultimately, however, we want to use this information to express properties of concrete execution traces. To do so, we need to formalise the link between worlds with their logical instrumentation and concrete states comprising variable stores and heaps. The first component that contributes to this link is a *region interpretation*, which specifies the implementation-dependent content of a shared region: for example, for a shared spin lock, the interpretation of the abstract shared region $\text{spin}_r^{\lambda}(x, l)$ is the view given by $x \mapsto l$, a single cell storing l at x .

Definition 3.18 (Region Interpretation). TaDA Live is parametrised by a *region interpretation function* $\mathcal{I}_t[\cdot, \cdot, \cdot]$: $\text{Rld} \times \text{Lvl} \times \text{AState} \rightarrow \text{View}_0$ for each $t \in \text{RType}$, such that, for every $r \in \text{Rld}$, $\lambda \in \text{Lvl}$, $a \in \text{AState}$, $\forall w \in \mathcal{I}_t[r, \lambda, a]$. $\forall r' \in \text{dom}(\theta_w) \setminus \{r\}$. $\theta_w(r') = \mathbf{0}$. We also require the interpretation to be λ -safe, a technical condition explained in Section 4.3 that is usually immediate to check (see Lemma 4.2). A region interpretation’s companion is the *syntactic region interpretation* $\mathcal{I}_t = (r, l, a, P)$ where $r, l, a \in \text{LVar}$, $\text{fv}(P) \subseteq \{r, l, a\}$, $\mathbf{0} \vDash P$ stable, and $\vdash_0 P[\lambda/l] \Rightarrow \text{emp}_{\text{Ob}}^{\text{Rld} \setminus \{r\}}$. We write $\mathcal{I}(t_{\mathbb{E}_1}^\lambda(\mathbb{E}_2))$ for $P[\mathbb{E}_1/r, \lambda/l, \mathbb{E}_2/a]$. We require that $\mathcal{I}_t[r, \lambda, a] = \mathcal{W}[\mathcal{I}(t_r^\lambda(a))]\mathbf{0}$; in practice, we will define region interpretations by writing syntactic interpretations and using the previous equation as a definition for the corresponding region interpretation functions.

It is important to understand that interpretations are not merely an indirect way of writing assertions. In our spin lock example, the crucial difference between the two assertions $\text{spin}^\lambda(x, l, \alpha)$ and $x \mapsto l$ is that the first is subjected to interference, while the latter expresses ownership of the cell at x . The requirement that the interpretation of some region with id r must imply $\text{emp}_{\text{Ob}}^{\text{Rld} \setminus \{r\}}$ forbids an interpretation to own local obligations of other regions. This is necessary for soundness: if we removed the restriction, we could fool ourselves into thinking that we fulfilled an obligation $[O]_r^l$ by creating another region with the obligation in its interpretation.

Remark 2 (On “opening” regions and levels). As in TaDA, the region interpretation is used to “open” a region: that is, import the region interpretation as local state in order to do a single atomic update. The idea is to obtain instantaneously the ownership of the content of the region for the atomic update, and to re-establish the region interpretation for the updated abstract state, before immediately relinquishing ownership by “closing” up the region. As in TaDA, this opening and closing mechanism depends on the level of the region, which is a device to avoid inconsistencies. With a specification at level λ , the rules enable a region to be opened at level $\lambda' < \lambda$ to obtain a resulting specification at level λ' . This means that, although a region can be shared ($\vdash t_r^\lambda(a) \Leftrightarrow t_r^{\lambda'}(a) * t_r^{\lambda'}(a)$), it cannot be *opened* twice, which would result in $\mathcal{I}(t_r^{\lambda'}(a)) * \mathcal{I}(t_r^{\lambda'}(a))$ with a potential contradictory duplication of non-duplicable resources.

The second component that expresses the link between the instrumented worlds and concrete states is the *reification* function. Reification has two main purposes. First, at level λ , all the regions with level lower than λ are closed, which means that the resources in their interpretation do not exist as far as the world is concerned. The concrete heap cells accounted for inside these interpretations will however correspond to cells in the concrete heap. To bridge this gap, the reification opens all closed regions importing the resources in their interpretation as local resources, obtaining a “collapsed” world. Second, all the “ghost” components of collapsed worlds (like regions, guards or obligations) do not have any representation in memory so reification erases them.

Definition 3.19 (Reification). Let $\lambda \in \text{Lvl}$ and let $\text{closed}(\lambda, w) \triangleq \{r \in \text{Rld} \mid \text{lvl}_w(r) < \lambda\}$. The *region collapse function*, $(\cdot) \downarrow^\lambda$: $\text{World}_{\mathcal{A}} \rightarrow \wp(\text{World}_{\mathcal{A}})$, is defined by:

$$w_0 \downarrow^\lambda \triangleq \left\{ w_0 \odot w_1 \odot \dots \odot w_n \left| \begin{array}{l} \text{closed}(\lambda, w_0) = \{r_1, \dots, r_n\}, \\ \rho_{w_0}(r_i) = (t_i, \lambda_i, a_i), w_i \in \mathcal{I}_{t_i}[r_i, \lambda_i, a_i], \\ \forall i \leq n. \forall r \in \text{dom}(\xi_{w_0 \odot \dots \odot w_i}). \xi_{w_0 \odot \dots \odot w_i}(r) = \mathbf{0} \end{array} \right. \right\}$$

The function $[w]_\lambda \triangleq \{h \in \text{Heap} \mid (h, _ \mapsto _ \mapsto _ \mapsto _) \in w \downarrow^\lambda\}$ is called the *world reification* of w at level λ . For any $p \in \text{World}_{\mathcal{A}}^\uparrow$, the function $\llbracket p \rrbracket_\lambda \triangleq \bigcup_{w \in p} [w]_\lambda$ is called *reification* of p at level λ .

To understand if a world w_1 can represent local resources consistent with some global heap h , we need to identify if there is a world w_2 representing the resources of the environment such that $h \in [w_1 \odot w_2]_\lambda$. That would mean that it is possible to factor h as $h = h_{w_1} \uplus h' \uplus h_{w_2}$ where h_{w_1} are

the cells fully owned by the local thread, h_{w_2} are the ones fully owned by the environment, and the cells in h' are the ones that are shared and come from opening the interpretations of closed regions in the world collapse. When collapsing, we are assuming, conceptually, that we are collapsing a world that represents *every* resource in the system. Correspondingly, the definitions that use reification —crucially, Definitions 3.20 and 3.22— always complete the local resources with some “global” frame before applying reification.

In addition to opening shared regions, the collapse function also checks that no environment obligations are assumed. To understand this, consider a world w_1 representing local resources, and a world w_2 completing it to a world $w = w_1 \odot w_2$ representing the global resources. The global world w cannot assert the existence of obligations in the environment: all those have been already accounted for in w_2 . The definition of collapse explicitly enforces this constraint by the condition on the environment obligation map. We explain why this condition is important in Section 3.7.

3.6 Frame preservation

Having established the link between worlds/views and concrete state, we can move to establishing a link between concrete *steps* in a trace and their logical justification in terms of logical state. The fundamental driver of this link is the notion of *frame-preserving update*, inspired by the frame-preserving update from [24], which represents the essence of the Rely/Guarantee reasoning in TaDA Live. The frame-preserving update looks at a specific concrete update from some h_1 to h_2 , and states under which conditions this logical update can be described as an update from logical state p_1 to logical state p_2 . The p_1 and p_2 are sets of worlds describing local resource, whereas the h_1 and h_2 are global concrete heaps. We therefore need to complete p_1 with some frame f and use reification to relate this logical state to h_1 : that is, $h_1 \in \llbracket p_1 * f \rrbracket_\lambda$. There will usually be more than one such f . The frame-preserving update requires that any such f that is a view should remain a valid frame after the update: that is, $h_2 \in \llbracket p_2 * f \rrbracket_\lambda$.

Definition 3.20 (Frame-Preserving Update). Given $h_1, h_2 \in \text{Heap}$, $p_1, p_2 \in \text{World}_{\mathcal{A}}^\dagger$ and $\lambda \in \text{Lvl}$, we define $(h_1, h_2) \vDash_{\lambda, \mathcal{A}} p_1 \rightarrow^* p_2$ to hold if and only if

$$\forall f \in \text{View}_{\mathcal{A}}. h_1 \in \llbracket p_1 * f \rrbracket_\lambda \Rightarrow h_2 \in \llbracket p_2 * f \rrbracket_\lambda.$$

TaDA Live implements the Rely/Guarantee proof principle by requiring every update to be frame-preserving. Views are resources that are preserved by protocol-compliant environment interference. The idea of a Rely, a set of allowed environment updates, is represented by assuming environment steps are frame-preserving updates on resources that are compatible with our current view. By frame preservation, any such update would preserve our view. Conversely, the idea of a Guarantee, an over-approximation of the effects of local steps under the assumption of Rely, is encoded by requiring every local step to be a frame-preserving update, and thus unable to disrupt any view held by the environment.

To see how this works more concretely, let us consider an example. We use the notation $\vDash_\lambda p \rightarrow^* q$ to mean $\forall h \in \llbracket p * \text{True} \rrbracket. \exists h'. (h, h') \vDash_\lambda p \rightarrow^* q$, that is, $\vDash_\lambda p \rightarrow^* q$ holds when p to q can be used to justify *some* concrete update.

Example 3.21. Assume we have a region type t with abstract states a, b, c, d , a single guard \mathbf{E} (with $\mathbf{E} \bullet \mathbf{E} = \perp$) and interference protocol consisting of transitions $\mathbf{E} : (a, \mathbf{0}) \rightsquigarrow (b, \mathbf{0})$ and $\mathbf{E} : (b, \mathbf{0}) \rightsquigarrow (c, \mathbf{0})$. We want to show that (for $\lambda < \lambda'$) $\vDash_{\lambda'} t_r^\lambda(a) * [\mathbf{E}]_r \rightarrow^* t_r^\lambda(c) * [\mathbf{E}]_r$ holds, but $\vDash_{\lambda'} t_r^\lambda(a) * [\mathbf{E}]_r \rightarrow^* t_r^\lambda(d) * [\mathbf{E}]_r$ and $\vDash_{\lambda'} t_r^\lambda(a) \rightarrow^* t_r^\lambda(b)$ do not. Consider any view f that is a frame of $t_r^\lambda(a) * [\mathbf{E}]_r$. The f cannot hold $[\mathbf{E}]_r$, because \mathbf{E} is not compatible with itself. As a consequence, since f is a view, it needs to be closed under world rely, which means that it is closed under the interference, which can transform a into b and b into c . For f to be compatible with $t_r^\lambda(a)$, it needs to contain

some world associating a to r ; to be a view, f needs to contain some other world associating c to r , which makes it compatible with $\mathfrak{t}_r^\lambda(c) * [\mathbf{E}]_r$. Therefore $\vDash_{\lambda'} \mathfrak{t}_r^\lambda(a) * [\mathbf{E}]_r \rightarrow^* \mathfrak{t}_r^\lambda(c) * [\mathbf{E}]_r$ holds.

Now, the view f above is not required to contain any world associating d to r . Such an f is a counterexample to $\vDash_{\lambda'} \mathfrak{t}_r^\lambda(a) * [\mathbf{E}]_r \rightarrow^* \mathfrak{t}_r^\lambda(d) * [\mathbf{E}]_r$ holding.

Finally, consider $\mathfrak{t}_r^\lambda(a)$; we can construct a frame f_a in which all worlds associate a to r and own the guard \mathbf{E} . Such set of worlds can be a view because owning \mathbf{E} disables the transition from a to b . However, f_a would be compatible with $\mathfrak{t}_r^\lambda(a)$ but not with $\mathfrak{t}_r^\lambda(b)$, which means $\vDash_{\lambda'} \mathfrak{t}_r^\lambda(a) \not\rightarrow^* \mathfrak{t}_r^\lambda(b)$ does not hold.

This definition of frame-preserving update simplifies drastically the semantics of TaDA specifications. For TaDA Live, however, we need to introduce the stronger notion of *atomic* frame-preserving update. To see the motivation behind the stronger condition, consider the region interference relation $\mathbf{E} : (a, \mathbf{k}) \rightsquigarrow (b, \mathbf{0})$ and $\mathbf{E} : (b, \mathbf{0}) \rightsquigarrow (c, \mathbf{k})$. The update from a to c via b is very different from a direct update from a to c . The intermediate step to b fulfils the obligation \mathbf{k} , which may be crucial information for the progress argument. We therefore want to enforce that if we are justifying a step as going from p to q , all the allowed transitions between region states need to match a *single transition* in the interference protocol.

Definition 3.22 (Atomic Frame-Preserving Update). Given $h_1, h_2 \in \text{Heap}$, $p_1, p_2 \in \text{World}_{\mathcal{A}}^\dagger$ and $\lambda \in \text{Lvl}$, we define $(h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_1 \rightarrow p_2$ to hold if and only if

$$\forall f \in \text{World}_{\mathcal{A}}^\dagger. h_1 \in \llbracket p_1 * f \rrbracket_\lambda \Rightarrow h_2 \in \llbracket p_2 * \mathbf{R}_{\mathcal{A}}^a(f) \rrbracket_\lambda$$

where the *atomic world rely relation*, $\mathbf{R}_{\mathcal{A}}^a$, is defined to be the smallest reflexive relation closed under the rules of Fig. 6, with the restriction that rules WR_1 and WR_2 can be applied at most once per region identifier. It is formally defined in Appendix E.1.

Intuitively, this says that if the environment has some resource f compatible with p_1 , it should expect that after a step, the resource f might be transformed into $\mathbf{R}_{\mathcal{A}}^a(f)$. When f is a view, one gets back Definition 3.20, as views are precisely the resources that cannot be invalidated by any number of updates of the environment. We will use atomic frame-preserving updates to check the safety of logical traces with respect to some specification in Definition 3.28.

3.7 Viewshifts and “classical” resources

Before moving to specifications, we define *viewshift*, a semantic generalisation of implication, which is a prime example of application of frame-preserving update, used in our CONS rule. They correspond to “purely logical” updates in that they update the ghost resources without affecting the concrete memory.

Definition 3.23 (Viewshift). Given $p_1, p_2 \in \text{World}_{\mathcal{A}}^\dagger$, the judgement $\lambda; \mathcal{A} \vDash p_1 \Rightarrow p_2$, holds if $\forall h \in \text{Heap}. (h, h) \vDash_{\lambda; \mathcal{A}} p_1 \rightarrow^* p_2$. For two assertions P, Q , the assertion P *viewshifts* to Q , written $\lambda; \mathcal{A} \vDash P \Rightarrow Q$, if and only if, $\forall \zeta : (\text{PVar} \uplus \text{LVar}) \rightarrow \text{AVal}, \lambda; \mathcal{A} \vDash \mathcal{W}\llbracket P \rrbracket_{\mathcal{A}}^\zeta \Rightarrow \mathcal{W}\llbracket Q \rrbracket_{\mathcal{A}}^\zeta$.

Viewshifts are typically employed to “allocate” a new region by sharing some local resource (a form of weakening). For example, assume $I(\mathfrak{t}_r(x, v)) \triangleq x \mapsto v * [\mathbf{A}]_r^\perp$. We have that $P_0 = (x \mapsto 0)$ viewshifts to $\exists r. \mathfrak{t}_r(x, 0)$: the underlying reification does not change, and any frame of P_0 with non-empty reification, must only have regions reifying to cells disjoint from x ; moreover, such frame will only have finitely many regions allocated, so it is always possible to draw a fresh r from the infinite set Rld to satisfy the existential quantification over r .

Viewshifts also ensure that obligation information is not updated inconsistently. For example, in the “region allocation” step above, we cannot viewshift P_0 to $P_1 = \exists r. \mathfrak{t}_r(x, 0) * [\mathbf{k}]_r^{\mathbf{E}}$, which would mean we are pretending there is an obligation \mathbf{k} in the environment without any evidence of that

being true. To show that the viewshift does not hold, we can choose $h = [x \mapsto v]$ and show that $(h, h) \vDash_{\lambda; \mathcal{A}} P_0 \rightarrow P_1$ is false. To see this, pick $f = \text{emp}$ as the global frame; h is in the reification of $P_0 * \text{emp}$ but the reification of $\exists r. \mathbf{t}_r(x, 0) * \lfloor \mathbf{K} \rfloor_r^E * \text{emp}$ is empty: the frame emp has empty local obligation map, so every world w considered by the region collapse of $P_1 * \text{emp}$ has $\xi_w \neq \mathbf{0}$. The idiomatic, and correct, pattern of creation of environment obligations would viewshift P_0 to, say, $\exists r. \mathbf{t}_r(x, 0) * \lfloor \mathbf{B} \rfloor_r^L$ for some relevant obligation \mathbf{B} compatible with \mathbf{A} , and then with implication obtain $\exists r. \mathbf{t}_r(x, 0) * \lfloor \mathbf{B} \rfloor_r^L * \lfloor \mathbf{B} \rfloor_r^E$: in this case the environment obligation has been created from the evidence of the existence of a corresponding local obligation.

It is important to note that, in a logic with the ability to share assertions, as regions in TaDA or invariants in Iris allow, having classical resources does not have the expected effect. By definition, a classical resource P cannot be “forgotten”, i.e. $P * Q \not\Rightarrow Q$. By using viewshift, however, it is possible to create a region with interpretation defined so it contains P , and immediately discard it (regions are not classical resources), obtaining $P * Q \Rightarrow Q$. This, for example, makes TaDA Live incapable of proving absence of memory leaks even if its heap assertions are classical. We however manage to avoid this issue for local obligation assertions $\lfloor O \rfloor_r^L$ because of their specific semantics. First, $\lfloor O \rfloor_r^L$ can only ever be part of the interpretation for the region r , as imposed by our restrictions on region interpretations. Second, the very notion of fulfilling the obligation is defined as transferring its ownership to the interpretation. Moreover, the region protocol constrains the loss of an obligation to happen only in correspondence with some region state change, so the only way to get rid of a local obligation is to induce the desired state change in the region and transfer the obligation to the interpretation. We need obligations to be classical resources in order for this to be the *only* way of losing them. We made heaps and guards behave classically for the sake of uniformity, but this is not essential.

The issue of having genuinely classical resources in a logic with regions/invariants has been tackled in [1], with the main use cases being proving absence of memory leaks. The techniques presented there could provide the basis for an alternative way of handling TaDA Live-style obligations.

3.8 Specification format

With all these definitions in place, we can now proceed to define TaDA Live specifications and their trace semantics. Most of the time, TaDA Live proofs manipulate triples of two forms:

$$m; \lambda; \mathcal{A} \vdash \forall x \in X \rightarrow_k X'. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle \quad (11a)$$

$$m; \lambda; \mathcal{A} \vdash \{P\} \mathbb{C} \{Q\} \quad (11b)$$

called *atomic triples* and *Hoare triples*, respectively. It is however possible for a command to manipulate some resources P_h non-atomically, and some other resources $P_a(x)$ atomically, at the same time. In general, specifications use *hybrid triples*:

$$m; \lambda; \mathcal{A} \vdash \forall x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle$$

a minor generalisation¹¹ of hybrid triple discussed in Section 2.1. Intuitively, the Hoare precondition P_h is a resource that is owned by the command and, as such, cannot be invalidated by actions of the environment. The command is allowed to manipulate this owned resource non-atomically, provided it satisfies the Hoare postcondition Q_h upon termination. The atomic precondition $P_a(x)$ represents the resource that can be shared between the command and the environment. The environment can update it, but only with the effect of going from $P_a(x)$ for some $x \in X$ to $P_a(x')$ for some $x' \in X$. The command is allowed to update it exactly once from $P_a(x)$ to perform its linearization point,

¹¹The difference is the $\exists y$, which is used in the uncommon case when the linearization point is non-deterministic *and* the Hoare postcondition depends on this non-deterministic choice.

transforming it to a resource satisfying the atomic postcondition $Q_a(x)$. The atomic postcondition only needs to be true *just after* the linearization point as the environment is allowed to update it immediately afterwards. The pseudo-quantified variable x has two important uses: it represents the “surface” of allowed interference by the environment; it is bound in the postcondition to the value of the parameter of the atomic precondition *just before* the linearization point.

The atomic and Hoare triples in (11a) and (11b) are then special cases of the hybrid triple:¹²

$$\forall \vec{v}_0. m; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X \rightarrow_k X'. \langle \vec{v}_0 \doteq \vec{v}_0 \mid P'(x) \rangle \mathbb{C} \exists \vec{v}_1. \langle \vec{v}_0 \doteq \vec{v}_0 \wedge \vec{v}_1 \doteq \vec{v}_1 \mid Q'(x) \rangle \quad (12a)$$

$$m; \lambda; \mathcal{A} \vdash \langle P \mid \text{emp} \rangle \mathbb{C} \langle Q \mid \text{emp} \rangle \quad (12b)$$

resp., where $\vec{v}_0 = \text{pv}(P(x))$, $\vec{v}_1 = \text{pv}(Q(x)) \setminus \vec{v}_0$, $P'(x) = P(x)[\vec{v}_0/\vec{v}_0]$ and $Q'(x) = Q(x)[\vec{v}_0/\vec{v}_0, \vec{v}_1/\vec{v}_1]$ (for technical reasons the atomic pre/post-conditions in the general triples cannot contain program variables). We omit the pseudo-quantifier from an atomic triple (as above) when the pseudo-quantified variable does not occur in the triple, and thus could be quantified as $\mathbb{W}x \in \{1\} \rightarrow_{\perp} \{1\}$. We also use the abbreviated form $\mathbb{W}x \in X$ when the liveness assumption is trivial, i.e. $\mathbb{W}x \in X \rightarrow_{\perp} X$.

Definition 3.24 (Specification). Specifications, $\mathbb{S} \in \text{Spec}$, have the form:

$$\mathbb{W}x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle_{m; \lambda; \mathcal{A}} \quad (\star)$$

where

- $m \in \mathcal{L}$, $\lambda \in \text{Lvl}$ and $\mathcal{A} \in \text{ACTxt}$;
- $x, y \in \text{LVar}$;
- $X' \subseteq X \subseteq \text{AVal}$ and $k \in \mathcal{L}$;
- $P_h, Q_h(v, v') \in \text{Stable}_{\mathcal{A}}$ for all $v \in X$ and $v' \in \text{AVal}$;
- $P_a(v), Q_a(v, v') \in \text{Assrt}$ for all $v \in X$ and $v' \in \text{AVal}$, and $\text{pv}(P_a) = \text{pv}(Q_a) = \emptyset$.
- $\forall x \in X. \vdash_{\mathcal{A}} P_a(x) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda}$.
- $\forall x \in X, y. \vdash_{\mathcal{A}} Q_a(x, y) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda}$.

In addition to the atomicity context \mathcal{A} , the *context of a specification* m, λ, \mathcal{A} consists also of a layer m , and a level λ . These components record information about the proof context of the judgement. The layer m indicates that we are in a context where we are forbidden from assuming as live obligations with layers $\geq m$ (or incomparable to m). The level λ indicates that the regions with level $\geq \lambda$ are open (and cannot be re-opened).

3.9 Trace Semantics of Specifications

Finally, we can define the semantics of a specification. The idea of the semantics is to collect all traces that are deemed as acceptable to a specification \mathbb{S} , so that we can later say a command satisfies \mathbb{S} if its traces are all accepted by the semantics of \mathbb{S} . The general principle in accepting a trace is the following: the local steps are only expected to be correctly implementing the functionality declared by \mathbb{S} if the environment satisfies the assumptions implied by the (safety and liveness) protocols and \mathbb{S} itself. If a trace has gone wrong as a consequence of the environment making moves outside of the assumptions, that trace is accepted, as the problem is not the responsibility of the local command itself. If a trace has gone wrong as a consequence of local steps, then the trace is rejected.

The semantics of a specification therefore traverses a trace to decide whether to accept or reject it, by determining who is to blame for failures. We decouple the traversal needed for checking the safety constraints, and the one checking the liveness ones. In terms of safety, a specification like (\star) (in Definition 3.24) expects that:

- the precondition holds: the starting resource satisfies $P_h * P_a(x)$, for some $x \in X$;

¹²We use the standard notation $a \doteq b$ to mean $a = b \wedge \text{emp}$.

- the interference precondition holds: every step of the environment, before the local linearization point takes place goes from a resource satisfying $P_a(x_1)$, for some $x_1 \in X$, to a resource satisfying $P_a(x_2)$, for some $x_2 \in X$.

If any of the above are violated, the blame is on the environment. In return, the local steps are expected to:

- respect atomicity: transform the resources of $P_a(x)$ exactly once to resources satisfying $Q_a(x, y)$, for any $x \in X$ and some y ;
- respect the pre-/postconditions: transform (in possibly many steps) the resources in P_h to resources satisfying $Q_h(x, y)$ at the end of the execution.

In this sense, the resources in $P_a(x)$ should be understood as shared: the environment can use them to change the value of x , and the local steps can use them atomically to perform the linearization point. Note that $Q_a(x, y)$ is only guaranteed to hold *immediately after* the linearization point.

Key idea of the liveness semantics. In terms of termination, the specification (\boxtimes) guarantees local termination *only if* the environment is *live*, i.e it satisfies the layered liveness invariants represented by the pseudo-quantifiers (of the specification and in \mathcal{A}) and the obligations. The idea is again to identify when non termination is caused by a bad environment or by bad local steps. Consider the case of liveness invariants encoded by obligations. Imagine we annotate each position of a trace indicating which obligations are held at that point by the environment and which are held locally. Now suppose the environment always eventually fulfils *every* obligation (i.e. for each obligation O there are infinitely many positions where O is not held by the environment). This environment is certainly live, so it cannot be blamed for non termination. The layer structure, however, allows the environment to fulfil obligations of layer k by relying on eventual fulfilment of obligations at layer $< k$. Therefore, if there is an obligation O that is locally held forever, the environment is still considered live if it never fulfills obligations at layer $> \text{lay}(O)$. In this scenario, the local steps are blamed for non termination: by holding O forever, the local steps are not allowed to rely on the environment being live at higher layers than $\text{lay}(O)$.

This scheme leads to the following semantic interpretation of layers. The local steps can blame the environment for non termination, by waiting for the fulfilment of some environment obligation O_1 indefinitely; in turn the environment can blame the local steps for the inability to fulfil O_1 by claiming to be waiting for the fulfilment of some local obligation O_2 with $\text{lay}(O_1) > \text{lay}(O_2)$; the local step can justify the indefinite postponement of the fulfilment of O_2 by shifting blame on the environment again, appealing to an environment obligation with even lower layer, and so on. This blame-shifting cannot be unbounded: every time the blame is shifted, the layers considered are strictly lower and, by well-foundedness of layers, this cannot happen ad libitum. Ultimately, the blame is unambiguously placed on the environment or the local steps and the trace is accepted or rejected accordingly.

This intuition about obligations extends to liveness assumptions attached to pseudo-quantifications in the triple and in the atomicity context. All these assumptions need to be layered to avoid unsound circularities, which is why the pseudo-quantifier carries a layer k . The specifications mention another layer, m , which represents a (strict) upper bound on the layers that we may consider live when proving some command satisfies the specification. An environment is still considered live by the specification, if it keeps an obligation of layer $\geq m$ forever unfulfilled.

We now define the formal semantics of specifications, as set of concrete traces that satisfy the specification. To check if a concrete trace τ satisfies a specification, the semantics first collects all the possible “logical” justifications of the trace in a set \mathbb{T} . To justify a trace means to instrument each step with sets of worlds that show how the trace respects the (safety) logical constraints of

the specification. The set \mathbb{T} is then further analysed to check that every instrumented trace where the environment satisfies the liveness assumptions is locally terminating.

We begin by defining the *trace safety judgement*, of the form $\tau \vDash_{\mathbb{S}} p_h, p_a, v : \mathbb{T}$, the purpose of which is to check the safety constraints implied by \mathbb{S} . The judgement formalises the idea of a specification \mathbb{S} as a trace acceptor, that is, an automaton reading a trace step by step, and either accepting or rejecting it. If we ignore \mathbb{T} for a moment, the trace safety judgement represents a snapshot of the state of this imaginary automaton at a point when some prefix of the trace has been already successfully processed, and τ is the suffix that remains to be processed. The automaton traverses the trace producing a guess for an instrumentation, i.e. logical resources corresponding to the concrete memory contents that explains why the trace is acceptable. The instrumentation needs to describe, for instance, when the linearization point is thought of taking place, what portions of the state are considered as shared and which owned. Let (σ, h) be the current concrete state, i.e. $\tau = (\sigma, h) \tau'$. The triple (p_h, p_a, v) in the judgement encode the automaton’s current state, representing the current guess for the instrumentation of (σ, h) . The resources currently considered as locally owned are represented by the view p_h ; the v can be either an abstract value, in which case the automaton thinks that we are still in the interference phase, or it can be a pair $\langle v_1, v_2 \rangle$ which means we are past the linearization point, which updated the abstract state from v_1 to v_2 ; the $p_a(v)$ is a set of worlds parametric on v and corresponds to the shared atomic resources if we are before the linearization point, or it is the empty resource if we are past it. The judgement assumes that $h \in \llbracket p_h * p_a(v) * f \rrbracket_{\lambda}$ for some frame f , i.e. the current concrete state is consistent with the current instrumentation guess. The initial state of the automaton will be chosen so that this holds at the beginning of the trace, and each transition of the automaton will by construction preserve this correspondence.

As it walks down a trace, the automaton updates p_h, p_a and v , trying to construct a consistent instrumentation for the whole trace. Such a sequence of automaton states constitute its *run* over the trace. *Specification traces* augment each state of a trace with the instrumentation from a run of the automaton.

Definition 3.25 (Specification Traces). Define $\text{AVal}' \triangleq \text{AVal} \uplus \{ \langle v_1, v_2 \rangle \mid v_1, v_2 \in \text{AVal} \}$, the set of *specification states* to be $\text{SState}_{\mathcal{A}} \triangleq \text{View}_{\mathcal{A}} \times (\text{AVal}' \rightarrow \text{World}_{\mathcal{A}}^{\uparrow}) \times \text{AVal}'$ and the set of *specification configurations* to be $\text{SConf}_{\mathcal{A}} \triangleq \text{Store} \times \text{Heap} \times \text{SState}_{\mathcal{A}}$. The set of *specification traces*, $\text{STrace}_{\mathcal{A}}$, is the set of infinite sequences of the form $\hat{c}_1 \pi_1 \hat{c}_2 \pi_2 \cdots$ where $\hat{c}_i \in \text{SConf}_{\mathcal{A}}$ and $\pi_i \in \{\text{loc}, \text{env}\}$. Given a set of specification traces $\mathbb{T} \subseteq \text{STrace}$, we write $\hat{c} \pi \mathbb{T}$ for the set $\{ \hat{c} \pi \hat{t} \mid \hat{t} \in \mathbb{T} \}$.

The trace safety judgement accumulates, as it traverses a trace, all the successful instrumentations of the trace in \mathbb{T} , which we can later check against liveness properties. Let us define the judgement formally, and then explain it in detail.

Definition 3.26 (Trace Safety). Let $\mathbb{S} \in \text{Spec}$ with components named as (\star) , $\tau \in \text{Trace}$, $\mathbb{T} \subseteq \text{STrace}$, and $(p_h, p_a, v) \in \text{SState}$ such that

$$p_a(x) = \mathcal{W}_a \llbracket P_a \rrbracket (x) \triangleq \begin{cases} \mathcal{W} \llbracket P_a(x) \wedge x \in X \rrbracket_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases}$$

The *trace safety judgement* is the relation $\tau \vDash_{\mathbb{S}} p_h, p_a, v : \mathbb{T}$ defined coinductively in Fig. 7.¹³ We write $\text{term}(\tau)$, if the trace τ contains no local steps.

The judgement $\tau \vDash_{\mathbb{S}} p_h, p_a, v : \mathbb{T}$ assumes the initial configuration (σ_0, h_0) of the trace τ satisfies $h_0 \in \llbracket p_h * p_a(v) * \text{True} \rrbracket_{\lambda}$. Rule **STUTTER** checks that any local step other than the linearization point

¹³Here τ ranges over subsequences of traces.

$$\begin{array}{c}
\frac{(h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_h * p_a(v) \rightarrow p'_h * p_a(v) \quad (\sigma_2, h_2) \tau \vDash_{\mathbb{S}} p'_h, p_a, v : \mathbb{T} \quad \text{term}(\tau) \Rightarrow v = \langle v_1, v_2 \rangle \wedge p'_h = \mathcal{W} \llbracket Q_h(v_1, v_2) \rrbracket_{\mathcal{A}}^{\sigma_2}}{(\sigma_1, h_1) \text{loc } (\sigma_2, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, v : ((\sigma_1, h_1, p_h, p_a, v) \text{loc } \mathbb{T})} \text{STUTTER} \\
\\
\frac{(h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_h * p_a(v) \rightarrow q'_h * \mathcal{W} \llbracket Q_a(v, v') \rrbracket_{\mathcal{A}} \quad \text{term}(\tau) \Rightarrow q'_h = \mathcal{W} \llbracket Q_h(v, v') \rrbracket_{\mathcal{A}}^{\sigma_2} \quad (\sigma_2, h_2) \tau \vDash_{\mathbb{S}} q'_h, p_a, \langle v, v' \rangle : \mathbb{T}}{(\sigma_1, h_1) \text{loc } (\sigma_2, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, v : ((\sigma_1, h_1, p_h, p_a, v) \text{loc } \mathbb{T})} \text{LINPT} \\
\\
\frac{\mathbb{T} = \bigcup \{ (\sigma, h_1, p_h, p_a, v) \text{env } \mathbb{T}_{v'} \mid v' \in X, E(v') \} \quad \forall v' \in X. E(v') \Rightarrow (\sigma, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, v' : \mathbb{T}_{v'} \quad v \in \text{AVal} \quad E(v') \triangleq (\exists p_e, p'_e. h_1 \in \llbracket p_h * p_a(v) * p_e \rrbracket_{\lambda} \wedge (h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_a(v) * p_e \rightarrow p_a(v') * p'_e)}{(\sigma, h_1) \text{env } (\sigma, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, v : \mathbb{T}} \text{ENV} \\
\\
\frac{\text{if } \exists p_e, p'_e. h_1 \in \llbracket p_h * p_e \rrbracket_{\lambda} \wedge (h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_e \rightarrow p'_e \text{ then } (\sigma, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, \langle v, v' \rangle : \mathbb{T} \text{ else } \mathbb{T} = \emptyset}{(\sigma, h_1) \text{env } (\sigma, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, \langle v, v' \rangle : ((\sigma_1, h_1, p_h, p_a, \langle v, v' \rangle) \text{env } \mathbb{T})} \text{ENV}' \\
\\
\frac{}{(\sigma, h) \text{env } \not\vdash \tau \vDash_{\mathbb{S}} p_h, p_a, v : \emptyset} \text{ENV} \not\vdash
\end{array}$$

Fig. 7. Safety Specification Semantics

updates the local Hoare view (to some p'_h) in a frame-preserving manner; this implies that, before the linearization point, the abstract state v needs to be preserved by such step. Rule **LINPT** checks that the linearization point is frame-preserving and consistent with the atomic postcondition Q_a . Both rules **STUTTER** and **LINPT** check that the Hoare postcondition is satisfied if we are considering the last local step of the trace (i.e. if $\text{term}(\tau)$ holds). Rule **ENV** checks whether the current environment step, assumed to happen before the linearization point, can be seen as a transition changing the abstract state from v to v' in a way that does not disrupt any frame (including p_h). If that is the case, the rest of the trace is checked for safety. Rule **ENV'** performs the same check but after the linearization point. In both cases, if the environment step cannot be seen as frame-preserving, then the trace is accepted since the environment did not respect the assumptions. Similarly, Rule **ENV** $\not\vdash$ accepts the trace after a fault caused by the environment.

As we briefly mentioned, Definition 3.26 is inspired by alternating automata [40]. The “alternation” aspect is necessary because of the angelic/demonic duality between local and environment steps: when processing an environment step, we need to be prepared to handle *every* possible interpretation of the update that took place; for local steps, we are allowed to pick *any* interpretation of the update. Note that these ambiguities arise purely from the fact that we are instrumenting the trace with “ghost” logical state: at each step there is no ambiguity in a trace about how the *concrete* state has been updated. This dual interpretation gives rise to the two kinds of transitions in an alternating automaton. An automata-based presentation of the trace safety judgement would use existentially branching transitions for local steps, and universally branching transitions for environment steps. We further mimic alternating automata in the way we factor safety and liveness constraints. Alternating automata impose safety constraints by constructing sets of runs that linearise the choices for the existential transitions and the branching due to universal transitions. In our setting these sets of runs correspond to \mathbb{T} . The liveness constraints can be then checked by, for example, requiring each run in the set to visit final states infinitely often, the usual Büchi-style acceptance condition. Here we also examine the instrumented traces of \mathbb{T} individually and impose a liveness acceptance condition; the condition in our case is more complex as it has to take into account

layers, pseudo-quantifiers and obligations. One key simplification introduced by this approach is that we can cleanly separate the branching (safety) aspect—the quantifier alternation due to duality environment/local steps—from the linear-time liveness aspect.

Building on trace safety, we can now define the semantics of a specification $\llbracket \mathbb{S} \rrbracket$ as the set of traces that are safe and that additionally satisfy the liveness constraints implied by the obligations and the liveness assumptions of \mathbb{S} . Conceptually, we want to require local termination, if the environment satisfies the layered liveness invariants represented by pseudo-quantifiers and obligations. To harmonise the pseudo-quantification and obligation-related liveness assumptions of a specification, \mathbb{S} , we collect all of them in a set of so-called *pseudo-obligations*:

$$\text{POb}^{\mathbb{S}} \triangleq \{ (r, O) \mid r \in \text{Rld}, O \in \text{AOB} \} \uplus \{ (r, \text{live}(\mathcal{A}, r)) \mid r \in \text{dom}(\mathcal{A}) \} \uplus \{ X \rightarrow_k X' \}$$

where \mathcal{A}, X, X' and k are taken from the specification.

We extend the layer function to $\text{lay}: \text{POb}^{\mathbb{S}} \rightarrow \mathcal{L}$ by setting $\text{lay}(a) = k$ if $a = (r, O)$ and $\text{lay}(O) = k$, or $a = (r, X \rightarrow_k X')$, or $a = (X \rightarrow_k X')$. Furthermore, define

$$\text{POb}_{<k}^{\mathbb{S}} \triangleq \{ \widehat{O} \in \text{POb}^{\mathbb{S}} \mid \text{lay}(\widehat{O}) < k \} \quad \text{AOB}_{<k} \triangleq \{ O \in \text{AOB} \mid \text{lay}(O) < k \}$$

Now we want to understand, for each position of a specification trace, which pseudo-obligations we are holding locally and which are held by the environment. This information is contained in the single worlds so as a first step we extract, from a specification trace, the set of traces of worlds that it represents.

Definition 3.27 (World Traces). Given an atomicity context, \mathcal{A} , we call *world traces*, $\text{WTrace}_{\mathcal{A}}$, ranged over by $\bar{\tau}, \bar{\tau}', \dots$, the infinite sequences of the form $(h_0, w_h^0, w_a^0, w_e^0, v^0) \pi_0 (h_1, w_h^1, w_a^1, w_e^1, v^1) \pi_1 \dots$, where, for all $i \in \mathbb{N}$, $h_i \in \text{Heap}$, $w_h^i, w_a^i, w_e^i \in \text{World}_{\mathcal{A}}$ and $v^i \in \text{AVal}'$. We define the function:

$$\mathcal{W}_{\lambda}(s, h, p_h, p_a, v) \triangleq \{ (h, w_h, w_a, w_e, v) \mid w_h \in p_h, w_a \in p_a(v), h \in \lfloor w_h \odot w_a \odot w_e \rfloor_{\lambda} \}$$

which we extend to specification traces by $\mathcal{W}_{\lambda}(\hat{c}_0 \pi_0 \hat{c}_1 \pi_1 \dots) \triangleq \{ c_0 \pi_0 c_1 \pi_1 \dots \mid \forall i. c_i \in \mathcal{W}_{\lambda}(\hat{c}_i) \}$. A world trace $(h_0, w_h^0, w_a^0, w_e^0, v^0) \pi_0 (h_1, w_h^1, w_a^1, w_e^1, v^1) \pi_1 \dots$ is $\mathbf{R}_{\mathcal{A}}^a$ -respecting if for all $i \in \mathbb{N}$:

$$\pi_i = \text{env} \Rightarrow w_h^i \mathbf{R}_{\mathcal{A}}^a w_h^{i+1} \quad \wedge \quad \pi_i = \text{loc} \Rightarrow w_e^i \mathbf{R}_{\mathcal{A}}^a w_e^{i+1}$$

Given specification trace $\hat{\tau} \in \text{STrace}_{\mathcal{A}}$, the set $\llbracket \hat{\tau} \rrbracket_{\lambda; \mathcal{A}}$ is the *set of world traces of $\hat{\tau}$* , defined by

$$\llbracket \hat{\tau} \rrbracket_{\lambda; \mathcal{A}} \triangleq \{ \bar{\tau} \in \mathcal{W}_{\lambda}(\hat{\tau}) \mid \bar{\tau} \text{ is } \mathbf{R}_{\mathcal{A}}^a\text{-respecting} \}$$

We lift $\llbracket \cdot \rrbracket_{\lambda; \mathcal{A}}$ to apply to sets of specification traces in the obvious way.

We can now define two predicates indicating when a pseudo-obligation is considered to be held by the environment (envheld_{λ}) or locally (locheld_{λ}) in a position of a world trace.

$$\text{envheld}_{\lambda}(\widehat{O}, (_, w_h, _, _, v)) \triangleq \begin{cases} \xi_{w_h}(r) \sqsupseteq \mathbf{O} \wedge \text{lvl}_{w_h}(r) < \lambda & \text{if } \widehat{O} = (r, \mathbf{O}) \\ \text{ast}_{w_h}(r) \notin X_2 \wedge \text{lvl}_{w_h}(r) < \lambda & \text{if } \widehat{O} = (r, X_1 \rightarrow_k X_2) \\ v \in X_1 \setminus X_2 & \text{if } \widehat{O} = (X_1 \rightarrow_k X_2) \end{cases}$$

$$\text{locheld}_{\lambda}((r, \mathbf{O}), (_, w_h, _, _, v)) \triangleq \theta_{w_h}(r) \sqsupseteq \mathbf{O} \wedge \text{lvl}_{w_h}(r) < \lambda$$

Equipped with these definitions, we can state the liveness constraints associated with a specification. The idea is that one can assign the “blame” for local non-termination either to the environment or to the local behaviour. If we deem the environment responsible for non-termination then the specification will classify the trace as acceptable, otherwise it will reject it. The idea behind this “blame” assignment is to examine the world traces justifying the safety of a trace and consider, for each position, which obligations are held by the environment and which are held locally. To understand the intuition, consider the case of liveness invariants encoded by obligations. Suppose

the environment always eventually fulfills *every* obligation, i.e. for each obligation O there are infinitely many positions where O is not held by the environment. This environment is certainly *live*, i.e. it respects the liveness assumptions and the local code is responsible for any non-terminating behaviour. But what if the environment itself is blocking on some locally held obligation, and as a consequence is not able to fulfill some O ? Whether the environment or the local code is to blame depends on the layers. The environment is to blame if, from some point in the trace, it never fulfills some O but the local steps always eventually fulfill every obligation of layer strictly lower than $\text{lay}(O)$. Conversely, an environment which keeps O unfulfilled because of some forever-unfulfilled obligation O' held locally with $\text{lay}(O) > \text{lay}(O')$ cannot be blamed for local non-termination.

This intuition about obligations extends to liveness assumptions attached to pseudo-quantifications in the triple and in the atomicity context. The liveEnv predicate given in Definition 3.28 formalises the above blame-assigning mechanism. A world trace which satisfies liveEnv is one where the environment cannot be blamed for local non-termination. The specification semantics then is the set of safe traces where, if liveEnv is satisfied, then the trace is locally terminating.

Definition 3.28 (Specification Semantics). Fix a specification $\mathbb{S} \in \text{Spec}$ with components named as in (\mathfrak{X}) . The $\text{liveEnv}_{\mathbb{S}}(\hat{\tau})$ predicate checks whether the environment is satisfying the liveness assumptions of the specification:

$$\text{liveEnv}_{\mathbb{S}}(\hat{\tau}) \triangleq \widehat{\forall} \hat{O} \in \text{POb}_{< m}^{\mathbb{S}} \text{ \textbf{if} } \forall r, O \in \text{AOb}_{\leq \text{lay}(\hat{O})}, \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_{\lambda}((r, O), \hat{\tau}(j)) \\ \text{ \textbf{then} } \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(\hat{O}, \hat{\tau}(j))$$

Let $p_h = \mathcal{W}[[P_h]]_{\mathcal{A}}^{\sigma_0}$, and $p_a = \mathcal{W}_a[[P_a]]$. We define the trace semantics $[\mathbb{S}] \subseteq \text{Trace}$ of specification \mathbb{S} as the set:

$$[\mathbb{S}] \triangleq \left\{ (\sigma_0, h_0) \tau \left| \begin{array}{l} \forall v_0 \in X. \text{ \textbf{if} } h_0 \in [p_h * p_a(v_0) * \text{True}]_{\lambda} \\ \text{ \textbf{then} } \exists \mathbb{T}. (\sigma_0, h_0) \tau \models_{\mathbb{S}} p_h, p_a, v_0 : \mathbb{T} \\ \wedge \forall \hat{\tau} \in \llbracket \mathbb{T} \rrbracket_{\lambda, \mathcal{A}}. \text{liveEnv}_{\mathbb{S}}(\hat{\tau}) \Rightarrow \text{lterm}((\sigma_0, h_0) \tau) \end{array} \right. \right\}$$

where λ and \mathcal{A} are the level and atomicity context from the specification \mathbb{S} .

The more precise intuition behind the specification semantics is as follows. Once it has been established that there is a way to instrument the trace to justify why the local steps satisfy the safety constraints of \mathbb{S} , we consider the set of valid instrumentations \mathbb{T} . First, we extract the set of world traces represented by the traces of \mathbb{T} . Each such world trace should either be locally terminating, in which case the trace is accepted or, if it is non-terminating, the non-termination should be due to the environment not satisfying the liveness assumptions of \mathbb{S} . The predicate $\text{liveEnv}_{\mathbb{S}}(\hat{\tau})$ holds for a specification trace $\hat{\tau}$ if the environment always eventually fulfills any pseudo-obligation with layer k and if *no obligation of layer $< k$ is constantly held by the local thread*. Blame for non-termination can be unambiguously assigned thanks to well-foundedness of layers: if there is a forever-unfulfilled local obligation O_0 , we can try to blame the environment by identifying a lower-layer obligation O_1 that is forever-unfulfilled by the environment; the environment can shift the blame back to the local steps if one can find a lower-layer local obligation O_2 that is forever-unfulfilled. Well-foundedness implies this blame-shifting game must be bounded in length and the ultimate culprit can always be identified. This effectively encodes the acyclicity of the layered termination argument.

3.10 The Semantic Judgement

We are now ready to define the semantic version of our judgements, $\models_{\Phi} \mathbb{C} : \mathbb{S}$, indexed by a function specification context, Φ , which, for each function, provides the arguments of the function and the specification of the function body.

Definition 3.29 (Function Specification Context). A function specification context, Φ , is a partial function $\Phi \in \text{FSpec} \triangleq \text{FName} \rightarrow (\text{PVar}^*, \text{Spec})$.

Definition 3.30 (Semantic Triple). Given $\varphi \in \text{Flmpl}$ and $\Phi \in \text{FSpec}$, a function implementation context φ is a correct implementation of Φ , written $\vDash \varphi : \Phi$, if and only if $\forall f, \vec{x}, \mathbb{S}. \Phi(f) = (\vec{x}, \mathbb{S}) \Rightarrow \exists \mathbb{C}. \varphi(f) = (\vec{x}, \mathbb{C}) \wedge \llbracket \mathbb{C} \rrbracket_\varphi \subseteq \llbracket \mathbb{S} \rrbracket$. The semantic triple $\vDash_\Phi \mathbb{C} : \mathbb{S}$, stating that command \mathbb{C} satisfies specification \mathbb{S} under any correct implementation of the functions specified in Φ , is defined by:

$$\vDash_\Phi \mathbb{C} : \mathbb{S} \quad \text{if and only if} \quad \forall \varphi. \vDash \varphi : \Phi \Rightarrow \llbracket \mathbb{C} \rrbracket_\varphi \subseteq \llbracket \mathbb{S} \rrbracket$$

Note that when \mathbb{C} has no free function names, the judgement $\vDash_\Phi \mathbb{C} : \mathbb{S}$ is equivalent to $\llbracket \mathbb{C} \rrbracket \subseteq \llbracket \mathbb{S} \rrbracket$.

Since the semantics of our triples is a complex conditional termination statement, it is useful to show when it corresponds to unconditional termination. Intuitively, to state facts about the behaviour of a command in an “empty” environment, we should be using a Hoare triple (no resource needs to be shared, no interference experienced) and there should be no assumption of obligations being owned by the environment. We characterise the preconditions that ensure this as the ones that always admit $\text{emp}_{\text{Ob}}^\lambda$ as a global frame.

Definition 3.31 (Grounded view). Fix an arbitrary level λ . We say $p \in \text{View}_\emptyset$ is λ -grounded if

$$\forall h. h \in \llbracket p * \text{True} \rrbracket_\lambda \Rightarrow h \in \llbracket p * \text{emp}_{\text{Ob}}^\lambda \rrbracket_\lambda$$

We say a stable assertion P is λ -grounded if, for all $\sigma \in \text{Store}$, the view $\mathcal{W}[P]_\emptyset^\sigma$ is λ -grounded.

Examples of grounded assertions are standard separation logic assertions like emp or $x \mapsto v$.

Unconditional termination applies to programs running in isolation. Note that, technically, we cannot consider traces without environment steps as the fairness constraint requires infinitely many of those. We therefore model the isolated executions of a command as the executions where the environment steps do not modify the current state. It is easy to check that, for each finite or infinite sequence of local steps of \mathbb{C} , there is a corresponding *fair* trace of \mathbb{C} with only identity environment steps, and vice versa.

Definition 3.32 (Closed-World Semantics). Given a command \mathbb{C} , its *closed-world semantics* $C[\mathbb{C}] \subseteq \text{Trace}$ is the subset of the open-world semantics $\llbracket \mathbb{C} \rrbracket$ of the traces where every environment step is an identity step, i.e. of the form $c \xrightarrow{\text{env}} c$. Additionally, for an assertion P , we define $C[\mathbb{C}](P)_\lambda \triangleq \{ (\sigma, h)\tau \in C[\mathbb{C}] \mid h \in \llbracket \mathcal{W}[P]_\emptyset^\sigma * \text{True} \rrbracket_\lambda \}$, that is the closed-world traces of \mathbb{C} that start from a state satisfying the precondition P .

THEOREM 3.33 (ADEQUACY). For every λ -grounded assertion P , if $m; \lambda; \emptyset \vDash \{P\} \mathbb{C} \{Q\}$ then all traces in $C[\mathbb{C}](P)_\lambda$ are locally terminating.

PROOF. Take a trace $(\sigma, h)\tau \in C[\mathbb{C}](P)_\lambda$ and let $p = \mathcal{W}[P]_\emptyset^\sigma$. From the semantic triple we know that $C[\mathbb{C}](P)_\lambda \subseteq \llbracket \mathbb{C} \rrbracket \subseteq \llbracket \{P\} \cdot \{Q\} \rrbracket$. We have $h \in \llbracket p * \text{True} \rrbracket_\lambda$ and, since P is λ -grounded, $h \in \llbracket p * \text{emp}_{\text{Ob}}^\lambda \rrbracket_\lambda$. By the definition of the specification trace semantics, we therefore know that $(\sigma, h) \tau \vDash_{\mathbb{S}} p, \text{emp}, 1 : \mathbb{T}$ for some \mathbb{T} . Note that a frame-preserving update on a grounded view keeps it grounded, and that identity environment steps can always be justified as a frame-preserving update that does not update the resources. In particular, from these facts we can deduce that there is some $\hat{\tau} \in \mathbb{T}$ such that at each point in time the global frame is $\text{emp}_{\text{Ob}}^\lambda$. From this we can extract a world-trace $\bar{\tau} \in \llbracket \mathbb{T} \rrbracket_{\lambda; \emptyset}$ such that $\forall \bar{O} \in \text{POb}_{< m}. \forall i \in \mathbb{N}. \neg \text{envheld}_\lambda(\bar{O}, \bar{\tau}(i))$ which implies $\text{liveEnv}(\bar{\tau})$. By the definition of the specification’s semantics this implies local termination of our concrete trace $(\sigma, h)\tau$. \square

As a corollary, we have that if $m; \lambda; \emptyset \vDash \{\text{emp}\} \mathbb{C} \{\text{True}\}$ holds, then every isolated execution of \mathbb{C} from the empty heap and arbitrary store terminates.

4 TADA LIVE RULES

We now introduce the rules of TaDA Live, summarised in Fig. 9, using a simple but tricky running example to motivate and explain them.

Example 4.1 (Distinguishing client). Consider the following client of a lock module:

lock(x);		var d = false in
[done] := true;		while(\neg d){
unlock(x);		lock(x); d := [done]; unlock(x);
		}

The code is interesting in that it can distinguish whether the lock implementation is a spin or CLH lock. Under weak fairness, when x is a spin lock, this client program does not always terminate. It is possible for the lock invocation of the left thread to be scheduled infinitely often but always in a state in which the lock is locked. As a result, `done` will never be set to `true`, making the while loop spin forever. The spin lock has been *starved* by the other thread. In contrast, when x is a CLH lock, this client program is guaranteed to terminate: a fair scheduler will eventually allow the left thread to enqueue itself in the internal queue of the lock; from then on, the thread on the right can only acquire the lock at most once; after unlocking, the next `lock(x)` call of the right thread would enqueue it after the left thread, which is now the only unblocked thread. The CLH lock is *starvation free*.

It is worth noting that none of the proof principles of [5, 8, 13, 14, 21, 32] are powerful enough to handle this example due to the blocking behaviour it displays. Even replacing locks with primitive locks, due to the mix of busy-waiting blocking and locks, the example cannot be handled by any of the proof systems of [3, 22, 26, 28]. Since LiLi does not have a rule for parallel, this client cannot be proven within the LiLi logic.

We show that the distinguishing client terminates with the CLH lock, by proving the Hoare triple $\top \vdash \{L(x, 0) * done \mapsto false\} C_\ell \parallel C_r \{True\}$, where C_ℓ and C_r are the left and right threads of the example, respectively. Since our triples are total, this triple immediately guarantees termination of the program. Our overall argument is as follows. The CLH specification guarantees termination of a call to `lock(x)` if the lock is always eventually unlocked by the environment. This is intuitively true for both threads: they always unlock the lock after having acquired it. The call to `lock(x)` will therefore terminate in both threads. The only other potentially non-terminating operation is the while loop in the right thread. The loop is implementing a busy-wait pattern on `done`, and needs the help of the left thread to terminate. We will be able to prove that since `done` is going to be eventually set to `true` (and never reset to `false`), the loop will terminate.

4.1 The Basics: Regions

Let us formalise the argument in TaDA Live, introducing the proof rules as they are needed. Recall the specifications of CLH lock:

$$\begin{aligned} \mathbf{1} &\vdash \forall l \in \{0, 1\} \rightarrow_0 \{0\}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle \\ \mathbf{0} &\vdash \langle L(x, 1) \rangle \text{unlock}(x) \langle L(x, 0) \rangle \end{aligned}$$

where we make explicit the previously omitted layers $\mathbf{1} > \mathbf{0}$ (we justify the choice of layers in the proof of CLH lock). These specifications will be available in the proof as “axioms” stored in a function specification context Φ parametrising every triple of the client proof; we omit the parametrisation to aid readability. The predicate $L(x, l)$ is given a definition in the proof of the lock module, and, in the spirit of CAP, the client proof should not be relying on the *definition* of the predicate but in its abstract properties. Here we rely on the fact that $L(x, _) * L(x, _)$ is false,

expressing that a lock is an exclusively owned resource —see Section 4.8 for the other properties of L exposed to the client.

The two threads of the distinguishing client both access the lock x and the heap cell $done$. Consider the precondition $L(x, 0) * done \mapsto \text{false}$. Both resources in the precondition are non-duplicable, so if we give them to \mathbb{C}_ℓ , the other thread would not be able to also have them. As we anticipated, shared state in TaDA is handled using regions. We therefore introduce a new region type \mathbf{dc} (for distinguishing client) which encapsulates the resources in the precondition: $\mathbf{dc}_r(x, done, l, d)$ is the shared resource encapsulating a lock at x with state l and a cell at $done$ storing the Boolean d . Although in this case the abstract state is not hiding any detail, since both l and d are visible, in general the abstraction of the contents is an essential mechanism for reasoning about *abstract* atomicity. In the proof of CLH lock, for example, to be able to see the operations as abstractly atomic, it is essential to hide the queue from the abstract state.

Assuming the lock region encapsulated by the L predicate has level λ , the lock specifications will have level $\lambda + 1$ in the context, indicating they consider the lock region closed. To allow \mathbf{dc} to encapsulate the lock region and use the lock specifications to derive updates to its own state, we let it have level $\lambda + 1$. The top-level triples for the distinguishing client have level $\lambda + 2$ as a consequence. We will elide all details about levels as they can be mechanically inferred from the applications of the **LIFTA** and **UPDREG** rules.

We now design the protocol of the region, with the intent of encoding the following safety invariants:

- (I1) the addresses of the lock and the flag never change;
- (I2) only the thread which acquired the lock can unlock it;
- (I3) only the left thread will ever modify $done$, and at most once from false to true ;

and the following liveness invariants:

- (I4) the lock will always eventually be unlocked;
- (I5) the value at $done$ will always eventually be true .

Note that, together, invariants (I3) and (I5) imply that eventually the value at $done$ will always be true . To encode invariants (I2) and (I3) we introduce a guard algebra for \mathbf{dc} generated from two guards \mathbf{k} (for **key**) and \mathbf{d} (for **Done**), with $\mathbf{k} \bullet \mathbf{k} = \perp$ and $\mathbf{d} \bullet \mathbf{d} = \perp$ to represent exclusivity of the permissions they give on the lock and flag respectively. We can reuse the same guard algebra for the obligation algebra associated with \mathbf{dc} : the atom obligations \mathbf{k} and \mathbf{d} will represent liveness invariants (I4) and (I5) respectively. The protocol $\mathcal{T}_{\mathbf{dc}}$ formalises all the invariants:

$$\mathbf{0} : ((x, done, 0, d), \mathbf{0}) \rightsquigarrow ((x, done, 1, d), \mathbf{k}) \quad (13)$$

$$\mathbf{k} : ((x, done, 1, d), \mathbf{k}) \rightsquigarrow ((x, done, 0, d), \mathbf{0}) \quad (14)$$

$$\mathbf{d} : ((x, done, l, \text{false}), \mathbf{d}) \rightsquigarrow ((x, done, l, \text{true}), \mathbf{0}) \quad (15)$$

The fact that no transition can change x and $done$ encodes (I1); this is such a common pattern that we adopt the convention to declare which are the *fixed* components of the abstract state and omit them from the protocol transitions completely. The choice for the guards of (14) and (15) reflects (I2) and (I3), respectively; we will give \mathbf{d} to the left-hand thread, and \mathbf{k} will be obtained by locking the lock. The obligation \mathbf{k} is obtained by locking and fulfilled by unlocking; the obligation \mathbf{d} is fulfilled by setting $done$ to true ; these facts encode (I4) and (I5).

We assign layers to the obligations to reflect the intuitive dependency: the lock needs to be assumed live in the process of fulfilling the obligation on the flag. We therefore set $\perp < \mathbf{0} = \text{lay}(\mathbf{k}) < \text{lay}(\mathbf{d}) = \mathbf{1} < \top$.

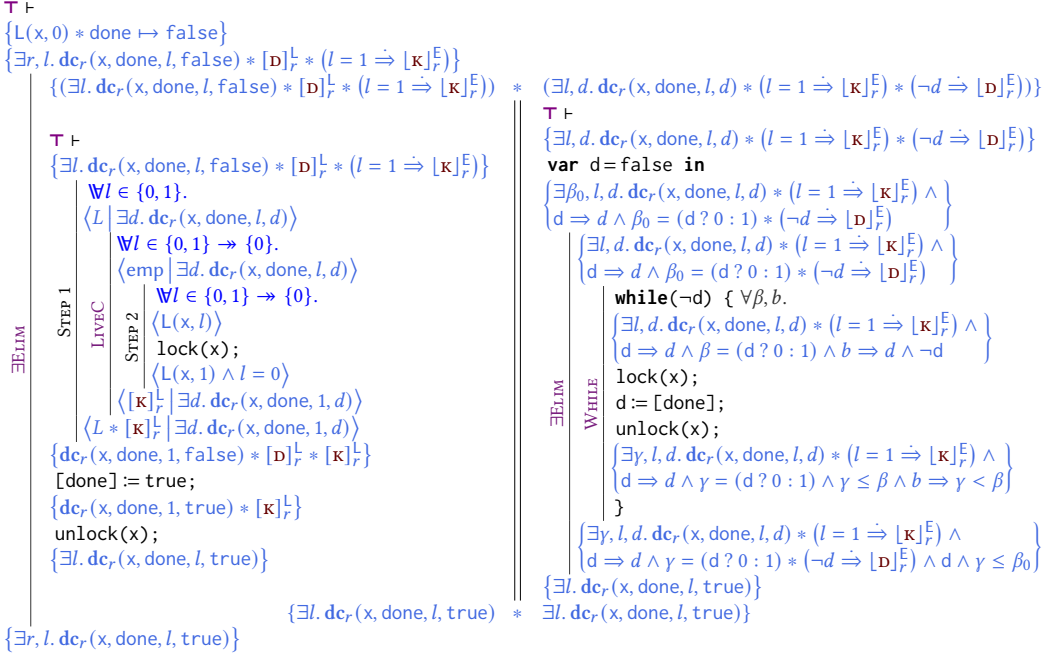


Fig. 8. Proof Sketch of the Distinguishing Client. Here $L = (\exists l, d. \text{dc}_r(x, \text{done}, l, d) * (l = 1 \dot{\Rightarrow} [\mathsf{K}]_r^E))$.

To complete the definition of shared region dc , we link its abstract state to the actual heap content that it encapsulates using the *region interpretation*:

$$\mathcal{I}(\text{dc}_r(x, \text{done}, l, d)) \triangleq L(x, l) * \text{done} \mapsto d * (l = 0 \dot{\Rightarrow} [\mathsf{K}]_r^L) * (d \dot{\Rightarrow} [\mathsf{D}]_r^L)$$

This assertion describes a portion of the heap being shared (the lock at x and the cell at done) and the linking of the ghost state (the guards and obligations) with the abstract state. The assertion $[\mathsf{K}]_r^L$ is an abbreviation for $[\mathsf{K}]_r * [\mathsf{K}]_r^L$, which indicates local ownership of the guard K and obligation K . The interpretation of a region establishes the invariant that, when $l = 0$, the guard and obligation K will be “owned” by the region (and by no thread as a consequence). Similar links are established between the value of d and D .

Now that we set the scene, we can proceed with the proof, outlined in Fig. 8. The first operation to do is to transform the precondition $L(x, 0) * \text{done} \mapsto \text{false}$ to an assertion about the $\text{dc}_r(x, \text{done}, l, d)$ region. We can do that by using the consequence rule:

$$\frac{\lambda; \mathcal{A} \vDash P \Rightarrow P' \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \{P'\} \mathbb{C} \{Q'\} \quad \lambda; \mathcal{A} \vDash Q' \Rightarrow Q}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P\} \mathbb{C} \{Q\}} \text{CONSH}$$

which allows the use of *viewshift* to logically manipulate the assertions. Since triples are only well-defined if the Hoare pre-/postconditions are stable, the rule asks to check stability of the assertions of the conclusion as this does not follow from stability of the assertions of the triple in the premise. An analogous rule, called **CONS**, holds for hybrid triples —with no stability checks on the atomic pre-/postconditions— so viewshifting is available at any point in a derivation.

In our example, we want to create the guards and obligations needed to match the interpretation of $\text{dc}_r(x, \text{done}, l, d)$ and create the region, replacing its interpretation. Here we might be tempted to

match the interpretation with $l = 0$ and $d = \text{false}$, $L(x, 0) * \text{done} \mapsto \text{false} * [\kappa]_r^l * [\mathbf{D}]_r^l$ to viewshift to $\mathbf{dc}_r(x, \text{done}, 0, \text{false}) * [\mathbf{D}]_r^l$. While this viewshift holds, there is an issue: in TaDA Live, all the assertions in Hoare triples (or in Hoare position in hybrid triples) need to be stable for the triple to have well-defined semantics. The proof system enforces the stability of these assertions, by inserting stability checks in crucial rules. This means that if we viewshift now to a non-stable assertion, then we would fail at some point in the derivation to satisfy a stability check. While $L(x, 0) * \text{done} \mapsto \text{false}$ is stable, as we own these resources, the assertion $\mathbf{dc}_r(x, \text{done}, 0, \text{false}) * [\mathbf{D}]_r^l$ is not stable: a region is subjected to the transitions of the protocol. Since we have the guard \mathbf{D} (from $[\mathbf{D}]_r^l$) the environment cannot own it, hence cannot fire the transitions guarded by \mathbf{D} ; this makes $d = \text{false}$ stable. The transitions changing the state of the lock, however, can affect the region. This leads us to¹⁴ $\exists r, l. \mathbf{dc}_r(x, \text{done}, l, \text{false}) * [\mathbf{D}]_r^l * (l = 1 \Rightarrow [\kappa]_r^E)$ where we also add $[\kappa]_r^E$ when $l = 1$, a stable fact.

4.2 The Parallel Rule

We now want to proceed with an application of the parallel rule:

$$\frac{m_1; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad \vdash_{\mathcal{A}} Q_1 \triangleright m_2 \leq m \quad m_2; \lambda; \mathcal{A} \vdash_{\Phi} \{P_2\} \mathbb{C}_2 \{Q_2\} \quad \vdash_{\mathcal{A}} Q_2 \triangleright m_1 \leq m}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}} \text{PAR}$$

The abbreviation $\vdash_{\mathcal{A}} P \triangleright k$ means $\forall r \in \text{Rld}. \vdash_{\mathcal{A}} P \Rightarrow r \triangleright k$, that is, all the obligations owned by P have layer $\geq k$. $\vdash_{\mathcal{A}} P \triangleright k \leq k'$ means $\vdash_{\mathcal{A}} P \triangleright k$ and $k \leq k'$. The intuition behind these constraints is as follows. The layer in the context of the triples indicates a strict upper bound on the layers that can be assumed live in the proofs of the triples. If thread 2 needs layers lower than m_2 , then if thread 1 has unfulfilled obligations by the end of its execution, these cannot conflict with the assumptions made by the proof of thread 2. It is not however sound to leave an obligation O_2 of layer $< m_2$ unfulfilled by thread 1: if thread 1 terminates first, leaving O_2 unfulfilled in its postcondition, thread 2 may be assuming O_2 live in a situation where it will never be fulfilled.

In our example we need to apply consequence again to massage the viewshifted precondition into an assertion of the form $P_1 * P_2$. The region assertion is duplicable, as is $l = 1 \Rightarrow [\kappa]_r^E$, but we want to give the non-duplicable resource $[\mathbf{D}]_r^l$ to the thread on the left, as it is the one that will fulfill the \mathbf{D} obligation. This has a side-effect though: since the \mathbf{D} guard is given to the left thread, the value of d from the right thread’s perspective is not stably false . Moreover, we want to know that the left thread has the obligation \mathbf{D} . So we use $[\mathbf{D}]_r^l \Leftrightarrow [\mathbf{D}]_r^l * [\mathbf{D}]_r^E$ to give $[\mathbf{D}]_r^E$ to the thread on the right, and we stabilise the assertion to $\neg d \Rightarrow [\mathbf{D}]_r^E$. All in all we obtain:

$$\begin{aligned} (\exists r, l. \mathbf{dc}_r(x, \text{done}, l, \text{false}) * [\mathbf{D}]_r^l * (l = 1 \Rightarrow [\kappa]_r^E)) &\Rightarrow \exists r. P_1 * P_2 \\ P_1 &\triangleq \exists l. \mathbf{dc}_r(x, \text{done}, l, \text{false}) * [\mathbf{D}]_r^l * (l = 1 \Rightarrow [\kappa]_r^E) \\ P_2 &\triangleq \exists l, d. \mathbf{dc}_r(x, \text{done}, l, d) * (l = 1 \Rightarrow [\kappa]_r^E) * (\neg d \Rightarrow [\mathbf{D}]_r^E) \end{aligned}$$

Which allows us to apply consequence and the standard $\exists\text{ELIM}$ rule to obtain a precondition in the form required by the **PAR** rule. For both threads we are aiming at postcondition $\exists l. \mathbf{dc}_r(x, \text{done}, l, \text{true})$ which has no obligation so it satisfies the layer conditions trivially.

To see why the layer conditions are important for soundness, imagine we forgot to unlock x in the left thread, obtaining a non-terminating program. We would obtain $[\kappa]_r^l$ in the postcondition of \mathbb{C}_l , but the check would fail as $0 = \text{lay}(\kappa) \not\geq \top$. Choosing 0 for the context layer of the triple for \mathbb{C}_r would not work: in its proof we need to assume \mathbf{D} live, and $\text{lay}(\mathbf{D}) = 1$.

¹⁴Recall that $\mathbb{B} \Rightarrow Q$ stands for $(\mathbb{B} \wedge Q) \vee (\neg \mathbb{B} \wedge \text{emp})$.

4.3 Handling a call to lock

Let us focus on the proof of the left-hand thread first. The difficult step is the execution of the first instruction, since this is the only potentially non-terminating instruction of the thread. If we let $L = \exists l, d. \text{dc}_r(x, \text{done}, l, d) * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)$, STEP 1 can be derived as follows:

$$\begin{array}{c}
 \frac{\mathbf{1} \vdash \mathbf{W}l \in \{0, 1\}. \langle L \mid \exists d. \text{dc}_r(x, \text{done}, l, d) \rangle \text{lock}(x) \langle L * \lfloor \mathbf{k} \rfloor_r^L \mid \exists d. \text{dc}_r(x, \text{done}, 1, d) \rangle}{\mathbf{1} \vdash \langle L \mid \exists l, d. \text{dc}_r(x, \text{done}, l, d) \rangle \text{lock}(x) \langle L * \lfloor \mathbf{k} \rfloor_r^L \mid \exists d. \text{dc}_r(x, \text{done}, 1, d) \rangle} \text{A}\exists\text{ELIM} \\
 \frac{\mathbf{1} \vdash \langle L \mid \exists l, d. \text{dc}_r(x, \text{done}, l, d) \rangle \text{lock}(x) \langle L * \lfloor \mathbf{k} \rfloor_r^L \mid \exists d. \text{dc}_r(x, \text{done}, 1, d) \rangle}{\mathbf{1} \vdash \{L * \exists l, d. \text{dc}_r(x, \text{done}, l, d)\} \text{lock}(x) \{L * \lfloor \mathbf{k} \rfloor_r^L * \exists d. \text{dc}_r(x, \text{done}, 1, d)\}} \text{ATOMW} \\
 \frac{\mathbf{1} \vdash \{L * \exists l, d. \text{dc}_r(x, \text{done}, l, d)\} \text{lock}(x) \{L * \lfloor \mathbf{k} \rfloor_r^L * \exists d. \text{dc}_r(x, \text{done}, 1, d)\}}{\mathbf{1} \vdash \{\exists l, d. \text{dc}_r(x, \text{done}, l, d) * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)\} \text{lock}(x) \{\exists d. \text{dc}_r(x, \text{done}, 1, d) * \lfloor \mathbf{k} \rfloor_r^L\}} \text{CONS} \\
 \frac{\mathbf{1} \vdash \{\exists l, d. \text{dc}_r(x, \text{done}, l, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)\} \text{lock}(x) \{\text{dc}_r(x, \text{done}, 1, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * \lfloor \mathbf{k} \rfloor_r^L\}}{\mathbf{1} \vdash \{\exists l, d. \text{dc}_r(x, \text{done}, l, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)\} \text{lock}(x) \{\text{dc}_r(x, \text{done}, 1, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * \lfloor \mathbf{k} \rfloor_r^L\}} \text{FRAMEH} \\
 \frac{\mathbf{1} \vdash \{\exists l, d. \text{dc}_r(x, \text{done}, l, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)\} \text{lock}(x) \{\text{dc}_r(x, \text{done}, 1, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * \lfloor \mathbf{k} \rfloor_r^L\}}{\mathbf{T} \vdash \{\exists l, d. \text{dc}_r(x, \text{done}, l, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E)\} \text{lock}(x) \{\text{dc}_r(x, \text{done}, 1, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L * \lfloor \mathbf{k} \rfloor_r^L\}} \text{LAYWH}
 \end{array}$$

Let us unpack the derivation. As a first step, we would like to frame the irrelevant resources, in this case $\lfloor \mathbf{D} \rfloor_r^L$. In TaDA Live, this step is more subtle and interesting than usual, because of the layer-related side-conditions of rule **FRAMEH** (a special case of rule **FRAME**):

$$\frac{\text{fv}(R) \cap \text{mod}(\mathbb{C}) = \emptyset \quad \mathbf{T} \vdash R \triangleright m \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \{P\} \mathbb{C} \{Q\} \quad \mathcal{A} \vDash R \text{ stable}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P * R\} \mathbb{C} \{Q * R\}} \text{FRAMEH}$$

With this rule, one can only frame obligations if they are of layer greater or equal the context layer. Here we can use consequence (omitted) to obtain a stable frame $R = \exists l, d. \text{dc}_r(x, \text{done}, l, \text{false}) * \lfloor \mathbf{D} \rfloor_r^L$ of the pre- and postconditions. We have $\mathbf{T} \vdash R \triangleright \mathbf{1}$ but $\not\vdash R \triangleright \mathbf{T}$; since the layer in the context of the goal is \mathbf{T} , before we can apply **FRAMEH** we need to artificially lower the layer using rule **LAYWH** before applying frame:

$$\frac{m_1 \leq m_2 \quad m_1; \lambda; \mathcal{A} \vdash_{\Phi} \{P\} \mathbb{C} \{Q\}}{m_2; \lambda; \mathcal{A} \vdash_{\Phi} \{P\} \mathbb{C} \{Q\}} \text{LAYWH}$$

Notice that lowering the layer in the context is always sound (even for hybrid triples): if we can prove the triple assuming live only layers $< k_1 \leq k_2$, then the proof is valid in contexts where layers up to k_2 can be assumed live.

The layer constraint of **FRAME** is crucial for soundness. Suppose we remove the constraint. Then we would be able to frame a locally held obligation O with layer $k < m$, i.e. one of the layers that might be assumed live by the proof. This would allow the proof to assume live environment obligations that have layer $\geq k$, the eventual fulfilment of which might depend on the eventual fulfilment of O . But since O is in the frame, it is constantly held and not fulfilled for the whole duration of the execution of the command we are proving. The frame condition forces us to record the “minimum” layer of the frame in the context, ruling out unsound circular reasoning.

After framing, we use the rule of consequence to massage the assertions to prepare them to the form required for the later application of **LIVEC**.

The rest of the derivation does not involve liveness reasoning, and follows a standard TaDA proof pattern. We use **A}\exists\text{ELIM}** and **ATOMW}** to turn the Hoare triple into an atomic triple:

$$\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X', z \in Z. \langle P_h \mid P_a(x, z) \rangle \mathbb{C} \langle Q_h \mid Q_a(x, z) \rangle}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X'. \langle P_h \mid \exists z \in Z. P_a(x, z) \rangle \mathbb{C} \langle Q_h \mid \exists z \in Z. Q_a(x, z) \rangle} \text{A}\exists\text{ELIM} \\
 \frac{\mathcal{A} \vDash P' \text{ stable} \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \langle P \mid P' \rangle \mathbb{C} \langle Q \mid Q' \rangle \quad \mathcal{A} \vDash Q' \text{ stable}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P * P'\} \mathbb{C} \{Q * Q'\}} \text{ATOMW}$$

Rule **A \exists ELIM** says that if one can prove the command is resilient to interference on z and does not affect the resource on z until its atomic update, then we can relax the specification to state that the command allows changes to z and might also affect z during the interference phase. Rule **ATOMW** says that if you prove a command is atomic, you can relax the specification not to insist on atomicity; this can be done provided the atomic pre- and postcondition are stable, as required for the Hoare triple to be well-defined.

The combination of **A \exists ELIM** and **ATOMW** simply states that if we can prove a command performs an update atomically, and the pre- and postconditions are stable, then we can prove that the command also performs the update non-atomically.

Now let us consider the derivation for **STEP 2**, which lifts the specification of CLH lock to the context of the client:

$$\frac{\frac{\frac{\frac{1 \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle}{1 \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}, d \in \text{Bool}. \langle L(x, l) \rangle \text{lock}(x) \langle L(x, 1) \wedge l = 0 \rangle} \text{SUBPQA}}{1 \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}, d \in \text{Bool}. \langle \mathcal{I}(dc_r(x, \text{done}, l, d)) \rangle \text{lock}(x) \langle \mathcal{I}(dc_r(x, \text{done}, 1, d)) * [\kappa]_r^\dagger \rangle} \text{FRAME}}{1 \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}, d \in \text{Bool}. \langle \text{emp} \mid dc_r(x, \text{done}, l, d) \rangle \text{lock}(x) \langle [\kappa]_r^\dagger \mid dc_r(x, \text{done}, 1, d) \rangle} \text{LIFTA}}{1 \vdash \mathbb{W}l \in \{0, 1\} \rightarrow \{0\}. \langle \text{emp} \mid \exists d. dc_r(x, \text{done}, l, d) \rangle \text{lock}(x) \langle [\kappa]_r^\dagger \mid \exists d. dc_r(x, \text{done}, 1, d) \rangle} \text{A}\exists\text{ELIM}}$$

Rule **SUBPQA** simply gives a way to manipulate the pseudo-quantified variable and its domain:

$$\frac{\begin{array}{l} f: X \rightarrow Y \quad Y' = f(X') \\ \forall x \in X. \vdash_{\mathcal{A}} P'(x) \Leftrightarrow P(f(x)) \quad \forall x \in X. \vdash_{\mathcal{A}} Q(f(x)) \Leftrightarrow Q'(x) \\ m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}y \in Y \rightarrow_k Y'. \langle P(y) \rangle \subset \langle Q(y) \rangle \end{array}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle P'(x) \rangle \subset \langle Q'(x) \rangle} \text{SUBPQA}$$

These are manipulations that would normally be carried out using consequence, but need to be done specially since the pseudo-quantifier is a component of triples and not of assertions. In our example we simply use it to remove the unused variable d , choosing $f: \{0, 1\} \times \text{Bool} \rightarrow \{0, 1\}$ to be the first projection.

The interesting step of the derivation of **STEP 2** is the application of rule **LIFTA**:

$$\frac{\begin{array}{l} r \in \text{dom}(\mathcal{A}) \Rightarrow R = \text{id} \quad \mathcal{A} \vDash P(x), Q(x, z) \text{ } \lambda\text{-safe} \\ \mathcal{A} \vDash P(x), Q(x, z) \text{ } \lambda+1\text{-obl. free} \quad \{((x, O_1), (z, O_2)) \mid x \in X \wedge R(x, z) \} \subset \mathcal{T}_r(G) \\ m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle \mathcal{I}(t_r^\dagger(x)) * [G]_r * P(x) * [O_1]_r^\dagger \rangle \subset \langle \exists z. \mathcal{I}(t_r^\dagger(z)) * Q(x, z) \wedge R(x, z) * [O_2]_r^\dagger \rangle \end{array}}{m; \lambda+1; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle [O_1]_r^\dagger \mid t_r^\dagger(x) * [G]_r * P(x) \rangle \subset \langle [O_2]_r^\dagger \mid \exists z. t_r^\dagger(z) * Q(x, z) \rangle} \text{LIFTA}$$

Let us unpack it. The purpose of the rule is to take an atomic specification that applies to some resource, and lift it to the effect the atomic update has on some region that contains that resource in its interpretation. In our example, it says: you have proven that the command locks the lock; the lock is part of the interpretation of $dc_r(s, x, \text{done}, l, d)$ and the update to the lock amounts to going from the interpretation with $l = 0$ to the interpretation with $l = 1$. The rule needs to make sure that the region update is among the ones permitted by the associated protocol. It does so by checking:

- (1) that there is a transition in the protocol matching the update;
- (2) that such transition is guarded by a guard that is owned;
- (3) that the local obligations are updated as the protocol mandates.

To check the first condition, the rule uses a relation R between abstract states of the region; by the fourth premise, R can only include updates that the owner of G is allowed to perform. The second condition is enforced by requiring the precondition to own G . The third condition is ensured by going from owning O_1 to owning O_2 , which, according to the fourth premise, is the expected update of obligations. In our example we have $O_1 = \mathbf{0}$ and $O_2 = \kappa$ and the update matches transition (13). The third premise uses the abbreviation ‘ $\mathcal{A} \vDash P \text{ } \lambda\text{-obl. free}$ ’ which stands for $\vdash_{\mathcal{A}} P \Rightarrow \text{emp}_{\text{Ob}}^\lambda$. This

implies that P and Q cannot own obligations of r , and so O_1 and O_2 capture the whole of the updated obligations. Note that because of the well-formedness restrictions on triples, in the conclusion of the rule the obligations are transferred to the Hoare pre/post-conditions: there they belong to a closed region. The first premise says: if the region we are updating is tracked by the atomicity context, this needs to be a trivial update, or else it would count as a linearisation point (which is instead handled using rule **UPDREG**). In our example $\mathcal{A} = \emptyset$ as we are not proving atomicity of the client, so we are allowed any protocol compliant update.

Finally, the second premise $\mathcal{A} \vDash Q(x, z)$ λ -safe requires Q to preserve its meaning at level $\lambda+1$. The formal definition of λ -safety is given in Appendix B.2.1; all the λ -safety conditions in our proofs can be immediately discharged by applications of the following lemma.

LEMMA 4.2. *The properties below hold, for arbitrary $\lambda \in \text{Lvl}$:*

- (1) $\text{emp}, \mathbb{E}_1 \mapsto \mathbb{E}_2$ and \mathbb{B} are λ -safe.
- (2) $\lceil G \rceil_r$ and $\lfloor O \rfloor_r^L$ are both λ -safe.
- (3) If $\lambda' < \lambda$, then $\mathbf{t}_r^{\lambda'}(a) * \lfloor O \rfloor_r^E$ is λ -safe.
- (4) If P, Q are both λ -safe, then so are $P \wedge Q, P \vee Q$, and $P * Q$.
- (5) If $P(v)$ is λ -safe for all $v \in \text{AVal}$, then $\exists x. P(x)$ is λ -safe.

4.4 The **LIVEC** rule

In a TaDA safety proof, the derivations of **STEP 1** and **STEP 2** could be plugged together: the safety specification of the lock operation does not contain the $\{0, 1\} \rightarrow \{0\}$ component, and the premise of **STEP 1** matches exactly the conclusion of **STEP 2** (modulo framing L , which would anyway not be used in a safety proof). In TaDA Live, the discrepancy between the two steps expresses the need for a termination argument for this call. What needs to be proven is the fact that, in the current context of the **dc** protocol, during the interference phase of this call to $\text{lock}(x)$, the environment will always eventually unlock the lock. The **LIVEC** rule allows to remove the liveness condition of the specification, in a context where the corresponding liveness invariant can be proven to hold:

$$\frac{n; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T \quad m \geq n \quad k \geq n \quad \forall x \in X. \vdash_{\lambda; \mathcal{A}} P_a(x) * T \Rightarrow x \in X' \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x) \rangle}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X. \langle P_h * L \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * L \mid Q_a(x) \rangle} \text{LIVEC}$$

The first premise $n; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T$ is called the *environment liveness condition*, and it roughly corresponds to checking $\square L \Rightarrow \square \diamond T$ (with M acting as a certificate of the property holding, explained later). Here, we pick:

$$L \triangleq \exists l \in \{0, 1\}, d. \text{dc}_r(x, \text{done}, l, d) * (l = 1 \Rightarrow \lfloor \mathbf{k} \rfloor_r^E) \quad (16)$$

$$T \triangleq \exists d. \text{dc}_r(x, \text{done}, 0, d) \quad (17)$$

and we can conclude $\square L \Rightarrow \square \diamond T$ because when T does not hold, i.e. $l = 1$, we know, from L , that $\lfloor \mathbf{k} \rfloor_r^E$ holds; if we can show \mathbf{k} is live, the protocol says that if the environment holds \mathbf{k} , it will eventually fulfil it; under the protocol the transition fulfilling it is setting $l = 0$ which brings us to T . The environment can always set $l = 1$ again after that, but the same argument then applies.

To show \mathbf{k} is live we have to look at the layers. Here we have $m = 1$ and $k = 0 = \text{lay}(\mathbf{k})$. Recall that $k \geq n$ holds if $\forall k' > k. k' \geq n$. We can therefore set $n = 1$: $0 \geq 1$ holds since $\forall k' > 0. k' \geq 1$. Since we don't own any obligation locally (\mathbf{d} has been framed, recording this fact in the context layer m) we can consider \mathbf{k} live when proving the environment liveness condition.

The environment liveness condition is the central component of both **LIVEC** and **WHILE**; we explain it in depth now, and then resume our proof of the distinguishing client.

$$\begin{array}{c}
\frac{n; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T \quad m \geq n \quad k \geq n \quad \forall x \in X. \vdash_{\lambda; \mathcal{A}} P_a(x) * T \Rightarrow x \in X' \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x) \rangle}{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X. \langle P_h * L \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * L \mid Q_a(x) \rangle} \text{LIVEC} \\
\\
\frac{\forall \beta \leq \beta_0. m(\beta); \lambda; \mathcal{A} \vdash L \xrightarrow{M} T(\beta) \quad \forall \beta \leq \beta_0. \vdash_{\mathcal{A}} P(\beta) \triangleright m(\beta) \leq m \quad \forall \alpha. \mathcal{A} \models \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \quad \text{pv}(T, L, M) \cap \text{mod}(\mathbb{C}) = \emptyset \quad \forall \beta \leq \beta_0. \forall b \in \text{Bool}. m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta) * (b \Rightarrow T(\beta)) \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma \leq \beta * (b \Rightarrow \gamma < \beta)\}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta_0) * L\} \text{while}(\mathbb{B})\{\mathbb{C}\} \{\exists \gamma. P(\gamma) * L \wedge \neg \mathbb{B} \wedge \gamma \leq \beta_0\}} \text{WHILE} \\
\\
\frac{\frac{m_1; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad \vdash_{\mathcal{A}} Q_1 \triangleright m_2 \leq m \quad m_1 \leq m_2}{m_2; \lambda; \mathcal{A} \vdash_{\Phi} \{P_2\} \mathbb{C}_2 \{Q_2\} \quad \vdash_{\mathcal{A}} Q_2 \triangleright m_1 \leq m} \text{PAR} \quad \frac{m_1 \leq m_2}{m_1; \lambda; \mathcal{A} \vdash_{\Phi} \{P\} \mathbb{C} \{Q\}} \text{LAYWH}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}} \text{LAWH} \\
\\
\frac{\forall x \in X. \vdash_{\mathcal{A}} R_h * R_a(x) \triangleright m \quad \text{pv}(R_h, R_a(x)) \cap \text{mod}(\mathbb{C}) = \emptyset \quad \mathcal{A} \models R_h \text{ stable} \quad \forall x \in X. \mathcal{A} \models R_a(x) \text{ stable}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle} \text{FRAME} \\
\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle P_h * R_h \mid P_a(x) * R_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * R_h \mid Q_a(x, y) * R_a(x) \rangle}{\mathcal{A} \models P' \text{ stable} \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \langle P \mid P' \rangle \mathbb{C} \langle Q \mid Q' \rangle \quad \mathcal{A} \models Q' \text{ stable}} \text{ATOMW} \\
\frac{\lambda < \lambda' \quad r \notin \text{dom}(\mathcal{A}) \quad \mathcal{A}' = \mathcal{A}[r \mapsto (X, k, X', T)] \quad T \subseteq \mathcal{T}_r(G) \quad R = \text{io}(T) \quad \forall x \in X. \mathcal{A} \models t_r^\lambda(x) * [G]_r \text{ stable}}{m; \lambda'; \mathcal{A}' \vdash_{\Phi} \{\exists x \in X. t_r^\lambda(x) * r \Rightarrow \blacklozenge\} \mathbb{C} \{\exists x, y. R(x, y) \wedge r \Rightarrow (x, y)\}} \text{MKATOM} \\
\frac{m; \lambda'; \mathcal{A}' \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle t_r^\lambda(x) * [G]_r \rangle \mathbb{C} \langle \exists y. t_r^\lambda(y) * [G]_r \wedge R(x, y) \rangle}{r \in \text{dom}(\mathcal{A}) \quad \mathcal{A}' = \mathcal{A}[r \mapsto \perp] \quad \mathcal{A} \models P(x), Q_1(x, y), Q_2(x, y) \text{ } \lambda\text{-safe} \quad \mathcal{A} \models P(x), Q_1(x, y), Q_2(x, y) \text{ } \lambda+1\text{-obl. free} \quad \{(x, O_1), (z, O_2(x)) \mid x \in X\} \subseteq \text{tr}(\mathcal{A}, r)} \text{UPDREG} \\
\frac{m; \lambda; \mathcal{A}' \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle [O_1]_r^\perp \mid t_r^\lambda(x) * P(x) * r \Rightarrow \blacklozenge \rangle \mathbb{C} \langle \exists z. [O_2(x)]_r^\perp * \left(\begin{array}{l} R(x, z) \wedge Q_1(x, z) \\ \vee x = z \wedge Q_2(x) \end{array} \right) \rangle}{m; \lambda+1; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle [O_1]_r^\perp \mid t_r^\lambda(x) * P(x) * r \Rightarrow \blacklozenge \rangle \mathbb{C} \langle [O_2(x)]_r^\perp \mid \exists z. t_r^\lambda(z) * \left(\begin{array}{l} Q_1(x, z) * r \Rightarrow (x, z) \\ \vee Q_2(x) * r \Rightarrow \blacklozenge \end{array} \right) \rangle} \\
\\
\frac{r \in \text{dom}(\mathcal{A}) \Rightarrow R = \text{id} \quad \mathcal{A} \models P(x), Q(x, z) \text{ } \lambda\text{-safe} \quad \mathcal{A} \models P(x), Q(x, z) \text{ } \lambda+1\text{-obl. free} \quad \{(x, O_1), (z, O_2)\} \mid x \in X \wedge R(x, z)\} \subseteq \mathcal{T}_r(G)}{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle [G]_r * P(x) * [O_1]_r^\perp \rangle \mathbb{C} \langle \exists z. [G]_r * Q(x, z) \wedge R(x, z) * [O_2]_r^\perp \rangle} \text{LIFTA} \\
\\
\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X', z \in Z. \langle P_h \mid P_a(x, z) \rangle \mathbb{C} \langle Q_h \mid Q_a(x, z) \rangle}{m; \lambda; \mathcal{A} \vdash_{\Phi} \mathbb{W}x \in X \rightarrow_k X'. \langle P_h \mid \exists z \in Z. P_a(x, z) \rangle \mathbb{C} \langle Q_h \mid \exists z \in Z. Q_a(x, z) \rangle} \text{A}\exists\text{ELIM}
\end{array}$$

Fig. 9. Key TaDA Live rules. Abbreviations:

$\vdash_{\mathcal{A}} P \triangleright k$ means $\forall r \in \text{Rld}. \vdash_{\mathcal{A}} P \Rightarrow r \triangleright k$;

$\mathcal{A} \models P$ λ -safe can be discharged using Lemma 4.2;

$\mathcal{A} \models P$ λ -obl. free means $\vdash_{\mathcal{A}} P \Rightarrow \text{emp}_{\text{Ob}}^\lambda$;

$k \geq n$ means $(\forall k' > k. k' \geq n)$.

$$\begin{array}{c}
\frac{\lambda; \mathcal{A} \vDash L \text{ stable} \quad \vdash_{\lambda; \mathcal{A}} L \Rightarrow L * \exists \alpha. M(\alpha)}{m; \lambda; \mathcal{A} \vdash L * M(\alpha) : L * M(\alpha) \rightarrow T} \text{ENVLIVE} \\
\frac{}{m; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T} \\
\frac{m; \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \rightarrow T \quad m; \lambda; \mathcal{A} \vdash L(\alpha) : L_2(\alpha) \rightarrow T}{m; \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \vee L_2(\alpha) \rightarrow T} \text{ECASE} \quad \frac{\forall x \in X. m; \lambda; \mathcal{A} \vdash L(\alpha) : L(x, \alpha) \rightarrow T}{m; \lambda; \mathcal{A} \vdash L(\alpha) : \exists x \in X. L(x, \alpha) \rightarrow T} \text{EQUANT} \\
\frac{\forall \alpha. \vdash_{\mathcal{A}} T'(\alpha) \Rightarrow T}{m; \lambda; \mathcal{A} \vdash L(\alpha) : T'(\alpha) \rightarrow T} \text{LIVET} \\
\frac{\text{impr}_{\mathcal{A}}(L', L, T) \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright \text{lay}(O(x)) \quad \lambda < \lambda' \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x. \mathfrak{t}_r^\lambda(x) * \lfloor O(x) \rfloor_r^E * \text{True} \wedge m > \text{lay}(O(x))}{m; \lambda'; \mathcal{A} \vdash L(\alpha) : L'(\alpha) \rightarrow T} \text{LIVEO} \\
\frac{\text{impr}_{\mathcal{A}}(L', L, T) \quad m > k \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright k \quad (X \rightarrow_k X') = \text{live}(\mathcal{A}, r) \quad \lambda < \lambda' \quad \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x \in X \setminus X'. \mathfrak{t}_r^\lambda(x) * r \Vdash \diamond * \text{True}}{m; \lambda'; \mathcal{A} \vdash L(\alpha) : L'(\alpha) \rightarrow T} \text{LIVEA}
\end{array}$$

Fig. 10. Environment Liveness Condition Rules

4.5 The Environment Liveness Condition

The essence of the termination argument is captured in **LIVEC** and **WHILE** by the conditions of the form $m; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T$. They establish “always eventually T holds” facts. The condition is parametrised by L , an assertion that holds at any point in the traces we are considering, an assertion T , characterising the so-called *target* states, and an assertion $M(\alpha)$ parametric on some ordinal α , which represents the environment progress measure. Intuitively, the condition states that, from any state satisfying $L * M(\alpha)$, for some α , we can find an environment transition that *must eventually* happen that would take us either to T , or to some state satisfying $L * M(\alpha')$ with $\alpha' < \alpha$. Additionally, any transition from L to L that *may* happen does not strictly increase the progress measure, unless they end in a target state. The transitions that must happen are characterised by being those that either: (1) fulfil some obligation known to be in the environment and with layer lower than the ones we may hold locally, or (2) fulfil some environment liveness assumption stored in \mathcal{A} with layer lower than the ones we may hold locally. This entails that, under an environment that always eventually fulfils the obligations we are assuming live, $\square L \Rightarrow \square \diamond T$ holds, as desired.

In the **WHILE** rule, an environment liveness condition is combined with the condition

$$\forall \alpha. \mathcal{A} \vDash \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable}$$

which requires us to prove that any protocol compliant step from a state satisfying $L * M(\alpha_0)$ for some α_0 will take us to a state satisfying $L * M(\alpha_1)$ for some $\alpha_1 \leq \alpha_0$. In other words, in traces satisfying $\square L$ the progress measure never increases. This, in conjunction with $m; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T$, entails $\square L \Rightarrow \square \diamond T$, as needed for soundness of rule **WHILE**.

Take the environment liveness condition required by the application of **LIVEC** in the proof of the distinguishing client. Given $m = 1$ and $M(\alpha) = (\alpha = 0)$ we have to prove:

$$m; \lambda; \mathcal{A} \vdash \exists l \in \{0, 1\}. \text{dc}_r(x, \text{done}, l, _) * (l = 1 \Rightarrow \lfloor \mathfrak{k} \rfloor_r^E) \xrightarrow{M} \text{dc}_r(x, \text{done}, 0, _)$$




Fig. 11. Illustration of rule **ENVLIVE** and the $\text{impr}_{\mathcal{A}}$ condition.

That is, during the interference phase, we know that at any point in time the lock will be in some state $l \in \{0, 1\}$; we want to prove that the environment will always eventually set l to 0. Here this is particularly easy to show: L states that when $l = 1$ the obligation κ is held by the environment; since $\text{lay}(\kappa) = 0 < 1 = m$ (and L does not hold obligations) we can assume the obligation will be eventually fulfilled; the only transition that can fulfil it, is the one that sets $l = 0$, so in exactly one such step we reach T . This justifies the trivial definition of M : we do not need to keep track of progress towards T as we reach it in exactly one of the steps that *must eventually* happen.

The environment liveness condition can be proven using the rules in Fig. 10. The only rule that applies directly is **ENVLIVE**, which checks that in a state satisfying L one can always measure progress (second premise), and then asks to discharge an auxiliary judgement of the form $m; \lambda; \mathcal{A} \vdash L(\alpha) : L(\alpha) \xrightarrow{M} T$ which is best explained with the help of the illustration in Fig. 11. The condition works under the hypothesis that the assertion L holds at any point of the traces under consideration, so in the picture we are considering infinite sequences of steps within the outer rectangle. The target states T describe some subset of L , which we want to show is visited infinitely often¹⁵ by any infinite trace that complies with the liveness rely as specified by the region protocols and pseudo-quantifiers. Rule **ECASE** allows the splitting of the invariant L into a disjunction of, say, L_1, L_2, L_3 and T as in the picture. We need to prove there is going to be eventual progress towards reaching T from each of these cases. If we start from T we already reached the target, and this case can be discharged by rule **LIVET**. The other cases are covered by Rule **LIVEO** which justifies progress by appealing to an environment-owned atomic obligation O which is live (premises two and three); and Rule **LIVEA** which justifies progress by appealing to a liveness assumption stored in the atomicity context. The **EQUANT** rule is a generalisation of rule **ECASE**.

To see how progress is justified, consider the trace of Fig. 11 starting from w_0 . Assume the progress measure at w_0 is α_1 (i.e. $L * M(\alpha_1)$ holds in w_0). Each case L_i can be discharged with either rule **LIVEO** or rule **LIVEA**. Imagine L_1 is discharged using **LIVEO**: the rule requires us to find an obligation A which, in every state of L_1 , is necessarily owned by the environment ($\llbracket A \rrbracket_r^E$) for some region $t_r(_)$. Then the $\text{impr}_{\mathcal{A}}$ condition checks that the progress measure will *improve* when the environment will fulfil A ; formally:

¹⁵Notice that “ T is visited infinitely often” is equivalent to $\square \diamond T$.

Definition 4.3 ($\text{impr}_{\mathcal{A}}$). Given assertions $L(\alpha)$, $L'(\alpha)$ and T , the condition $\text{impr}_{\mathcal{A}}(L', L, T)$ holds if and only if, for arbitrary $\sigma \in \text{Store}$, letting

$$l(\alpha) = \mathcal{W}[[L(\alpha)]]_{\mathcal{A}}^{\sigma} \quad l'(\alpha) = \mathcal{W}[[L'(\alpha)]]_{\mathcal{A}}^{\sigma} \quad t = \mathcal{W}[[T * \text{True}]]_{\mathcal{A}}^{\sigma}$$

the following holds:

$$\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(l'(\alpha_1)) \cap l(\alpha_2) \subseteq l'(\alpha_1) \cup t$$

Intuitively, the $\text{impr}_{\mathcal{A}}$ condition considers an arbitrary transition (w_1, w_2) from the current case L' to L , obeying the atomic rely (thus allowed by the safety constraints of the protocols). It then compares the progress measure α_1 and α_2 , taken before and after the transition, checking that:

- (1) the measure strictly improved ($\alpha_1 > \alpha_2$); or
- (2) the measure stalled ($\alpha_1 = \alpha_2$) but we remained within case L' , and thus the pending obligation O /liveness assumption are still pending; or
- (3) we reached T (allowing the measure to vary arbitrarily)

Examine the trace from w_0 in Fig. 11. While the trace stays within L_1 the environment obligation \mathbf{A} stays unfulfilled, (steps are labelled with the obligation they fulfil, if any) and $\text{impr}_{\mathcal{A}}$ requires the measure α_1 to decrease, or in the worst case stay constant. Since \mathbf{A} is live, the environment will eventually fulfill it, thus taking us outside of L_1 . If such transition takes us to another case, L_2 in the illustration, $\text{impr}_{\mathcal{A}}$ requires the measure to strictly decrease to some $\alpha_2 < \alpha_1$. This process cannot repeat *ad libitum*: the progress measure is an ordinal and hence well-founded. The effect is that eventually, the only option is to reach the target. Note that transitions that end in the target are allowed by $\text{impr}_{\mathcal{A}}$ to increase the progress measure: in the picture the transition reaching T increases the measure from α_2 to α_3 . This allows the “reset” of the measure so that the trace can go outside of T and the whole process of reaching T again can be repeated an unbounded number of times.

The idea behind Rule **LIVEA** is analogous to the above description, but progress is justified by appealing to an environment liveness assumption stored in \mathcal{A} . The layer of the assumption needs to be lower than any layer we may be holding. Since the environment liveness assumptions only hold in the interference phase of an update, the rule needs evidence that the linearisation point on r has not occurred yet, which is provided by $r \mapsto \diamond$.

In the proof of the distinguishing client, the environment liveness condition for the application of rule **LIVEC** between **STEP 1** and **STEP 2**, is proved by:

$$\frac{\frac{\forall \alpha. \mathbf{1} \vdash_0 L_0(\alpha) \Rightarrow T}{\mathbf{1}; \emptyset \vdash L(\alpha) : L_0(\alpha) \longrightarrow T} \text{LIVET} \quad \frac{\text{impr}_{\emptyset}(L_1, L, T)}{\mathbf{1}; \emptyset \vdash L(\alpha) : L_1(\alpha) \longrightarrow T} \text{LIVEO}}{\mathbf{1}; \emptyset \vdash L(\alpha) : L_0(\alpha) \vee L_1(\alpha) \longrightarrow T} \text{ECASE}}{\mathbf{1}; \emptyset \vdash L \xrightarrow{M} T} \text{ENVLIVE}$$

where L and T are defined in (16) and (17), and $M(\alpha) = (\alpha = 0)$. Since L trivially implies $L * \exists \alpha. M(\alpha)$, we can apply **ENVLIVE**, setting $L(\alpha) = (L \wedge \alpha = 0)$. Then we apply **ECASE** to split on the value of l : $L(\alpha) = L_0(\alpha) \vee L_1(\alpha)$ where $L_0(\alpha) = \mathbf{dc}_r(x, \text{done}, 0, _) \wedge \alpha = 0$ and $L_1(\alpha) = \mathbf{dc}_r(x, \text{done}, 1, _) * \lfloor \mathbf{k} \rfloor_r^E \wedge \alpha = 0$. If $l = 0$ we can apply **LIVET** as we are already in T ; if $l = 1$, L_1 entails $\lfloor \mathbf{k} \rfloor_r^E$ so we can apply **LIVEO** with $\mathbf{t}_r = \mathbf{dc}_r$ and $O = \mathbf{k}$. The atomic obligation \mathbf{k} is live: $\mathbf{1} > \text{lay}(\mathbf{k}) = \mathbf{0}$, and L_1 holds no obligations. To check that the $\text{impr}_{\mathcal{A}}$ condition is satisfied, we need to consider the transitions allowed by the protocol \mathbf{dc} :

- $l = 1$ to $l = 1$ keeps the measure constant but keeps us in L_1 ,
- $l = 1$ to $l = 0$ brings us directly in T .

Although in this case the progress measure is trivial and the proof of the environment liveness condition simple, the generality provided by non-trivial progress measures is needed for more interesting examples. For instance, our proofs of spin lock (Section 5.1) and CLH lock (Section 5.2) do make use of this added generality.

We chose to state the $\text{impr}_{\mathcal{A}}$ condition as a semantic check; while this achieves full generality, in typical proofs the environment liveness condition only involves a single region, and $\text{impr}_{\mathcal{A}}$ can be checked by examining the region’s protocol.

4.6 The while rule

By using the rules we described so far, one can justify most of the proof outline of the distinguishing client in Fig. 8. For instance, the proof of $\text{lock}(x)$ for the left thread can be reused as is to prove the $\text{lock}(x)$ call in the body of the loop of the right-hand thread.

The main missing step is the application of the **WHILE** rule:

$$\frac{\begin{array}{l} \forall \beta \leq \beta_0. m(\beta); \lambda; \mathcal{A} \vdash L \xrightarrow{M} T(\beta) \quad \forall \beta \leq \beta_0. \vdash_{\mathcal{A}} P(\beta) \triangleright m(\beta) \leq m \\ \forall \alpha. \mathcal{A} \vDash \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \quad \text{pv}(T, L, M) \cap \text{mod}(\mathbb{C}) = \emptyset \\ \forall \beta \leq \beta_0. \forall b \in \text{Bool}. m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta) * (b \dot{\Rightarrow} T(\beta)) \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma \leq \beta * (b \dot{\Rightarrow} \gamma < \beta)\} \end{array}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta_0) * L\} \text{while}(\mathbb{B})\{\mathbb{C}\} \{\exists \gamma. P(\gamma) * L \wedge \neg \mathbb{B} \wedge \gamma \leq \beta_0\}} \text{WHILE}$$

Let us review the main differences with the simplified **WHILEB** rule presented in Section 2. First, the two triples in the premises of **WHILEB** (corresponding to the blocked and unblocked case) are compressed here in a single triple: this is convenient in proofs as the proof of the two triples only differs on the treatment of the variant. When $b = \text{false}$, $(b \dot{\Rightarrow} T) = \text{emp} = (b \dot{\Rightarrow} \gamma < \beta)$, obtaining the first triple of **WHILEB**, when $b = \text{true}$ we obtain the other triple. Second, the target states T are parametrised over the variant β : each value of the variant may represent a different “phase” of the local progress of the while loop; in each of these phases the loop may be blocked waiting for a different set of target states to be reached. Third, as anticipated, the $\square L \Rightarrow \diamond \square T$ condition is expressed as the conjunction of the first and third premise.

There are two additional side-conditions. Since T, L and M assert facts about arbitrary intermediate states of an iteration, they cannot refer to any local variable that may be modified by the body of the loop, hence the fourth premise.

The most important addition is the layer condition of the second premise. The idea is that we should be forbidden from constantly owning obligations of layers that we might assume live in the proof of the environment liveness condition. By requiring $P(\beta) \triangleright m(\beta)$ we make sure that the loop invariant only owns obligations of layer higher than $m(\beta)$, and the $m(\beta)$ in the context of the environment liveness condition indicates that only layers lower than that may be assumed live. The layer m in the context of the triple in the conclusion is an upper bound for any layer that may be assumed live in the proof of the loop.

Consider the application of **WHILE** in the proof of the distinguishing client. The while loop of the right-hand thread is busy-waiting until done is set to true. The target states are therefore $T \triangleq \text{dc}_r(x, \text{done}, _, \text{true})$. In this example, the target states do not depend on the variant β , which itself is quite trivial: when the loop is finally unblocked, it needs at most one iteration to terminate. The local variant can simply be $\beta = (d ? 0 : 1)$, i.e. when d is false there needs to be one unblocked iteration to terminate, and when d is finally true the loop will take no more iterations. The loop invariant is

$$P(\beta) \triangleq \exists l, d. \text{dc}_r(x, \text{done}, l, d) * (l = 1 \dot{\Rightarrow} \lfloor \kappa \rfloor_r^E) \wedge d \Rightarrow d \wedge \beta = (d ? 0 : 1)$$

on which we can frame the stable assertion

$$L \triangleq T \vee \mathbf{dc}_r(x, \text{done}, _, \text{false}) * [\mathbf{D}]_r^E$$

Since the loop invariant owns no obligations, we can set $m(\beta) = \top = m$, and we need to prove the environment liveness condition $\top; \emptyset \vdash L \xrightarrow{M} T$; here, as for the application of **LIVEC**, with the fulfilment of the environment obligation \mathbf{D} we immediately reach the target, so M can be trivial ($M(\alpha) = (\alpha = 0)$). The derivation is as follows:

$$\frac{\frac{\frac{}{\top; \emptyset \vdash L(\alpha) : T \rightarrow T} \text{LIVET} \quad \frac{\text{impr}_{\mathcal{A}}(\mathbf{dc}_r(x, \text{done}, _, \text{false}) * [\mathbf{D}]_r^E, L, T)}{\top; \emptyset \vdash L(\alpha) : \mathbf{dc}_r(x, \text{done}, _, \text{false}) * [\mathbf{D}]_r^E \rightarrow T} \text{LIVEO}}{\top; \emptyset \vdash L(\alpha) : L(\alpha) \rightarrow T} \text{ECASE}}{\top; \emptyset \vdash L \xrightarrow{M} T} \text{ENVLIVE}$$

where $L(\alpha) = L \wedge \alpha = 0$. We split L into two cases using **ECASE**. In the first case T holds, so **LIVET** applies. In the second, $d = \text{false}$ and since $\text{lay}(\mathbf{D}) = \mathbf{1} < \top$, \mathbf{D} is a live obligation. The $\text{impr}_{\mathcal{A}}$ condition is satisfied: the allowed transitions either keep d constant or set it to false , taking us directly to T .

The stability of $\exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha$ holds trivially as α is constantly 0. The condition is this trivial in this case because it checks that transitions to and from T are not resetting the progress measure; here, once done is set to true, it will not be set to false any more, so once T is reached there is no way to leave it.

Non termination of distinguishing client with spin lock. If the lock at x is implemented as a spin lock, the distinguishing client may not terminate. Indeed, there is no TaDA Live proof for the distinguishing client if one assumes the spin lock specifications: in the precondition we need to specify an impedance budget α for the lock $L(x, 0, \alpha)$; whatever ordinal we may choose for α , there is no way to consume some budget at every potential iteration of the loop of \mathbb{C}_r and never exhaust the budget, as the number of iterations is effectively unbounded.

4.7 Other rules

The rules in Fig. 9 are the most important TaDA Live-specific rules. We have omitted standard rules like the axioms for primitive atomic commands, the rules handling sequencing, function calls (recall that for simplicity we restrict to non-recursive function definitions) and structural manipulations. They are reproduced in full in the Appendix.

Let us conclude with an explanation of the two TaDA Live-specific rules of Fig. 9 which are not illustrated by the proof of distinguishing client: rules **UPDREG** and **MkATOM**. While the goal of **LIFTA** is to lift an atomic update on a resource inside the interpretation of a region to the corresponding update on the region itself, **UPDREG** obtains the same effect but on a region r that is supposed to be updated once atomically (i.e. $r \in \text{dom}(\mathcal{A})$). While **LIFTA** applies to regions with $r \in \text{dom}(\mathcal{A})$, the update allowed in that case needs to be an identity step from the point of view of the abstract state of the region. A genuine update to the region needs to be recorded as the unique linearization point on that region; this is precisely the purpose of **UPDREG**. Most of the premises of **UPDREG** have the same function as in **LIFTA**: checking that the update of the abstract state *and* the obligations comply with the protocol. The difference is that here the update expected to take place in the linearization point is recorded in $\text{tr}(\mathcal{A}, r)$ (i.e. the components of $\mathcal{A}(r)$ recording the expected update to abstract state and obligations of r). To be able to claim the linearization point took place exactly once, the precondition of the triple requires the $r \Rightarrow \blacklozenge$ resource which represents the unique permission to perform the linearization point. The postcondition allows for two cases:

either the update was successful, in which case the atomicity tracking component is recording the update (x, z) ; or the update was not performed ($x = z$) and the $r \Rightarrow \blacklozenge$ resource is still available for future updates.

Rule **MkATOM** is another crucial rule for proving abstract atomicity: it states that a Hoare triple can be promoted to an atomic triple if it contains a “certificate”, in the form of $r \Rightarrow \blacklozenge$ being updated to $r \Rightarrow (x, y)$, that the region in question was updated atomically exactly once, with the expected update. The expected update, and the additional interference assumptions given by the pseudo-quantifiers need to be stored in the atomicity context so that the triple in the premise can make use of the interference precondition assumptions, and certificate the right update took place. Any expected update must be protocol compliant ($T \subseteq \mathcal{T}_t(G)$). Notice how the atomicity context records the liveness assumption expressed by the pseudo-quantifier, so that it is available for termination proofs in the proof of the triple in the premise; in particular they can be used by applications of **LIVEA**. The proof of spin lock and CLH lock in Section 5 illustrate applications of **MkATOM** and **UPDREG**.

4.8 Abstract Predicates

In the spirit of CAP, abstract resources provided by a library should be presented to clients by only exposing their abstract properties, and not their definition.

In our example, the $L(x, l)$ predicate is defined internally to the proof of the lock module, say using internal regions (of some maximum level λ) expressing the internal protocols of the module.

The proof of our distinguishing client only relies on the following abstract properties:

- (L1) $L(x, _) * L(x, _)$ is false, expressing that a lock is an exclusively owned resource;
- (L2) $L(x, l)$ is stable for all l ;
- (L3) $L(x, l)$ is λ -safe for all l ;
- (L4) $L(x, l)$ is obligation-free, i.e. $L(x, l) \Rightarrow \text{emp}_{\text{Ob}}^{\text{Rld}}$
(which also entails $L(x, l) \Rightarrow r \triangleright m$ for all $r \in \text{Rld}$ and $m \in \mathcal{L}$).

For instance, the interpretation $\mathcal{I}(\text{dc}_r(x, \text{done}, l, d))$ is well-formed thanks to properties (L2) and (L4). The proof also involves side conditions on layers, stability and λ' -safety which can be discharged by appealing to (L2) (L3) and (L4).

More generally, a module would typically expose to clients viewshifts representing separation properties of the abstract predicates (e.g. duplicability), stability properties, λ -safety, obligation freedom and relevant $P \triangleright m$ facts.

4.9 What is leaked by TaDA Live specifications?

TaDA Live’s triples are rather expressive: they support strong specifications of updates via logical atomicity, and conditional termination properties via liveness assumptions. It is natural to ask whether our triples force the leak of any unnecessary detail about the implementation. In particular, there are three components of the proof system that have a “global” flavour: the level and layer in the context of the judgment, and the layer decorating the liveness assumption of the pseudo-quantification. Although necessary for soundness, the management of levels is tedious but relatively straightforward. Iris introduced namespaces for invariants to ease the management of so-called masks, which serve essentially the same function as levels in TaDA. A similar construction could be used to ease management of layers. Here we keep it simple and require proofs of clients to use layers high enough to be able to reuse the libraries specifications.

The layers decorating a triple, on the other hand, are a more delicate matter. The main complication arises from the choice of parametrising TaDA Live with a global layer structure. If a specification insists on the use of a specific subset of layers, that could seem like an unnecessary

leak of implementation details. For example, there could be two valid implementations of a module that use wildly different internal layer structures to justify their internal blocking behaviour. Should the abstract specification of the module insist on a specific layer structure for the internal layers, that would rule out valid implementations for no good reason. In TaDA Live modularity of the layers can be achieved by exploiting a crucial property of derivations: their validity is invariant under a strict-ordering-preserving remapping of layers. This allows a style of specification which generalises the one we have seen in our examples until now, where the layer structure relevant for the proof of an implementation is parametrised over a client-provided remapping of layers. To avoid cluttering the proofs we do not explicitly parametrise the proofs in Section 5. In Section 5.5, where the construction becomes relevant and used in a non-trivial way, we explain how to convert a proof so that it is parametric on the layer remapping.

In terms of behaviour, TaDA Live’s specifications are able to hide internal blocking, as showed by the blocking counter example of Section 2.1 (formalised in Section 5.3). There is, in fact, one progress property leaked by the specification layers that is currently not exploited by TaDA Live. In the special case when the layer in the context is the *globally* smallest layer \perp , the proof of the triple cannot rely on any liveness assumption at all. This can be used to differentiate, for example, a wait-free counter implemented as an hardware-atomic fetch-and-add (which admits a proof with \perp in the context) and a blocking counter (which only admits proofs with layer $> \perp$). This is a useful distinction: wait-freedom is an important progress property, asserting termination without assumptions on liveness of other threads and *without fairness assumptions on the scheduler* [19]. Currently, however, TaDA Live’s semantics does not support deriving wait-freedom as a consequence of \perp as the context layer: the current triple semantics only implies termination of the *fair* traces. Extending TaDA Live’s semantics to encompass wait-freedom is left as future work.

4.10 Soundness

We have proven soundness of TaDA Live rules against the semantic judgement of Definition 3.30.

THEOREM 4.4 (SOUNDNESS). *If $\vdash_{\Phi} \mathbb{C} : \mathbb{S}$ then $\vDash_{\Phi} \mathbb{C} : \mathbb{S}$.*

The detailed proofs of the liveness-related rules are produced in Appendix E. The soundness of most rules is an adaptation of the soundness arguments of the corresponding TaDA rules. The rules that drive the liveness argument are rule **PAR**, rule **LIVEC** and rule **WHILE**.

The soundness of the parallel rule follows from the layered liveness invariants semantics explained in Section 3.9. The argument is roughly as follows. There are two possible ways the parallel composition $\mathbb{C}_1 \parallel \mathbb{C}_2$ may fail to terminate: either one thread terminates and the other does not, or they both do not terminate. In the first case, when the terminating thread, say \mathbb{C}_1 , terminated, we are in a state where thread 1 does not own any obligation of layers that may be assumed live by \mathbb{C}_2 (this is from the conditions on the layers of the postcondition of \mathbb{C}_1). By the triple about \mathbb{C}_2 in the premises, \mathbb{C}_2 is only allowed not to terminate if the environment is constantly owning an obligation O of layer lower than m_2 . Since \mathbb{C}_1 cannot do that, we obtain that said O must be owned by the overall environment of the parallel composition. In such case the triple of the conclusion allows the program to diverge.

In the second case, both threads are not terminating. Each of the threads, say 1, is allowed to keep an obligation constantly unfulfilled, as long as it can blame thread 2 by showing an obligation of strictly lower layer that is kept constantly unfulfilled by 2. Since layers are well-founded there needs to be some thread that will not be justified in not fulfilling some of its obligations. This cannot be as we were able to prove the two triples in the premises.

The soundness of rule **WHILE** considers the worst-case scenario for progress: an infinite sequence of iterations, all of which do not start from a target state in $T(\beta)$, and therefore do not decrease the

variant β . In such case we know that the assertion L holds at every point of the trace: it has been framed so the local steps and the environment steps must preserve it, and it is stable (as checked by `ENVLIVE`). We are thus within the hypothesis of the environment liveness condition, which proves, together with the premise asking the progress measure α to never increase, that eventually the target states will be reached. Although they may be reached in the middle of an iteration, instead of at the beginning as it would be required to invoke the triple that decreases the variant β , in the worst case this can happen boundedly many times (the progress measure is well-founded and must always decrease). Therefore we eventually reach $T(\beta)$ and not leave it until the next iteration starts from a state satisfying $P(\beta) * T(\beta) \wedge \mathbb{B}$, which matches the premise that ensures the variant decreases. This can only happen boundedly many times as the variant is well-founded.

Rule `LIVEC`’s soundness argument is a variation of the one for `WHILE`.

As a simple corollary of soundness and Theorem 3.33, if we can prove $m; \lambda; \emptyset \vdash \{\text{emp}\} \mathbb{C} \{\text{True}\}$ then `C` run in isolation terminates from the empty heap. For our distinguishing client (Example 4.1) for instance, we can wrap up the proof by initialising the state and prove

$$\top \vdash \{\text{emp}\} \text{ var done} = \text{false}, x \text{ in } x := \text{makeLock}(); (\mathbb{C}_\ell \parallel \mathbb{C}_r) \{\text{True}\}$$

which implies termination of the program.

5 EVALUATION

In the previous section, we introduced the TaDA Live proof system, explaining the rules on the distinguishing client, which showcases in a simple setting the proof mechanics of the logic.

In this section we consider more challenging case studies, to demonstrate how TaDA Live achieves proof scalability and reuse in practice.

We start by proving correctness of the spin and CLH lock implementations against the specifications we discussed in Section 2. The proof of spin lock highlights the use of the liveness assumption of a pseudo-quantifier in a proof, and the handling of impedance through the impedance budget. The proof of CLH has a number of interesting features. The CLH code exhibits both *internal blocking*, i.e. blocking that is resolved internally and does not leak to the client, and *external blocking*, i.e. blocking that has to be resolved by the client and thus leaks in the liveness assumption of the pseudo-quantifier. As a consequence, the termination argument requires using a combination of obligations (for internal blocking) and the liveness assumption of the pseudo-quantifier (for external blocking). Moreover, the obligations (and their layers) are not simple tokens like the ones for the simple examples of Sections 2 and 4, but form an infinite set. This reflects the unboundedness of the internal queue of threads.

The two lock examples demonstrate TaDA Live’s ability to abstract from implementation details, and only leak to the client the parts of the termination argument which depend on the choices of the clients. In the same vein, we will follow this with a counter module using a spin lock to protect access to a cell holding the value of the counter. Interestingly, since the blocking due to the use of a lock is internal, the specification of the counter will not be blocking. The impedance suffered by the internal spin lock does however leak to the interface for the counter: the counter will have its own impedance budget which will be internally spent to call operations on the lock.

To exhibit TaDA Live’s ability to reason about liveness locally, we will verify a double blocking counter, showing that for simple common programming patterns, the layer system leads to natural and modular client proofs.

Finally, we comment on a proof of a lock-coupling set, produced in full in Appendix C. The example considers a data structure implemented as a linked list with CLH locks guarding the single cells. The example is challenging for the presence of a dynamic number of locks. At first sight it

```

def makeLock() {
  ret := alloc(1);
  [ret] := 0;
}

def lock(x) {
  var d=0 in
  while(d=0){
    d := CAS(x,0,1);
  }
}

def unlock(x) {
  [x] := 0;
}

```

Fig. 12. Code of spin lock operations.

might seem it is impossible to represent this using the static association of layers to obligations of TaDA Live.

Obligations however, as demonstrated in this case study, are a very general form of ghost state and can easily represent dynamic properties of state.

Other case studies. Ticket lock and MCS lock [18] are alternative implementations of starvation-free locks; they can be given the same specification as the CLH lock, and their liveness argument can be carried out in the same way as the one we present for CLH.¹⁶ A paradigmatic example of fine-grained data structure is the Treiber stack [18] which, in its standard form, is non-blocking and has been proven in Total TaDA already. It is easy to adapt the code to have a pop operation which blocks on an empty stack. Such operation would be blocking and suffers impedance. Its specification and proof mirrors closely the proof of the spin lock. Challenging variants of the lock-coupling set are the “optimistic” and “lazy” sets. The proof of optimistic set uses a combination of the proof of the lock-coupling set and the impedance budget technique (optimistic set operations impede each other).

These case studies cover all the proof patterns needed to prove all the examples of the LiLi papers [30, 31]. Notably, proofs in LiLi involving modules that use locks, require in-lining some non-atomic implementation of the lock operations in the client, resulting in non-modular proofs and unnecessarily intertwined termination arguments.

5.1 Spin lock

Code. The spin lock module implements a lock by storing a single bit in a heap cell; locking is implemented by trying to CAS the heap cell from 0 to 1 until the CAS succeeded; unlocking simply sets the cell back to 0. In Fig. 12 we give all the operations of a spin lock module.

Specifications. We will prove the module satisfies the following specifications:

$$\begin{aligned}
\forall \alpha. \mathbf{0} \vdash \{\text{emp}\} \text{makeLock}() \{ \exists r. L(r, \text{ret}, 0, \alpha) \} \\
\forall \phi. \mathbf{1} \vdash \forall l \in \{0, 1\} \rightarrow_0 \{0\}, \alpha. \langle L(r, x, l, \alpha) \wedge \alpha > \phi(\alpha) \rangle \text{lock}(x) \langle L(r, x, 1, \phi(\alpha)) \wedge l = 0 \rangle \\
\mathbf{0} \vdash \langle L(r, x, 1, \alpha) \rangle \text{unlock}(x) \langle L(r, x, 0, \alpha) \rangle
\end{aligned}$$

where $L(r, x, l, \alpha)$ abstractly represents the lock resource at abstract location r (omitted for readability in Section 2) and concrete address x , with abstract state $l \in \{0, 1\}$ and impedance budget α (an ordinal). The purpose of the impedance budget, as described in Section 2, is to prevent the environment from taking possession of the lock an unbounded number of times. Without this bound, the CAS operation in the implementation of `lock` could be indefinitely preempted by the environment locking the lock, preventing it from ever taking its possession and terminating, even

¹⁶The proof of ticket lock requires some minor ghost code to side-step the lack of support for helping.

if the environment always unlocks the lock when it is locked. This is enforced by requiring the lock operation to strictly decrease the impedance budget using $\phi: \mathbb{O} \rightarrow \mathbb{O}$, a function that can be freely instantiated by the client upon usage of the specification, which indicates precisely how much the budget will decrease after this call (which is client dependent information). The specification of `makeLock` then allows the client to pick an arbitrary ordinal as the initial budget.

Shared Region. The abstract shared lock resource will be represented by a region $\mathbf{spin}_r(x, l, \alpha)$ where $x \in \text{Addr}$, $l \in \{0, 1\}$, $\alpha \in \mathbb{O}$. Here x is a fixed parameter of the region.

Convention 1. An exclusive guard, \mathbf{E} , is very commonly used to express some exclusive permission on some shared resource, which cannot be composed with itself: i.e. $\mathbf{E} \bullet \mathbf{E} = \perp$. Local ownership of \mathbf{E} is exclusive in that no other thread can at the same time assert ownership of \mathbf{E} . A ubiquitous use of this guard is in representing the resource offered by a module.

Take for example the current spin lock module. Since this is a concurrent module it uses internally shared resources. We therefore have a region $\mathbf{spin}_r(x, l, \alpha)$ encapsulating the shared internal resources of the counter. From the perspective of the client, however, at the moment of creation of a lock by, say, a `makeLock()` operation, the lock is exclusively owned by the client. This, for example, is reflected in the fact that, until the client shares the lock or invokes operations on it, it remains unlocked. To represent this fact, one typically defines an exclusive guard \mathbf{E} guarding each transition of the region interference: e.g. $\mathbf{E} : (0, O_1) \rightsquigarrow (1, O_2)$, $\mathbf{E} : (1, O_1) \rightsquigarrow (0, O_2)$. Then the `makeLock()` operation can be given the specification above, which gives to the client the stable assertion $\mathbf{spin}_r(\text{ret}, 0, \alpha) * [\mathbf{E}]_r$, wrapped in the predicate $L(r, \text{ret}, 0, \alpha)$. (Note how $\mathbf{spin}_r(\text{ret}, 0, \alpha)$ is not stable on its own.) To re-share the lock, the client will create its own region encoding the invariants governing the interaction over the lock (and the other resources of the client) the interpretation of which will contain the guard $[\mathbf{E}]_r$.

Note that assertions have very different meanings if occurring in the *atomic* precondition of a triple, as opposed to the *Hoare* precondition: the resources in the atomic precondition are not owned by the local thread, but only acquired instantaneously at the linearization point. For example, in the triple

$$\forall \phi. \vdash \forall l \in \{0, 1\} \rightarrow_0 \{0\}, \alpha. \langle \mathbf{spin}_r(x, l, \alpha) * [\mathbf{E}]_r \wedge \alpha > \phi(\alpha) \rangle \text{lock}(x) \langle \mathbf{spin}_r(x, 1, \phi(\alpha)) * [\mathbf{E}]_r \wedge l=0 \rangle$$

the exclusivity of \mathbf{E} is only granted *instantaneously* to the thread acting on it atomically, i.e. either the environment during the interference phase as allowed by the pseudo-quantifier, or the local thread at the linearization point.

Since this pattern is ubiquitous, we reserve the \mathbf{E} guard constructor for this use, and will omit the $\mathbf{E} \bullet \mathbf{E} = \perp$ axiom when specifying guard algebras.

Guards and Obligations. For the \mathbf{spin} region we only have the exclusive guard \mathbf{E} , and no obligation constructors, as the implementation has no internal blocking. All the blocking behaviour is represented by the liveness assumption in the pseudo-quantifier of the specification of lock. Note that without the exclusive guard, the specification of `makeLock` would not hold as the lock would not be stably unlocked.

Region protocol. The interference protocol for \mathbf{spin} is very simple:

$$\begin{aligned} \mathbf{E} : ((0, \alpha), \mathbf{0}) &\rightsquigarrow ((1, \beta), \mathbf{0}) \text{ only if } \beta < \alpha \\ \mathbf{E} : ((1, \alpha), \mathbf{0}) &\rightsquigarrow ((0, \alpha), \mathbf{0}) \end{aligned}$$

It states that whoever owns \mathbf{E} can freely acquire or release the lock, provided that at each acquisition, some budget is spent ($\beta < \alpha$), preventing the lock from being locked an unbounded number of times.

$$\begin{array}{c}
\forall \phi. \mathbf{1}; \mathbf{0} \vdash \mathbf{W}l \in \{0, 1\} \rightarrow_0 \{0\}, \alpha. \\
\langle L(r, x, l, \alpha) \wedge \alpha > \phi(\alpha) \rangle \\
\left(\text{spin}_r(x, l, \alpha) * [\mathbf{E}]_r \wedge \alpha > \phi(\alpha) \right) \\
\mathbf{1}; [r \mapsto (\{0, 1\} \times \mathbf{0}, \mathbf{0}, \{0\} \times \mathbf{0}, ((0, \alpha), \mathbf{0}) \rightsquigarrow ((1, \phi(\alpha)), \mathbf{0}))] \vdash \\
\{ \exists l, \alpha. \text{spin}_r(x, l, \alpha) \wedge \alpha > \phi(\alpha) * r \Rightarrow \blacklozenge \} \\
\mathbf{var} \ d = \mathbf{0} \ \mathbf{in} \\
\{ \exists l, \alpha. \text{spin}_r(x, l, \alpha) * (d = 0 \wedge r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha)) \} \\
\{ P(\beta_0) \} \\
\text{while}(d=0)\{ \\
\quad \forall b \in \text{Bool}, \beta. \\
\quad \{ P(\beta) * b \Rightarrow T(\beta) \wedge d = 0 \} \\
\quad d := \mathbf{CAS}(x, \mathbf{0}, \mathbf{1}); \\
\quad \{ \exists \gamma. P(\gamma) \wedge \beta \geq \gamma \wedge b \Rightarrow \gamma < \beta \} \\
\} \\
\{ \exists \gamma. P(\gamma) \wedge \beta_0 \geq \gamma \wedge d \neq 0 \} \\
\{ \exists l, \alpha. r \mapsto ((l, \alpha), (1, \phi(\alpha))) \wedge l = 0 \} \\
\langle \text{spin}_r(x, 1, \phi(\alpha)) * [\mathbf{E}]_r \wedge l = 0 \rangle \\
\langle L(r, x, 1, \phi(\alpha)) \wedge l = 0 \rangle
\end{array}$$

Fig. 13. Spin lock: proof of lock.

Region interpretation. The implementation uses a single cell stored in the heap, and we have no non-trivial guards/obligations; the interpretation is thus straightforward:

$$\mathcal{I}(\text{spin}_r(x, l, \alpha)) \triangleq x \mapsto l$$

Note how α is pure ghost state in that it is not linked to any information in the concrete memory.

Predicates. The lock resource is abstractly represented by the predicate

$$L(r, x, l, \alpha) \triangleq \text{spin}_r(x, l, \alpha) * [\mathbf{E}]_r$$

Verification of lock. Figure 13 is the proof of the lock operation. The only step that involves reasoning about liveness is the application of the **WHILE** rule. To apply this rule, we must first define the loop invariant, $P(\beta)$, the target states, $T(\beta)$, the persistent loop invariant, L , $m(\beta)$, and the environmental progress measure, $M(\alpha)$.

The loop invariant is

$$P(\beta) \triangleq \exists l, \alpha. \text{spin}_r(x, l, \alpha) \wedge \beta \geq \alpha * \left(\begin{array}{l} (d = 0 \wedge r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha)) \\ \vee (\exists l', \alpha'. d = 1 \wedge r \Rightarrow ((l', \alpha'), (1, \phi(\alpha')))) \wedge l' = 0 \end{array} \right)$$

which contains:

- the safety information to prove the uniqueness of the linearization point, namely that if the **CAS** failed, i.e. $d = 0$, then we have not touched the resource yet and we still have permission to perform the linearization point ($r \Rightarrow \blacklozenge$); whereas if the **CAS** succeeded, i.e. $d = 1$, then we did perform the linearization point with the expected effect.
- the definition of the local variant β as an upper bound on the impedance budget α .

Indeed, whenever some budget is spent, the loop approaches termination as eventually, the exhaustion of the budget prevents further interference, allowing the CAS operation to succeed and the loop to terminate. Therefore decreasing the upper bound to the interference budget corresponds to progress for the while operation. Without additional information however, we cannot show the local variant must eventually strictly decrease, indeed, in the case $l = 1$ we cannot exit the loop and the environment is not forced to spend budget. Therefore, the termination argument will need the assumption that the environment always eventually unlocks the lock to allow the

termination of the while loop or further decrease of the variant due to the environment locking the lock. This guarantee is given by the atomicity context $\mathcal{A} = [r \mapsto (\{0, 1\} \times \mathbb{O}, \mathbf{0}, \{0\} \times \mathbb{O}, R)]$ with $R = ((0, \alpha), \mathbf{0}) \rightsquigarrow ((1, \phi(\alpha)), \mathbf{0})$.

The target states, T , must clearly include unlocked states, where $l = 0$, but, as it must eventually be stable, this is insufficient, since once the lock is unlocked, the environment can lock it again. However, when the lock is unlocked, if the environment takes possession of it, the environment must also simultaneously decrease the impedance budget, i.e. $\beta > \alpha$.

The argument that T is always eventually true, relies on the assumption from the atomicity context that the environment will always eventually unlock the lock. However, this assumption only holds before the linearization point. In particular, as the loop variant must contain $r \mapsto \blacklozenge$, since the loop body may perform the linearization point, the persistent loop invariant cannot, and therefore T must also contain a disjunct where the linearization point has occurred and T holds a witness $r \mapsto (_, _)$.

We therefore declare the target states as the ones where, either the linearization point has been performed, or the lock is unlocked, or some budget was spent:

$$T(\beta) \triangleq \exists l, \alpha. \mathbf{spin}_r(x, l, \alpha) \wedge (r \mapsto (_, _) \vee l = 0 \vee \beta > \alpha)$$

The persistent loop invariant here is simply $L = \mathbf{spin}_r(x, _, _)$, which is a valid stable frame of the loop.

To apply **WHILE** we also need to specify $m(\beta)$, which in this case is simply $\mathbf{1}$, which satisfies the layer constraints of the rule; and the environment progress measure M :

$$M(\alpha_e) \triangleq \exists l, \alpha. \mathbf{spin}_r(x, l, \alpha) \wedge \alpha_e = 2\alpha + l$$

(here we use the variable α_e for the environment progress measure variable, to avoid clashes with the impedance budget α .) This environmental progress measure is decreased by both the environment locking and unlocking the lock:

- Unlocking the lock decreases l from 1 to 0, so as $2\alpha + 1 > 2\alpha + 0$, the environmental progress measure decreases.
- Locking the lock decreases the impedance budget from α to $\alpha' < \alpha$, while also increasing l from 0 to 1. Since $\alpha' < \alpha$ implies $\alpha' + 1 \leq \alpha$, $2\alpha + 0 \geq 2\alpha' + 2 > 2\alpha' + 1$, the environmental progress measure decreases.

Given these parameters, the proof first establishes the loop invariant holds at the beginning for some β_0 , by applying **CONS**:

$$\exists l, \alpha. \mathbf{spin}_r(x, l, \alpha) \wedge \alpha > \phi(\alpha) * r \mapsto \blacklozenge \wedge d = 0 \implies \exists \beta_0. P(\beta_0) * L$$

$$\exists \beta_0, \beta. P(\beta) * L \wedge d \neq 0 \wedge \beta_0 \geq \beta \implies \exists \alpha. \mathbf{spin}_r(x, _, _) * r \mapsto ((0, \alpha), (1, \phi(\alpha))) \wedge l = 0$$

Note that we will often implicitly apply the **CONS** rule in proofs, only detailing the application when emphasis is desired. Next, **EXELIM** on β_0 gets rid of the existential quantification, so we are ready to apply **WHILE**.

To complete the application of the rule we need to show

$$\mathbf{1}; \mathcal{A} \vdash L \xrightarrow{M} T(\beta) \tag{18}$$

$$\forall \alpha. \mathcal{A} \vDash \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \tag{19}$$

$$\text{pv}(T, L, M) \cap \text{mod}(\mathbb{C}) = \emptyset \tag{20}$$

Condition (19) is easily seen to hold, as we showed above, all possible environmental interference on the region decreases the environmental progress metric, which is sufficient for this to hold.

Condition (20) is also easily seen to hold as the only program variable predicated over in T , L and M is x , which is not modified by the body of the loop.

Finally, condition (18) is proven as follows. We observe that:

$$\begin{aligned} L(\alpha_e) = L * M(\alpha_e) &\equiv (\exists l, \alpha. \mathbf{spin}_r(x, l, \alpha) * (r \Rightarrow (_, _) \vee l \doteq 0) \wedge \alpha_e = 2\alpha + l) & (L_1(\alpha_e)) \\ &\vee (\exists \alpha. \mathbf{spin}_r(x, 1, \alpha) * r \Rightarrow \diamond \wedge \alpha_e = 2\alpha + 1) & (L_2(\alpha_e)) \end{aligned}$$

We can then derive the environment liveness condition:

$$\frac{\frac{\forall \alpha_e. \vdash_{\mathcal{A}} L_1(\alpha_e) \Rightarrow T(\beta)}{\forall \alpha_e. \mathbf{1}; \mathcal{A} \vdash L(\alpha_e) : L_1(\alpha_e) \rightarrow T(\beta)} \text{LIVET} \quad \frac{\text{impr}_{\mathcal{A}}(L_2, L, T(\beta))}{\forall \alpha_e. \mathbf{1}; \mathcal{A} \vdash L(\alpha_e) : L_2(\alpha_e) \rightarrow T(\beta)} \text{LIVEA}}{\frac{\forall \alpha_e. \mathbf{1}; \mathcal{A} \vdash L(\alpha_e) : L(\alpha_e) \rightarrow T(\beta)}{\mathbf{1}; \mathcal{A} \vdash L \xrightarrow{M} T(\beta)} \text{ECASE}} \text{ENVLIVE}$$

Formally, the application of **ENVLIVE** requires us to prove $\vdash_{\mathcal{A}} L \Rightarrow L * \exists \alpha_e. M(\alpha_e)$ which is trivial. An application of the **ECASE** rule then splits between the cases where L_1 and L_2 hold. Intuitively, L_1 represents the case where we performed the linearization point or the lock is unlocked, while L_2 the case where we still have not performed the linearization point and the lock is locked. If L_1 holds, then T holds, so no progress of the environment is required, therefore, this case can be discharged via an application of rule **LIVET**. In the case where L_2 holds we can apply rule **LIVEA** to invoke the liveness assumption stored in \mathcal{A} : if the lock is unlocked, the metric strictly decreases.

To show that the liveness assumption encoded in the atomicity context for the region \mathbf{spin}_r , $\text{live}(\mathcal{A}, r) = \{0, 1\} \times \mathbb{O} \rightarrow_k \{0\} \times \mathbb{O}$ is active, the **LIVEA** rule requires that in the current case:

- The abstractly atomic update being tracked on r has yet to occur:

$$\forall \alpha_e. \vdash_{\mathcal{A}} L_2(\alpha_e) \Rightarrow \exists (l, \alpha) \in (\{0, 1\} \times \mathbb{O}) \setminus (\{0\} \times \mathbb{O}). \mathbf{spin}_r(x, l, \alpha) * r \Rightarrow \diamond * \text{True}$$

- No obligations of layer less than or equal to k is continuously held locally:

$$\begin{aligned} m &> k \\ \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) &\triangleright k \end{aligned}$$

If these hold, then the $\text{impr}_{\mathcal{A}}(L_2, L, T(\beta))$ predicate shows that discharging the liveness invariant will strictly decrease α_e . To show this holds, taking $\sigma \in \text{Store}$ arbitrary and letting

$$l(\alpha) = \mathcal{W}[[L(\alpha)]]_{\mathcal{A}}^{\sigma} \quad l'(\alpha) = \mathcal{W}[[L_2(\alpha)]]_{\mathcal{A}}^{\sigma} \quad t = \mathcal{W}[[T(\beta) * \text{True}]]_{\mathcal{A}}^{\sigma}$$

we need to show

$$\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(l'(\alpha_1)) \cap l(\alpha_2) \subseteq l'(\alpha_1) \cup t$$

This holds, as, given an arbitrary $\alpha_1 \in \mathbb{O}$, any step taken from $l'(\alpha_1)$ by the atomic world rely relation either leaves the state of the region \mathbf{spin}_r unchanged, preserving the state $l'(\alpha_1)$, or releases the lock, decreasing the metric. Therefore, for any $\alpha_2 \geq \alpha_1$, $\mathbf{R}_{\mathcal{A}}^a(l'(\alpha_1)) \cap l(\alpha_2) \subseteq l'(\alpha_1)$ holds, which implies the goal.

To conclude the argument, we briefly comment on the proof of the body of the while loop. The full proof of the body can be found in figure 14. The applications of rules **UPDREG** and **FRAME** lift the concrete atomic **CAS** to a (potential) update to the \mathbf{spin}_r region. An application of **CONS** allows us to introduce γ as an upper bound to the impedance budget, initially δ after the linearization point.

Then, we apply rule **AΞELIM** to remove the pseudo-quantification on l and α . At this point, the abstract state l, α of the region \mathbf{spin}_r in the postcondition is weakened to any state that might be reached before or after the linearization point. However, we keep record of what happened exactly at the linearization point because of the $r \Rightarrow _$ assertions. The later application of **MkATOM** will be

$$\begin{array}{l}
\forall b \in \{0, 1\}, \beta. \\
\left\{ \exists l, \alpha. \text{spin}_r(x, l, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha) \wedge \beta \geq \alpha \right\} \\
\left\{ \wedge b \Rightarrow (l = 0 \vee \beta > \alpha) \wedge d = 0 \right\} \\
\left. \begin{array}{l}
\text{ATOMW; A}\exists\text{ELIM} \\
\text{CONS} \\
\text{UPDREG} \\
\text{FRAME}
\end{array} \right\} \\
\left\{ \begin{array}{l}
\forall l \in \{0, 1\}, \alpha. \\
\left\{ \text{spin}_r(x, l, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha) \wedge \beta \geq \alpha \wedge \right\} \\
\left\{ b \Rightarrow (l = 0 \vee \beta > \alpha) \wedge d = 0 \right\} \\
\left\{ \text{spin}_r(x, l, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha) \wedge \beta \geq \alpha \wedge \right\} \\
\left\{ b \Rightarrow (l = 0 \vee \beta > \alpha) \wedge d = 0 \right\} \\
\left\{ x \mapsto l \wedge \alpha > \phi(\alpha) \wedge \beta \geq \alpha \wedge \right\} \\
\left\{ b \Rightarrow (l = 0 \vee \beta > \alpha) \right\} \\
\left\{ x \mapsto l \right\} \\
d := \text{CAS}(x, 0, 1); \\
\left\{ x \mapsto 1 \wedge ((d = 0 \wedge l = 1) \vee (d = 1 \wedge l = 0)) \right\} \\
\left\{ \exists \delta. x \mapsto 1 \wedge \left(\begin{array}{l}
(d = 1 \wedge l = 0 \wedge \delta = \phi(\alpha) \wedge \beta > \phi(\alpha)) \\
\vee (d = 0 \wedge l = 1 \wedge \delta = \alpha \wedge \alpha > \phi(\alpha) \wedge b \Rightarrow \beta > \alpha)
\end{array} \right) \wedge \beta \geq \alpha \right\} \\
\left\{ \exists \delta. \text{spin}_r(x, 1, \delta) * \left(\begin{array}{l}
(d = 1 \wedge l = 0 \wedge \beta > \delta \wedge r \Rightarrow ((l, \alpha), (1, \phi(\alpha)))) \\
\vee (d = 0 \wedge l = 1 \wedge \delta > \phi(\delta) \wedge b \Rightarrow \beta > \delta \wedge r \Rightarrow \blacklozenge)
\end{array} \right) \wedge \beta \geq \delta \right\} \\
\left\{ \exists \gamma, \delta. \text{spin}_r(x, 1, \delta) \wedge \beta \geq \gamma \geq \delta \wedge b \Rightarrow \beta > \gamma \right. \\
\left. * \left(\begin{array}{l}
(d = 0 \wedge \delta > \phi(\delta) \wedge r \Rightarrow \blacklozenge) \\
\vee (d = 1 \wedge l = 0 \wedge r \Rightarrow ((l, \alpha), (1, \phi(\alpha))))
\end{array} \right) \right\} \\
\left\{ \exists l, \alpha, \gamma. \text{spin}_r(x, l, \alpha) \wedge \beta \geq \gamma \geq \alpha \wedge b \Rightarrow \beta > \gamma \right. \\
\left. * \left(\begin{array}{l}
(d = 0 \wedge \alpha > \phi(\alpha) \wedge r \Rightarrow \blacklozenge) \\
\vee (\exists l', \alpha'. d = 1 \wedge r \Rightarrow ((l', \alpha'), (1, \phi(\alpha')))) \wedge l' = 0
\end{array} \right) \right\}
\end{array} \right\}
\end{array}$$

Fig. 14. Spin lock: Proof of **while** loop body.

<p style="margin: 0;">PROOF OF makeLock():</p> <p style="margin: 0;">$0; 0 \vdash$</p> <p style="margin: 0;">{emp}</p> <p style="margin: 0;">{emp}</p> <p style="margin: 0;">CONS</p> <p style="margin: 0;">ret := alloc(1);</p> <p style="margin: 0;">[ret] := 0;</p> <p style="margin: 0;">{ret \mapsto 0}</p> <p style="margin: 0;">{$\exists r. l(r, \text{ret}, 0, \alpha)$}</p>	<p style="margin: 0;">PROOF OF unlock(x):</p> <p style="margin: 0;">$0; 0 \vdash$</p> <p style="margin: 0;">$\langle l(r, x, 1, \alpha) \rangle$</p> <p style="margin: 0;">CONS</p> <p style="margin: 0;">$\langle \text{spin}_r(x, 1, \alpha) * [E]_r \rangle$</p> <p style="margin: 0;">STEP 3</p> <p style="margin: 0;">$\langle x \mapsto 1 \rangle$</p> <p style="margin: 0;">[x] := 0;</p> <p style="margin: 0;">$\langle x \mapsto 0 \rangle$</p> <p style="margin: 0;">$\langle \text{spin}_r(x, 0, \alpha) * [E]_r \rangle$</p> <p style="margin: 0;">$\langle l(r, x, 0, \alpha) \rangle$</p>
--	--

Fig. 15. Spin lock: proof of makeLock and unlock. Here STEP 3 is LIFTA, FRAME, SUBPQ.

able to fetch the atomic update witness $r \Rightarrow ((l, \alpha), (1, \phi(\alpha)))$ and declare the appropriate atomic update in the overall specification. Note that the overall Hoare postcondition after the application of **ATOMW** is stable.

Finally, Fig. 15 shows the proof outlines for the makeLock and unlock operations. The only notable step of the proof of makeLock is the last application of **CONS** to viewshift the postcondition from $\text{ret} \mapsto 0$ to $\exists r. \text{spin}_r(x, 0, \alpha) * [E]_r$, which is possible because the interpretation of the region matches with this resource, so the reifications of the two assertions coincide.

The proof of unlock is a straightforward lifting of the atomic reset of the cell at x to the region spin_r . Neither proof involves a liveness argument.

5.2 CLH lock

Code. A CLH lock is an implementation of a fair lock module that guarantees fairness by queuing the threads that are waiting to take its possession. Its implementation is shown in Fig. 17.




Fig. 16. Illustration of the memory layout of CLH lock.

```

1 def makeLock() {
2   var x, h in
3   h := alloc(1); [h] := 0;
4   x := alloc(2);
5   [x] := h;
6   [x + 1] := h;
7   ret := x;
8 }

1 def lock(x) {
2   var c, p, v in
3   c := alloc(1); [c] := 1;
4   p := FAS(x + 1, c);
5   v := [p];
6   while(v ≠ 0) { v := [p]; }
7   [x] := c
8   dealloc(p)
9 }

1 def unlock(x) {
2   var h in
3   h := [x];
4   [h] := 0;
5 }

```

Fig. 17. Code of CLH lock operation.

The diagram in Fig. 16 describes the state of the queued threads, t_1, t_2, \dots, t_n , waiting to take possession of the lock, as well as the module's head and tail pointers into the queue.

As described in Section 2, this queue is represented by associating each of the n threads queuing on the lock with the heap cells $cell_1, cell_2, \dots, cell_{n-1}, cell_n$ in memory. Each thread executing the lock operation to take possession of the lock then holds in its local state the address of its cell and that of its predecessor's cell. These are held in the program variables c and p respectively in the implementation of `lock`. The local instance of these program variables for each queued threads and the cells they are pointing to can be seen in Fig. 16.

The thread associated with the cell at the head of the queue is said to hold the lock, and the value stored in its cell determines the state of the lock, l . When a thread first takes possession of the lock, the lock will be locked. Therefore, the initial value in these cells, when the associated threads join the queue, is 1. This can be seen in the implementation of the `lock` operation which allocates and sets its associated cell to value 1 on line 3 before enqueueing itself. Once the thread holding the lock wishes to release it, it can do so by setting the value of its cell to 0, unlocking the lock and signalling to the next thread in the queue that it can now take possession of the lock. This can be seen in the implementation of the `unlock` operation which fetches the address of the cell associated with the lock's owner from the queue's head pointer and then sets its value to 0.

In Fig. 16 the thread t_1 is at the head of the queue, waiting for the lock to be released. If the lock is released by its owner, t_1 then gains the exclusive permission to take possession of the lock by setting the value of the module's head pointer to the address of its associated cell. t_1 detects the

lock has been released by repeatedly reading the value of its predecessor’s cell in the while loop on line 6 and then sets the head pointer to the address of its cell, c , on line 7.

Once the lock is released, only the thread at the head of the queue (if any) has the permission to take possession of the lock next. Due to this, if the owners of the lock continuously eventually release it, the threads waiting on the lock take possession of it in the order they are enqueued.

To enqueue itself, the lock operation performs a FAS operation on the tail pointer, placing the cell it has allocated with value 1 at the tail of the queue, and writing the address of its predecessor to the p program variable. The order in which the lock operations are enqueued is then the order in which they executed line 4. Any weakly fair scheduler will eventually give each thread executing the lock operation the opportunity to execute this FAS operation, allowing it to enqueue itself.

As long as the client then guarantees that every thread holding the lock eventually releases it, the thread will eventually take possession of the lock once it reaches the front of the queue and the lock operation will terminate, guaranteeing fairness.

To be able to provide the same guarantee, that every thread requesting the lock will eventually be able to take its possession as long as the lock is always eventually released, the spin lock requires that its client only call the lock operation concurrently a finite number of times. This is exposed in the spin lock specification via ordinals bounding the impedance on the lock.

An interesting aspect of this example is that it features a combination of internal and external blocking: the client needs to always eventually unlock the lock —external blocking, requiring the client to provide a guarantee— and the lock operation needs to eventually take possession of the lock once the previous thread signals its release —internal blocking, guaranteed by the implementation. This second guarantee will be enforced using obligations not exposed in the specification. The proof will therefore involve an environment liveness condition discharged using both **LIVEO** and **LIVEA**.

Specifications. We will prove the following fair lock module specifications:

$$\begin{aligned} 1 &\vdash \forall l \in \{0, 1\} \rightarrow_0 \{0\}. \langle L(s, x, l) \rangle \text{lock}(x) \langle L(s, x, 1) \wedge l = 0 \rangle \\ 0 &\vdash \langle L(s, x, 1) \rangle \text{unlock}(x) \langle L(s, x, 0) \rangle \end{aligned}$$

where $L(s, x, l)$ abstractly represents the lock resource at abstract location s (omitted for readability in Section 2) and concrete address x , with abstract state $l \in \{0, 1\}$.

To abstract the representation of a thread’s position in the queue, we will associate, through ghost state, to each thread requesting the lock, a *ticket number* $t \in \mathbb{N}$ which corresponds to the order of arrival of the lock implementation at line 4. Every time a thread joins the queue, it gets assigned the next available ticket.

This example shows a common proof pattern of TaDA Live: there is an inner region that exposes all the information needed for the termination argument (here the value of the next ticket to be handed out, t , so that individual threads can reason about the threads queuing on the lock) and an outer one that hides enough information to make the operation abstractly atomic. This pattern nicely separates the concerns in the proof: proving atomicity is done via the outer region, termination via the inner one. Because of this, the abstract location of the lock s will consist of the pair of inner and outer region identifiers. This is not a concern for modularity however: the type of s can be made opaque to the client, which just threads it through the proof unmodified.

Shared Regions. The abstract shared lock resource will be represented by a region $\text{clh}_r(r', x, h, l, o)$ where $r' \in \text{RId}$, $x, h \in \text{Addr}$, $l \in \{0, 1\}$, $o \in \mathbb{N}$. Here r' , the region identifier of the inner region and x , the address of the lock, are the fixed parameters of the region. The abstract state of the region

includes l , which represents the lock's state, o , which is the ticket number of the thread holding the lock, and h is the address of the cell associated with the owner.

Once a lock operation has enqueued itself, the difference between the ticket of the lock's owner, o and the operation's ticket, t , $t - o$, corresponds to the thread's current position in the queue.

The internal region $\mathbf{lclh}_{r'}(x, h, l, o, t)$ also exposes the next ticket to be handed to the next thread queuing on the lock, $t \in \mathbb{N}$.

Notation. Lists will frequently be used in the ghost state for the proof of the CLH lock. We introduce notation to manipulate lists so as to simplify the exposition of the reasoning. Given $n \in X$ and $ns, ns' \in X^*$ lists of elements of X , we write $n \oplus ns$, $ns \oplus n$, and $ns \oplus ns'$ for prepend, append, and concatenation, respectively; $|ns|$ is the length of ns , and $ns(i) = n$ states that the i -th element (from 0) in ns is n and $i < |ns|$; $\text{fst}(ns)$ and $\text{last}(ns)$ are the first and the last element of ns , respectively and $\text{tail}(ns)$ represents the list ns without the first element when ns is non empty.

Guard algebra: Take $p, c \in \text{Addr}$, $ns \in \text{Addr}^*$, $o, t \in \mathbb{N}$ arbitrary. For this proof, two guards will be necessary. First $\mathbf{T}(p, c, t)$, which encodes the current thread's ticket, t , once it has joined the queue, as well as $p, c \in \text{Addr}$, pointers to the thread's predecessor's cell in the queue and its own respectively. The second guard we require is $\mathbf{Q}(ns, o)$, which is used to track the overall queue, by tracking the cells associated with enqueued threads, $ns \in \text{Addr}^*$, and the ticket number of the current owner, $o \in \mathbb{N}$.

To use this as intended, a few axioms on the guard algebra will be required. First, an axiom to create new tickets, adding a new cell to the queue and associating a new, unique ticket number to the thread:

$$\mathbf{Q}(ns \oplus [p], o) = \mathbf{Q}(ns \oplus [p, c], o) \bullet \mathbf{T}(p, c, o + |ns| + 1)$$

This will be used to create the relevant guard resources \mathbf{T} , when a lock operation enqueues itself on line 4. Similarly, an axiom to remove a thread's predecessor from the queue once it can take possession of the lock:

$$\mathbf{Q}([p, c] \oplus ns, o) \bullet \mathbf{T}(p, c, o + 1) = \mathbf{Q}([c] \oplus ns, o + 1)$$

This will be used to update the relevant guard resources \mathbf{Q} with the relevant \mathbf{T} , when a lock operation takes possession of the lock on line 7, placing its associated cell, c , at the head of the queue. Finally, an axiom to guarantee that a ticket guard, \mathbf{T} is well-formed with respect to the queue in a guard \mathbf{Q} :

$$\mathbf{Q}(ns, o) \bullet \mathbf{T}(p, c, t) \neq \perp \Leftrightarrow ns(t - o - 1) = p \wedge ns(t - o) = c$$

Obligation algebra: Take $o, o', t, t' \in \mathbb{N}$ arbitrary. As mentioned above, to verify the totality of the CLH lock operation, once a thread is enqueued, if its predecessor relinquishes possession of the lock, it must eventually take its possession. Otherwise, although the lock will be permanently unlocked, no other thread waiting for the lock can take its possession, as they are not at the head of the queue.

To encode this liveness invariant which must be fulfilled, we associate an atom obligation $\mathbf{P}(t)$ with the ownership of the ticket $t \in \mathbb{N}$. The CLH lock's transition system will then require that this obligation be discharged by taking possession of the lock once it is unlocked by the thread with ticket $t - 1$.

The layer associated with $\mathbf{P}(t)$ is then t , so that these obligations are resolved in the order the associated threads are enqueued. Finally, as with the guard algebra, we have an obligation $\mathbf{O}(o, t)$, which will remain in the shared region's state and track the owner's ticket, o , and the next ticket to

be handed out, t , associated with the obligation \mathbf{P} via the obvious axioms.

$$\begin{aligned} \mathbf{o}(o, t) &= \mathbf{o}(o, t + 1) \bullet \mathbf{P}(t) & \mathbf{o}(o + 1, t) &= \mathbf{o}(o, t) \bullet \mathbf{P}(o + 1) \\ \mathbf{o}(o, t) \bullet \mathbf{P}(t') &\neq \perp & \Leftrightarrow & o \leq t' < t \\ \mathcal{L} &\triangleq \mathbb{N} \cup \{1, \mathbf{0}\} & \forall i \in \mathbb{N}. 1 > i > \mathbf{0} & \quad \text{lay}(\mathbf{o}(o, t)) = 0 \quad \text{lay}(\mathbf{P}(t)) = t \end{aligned}$$

Region protocols. The interference protocol for the **lclh** region is as follows:

$$\begin{aligned} \mathbf{E} &: ((h, l, o, t), \mathbf{0}) \rightsquigarrow ((h, l, o, t + 1), \mathbf{P}(t)) \\ \mathbf{E} &: ((h, 0, o, t), \mathbf{P}(o + 1)) \rightsquigarrow ((h', 1, o + 1, t), \mathbf{0}) \\ \mathbf{E} &: ((h, 1, o, t), \mathbf{0}) \rightsquigarrow ((h, 0, o, t), \mathbf{0}) \end{aligned}$$

The first transition allows a thread to place itself in the queue waiting to obtain the CLH lock, updating the next ticket to be handed out from t to $t + 1$. While doing so, the thread acquires an obligation, $\mathbf{P}(t)$, requiring it to eventually take possession of the lock once it is at the head of the queue. The second, allows the thread at the head of the queue to take possession of the lock, by changing the state, l , incrementing the owner ticket, o , to its own (tracked by the thread’s obligation) and changing the owner pointer of the lock to that of its own associated cell. This discharges the obligation $\mathbf{P}(o + 1)$, as the thread then leaves the queue, to take possession of the lock. Finally, the third transition allows the lock to be unlocked.

The interference protocol for the **clh** region is then:

$$\begin{aligned} \mathbf{E} &: ((h, l, o), \mathbf{0}) \rightsquigarrow ((h, l, o), \mathbf{0}) \\ \mathbf{E} &: ((h, 0, o), \mathbf{0}) \rightsquigarrow ((h', 1, o + 1), \mathbf{0}) \\ \mathbf{E} &: ((h, 1, o), \mathbf{0}) \rightsquigarrow ((h, 0, o), \mathbf{0}) \end{aligned}$$

This hides the enqueueing step of the lock operation, allowing the operation to appear atomic.

Region interpretation. As explained above, the CLH lock associates a cell with each thread queuing on it, as well as its owner. The list of each of these cells in the order in which the associated threads are queued, with the owner’s cell as the head, will be denoted ns . $\text{tail}(ns)$ is then the list of cells queuing on the lock. While threads are queuing, the associated cells must have value 1; this is represented using the predicate ones:

$$\text{ones}(ns) \triangleq ns(1) \mapsto 1 * \dots * ns(|ns| - 1) \mapsto 1$$

The inner shared region, **lclh**, holds the cells associated with each queued thread, this is represented by the resource $\text{ones}(ns)$ in the region interpretation.

The shared region also holds a pointer to the tail of the queue, ns , as well as a pointer to its owner’s cell, whose value is the state of the lock, l , as described above. This is represented by the resource:

$$x \mapsto h, \text{last}(ns) * h \mapsto l$$

The shared region’s ghost state is then comprised of:

- $[\mathbf{Q}(ns, o)]_{r'}$, the guard keeping track of the list of cells, $ns \in \text{Addr}^*$ and the current owner of the lock, $o \in \mathbb{N}$.
- $[\mathbf{o}(o, t)]_{r'}$, the obligation keeping track of the next ticket to hand out, $t \in \mathbb{N}$, and the current owner’s ticket, $o \in \mathbb{N}$.

Finally, the invariant $t - o = |ns|$ is used to guarantee that each thread that holds a ticket is associated with a cell in the queue ns and $h = ns(0)$, associates the head of ns and the address of the owner’s

$$\begin{array}{l}
\mathbf{1}; \emptyset \vdash \mathbb{W}l \in \{0, 1\} \rightarrow_0 \{0\}. \\
\langle L(s, x, l) \rangle \quad s = (r, r') \\
\mathbf{1}; \emptyset \vdash \\
\mathbb{W}l \in \{0, 1\} \rightarrow_0 \{0\}, o, h. \\
\langle \text{clh}_r(r', x, h, l, o) * [\mathbb{E}]_{r'} \rangle \\
\mathbf{1}; [r \mapsto (X_1, 0, X_2, R)] \vdash \\
\{ \exists l, o, h. \text{clh}_r(r', x, h, l, o) * r \Rightarrow \blacklozenge \} \\
c := \text{alloc}(1); [c] := 1; \\
\{ \exists l, o, h. \text{clh}_r(r', x, h, l, o) * r \Rightarrow \blacklozenge * c \mapsto 1 \} \\
p := \text{FAS}(x + 1, c); \\
v := [p]; \\
\{ \exists \beta, t'. P(\beta) * L \} \\
\text{while}(v \neq 0) \{ v := [p]; \} \\
\{ \exists t'. P(0) * L \wedge v = 0 \} \\
[x] := c; \\
\{ \exists o \in \mathbb{N}, h, h'. r \Rightarrow ((h, 0, o), (h', 1, o + 1)) * p \mapsto 0 \wedge l = 0 \} \\
\text{dealloc}(p) \\
\{ \exists o. r \Rightarrow ((_, 0, o), (_, 1, o + 1)) \wedge l = 0 \} \\
\langle \exists h'. \text{clh}_r(r', x, h', 1, o + 1) * [\mathbb{E}]_{r'} \wedge l = 0 \rangle \\
\langle L(s, x, 1) \wedge l = 0 \rangle
\end{array}$$

$$\begin{array}{l}
X_1 \triangleq \{(h, l, o) \mid h \in \text{Addr}, l \in \{0, 1\}, o \in \mathbb{N}\} \\
X_2 \triangleq \{(h, 0, o) \mid h \in \text{Addr}, o \in \mathbb{N}\} \\
R \triangleq \{((h, 0, o), (h', 1, o)) \mid h, h' \in \text{Addr}, o \in \mathbb{N}\} \\
P(\beta) \triangleq \exists l, o, h. \text{clh}_r(r', x, h, l, o) * [\mathbb{T}(p, c, t')]_{r'} \wedge \\
(v = 0 \Rightarrow (t' = o + 1 \wedge l = 0 \wedge h = p)) \wedge \\
o < t' \wedge \beta = v \\
L = \exists l \in \{0, 1\}, o, h. \text{clh}_r(r', x, h, l, o) * [p(t')]_{r'}^{\perp} * \\
*_{i=o+1}^{t'-1} [p(i)]_{r'}^{\mathbb{E}} * r \Rightarrow \blacklozenge \wedge o < t'
\end{array}$$

$$\begin{array}{l}
X_1 \triangleq \{(h, l, o) \mid h \in \text{Addr}, l \in \{0, 1\}, o \in \mathbb{N}\} \\
X_2 \triangleq \{(h, 0, o) \mid h \in \text{Addr}, o \in \mathbb{N}\} \\
R \triangleq \{((h, 0, o), (h', 1, o)) \mid h, h' \in \text{Addr}, o \in \mathbb{N}\} \\
P(\beta) \triangleq \exists l, o, h. \text{clh}_r(r', x, h, l, o) * [\mathbb{T}(p, c, t')]_{r'} \wedge \\
(v = 0 \Rightarrow (t' = o + 1 \wedge l = 0 \wedge h = p)) \wedge \\
o < t' \wedge \beta = v \\
L = \exists l \in \{0, 1\}, o, h. \text{clh}_r(r', x, h, l, o) * [p(t')]_{r'}^{\perp} * \\
*_{i=o+1}^{t'-1} [p(i)]_{r'}^{\mathbb{E}} * r \Rightarrow \blacklozenge \wedge o < t'
\end{array}$$

Fig. 18. Outline of CLH lock proof.

cell. All of this ties together to give the following region interpretation:

$$\begin{aligned}
\mathcal{I}(\text{lclh}_{r'}(x, h, l, o, t)) &\triangleq \exists ns \in \text{Addr}^*. x \mapsto h, \text{last}(ns) * \\
&h \mapsto l * \text{ones}(ns) * [\mathbb{Q}(ns, o)]_{r'} * [o(o, t)]_{r'}^{\perp} \wedge t - o = |ns| \wedge ns(0) = h
\end{aligned}$$

The outer shared region then holds full permission to update the inner region, $[\mathbb{E}]_{r'}$, and asserts that each thread queuing on the lock, with tickets $o + 1$ to $t - 1$, holds an obligation to take possession of the lock once their predecessor releases it, $*_{i=o+1}^{t-1} [p(i)]_{r'}^{\mathbb{E}}$, where r' is the identifier of the inner region:

$$\mathcal{I}(\text{clh}_r(r', x, h, l, o)) \triangleq \exists t \in \mathbb{N}. \text{lclh}_{r'}(x, h, l, o, t) * [\mathbb{E}]_{r'} * *_{i=o+1}^{t-1} [p(i)]_{r'}^{\mathbb{E}}$$

Predicates. The lock resource is then abstractly represented by the predicate:

$$L(s, x, l) \triangleq \exists r, r'. s = (r, r') \wedge \exists o \in \mathbb{N}. \exists h \in \text{Addr}. \text{clh}_r(r', x, h, l, o) * [\mathbb{E}]_{r'}$$

which abstracts away the CLH lock's implementation details: the ticket and cell address associated with the lock's current owner.

Proof of lock. Figure 18 gives an outline of the proof of the `clh` lock operation implementation, the definition of the loop invariant $P(\beta)$ will be given later. The steps involving liveness are the **FAS** operation, which enqueues the thread, hence obtaining the obligation to take possession of the lock once the previous thread relinquishes possession of it, the while loop which waits for the previous thread to release the lock, whose liveness depends on the previous threads in the queue taking possession and then releasing the lock in turn and the write operation at line 7 which takes possession of the lock. We begin with the details of the **FAS** operation's proof, shown in Fig. 19.

There, **STEP 4** is composed of the rules: **FRAMEH**, **ATOMW**, **A3ELIM**, **LIFTA**, **A3ELIM**, **LIFTA**, **A3ELIM**. The application of the **FRAMEH** rule frames off the view $r \Rightarrow \blacklozenge$, the **ATOMW** rule transfers all the remaining resources to the atomic precondition and postcondition, the **A3ELIM** rule pseudo-quantifies l, o and h , **LIFTA** then opens the region clh_r , the applications of **A3ELIM** and **LIFTA**

$$\begin{array}{l}
1; [r \mapsto (X_1, 0, X_2, R)] \vdash \\
\{ \exists l \in \{0, 1\}, o \in \mathbb{N}, h \in \text{Addr. } \text{clh}_r(r', x, h, l, o) * r \Rightarrow \blacklozenge * c \mapsto 1 \} \\
\begin{array}{l}
1; [r \mapsto (X_1, 0, X_2, R)] \vdash \\
\text{WI} \in \{0, 1\}, o, t \in \mathbb{N}, h \in \text{Addr}, ns \in \text{Addr}^*. \\
\left\langle \begin{array}{l}
x \mapsto h, \text{last}(ns) * h \mapsto l * \text{ones}(ns) * [\text{Q}(ns, o)]_{r'} * [\text{O}(o, t)]_{r'}^L * \\
*_{i=o+1}^{t-1} [\text{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge ns(0) = h * c \mapsto 1 \\
x \mapsto h, \text{last}(ns) * h \mapsto l * \text{ones}(ns) * [\text{Q}(ns, o)]_{r'} * [\text{O}(o, t)]_{r'}^L * \\
*_{i=o+1}^{t-1} [\text{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge ns(0) = h * c \mapsto 1
\end{array} \right\rangle \\
\text{STEP 4} \quad \left\langle \begin{array}{l}
0; [r \mapsto (X_1, 0, X_2, R)] \vdash \\
\langle x + 1 \mapsto \text{last}(ns) \rangle \\
\text{p} := \text{FAS}(x + 1, c); \\
\langle x + 1 \mapsto c \wedge \text{p} = \text{last}(ns) \rangle \\
x \mapsto h, c * h \mapsto l * \text{ones}(ns) * c \mapsto 1 * [\text{Q}(ns, o)]_{r'} * [\text{O}(o, t)]_{r'}^L * \\
*_{i=o+1}^{t-1} [\text{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge ns(0) = h \\
\exists ns' \in \text{Addr}^*. x \mapsto h, \text{last}(ns') * h \mapsto l * \text{ones}(ns') * [\text{Q}(ns', o)]_{r'} * [\text{O}(o, t + 1)]_{r'}^L * \\
*_{i=o+1}^t [\text{P}(i)]_{r'}^E \wedge (t + 1) - o = |ns'| \wedge ns(0) = h \wedge ns' = ns \oplus c * (\text{T}(p, c, t)]_{r'} * \\
[\text{P}(t)]_{r'}^L * *_{i=o+1}^{t-1} [\text{P}(i)]_{r'}^E \wedge o < t
\end{array} \right\rangle \\
\{ \exists l \in \{0, 1\}, o, t' \in \mathbb{N}, h \in \text{Addr. } \text{clh}_r(r', x, h, l, o) * r \Rightarrow \blacklozenge * [\text{T}(p, c, t')]_{r'} * [\text{P}(t')]_{r'}^L * *_{i=o+1}^{t'-1} [\text{P}(i)]_{r'}^E \wedge o < t' \}
\end{array}
\end{array}$$

Fig. 19. Proof outline of the FAS call of CLH lock.

then pseudo-quantify t and open the region `lclh` and the final application of `AΞELIM` rule pseudo-quantifies ns .

After using `LAYWH` to decrease the level of the assertion to `0` and `FRAME` to frame off everything except the region interpretation’s tail pointer, the `FAS` operation atomically updates it. After everything is framed back on, the consequence rule is then applied to the postcondition so as to re-establish the invariant. The axioms

$$\begin{aligned}
\text{Q}(ns \oplus [p], o) &= \text{Q}(ns \oplus [p, c], o) \bullet \text{T}(p, c, o + |ns| + 1) \\
\text{O}(o, t) &= \text{O}(o, t + 1) \bullet \text{P}(t)
\end{aligned}$$

are used to update the queue ns , by enqueueing c —the local thread’s cell— at its tail, and updating the next ticket to $t' + 1$. While doing so, the thread acquires the guard $\text{T}(p, c, t')$, the obligation, $\text{P}(t')$, which represent the thread’s position in the queue and its obligation to take possession of the lock once its predecessor relinquishes it respectively.

As environmental obligations can always be duplicated, the thread also obtains $*_{i=o+1}^{t'-1} [\text{P}(i)]_{r'}^E$ locally. These environmental assertions will be necessary for the application of the `WHILE` rule. To finish reestablishing the invariant, as the thread is retaining $[\text{P}(t')]_{r'}^L$ locally, it can leave $[\text{P}(t')]_{r'}^E$ in the region invariant. Finally, using the axiom

$$\text{O}(o, t) \bullet \text{P}(t') \neq \perp \Leftrightarrow o \leq t' < t$$

as we hold $\text{P}(t')$ locally, the assertion $o < t'$ holds stably.

Next, consider the proof of the `while` loop. The loop invariant is:

$$\begin{aligned}
P(\beta) &\triangleq \exists l, o, t', h. \text{clh}_r(r', x, h, l, o) * [\text{T}(p, c, t')]_{r'} \wedge o < t' \\
&\wedge (v = 0 \Rightarrow (t' = o + 1 \wedge l = 0 \wedge h = p)) \wedge \beta = v
\end{aligned}$$

which asserts that:

- $[\text{T}(p, c, t')]_{r'}$, the local thread is queueing for the lock with ticket t' and with the address of the predecessor’s cell and the current thread’s cell in p and c respectively.

- $o < t'$, the current owner must come before the local thread with ticket t' . This is stable due to the τ guard.
- $v = 0 \Rightarrow (t' = o + 1 \wedge l = 0 \wedge h = p)$, if v , the last read of the value of the predecessor cell, is 0, then the owner is the predecessor of the current thread has unlocked the lock, as only then can it set its cell to 0. Therefore $t' = o + 1$, and, consequently the lock is unlocked, $l = 0$. The owner's cell, h , will also take the value of that of the predecessor.
- $\beta = v$ which asserts that $\beta = 0$ once the thread has observed that its predecessor has taken possession of and then unlocked the lock (by reading the cell at address p into v). β will have value 1 otherwise.

A thread with ticket t' can take possession of a CLH lock once its predecessor has taken possession of and relinquished the lock. Once the lock reaches this state, $o = t' - 1$ and $l = 0$ hold stably as all transitions from this state would set $o \geq t'$, however we know that, $o < t'$.

The intent of this loop is to wait till this occurs, allowing the thread to safely take possession of the lock once the loop terminates. Hence, the goal state is:

$$T = \exists l \in \{0, 1\}, o \in \mathbb{N}, h \in \text{Addr}. \text{clh}_r(r', x, h, l, o) \wedge t' = o + 1 \wedge l = 0 \wedge h = p$$

Once the lock reaches this state, a subsequent iteration of this **while** loop will terminate with $v = 0$, breaking the loop. To reach the goal state, threads that come before the current thread must both take possession and then unlock the lock. The first is guaranteed due to obligations $\mathbf{p}(t')$ for $t' < t$ and the second due to the pseudo-quantifier, guaranteeing that the lock must always eventually be released. The progress measure

$$M(\alpha) = \exists l \in \{0, 1\}, o \in \mathbb{N}, h \in \text{Addr}. \text{clh}_r(r', x, h, l, o) \wedge \alpha = 2(t' - o - 1) + l$$

is decreased by both of these actions, and as $t' > o$ implies $2(t' - o - 1) + l \geq 0$, the progress measure, α , is a natural number, and therefore well-founded.

The use of the difference between t' , the local thread's ticket and the owner's ticket, o , to bound the number of threads that can take possession of the lock before the local thread removes the necessity for the impedance bound, α , required in the proof of the spin lock module, and that must leak in the associated specification (as it imposes a restriction on any client).

To support this argument, the persistent loop invariant, L , must contain the resource $r \Rightarrow \blacklozenge$ to make use of the liveness assumptions of the pseudo-quantifier, guaranteeing that the lock is always eventually unlocked, and the relevant environmental liveness assertions guaranteeing the threads queued before the current thread will take possession of it once their predecessor relinquishes it:

$$L = \exists l \in \{0, 1\}, o, t' \in \mathbb{N}, h \in \text{Addr}. \text{clh}_r(r', x, h, l, o) * [\mathbf{p}(t')]_{r'}^L * \bigstar_{i=o+1}^{t'-1} [\mathbf{p}(i)]_{r'}^E * r \Rightarrow \blacklozenge \wedge o < t'$$

The **WHILE** rule is applied as in Fig. 20. The rule $\exists\text{ELIM}$ is applied to quantify t and β_0 over the antecedent. To complete the application of the rule we need to show

$$\mathbf{1}; \mathcal{A} \vdash L \xrightarrow{M} T(\beta) \quad (21)$$

$$\forall \alpha. \mathcal{A} \vDash \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \quad (22)$$

Condition (22) holds trivially, as seen above all the possible operations on the module decrease the environmental metric.

$$\begin{array}{l}
\mathbf{1}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\{ \exists t' \in \mathbb{N}. \exists \beta_0. P(\beta_0) * L \} \\
\left| \begin{array}{l}
\forall \beta_0, t \in \mathbb{N}. \\
\mathbf{1}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\{ P(\beta_0) * L \} \\
\text{CONS; } \exists\text{ELIM} \left| \begin{array}{l}
\text{WHILE} \left\{ \begin{array}{l}
\text{while}(v \neq 0) \{ \\
\forall \beta \leq \beta_0, b \in \mathbb{B}. \\
\{ P(\beta) * b \Rightarrow T(\beta) \wedge v \neq 0 \} \\
v := [p]; \\
\{ \exists \gamma. P(\gamma) \wedge \gamma \leq \beta \wedge b \Rightarrow \gamma < \beta \} \\
\} \\
\{ \exists \gamma. P(\gamma) * L \wedge \gamma \leq \beta_0 \wedge v = 0 \} \\
\{ \exists o \in \mathbb{N}. \text{clh}_r(r', x, p, 0, o) * r \Rightarrow \blacklozenge * [\mathbf{T}(p, c, o+1)]_{r'} * \lfloor \mathbf{P}(o+1) \rfloor_{r'}^{\perp} \}
\end{array} \right. \\
\end{array} \right. \\
\end{array}
\end{array}$$

$$\begin{array}{l}
X_1 \triangleq \{ (h, l, o) \mid h \in \text{Addr}, l \in \{0, 1\}, o \in \mathbb{N} \} \\
X_2 \triangleq \{ (h, 0, o) \mid h \in \text{Addr}, o \in \mathbb{N} \} \\
R \triangleq \{ ((h, 0, o), (h', 1, o)) \mid h, h' \in \text{Addr}, o \in \mathbb{N} \}
\end{array}$$

Fig. 20. Application of **WHILE** in the CLH lock proof.

$$\forall \beta_0, t \in \mathbb{N}, \beta, b \in \mathbb{B}.$$

$$\begin{array}{l}
\mathbf{1}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\left\{ \begin{array}{l}
(\exists l \in \{0, 1\}, o \in \mathbb{N}, h \in \text{Addr}. \text{clh}_r(r', x, h, l, o) * [\mathbf{T}(p, c, t)]_{r'} \wedge \\
o < t \wedge (v = 0 \Rightarrow (t = o + 1 \wedge l = 0 \wedge h = p)) \wedge \beta = v \wedge b \Rightarrow (t = o + 1 \wedge l = 0 \wedge h = p) \wedge (v \neq 0))
\end{array} \right\} \\
\text{CONS; } \exists\text{ELIM; } \text{LIFTA; } \text{A}\exists\text{ELIM; } \text{LIFTA; } \text{A}\exists\text{ELIM; } \text{CONS} \left| \begin{array}{l}
\mathbf{1}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\forall l \in \{0, 1\}, h \in \text{Addr}, ns \in \text{Addr}^*, o, nt \in \mathbb{N}. \\
\left\langle \begin{array}{l}
x \mapsto h, \text{last}(ns) * h \mapsto l * \text{ones}(ns) * [\mathbf{Q}(ns, o)]_{r'} * \lfloor \mathbf{O}(o, nt) \rfloor_{r'}^{\perp} * \blacklozenge_{i=o+1}^{nt-1} [\mathbf{P}(i)]_{r'}^E \wedge nt - o = |ns| \wedge \\
ns(0) = h * ([\mathbf{T}(p, c, t)]_{r'} \wedge o < t \wedge \beta \geq 1 \wedge b \Rightarrow (t = o + 1 \wedge l = 0 \wedge h = p) \wedge p \in ns)
\end{array} \right\rangle \\
\text{LAWWH; FRAME} \left| \begin{array}{l}
\mathbf{0}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\forall v \in \{0, 1\}. \\
\langle p \mapsto v \rangle \\
v := [p]; \\
\langle p \mapsto v \wedge v = v \rangle
\end{array} \right. \\
\left\langle \begin{array}{l}
x \mapsto h, \text{last}(ns) * h \mapsto l * \text{ones}(ns) * [\mathbf{Q}(ns, o)]_{r'} * \lfloor \mathbf{O}(o, nt) \rfloor_{r'}^{\perp} * \blacklozenge_{i=o+1}^{nt-1} [\mathbf{P}(i)]_{r'}^E \wedge nt - o = |ns| \wedge \\
ns(0) = h * ([\mathbf{T}(p, c, t)]_{r'} \wedge \beta = 1 \wedge \\
\exists v \in \{0, 1\}. v = v \wedge b \Rightarrow v = 0 \wedge (v = 0 \Rightarrow (t = o + 1 \wedge l = 0 \wedge h = p)))
\end{array} \right\rangle \\
\left\{ \begin{array}{l}
\exists l \in \{0, 1\}, o \in \mathbb{N}, h \in \text{Addr}, \gamma. \text{clh}_r(r', x, h, l, o) * [\mathbf{T}(p, c, t)]_{r'} \wedge \\
o < t \wedge (v = 0 \Rightarrow (t = o + 1 \wedge l = 0 \wedge h = p)) \wedge \gamma = v \wedge \gamma \leq \beta \wedge b \Rightarrow \gamma = 0
\end{array} \right\}
\end{array}
\end{array}$$

Fig. 21. Proof outline of the CLH lock’s loop body.

To prove (21), take

$$\begin{aligned}
L''_o(\alpha) &= \left(\begin{array}{l} \exists l \in \{0, 1\}, h. \text{clh}_r(r', x, h, l, o) * \lfloor \mathbf{P}(t') \rfloor_{r'}^{\perp} * \\ \blacklozenge_{i=o+2}^{t'-1} [\mathbf{P}(i)]_{r'}^E * r \Rightarrow \blacklozenge \wedge l = 0 \wedge o + 1 < t' \end{array} \right) * M(\alpha) \\
L'_0(\alpha) &= \left(\begin{array}{l} \exists l \in \{0, 1\}, o, h. \text{clh}_r(r', x, h, l, o) * \lfloor \mathbf{P}(t') \rfloor_{r'}^{\perp} * \\ \blacklozenge_{i=o+2}^{t'-1} [\mathbf{P}(i)]_{r'}^E * r \Rightarrow \blacklozenge \wedge l = 0 \wedge o + 1 < t' \end{array} \right) * M(\alpha) \\
L'_1(\alpha) &= \left(\begin{array}{l} \exists l \in \{0, 1\}, o, h. \text{clh}_r(r', x, h, l, o) * \lfloor \mathbf{P}(t') \rfloor_{r'}^{\perp} * \\ \blacklozenge_{i=o+1}^{t'-1} [\mathbf{P}(i)]_{r'}^E * r \Rightarrow \blacklozenge \wedge l = 1 \end{array} \right) * M(\alpha) \\
L(\alpha) &= L * M(\alpha)
\end{aligned}$$

First split on $\alpha = 0 \vee \alpha > 0$:

$$\frac{\frac{\forall \alpha. \vdash_{\mathcal{A}} L(\alpha) \wedge \alpha = 0 \Rightarrow T}{\mathbf{1}; \mathcal{A} \vdash L(\alpha) : L(\alpha) \wedge \alpha = 0 \rightarrow T} \text{LIVET} \quad \frac{(23) \quad (24)}{\mathbf{1}; \mathcal{A} \vdash L(\alpha) : (L'_0(\alpha) \vee L'_1(\alpha)) \wedge \alpha > 0 \rightarrow T} \text{ECASE}}{\mathbf{1}; \mathcal{A} \vdash L(\alpha) : L(\alpha) \rightarrow T} \text{ECASE} \quad \text{ENVLIVE}$$

$$\frac{\mathbf{1}; \mathcal{A} \vdash L \xrightarrow{M} T}{\mathbf{1}; \mathcal{A} \vdash L \xrightarrow{M} T} \text{ENVLIVE}$$

In the case $\alpha = 0$, the rule **LIVET** applies directly. To show $\mathbf{1}; \mathcal{A} \vdash L(\alpha) : L'(\alpha) \wedge \alpha > 0 \xrightarrow{M} T$ holds, split on the state of the lock, $l = 0 \vee l = 1$.

In the case $l = 0$, for each $o \in \mathbb{N}$, the ticket of the current owner of the lock, the environment is guaranteed to eventually take possession of the lock due to the environmental obligation assertion $[\mathbf{P}(o+1)]_{r'}^E$. To consider each case for $o \in \mathbb{N}$, we first apply the rule **EQUANT** and then the **LIVEO** rule:

$$\frac{\text{impr}_{\mathcal{A}}(L''_o, L, T) \quad \forall \alpha. \vdash_{\mathcal{A}} L''_o(\alpha) \Rightarrow \text{clh}_r(_, _, o) * [\mathbf{P}(o+1)]_{r'}^E * \text{True}}{\frac{\forall \alpha. \vdash_{\mathcal{A}} L''_o(\alpha) \triangleright \text{lay}(\mathbf{P}(o+1)) \quad \mathbf{1} > \text{lay}(\mathbf{P}(o+1))}{\forall o \in \mathbb{N}. \mathbf{1}; \mathcal{A} \vdash L(\alpha) : L''_o(\alpha) \rightarrow T} \text{LIVEO}}{\mathbf{1}; \mathcal{A} \vdash L(\alpha) : \exists o \in \mathbb{N}. L''_o(\alpha) \rightarrow T} \text{EQUANT} \quad (23)$$

With the exception of $\text{impr}_{\mathcal{A}}(L''_o, L, T)$, all of these conditions hold trivially. This last condition holds as, given $\alpha_0 \in \mathbb{O}$, all possible transitions either preserve $L''_o(\alpha)$ or decrease the metric.

In the case $l = 1$, progress is guaranteed due to the assumptions in the atomicity context, \mathcal{A} , that eventually, the lock must be released, so the **LIVEA** rule is applied:

$$\frac{\text{impr}_{\mathcal{A}}(L'_1, L, T) \quad \mathbf{1} > \mathbf{0} \quad \forall \alpha. \vdash_{\mathcal{A}} L'_1(\alpha) \triangleright k}{(X_1 \rightarrow_0 X_2) = \text{live}(\mathcal{A}, r) \quad \vdash_{\mathcal{A}} L'_1(\alpha) \Rightarrow \exists x \in X_1 \setminus X_2. \text{clh}_r(r', x, x) * r \Rightarrow \diamond * \text{True}} \text{LIVEA}}{\mathbf{1}; \mathcal{A} \vdash L(\alpha) : L'_1(\alpha) \rightarrow T} \quad (24)$$

Once again, with the exception of $\text{impr}_{\mathcal{A}}(L'_1, L, T)$, all of these conditions hold trivially. This last condition holds as, given $\alpha_0 \in \mathbb{O}$, all possible transitions either preserve $L'_1(\alpha)$ or decrease the metric.

To conclude the proof of lock, the argument for the body of the **while** loop's proof is purely a safety argument, the full proof is in Fig. 21.

The key step uses the axiom

$$\mathbf{Q}(ns, o) \bullet \mathbf{T}(p, c, t) \neq \perp \Leftrightarrow ns(t - o - 1) = p \wedge ns(t - o) = c$$

Since we hold the guard $\mathbf{T}(p, c, t)$, we can infer $p \in ns$. Then, after the value of the cell at p has been read, if the value, v , is 0, then, since only the thread holding the lock can change the value of their associated cell to 0, then, $t = o + 1 \wedge l = 0 \wedge h = p$. As a consequence, if b holds initially, then $v = 0$ after the body of the loop is executed, therefore the loop variant in the postcondition, $\gamma = 0$. As initially, we know $v \neq 0$ from the loop condition, $\beta = 1$, therefore $\gamma < \beta$.

Finally, in Fig. 22 we consider the details of the linearization point, when the lock operation takes possession of the lock. First **EXELIM** rule is applied to quantify the ticket of the current owner, o , (the predecessor of the current thread) over the antecedent. Then the **ATOMW** and **UPDREG** rules are applied to atomically update the region state by acting on its interpretation. The rules **AEXELIM**, **LIFTA** and **AEXELIM** are then applied to pseudo-quantify t and ns , the two variables that are existentially quantified within the region invariants and open the region **lclh**. Finally the **CONS** rule is applied to re-establish the invariant in the postcondition by adjusting the ghost state. Specifically,

$$\begin{array}{l}
\mathbf{1}; [r \mapsto (X_1, \mathbf{0}, X_2, R)] \vdash \\
\{ \exists o \in \mathbb{N}. \mathbf{clh}_r(r', x, p, \mathbf{0}, o) * r \Rightarrow \blacklozenge * [\mathbf{T}(p, c, o + 1)]_{r'} * [\mathbf{P}(o + 1)]_{r'}^L \} \\
\mathbf{1}; \emptyset \vdash \\
\forall t \in \mathbb{N}, ns \in \text{Addr}^*. \\
\left\langle \begin{array}{l}
x \mapsto p, \text{last}(ns) * p \mapsto \mathbf{0} * \text{ones}(ns) * [\mathbf{Q}(ns, o)]_{r'} * [\mathbf{O}(o, t)]_{r'}^L * \bigstar_{i=o+1}^{t-1} [\mathbf{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge \\
ns(0) = p * ([\mathbf{T}(p, c, o + 1)]_{r'} * [\mathbf{P}(o + 1)]_{r'}^L \wedge ns(1) = c)
\end{array} \right\rangle \\
\left\langle \begin{array}{l}
x \mapsto p, \text{last}(ns) * p \mapsto \mathbf{0} * \text{ones}(ns) * [\mathbf{Q}(ns, o)]_{r'} * [\mathbf{O}(o, t)]_{r'}^L * \bigstar_{i=o+1}^{t-1} [\mathbf{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge \\
ns(0) = p * ([\mathbf{T}(p, c, o + 1)]_{r'} * [\mathbf{P}(o + 1)]_{r'}^L \wedge ns(1) = c)
\end{array} \right\rangle \\
\text{STEP 5} \quad \left\langle \begin{array}{l}
\mathbf{0}; \emptyset \vdash \\
\langle x \mapsto p \rangle \\
[x] := c; \\
\langle x \mapsto c \rangle \\
\text{CONS} \quad \left\langle \begin{array}{l}
x \mapsto c, \text{last}(ns) * p \mapsto \mathbf{0} * \text{ones}(ns) * [\mathbf{Q}(ns, o)]_{r'} * [\mathbf{O}(o, t)]_{r'}^L * \\
\bigstar_{i=o+1}^{t-1} [\mathbf{P}(i)]_{r'}^E \wedge t - o = |ns| \wedge ns(0) = p * ([\mathbf{T}(p, c, o + 1)]_{r'} * [\mathbf{P}(o + 1)]_{r'}^L \wedge ns(1) = c)
\end{array} \right\rangle \\
\left\langle \begin{array}{l}
\exists ns' \in \text{Addr}^*. x \mapsto c, \text{last}(ns') * c \mapsto \mathbf{1} * \text{ones}(ns') * [\mathbf{Q}(ns', o + 1)]_{r'} * [\mathbf{O}(o + 1, t)]_{r'}^L * \\
\bigstar_{i=o+2}^{t-1} [\mathbf{P}(i)]_{r'}^E \wedge t - (o + 1) = |ns'| \wedge ns'(0) = c * (p \mapsto \mathbf{0} * ns = p \oplus ns')
\end{array} \right\rangle \\
\{ \exists o \in \mathbb{N}, h, h' \in \text{Addr}. r \Rightarrow ((h, \mathbf{0}, o), (h', \mathbf{1}, o + 1)) * p \mapsto \mathbf{0} \wedge l = \mathbf{0} \}
\end{array}
\end{array}$$

Fig. 22. Proof outline for the linearization point of CLH lock.

STEP 5 is CONS, \exists ELIM, ATOMW, UPDREG, $\mathbf{A}\exists$ ELIM, LIFTA, $\mathbf{A}\exists$ ELIM, CONS

the guard \mathbf{T} and the obligation \mathbf{P} are reabsorbed into \mathbf{Q} and \mathbf{O} respectively, to update the list of threads waiting on the lock and increment the owner. This is done using the axioms:

$$\begin{aligned}
\mathbf{Q}([p, c] \oplus ns, o) \bullet \mathbf{T}(p, c, o + 1) &= \mathbf{Q}(c \oplus ns, o + 1) \\
\mathbf{O}(o, t) \bullet \mathbf{P}(o + 1) &= \mathbf{O}(o + 1, t)
\end{aligned}$$

The inner part of the proof then decreases the layer and frames off unnecessary resources to apply the update. Note that this step of the proof discharges the obligation $\mathbf{P}(t')$. This concludes the verification of the lock operation.

The CLH lock proof is able to internally encode the impedance bound enforced by thread queuing using ghost state: the local ticket numbers of each thread queuing for the lock and the owner’s ticket number which is visible in the abstract state of the region \mathbf{clh} , but hidden from the client.

Proof of unlock. Let $X = \{(h, \mathbf{1}, o) \mid h \in \text{Addr}, o \in \mathbb{N}\}$ and $R = \{((h, \mathbf{1}, o), (h, \mathbf{0}, o)) \mid h \in \text{Addr}, o \in \mathbb{N}\}$. The proof of the unlock operation is as follows:

$$\begin{array}{l}
\mathbf{0}; \emptyset \vdash \\
\langle \mathbf{L}(r, x, \mathbf{1}) \rangle \\
\left\langle \begin{array}{l}
\mathbf{0}; [r \mapsto (X, \mathbf{0}, X, R)] \vdash \\
\{ \exists o \in \mathbb{N}, h \in \text{Addr}. \mathbf{clh}_r(x, h, \mathbf{1}, o) * r \Rightarrow \blacklozenge \} \\
h := [x]; \\
\{ \exists o \in \mathbb{N}. \mathbf{clh}_r(x, h, \mathbf{1}, o) * r \Rightarrow \blacklozenge \} \\
[h] := \emptyset; \\
\{ \exists o \in \mathbb{N}. r \Rightarrow ((h, \mathbf{1}, o), (h, \mathbf{0}, o)) \}
\end{array} \right\rangle \\
\langle \mathbf{L}(r, x, \mathbf{0}) \rangle
\end{array}$$

5.3 Blocking Counter

We sketch the proof of a blocking counter module: a single cell storing a natural number that can be incremented, guarded by a non-fair lock for concurrent access. The example illustrates how the TaDA Live specifications and proofs neatly support hiding blocking when it is unobservable by the client, while still leaking the requirement of bounded impedance from the lock. This requires any client to only call operations making use of the lock (in this case the `incr` operation) a bounded number of times.

Code. The implementation of the module's operations is:

```

1 def makeCounter(x) {
2   var x, l in
3   x := alloc(2);
4   l := makeLock();
5   [x] := 1;
6   [x + 1] := 0;
7   ret := x;
8 }
1 def incr(x) {
2   var l in
3   l := [x];
4   lock(l);
5   v := [x + 1];
6   [x + 1] := v + 1;
7   unlock(l);
8   ret := v;
9 }
1 def read(x) {
2   var l in
3   l := [x];
4   lock(l);
5   ret := [x + 1];
6   unlock(l);
7 }

```

Specifications. The abstract predicate $C(s, x, n, \alpha)$ represents a blocking counter at address x with value n and impedance bound α .

$$\begin{aligned}
& \forall \alpha. \mathbf{1} \vdash \{\text{emp}\} \text{makeCounter}() \{\exists s. C(s, \text{ret}, 0, \alpha)\} \\
& \forall \phi. \mathbf{1} \vdash \forall n \in \mathbb{N}, \alpha. \langle \text{emp} \mid C(s, x, n, \alpha) \wedge \alpha > \phi(\alpha) \rangle \text{incr}(x) \langle \text{ret} = n \mid C(s, x, n + 1, \phi(\alpha)) \rangle \\
& \mathbf{1} \vdash \forall n \in \mathbb{N}, \alpha. \langle \text{emp} \mid C(s, x, n, \alpha) \rangle \text{read}(x) \langle \text{ret} = n \mid C(s, x, n, \alpha) \rangle
\end{aligned}$$

Shared Regions. This proof will use two region types: $\mathbf{cnt}_r(r', x, s, la, n, \alpha)$ and $\mathbf{lcnt}_{r'}(x, s, la, l, n, \alpha)$ where $r, r' \in \text{RId}$, $x, la \in \text{Addr}$, $l \in \{0, 1\}$, $n \in \mathbb{N}$, $\alpha \in \mathbb{O}$ and s is the abstract location of the lock guarding the counter resource. Here r', x, s and la are the fixed parameters of the regions, representing respectively the region identifier of the inner region, the address of the blocking counter and the abstract location and address of the associated lock.

As in the CLH lock example, we will use two nested regions. The region type \mathbf{lcnt} will be used as an inner region revealing sufficient information to prove desired liveness properties, in particular, exposing the state of the lock, l . The region type \mathbf{cnt} will be used to prove linearizability of our operations; to this end, it only exposes the value of the blocking counter n , and the lock's impedance bound α .

Guards and Obligations. We associate the exclusive guard \mathbf{E} with both \mathbf{cnt} and \mathbf{lcnt} . Besides this, this proof will also require the guards \mathbf{U} , $\mathbf{L}(n, n')$ and $\mathbf{K}(n, n')$, where $n, n' \in \mathbb{N}$, for the latter region. These guards will be used to record the update to the value of the counter that will occur at the moment the module's lock is locked in the proof of `incr`. Since other threads cannot observe the value of the counter without first holding the lock, performing this abstract update on the state of the outer region, \mathbf{cnt} , and then updating the concrete state of the counter before releasing the lock results in a linearizable implementation.

To allow this, once the lock is locked, the concrete value of the counter, $n' \in \mathbb{N}$, and the updated value of the counter, $n \in \mathbb{N}$, are stored in the guard $\mathbf{L}(n, n')$ within the region \mathbf{cnt} . The thread

holding the lock then holds the guard $\mathbf{k}(n, n')$, which keeps a local record of the concrete and updated counter values; the values are required to match with those stored in $\mathbf{l}(n, n')$ within the region by the axiom:

$$\mathbf{l}(n, n') \bullet \mathbf{k}(m, m') \text{ is defined } \Leftrightarrow n = m \wedge n' = m'$$

When the lock is unlocked, the guard \mathbf{u} is stored within the region \mathbf{cnt} . When a thread takes possession of the lock, it can be split into the guards $\mathbf{l}(n, n')$ and $\mathbf{k}(n, n')$ using the axiom:

$$\mathbf{u} = \mathbf{l}(n, n') \bullet \mathbf{k}(n, n')$$

Finally, if a thread holds the guard $\mathbf{k}(n, n')$, it holds the lock, which can be inferred from the axiom:

$$\mathbf{u} \bullet \mathbf{k}(n, n') \text{ is undefined.}$$

This pattern of three guards is often used as a TaDA pattern to encode mutual exclusion on some resource when a thread has possession of a shared lock.

We also associate a single atom obligation \mathbf{k} with the region type \mathbf{lcnt} . This obligation encodes ownership of the blocking counter's lock, as well as the obligation to unlock it. We set $\text{lay}(\mathbf{k}) = \mathbf{0}$.

Region Protocols. The guard-labelled transition system of the region \mathbf{cnt} is:

$$\mathbf{E} : ((n, \alpha), \mathbf{0}) \rightsquigarrow ((n + 1, \beta), \mathbf{0}) \quad \alpha > \beta$$

and the guard-labelled transition system of the region \mathbf{lcnt} is:

$$\begin{aligned} \mathbf{E} &: ((0, n, \alpha), \mathbf{0}) \rightsquigarrow ((1, n, \beta), \mathbf{k}) & \alpha > \beta \\ \mathbf{E} &: ((1, n, \alpha), \mathbf{0}) \rightsquigarrow ((1, n + 1, \alpha), \mathbf{0}) \\ \mathbf{E} &: ((1, n, \alpha), \mathbf{k}) \rightsquigarrow ((0, n, \alpha), \mathbf{0}) \end{aligned}$$

Region Interpretations. The interpretation of the locked counter region \mathbf{lcnt} links the state of the lock and counter to the abstract state of the region and the ownership of \mathbf{k} .

The region \mathbf{cnt} is a wrapper around the \mathbf{lcnt} region that hides the state of the lock and allows the counter value of the region \mathbf{lcnt} to be disconnected from that of the outer region when the lock is locked.

$$\begin{aligned} \mathcal{I}(\mathbf{lcnt}_r(x, s, la, l, n, \alpha)) &\triangleq x \mapsto la, n * \mathcal{L}(s, la, l, \alpha) * (l = 0 \Rightarrow [\mathbf{k}]_r^l) \\ \mathcal{I}(\mathbf{cnt}_r(r', x, s, la, n, \alpha)) &\triangleq \exists n' \in \mathbb{N}, l \in \{0, 1\}. \mathbf{lcnt}_{r'}(x, s, la, l, n', \alpha) * [\mathbf{E}]_{r'} \\ &\quad * ((l = 0 \wedge n = n' \wedge [\mathbf{U}]_r) \vee (l = 1 \wedge [\mathbf{l}(n, n')]_r * [\mathbf{k}]_{r'}^E)) \end{aligned}$$

Predicates. The counter resource is abstractly represented by the predicate

$$\mathbf{C}((r, r', s, la), x, n, \alpha) \triangleq \mathbf{cnt}_r(r', x, s, la, n, \alpha) * [\mathbf{E}]_r$$

Verification of incr. The proof of incr can be found in Fig. 24. The only step requiring liveness reasoning is the call $\text{lock}(x)$, which is handled very similarly to the same call in the left thread of the distinguishing client where the environment liveness condition of the LIVEC rule application is discharged using the fact that when $l = 1$ holds, then $[\mathbf{k}]_{r'}^E$, which, in this case, is obtained from the interpretation of the outer region, \mathbf{cnt} . The details of the proof of the lock operation can be found in Fig. 23.

$$\begin{array}{l}
\{ \exists n, \alpha, l. \text{cnt}_r(r', x, s', 1, n, \alpha) * r \Rightarrow \blacklozenge \text{lcnt}_{r'}(x, s', 1, l, _ _) * l = 1 \Rightarrow [\mathbf{K}]_{r'}^E \wedge \alpha > \phi(\alpha) \} \\
\left\langle \begin{array}{l}
\mathbb{W}n \in \mathbb{N}, \alpha. \\
\langle \exists l. \text{lcnt}_{r'}(x, s', 1, l, _ _) * l = 1 \Rightarrow [\mathbf{K}]_{r'}^E \mid \text{cnt}_r(r', x, s', 1, n, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha) \rangle \\
\left\langle \exists l. \text{lcnt}_{r'}(x, s', 1, l, _ _) * l = 1 \Rightarrow [\mathbf{K}]_{r'}^E \mid \left(\begin{array}{l}
\exists n', l. \text{lcnt}_{r'}(x, s', 1, l, n', \alpha) * [\mathbf{E}]_{r'} * \\
(l = 0 \wedge n = n' \wedge [\mathbf{U}]_r) \vee \\
(l = 1 \wedge [\mathbf{L}(n, n')]_r * [\mathbf{K}]_{r'}^E) \end{array} \right) \wedge \alpha > \phi(\alpha) \right\rangle \\
\left\langle \mathbb{W}n, n' \in \mathbb{N}, l \in \{0, 1\}, \alpha. \right. \\
\left. \left\langle \exists l. \text{lcnt}_{r'}(x, s', 1, l, _ _) * l = 1 \Rightarrow [\mathbf{K}]_{r'}^E \mid \left(\begin{array}{l}
\text{lcnt}_{r'}(x, s', 1, l, n', \alpha) * [\mathbf{E}]_{r'} * \\
(l = 0 \wedge n = n' \wedge [\mathbf{U}]_r) \vee \\
(l = 1 \wedge [\mathbf{L}(n, n')]_r * [\mathbf{K}]_{r'}^E) \end{array} \right) \wedge \alpha > \phi(\alpha) \right\rangle \right. \\
\left. \left\langle \mathbb{W}n, n' \in \mathbb{N}, l \in \{0, 1\} \rightarrow_0 \{0\}, \alpha. \right. \right. \\
\left. \left\langle \text{lcnt}_{r'}(x, s', 1, l, n', \alpha) * [\mathbf{E}]_{r'} * \left(\begin{array}{l}
(l = 0 \wedge n = n' \wedge [\mathbf{U}]_r) \vee \\
(l = 1 \wedge [\mathbf{L}(n, n')]_r \wedge [\mathbf{K}]_{r'}^E) \end{array} \right) \wedge \alpha > \phi(\alpha) \right\rangle \right. \\
\left. \left\langle \mathbb{W}l \in \{0, 1\} \rightarrow_0 \{0\}, \alpha. \right. \right. \\
\left. \left\langle \mathbf{L}(s', 1, l, \alpha) \wedge \alpha > \phi(\alpha) \right\rangle \right. \\
\left. \text{lock}(1); \right. \\
\left. \left\langle \mathbf{L}(s', 1, 1, \alpha) \wedge l = 0 \right\rangle \right. \\
\left. \left\langle [\mathbf{K}]_{r'}^L \mid \text{lcnt}_{r'}(x, s', 1, 1, n', \alpha) * [\mathbf{E}]_{r'} * \left(\begin{array}{l}
(l = 0 \wedge n = n' \wedge [\mathbf{U}]_r) \vee \\
(l = 1 \wedge [\mathbf{L}(n, n')]_r \wedge [\mathbf{K}]_{r'}^E) \end{array} \right) \wedge l = 0 \right\rangle \right. \\
\left. \left\langle [\mathbf{K}]_{r'}^L \mid \text{lcnt}_{r'}(x, s', 1, 1, n', \alpha) * [\mathbf{E}]_{r'} * [\mathbf{U}]_r \wedge n = n' \right\rangle \right. \\
\left. \left\langle [\mathbf{K}]_{r'}^L \mid \exists n', l. \text{lcnt}_{r'}(x, s', 1, 1, n', \alpha) * [\mathbf{E}]_{r'} * \left(\begin{array}{l}
(l = 0 \wedge n + 1 = n' \wedge [\mathbf{U}]_r) \vee \\
(l = 1 \wedge [\mathbf{L}(n + 1, n')]_r \wedge [\mathbf{K}]_{r'}^E) \end{array} \right) * [\mathbf{K}(n + 1, n)]_r \right\rangle \right. \\
\left. \langle [\mathbf{K}]_{r'}^L \mid \text{cnt}_r(r', x, s', 1, n + 1, \alpha) * r \Rightarrow \langle (n, \alpha), (n + 1, \phi(\alpha)) \rangle * [\mathbf{K}(n + 1, n)]_r \right\rangle \\
\{ \exists n, \alpha. \text{cnt}_r(r', x, s', 1, n + 1, \phi(\alpha)) * r \Rightarrow \langle (n, \alpha), (n + 1, \phi(\alpha)) \rangle * [\mathbf{K}]_{r'} * [\mathbf{K}(n + 1, n)]_r \}
\end{array} \right.
\end{array}$$

Fig. 23. Details of the proof of the lock(1) call of incr. STEP 6 is LIFTA, FRAME.

Verification of the makeCounter and read operations. The proof of makeCounter proceeds, using standard steps on Hoare triples, by establishing the postcondition $\exists x, la, lr, \alpha. x \mapsto la, 0 * \mathbf{L}(lr, la, 0, \alpha)$ which can be viewshifted to $\exists x, la, r, r', lr, \alpha. \mathbf{C}((r, r', lr, la), x, 0, \alpha)$.

The proof of read is almost identical to the proof in Fig. 24. The reader might wonder if the lock acquisition in the code is strictly necessary. Indeed, it is not given the current set of operations available to the client. To prove the version where read does not acquire the lock, however, we would need to change the region's protocol to encode the fact that while holding a lock a single write to it is possible. Since one would conceivably want to extend the module with other operations that write to the counter multiple times while holding the lock, we formalised the more general protocol. In the presence of such additional operations, read would need to acquire the lock to be correct.

5.4 Double Blocking Counter

We now develop the proof of a double blocking counter module, that is, a module encapsulating two integers each protected by a fair lock. The module offers linearizable operations to increment/read each counter in isolation and an incrBoth operation to atomically increment both. The implementation of incrBoth needs to deal with the ubiquitous pattern of locking multiple locks in a nested fashion, which is one of the most common sources of deadlocks in coarse-grained concurrent programs. The example illustrates how the specification format and layer system of TaDA Live allow for modular proofs of deadlock-freedom. In particular, verifying the example in LiLi would require: (i) replacing the calls to the lock operations with some non-atomic abstract code (ii) building a termination argument that talks about the queues of the two fair locks; in particular the variant argument would need to consider both queues at the same time and argue about all the possible ways the threads in the environment may enter and exit both queues. We

 PROOF OF $\text{incr}(x)$:

$$\begin{array}{l}
 \forall \phi. \mathbf{1}; \mathbf{0} \vdash \forall n \in \mathbb{N}, \alpha. \\
 \langle \text{emp} \mid C(s, x, n, \alpha) \wedge \alpha > \phi(\alpha) \rangle \\
 \left\{ \begin{array}{l}
 \langle \text{emp} \mid \text{cnt}_r(r', x, s', la, n, \alpha) * [\mathbf{E}]_r \wedge \alpha > \phi(\alpha) \rangle \\
 \mathbf{1}; [r \mapsto (\mathbb{N} \times \mathbf{0}, \mathbf{0}, \mathbb{N} \times \mathbf{0}, \{(n, \alpha), (n+1, \beta)\} \mid n \in \mathbb{N}, \alpha, \beta \in \mathbf{0}, \alpha > \beta)] \vdash \\
 \{ \exists n, \alpha. \text{cnt}_r(r', x, s', la, n, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \phi(\alpha) \} \\
 l := [x]; \\
 \{ \exists n, \alpha, l. \text{cnt}_r(r', x, s', l, n, \alpha) * r \Rightarrow \blacklozenge * \text{lcnt}_{r'}(x, s', l, l, _ , _) * l = 1 \Rightarrow [\mathbf{K}]_{r'}^E \wedge \alpha > \phi(\alpha) \} \\
 \text{lock}(1); \\
 \{ \exists n, \alpha. \text{cnt}_r(r', x, s', l, n+1, \phi(\alpha)) * r \Rightarrow \langle (n, \alpha), (n+1, \phi(\alpha)) \rangle * [\mathbf{K}]_{r'}^L * [\mathbf{K}(n+1, n)]_r \} \\
 v := [x + 1]; \\
 \{ \exists n, \alpha. \text{cnt}_r(r', x, s', l, n+1, \phi(\alpha)) * r \Rightarrow \langle (n, \alpha), (n+1, \phi(\alpha)) \rangle * [\mathbf{K}]_{r'}^L * [\mathbf{K}(n+1, n)]_r \wedge v = n \} \\
 [x + 1] := v + 1; \\
 \{ \exists n, \alpha. \text{cnt}_r(r', x, s', l, n+1, \phi(\alpha)) * r \Rightarrow \langle (n, \alpha), (n+1, \phi(\alpha)) \rangle * [\mathbf{K}]_{r'}^L * [\mathbf{K}(n+1, n+1)]_r \wedge v = n \} \\
 \text{unlock}(1); \\
 \{ \exists n, \alpha. r \Rightarrow \langle (n, \alpha), (n+1, \phi(\alpha)) \rangle \wedge v = n \} \\
 \text{ret} := v; \\
 \{ \exists n, \alpha. r \Rightarrow \langle (n, \alpha), (n+1, \phi(\alpha)) \rangle \wedge \text{ret} = n \} \\
 \langle \text{ret} = n \mid \text{cnt}_r(r', x, s', la, n+1, \phi(\alpha)) * [\mathbf{E}]_r \rangle \\
 \langle \text{ret} = n \mid C(s, x, n+1, \phi(\alpha)) \rangle
 \end{array} \right. \\
 \text{CONS}; s = (r, r', s', la) \\
 \text{MkATOM}
 \end{array}$$

Fig. 24. Blocking counter: proof of incr .

avoid these complications by: (i) reusing the (fair) lock specifications which are truly atomic and properly hide the queues (ii) arguing about termination by means of two obligations with layers the order of which reflect the order of acquisition of locks. These obligations only represent the liveness invariant that each lock is always eventually released, the layers represent the dependency between the two locks. The proof requires no detail about why, thanks to the internal queues, this is sufficient to ensure global progress: that part of the argument has already been made in proving the lock specifications!

Code. The implementation of the module’s operations is in Fig. 25 using the following abbreviations for readability:

$$x.\text{lock1} \triangleq [x] \quad x.\text{lock2} \triangleq [x+1] \quad x.\text{cnt1} \triangleq [x+2] \quad x.\text{cnt2} \triangleq [x+3]$$

Specifications. The fair lock module specifications assumed in this example are

$$\begin{array}{l}
 \mathbf{1}_r \vdash \forall l \in \{0, 1\} \rightarrow_{\mathbf{0}_r} \{0\}. \langle L(r, x, l) \rangle \text{lock}(x) \langle L(r, x, 1) \wedge l = 0 \rangle \\
 \mathbf{0}_r \vdash \langle L(r, x, 1) \rangle \text{unlock}(x) \langle L(r, x, 0) \rangle
 \end{array}$$

where $\mathbf{1}_r$ and $\mathbf{0}_r$ are layers parametrised on the region identifier r of the shared lock. It is a common TaDA Live pattern to parametrise the layers of specifications so that they can be instantiated with differently for each instance of the module. In Section 5.5 we explain this parametrisation in general, and how to parametrise the implementation proof accordingly.

The abstract predicate $\text{DC}(t, x, n, m)$ represents a double counter at address x with abstract location t and values n and m respectively. We wish to show the implementations of the


```

1 def makeDCounter() {      1 def incr1(x) {          1 def incr2(x) {          1 def incrBoth(x) {
2   var x,l1,l2 in        2   var l,v in          2   var l,v in          2   var l1,l2,v in
3   x := alloc(4);        3   l := x.lock1;      3   l := x.lock2;      3   l1 := x.lock1;
4   l1 := makeLock();    4   lock(l);           4   lock(l);           4   l2 := x.lock2;
5   l2 := makeLock();    5   v := x.cnt1;       5   v := x.cnt2;       5   lock(l1);
6   x.lock1 := l1;       6   x.cnt1 := v + 1;   6   x.cnt2 := v + 1;   6   lock(l2);
7   x.lock2 := l2;       7   unlock(l);         7   unlock(l);         7   v := x.cnt1;
8   x.cnt1 := 0;         8   ret := v           8   ret := v           8   x.cnt1 := v + 1;
9   x.cnt2 := 0;         9   }                  9   }                  9   v := x.cnt2;
10  ret := x              10                      10                      10  x.cnt2 := v + 1;
11 }                      11                      11                      11  unlock(l2);
                          12                      12                      12  unlock(l1);
                          13                      13                      13 }

```

Fig. 25. Code of the double blocking counter operations.

module's operations satisfy the following specifications:

$$\begin{aligned}
& 1 \vdash \{\text{emp}\} \text{makeDCounter}() \{\exists t. \text{DC}(t, \text{ret}, 0, 0)\} \\
& 1 \vdash \forall n, m \in \mathbb{N}. \langle \text{DC}(t, x, n, m) \rangle \text{incrBoth}(x) \langle \text{DC}(t, x, n+1, m+1) \rangle \\
& 1 \vdash \forall n, m \in \mathbb{N}. \langle \text{emp} \mid \text{DC}(t, x, n, m) \rangle \text{incr1}(x) \langle \text{ret} = n \mid \text{DC}(t, x, n+1, m) \rangle \\
& 1 \vdash \forall n, m \in \mathbb{N}. \langle \text{emp} \mid \text{DC}(t, x, n, m) \rangle \text{incr2}(x) \langle \text{ret} = m \mid \text{DC}(t, x, n, m+1) \rangle
\end{aligned}$$

It is important to note here that we are making explicit the parametrisation of the layers in the region identifiers s , because we will need to associate different layers with the two instances of the lock. As we will see later, we will have two region identifiers s_1 and s_2 , one per lock, with associated layers $1_{s_1}, 0_{s_1}, 1_{s_2}, 0_{s_2}$. The lock specifications themselves only require $1_{s_1} > 0_{s_1}$ and $1_{s_2} > 0_{s_2}$ but we will additionally impose, for this client proof, $0_{s_1} > 1_{s_2}$. This represents the fact that, in this client, the release of lock 1 will depend on the acquisition of lock 2.

Shared Regions. Like for the single counter example, we need two nested regions, one to prove the atomicity of the operation (**dcnt**) and an inner one to prove termination (**ldcnt**). They differ in that **dcnt** only records the abstract states of the counters, while **ldcnt** includes the abstract states of the locks. Formally: $\text{dcnt}_{r_1}((r_0, t_0), x, n, m)$ and $\text{ldcnt}_{r_0}(t_0, x, l_1, l_2, n, m)$ where $r_0, r_1 \in \text{Rld}$, $x \in \text{Addr}$, $l_1, l_2 \in \{0, 1\}$ and $n, m \in \mathbb{N}$, and t_0 is a tuple (la_1, la_2, s_1, s_2) with $la_1, la_2 \in \text{Addr}$ and $s_1, s_2 \in \text{Rld}$. Here $(r_0, t_0), x$, and t_0, x are the fixed parameters of the two regions respectively. The double blocking counter resource is abstractly represented by the predicate $\text{DC}((r_1, t_1), x, n, m) \triangleq \text{dcnt}_{r_1}(t_1, x, n, m) * \lceil \text{E} \rceil_{r_1}$.

Guards and Obligations. We introduce the guard constructors \mathbf{B}_i , \mathbf{C}_i , and \mathbf{W}_i , for $i \in \{1, 2\}$, for bookkeeping of the value of the counters. We need this ghost state because in **incrBoth** there is an intermediate state where one counter has been updated but the other hasn't; we cannot update the abstract state in two steps because we are proving atomicity of the operation, so we need to update both counter values in the abstract state in one go. We record the intermediate concrete state in these guards so the information is there locally without affecting the shared abstract state prematurely. The guard composition satisfies the axioms

$$\mathbf{B}_1 = \mathbf{C}_1(n, n') \bullet \mathbf{W}_1(n, n') \qquad \mathbf{B}_2 = \mathbf{C}_2(n, n') \bullet \mathbf{W}_2(n, n')$$

Here $c_i(n, n')$ tracks the reference value (left in the region interpretation) for the i -th counter’s abstract (n) and concrete (n') value and w_i is a local “witness” for the same information about the i -th counter, which can only be obtained when locking the i -th lock (otherwise it would not be stable information). This is enforced by the interpretation given later.

We associate two atom obligations \mathbf{k}_1 and \mathbf{k}_2 with the region type \mathbf{ldcnt} , encoding ownership of the double counter’s locks respectively, as well as the obligation to unlock them.

As anticipated, we choose the layers of the lock specifications in a way that represents the dependency between the two locks. We have a (double-counter-local) top ($\mathbf{1}$) and a bottom ($\mathbf{0}$) layer, and intermediate layers for the locks:¹⁷

$$\mathbf{0} = \mathbf{0}_{s_2} = \text{lay}(\mathbf{k}_2) < \mathbf{1}_{s_2} < \mathbf{0}_{s_1} = \text{lay}(\mathbf{k}_1) < \mathbf{1}_{s_1} = \mathbf{1}$$

Region Protocols. The interference protocol of the region \mathbf{dcnt} trivially allows for any change to the counter values:

$$\mathbf{E} : ((n, m), \mathbf{0}) \rightsquigarrow ((n', m'), \mathbf{0})$$

The interference protocol of the region \mathbf{ldcnt} encodes the constraint that we can update a counter only by holding the corresponding lock:

$$\begin{array}{ll} \mathbf{E} : ((0, l, n, m), \mathbf{0}) \rightsquigarrow ((1, l, n, m), \mathbf{k}_1) & \mathbf{E} : ((l, 0, n, m), \mathbf{0}) \rightsquigarrow ((l, 1, n, m), \mathbf{k}_2) \\ \mathbf{E} : ((1, l, n, m), \mathbf{k}_1) \rightsquigarrow ((0, l, n, m), \mathbf{0}) & \mathbf{E} : ((l, 1, n, m), \mathbf{k}_2) \rightsquigarrow ((l, 0, n, m), \mathbf{0}) \\ \mathbf{E} : ((1, l, n, m), \mathbf{k}_1) \rightsquigarrow ((1, l, n', m), \mathbf{k}_1) & \mathbf{E} : ((l, 1, n, m), \mathbf{k}_2) \rightsquigarrow ((l, 1, n, m'), \mathbf{k}_2) \end{array}$$

Region Interpretations. The interpretation of \mathbf{dcnt} formalises the fact that the outer region simply hides the state of the locks for the atomicity argument, while the actual internal protocol of the module is encoded in the interpretation of the inner region \mathbf{ldcnt} :

$$\begin{aligned} \mathcal{I}(\mathbf{dcnt}_{r_1}((r_0, t_0), x, n, m)) &\triangleq \exists l_1, l_2 \in \{0, 1\}. \mathbf{ldcnt}_{r_0}(t_0, x, l_1, l_2, n, m) * \lceil \mathbf{E} \rceil_{r_0} * \\ &\quad l_1 = 1 \Rightarrow \lceil \mathbf{k}_1 \rceil_{r_0}^E * l_2 = 1 \Rightarrow \lceil \mathbf{k}_2 \rceil_{r_0}^E \\ \mathcal{I}(\mathbf{ldcnt}_{r_0}((la_1, la_2, s_1, s_2), x, l_1, l_2, n, m)) &\triangleq \exists n', m' \in \mathbb{N}. \\ &\quad x \mapsto la_1, la_2, n', m' * \mathbf{L}(s_1, la_1, l_1) * \mathbf{L}(s_2, la_2, l_2) \\ &\quad * \left(\begin{array}{l} (l_1 = 0 \wedge \lceil \mathbf{k}_1 \rceil_{r_0}^L * \lceil \mathbf{B}_1 \rceil_{r_0} \wedge n = n') \\ \vee (l_1 = 1 \wedge \lceil \mathbf{c}_1(n, n') \rceil_{r_0}) \end{array} \right) \\ &\quad * \left(\begin{array}{l} (l_2 = 0 \wedge \lceil \mathbf{k}_2 \rceil_{r_0}^L * \lceil \mathbf{B}_2 \rceil_{r_0} \wedge m = m') \\ \vee (l_2 = 1 \wedge \lceil \mathbf{c}_2(m, m') \rceil_{r_0}) \end{array} \right) \end{aligned}$$

Proof of incrBoth. The proof outline of incrBoth is reproduced in Fig. 26. Most of the proof is routine; the derivation for the acquisition of the first lock follows closely the pattern we already explained in Sections 4 and 5.3. We show the proof of the acquisition of the second lock in more detail, to show the interplay between the layers. At that point we are continuously holding the obligation of the first lock, with layer greater than $\mathbf{1}_{s_2}$, so apply LAYWH to lower the layer to $\mathbf{1}_{s_2}$ enabling the application of FRAME to frame $r_1 \Rightarrow \blacklozenge * \lceil \mathbf{k}_1 \rceil_{r_0}^L * \lceil \mathbf{w}_1(n, n) \rceil_{r_0}$. The obligation \mathbf{k}_2 has layer lower than $\mathbf{1}_{s_2}$ so we are allowed to invoke it to discharge the environment liveness condition of the LIVEC application, in a way that is analogous to the derivations of the distinguishing client and Section 5.3.

¹⁷The proof works with $\mathbf{1}_{s_2} = \mathbf{0}_{s_1}$ too, but the ordered version better emphasizes the dependency between the locks.

PROOF OF $\text{incrBoth}(x)$:

$1; 0 \vdash \forall n, m \in \mathbb{N}.$ $\langle \text{emp} \mid \text{DC}(t, x, n, m) \rangle$	$\langle \text{dcnt}_{r_1}(t_1, x, n, m) * \lceil \mathbb{E} \rceil_{r_1} \rangle$ $1; \mathcal{A} \triangleq [r_1 \mapsto (\mathbb{N}^2, 0, \mathbb{N}^2, \{(n, m), (n+1, m+1)\} \mid n, m \in \mathbb{N})] \vdash$ $\{ \exists n, m. \text{dcnt}_{r_1}(t_1, x, n, m) * r_1 \Rightarrow \blacklozenge * \}$ $l1 := [x];$ $l2 := [x + 1];$ $// t'_1 \triangleq (r_0, t'_0), t'_0 \triangleq (l1, l2, s_1, s_2)$ $\{ \exists n, m. \text{dcnt}_{r_1}(t'_1, x, n, m) * r_1 \Rightarrow \blacklozenge * \}$ $\{ \exists l_1, l_2. \text{ldcnt}_{r_0}(t'_0, x, l_1, l_2, _ _) * l_1 = 1 \Rightarrow \lfloor \mathbb{K}_1 \rfloor_{r_0}^E * l_2 = 1 \Rightarrow \lfloor \mathbb{K}_2 \rfloor_{r_0}^E \}$ $\text{lock}(l1);$ $\{ \exists n, m. \text{dcnt}_{r_1}(t'_1, x, n, m) * r_1 \Rightarrow \blacklozenge * \lfloor \mathbb{K}_1 \rfloor_{r_0}^L * \lceil \mathbb{W}_1(n, n) \rceil_{r_0} * \}$ $\{ \exists l_2. \text{ldcnt}_{r_0}(t'_0, x, _ _, l_2, _ _) * l_2 = 1 \Rightarrow \lfloor \mathbb{K}_2 \rfloor_{r_0}^E \}$
$\text{CONS}; \text{Sub } t = (r_1, t_1), t_1 = (r_0, t_0), t_0 = (l_{a1}, l_{a2}, s_1, s_2)$	$1_{s_2}; \mathcal{A} \vdash$ $\forall n, m \in \mathbb{N}.$ $\langle \exists l_2. \text{ldcnt}_{r_0}(t'_0, x, _ _, l_2, _ _) * l_2 = 1 \Rightarrow \lfloor \mathbb{K}_2 \rfloor_{r_0}^E \mid \text{dcnt}_{r_1}(t'_1, x, n, m) \rangle$ $\forall l_1, l_2 \in \{0, 1\}.$ $\langle \exists l_2. \text{ldcnt}_{r_0}(t'_0, x, _ _, l_2, _ _) * l_2 = 1 \Rightarrow \lfloor \mathbb{K}_2 \rfloor_{r_0}^E \mid \text{ldcnt}_{r_0}(t'_0, x, l_1, l_2, n, m) * \lceil \mathbb{E} \rceil_{r_0} \rangle$
MKATOM	$1_{s_2}; \mathcal{A} \vdash$ $\forall n, m \in \mathbb{N}, l_1 \in \{0, 1\}, l_2 \in \{0, 1\} \rightarrow_{0, s_2} \{0\}.$ $\langle \text{ldcnt}_{r_0}(t'_0, x, l_1, l_2, n, m) * \lceil \mathbb{E} \rceil_{r_0} * l_2 = 1 \Rightarrow \lfloor \mathbb{K}_2 \rfloor_{r_0}^E \rangle$
$\text{LAWWH}; \text{FRAME}; \text{ATOMW}; \text{A} \in \text{ELIM}$	$1_{s_2}; \mathcal{A} \vdash$ $\forall l_2 \in \{0, 1\} \rightarrow_{0, s_2} \{0\}.$ $\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{ldcnt}_{r_0}(t'_0, x, l_1, 1, n, m) * \lceil \mathbb{E} \rceil_{r_0} * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$ $\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{ldcnt}_{r_0}(t'_0, x, l_1, 1, n, m) * \lceil \mathbb{E} \rceil_{r_0} * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$ $\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{dcnt}_{r_1}(t'_1, x, n, m) * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$
$\text{LIFTA}; \text{A} \in \text{ELIM}; \text{FRAME}$	$\{ \exists n, m. \text{dcnt}_{r_1}(t'_1, x, n, m) * r_1 \Rightarrow \blacklozenge * \}$ $\{ \lfloor \mathbb{K}_1 \rfloor_{r_0}^L * \lceil \mathbb{W}_1(n, n) \rceil_{r_0} * \lfloor \mathbb{K}_2 \rfloor_{r_0}^L * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \}$
LIVEC	$v := x.\text{cnt1}; x.\text{cnt1} := v + 1; v := x.\text{cnt2}; x.\text{cnt2} := v + 1;$ $\{ \exists n, m. \text{dcnt}_{r_1}(t'_1, x, n, m) * r_1 \Rightarrow \blacklozenge * \}$ $\{ \lfloor \mathbb{K}_1 \rfloor_{r_0}^L * \lceil \mathbb{W}_1(n, n+1) \rceil_{r_0} * \lfloor \mathbb{K}_2 \rfloor_{r_0}^L * \lceil \mathbb{W}_2(m, m+1) \rceil_{r_0} \}$
STEP 7	$\text{lock}(l2);$ $\langle \mathbb{L}(s_2, l2, l_2) \rangle$ $\langle \mathbb{L}(s_2, l2, 1) \wedge l_2 = 0 \rangle$
FRAME	$\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{ldcnt}_{r_0}(t'_0, x, l_1, 1, n, m) * \lceil \mathbb{E} \rceil_{r_0} * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$ $\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{ldcnt}_{r_0}(t'_0, x, l_1, 1, n, m) * \lceil \mathbb{E} \rceil_{r_0} * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$ $\langle \lfloor \mathbb{K}_2 \rfloor_{r_0}^L \mid \text{dcnt}_{r_1}(t'_1, x, n, m) * \lceil \mathbb{W}_2(m, m) \rceil_{r_0} \rangle$
STEP 7	$\text{unlock}(l2);$ $\{ \exists n, m. \text{dcnt}_{r_1}(t'_1, x, n, m) * r_1 \Rightarrow ((n, m), (n+1, m+1)) * \}$ $\{ \lfloor \mathbb{K}_1 \rfloor_{r_0}^L * \lceil \mathbb{W}_1(n+1, n+1) \rceil_{r_0} * \lfloor \mathbb{K}_2 \rfloor_{r_0}^L * \lceil \mathbb{W}_2(m, m+1) \rceil_{r_0} \}$
FRAME	$\text{unlock}(l1);$ $\{ \exists n, m. r_1 \Rightarrow ((n, m), (n+1, m+1)) \}$
STEP 7	$\langle \text{dcnt}_{r_1}(t_1, x, n+1, m+1) * \lceil \mathbb{E} \rceil_{r_1} \rangle$
STEP 7	$\langle \text{DC}(t, x, n+1, m+1) \rangle$

Fig. 26. Double blocking counter: proof of incrBoth .

STEP 7 is LIFTA, FRAME.

A comparison with LiLi. As we have seen in Section 2 (Innovation 3), the call of a CLH lock in LiLi involves two distinct atomic actions: requesting the lock, and acquiring it. Requesting a lock x is a non-blocking action as it just enqueues the current thread in the (concrete) queue for x , but the acquisition is represented with a (primitive) blocking operation that waits until the current thread is at the head of the lock’s queue, and the lock is unlocked. When proving the call to `lock(l1)` in `incrBoth`, the LiLi proof would require arguing about termination of acquisition by appealing to progress of the threads in the environment.

To do so, in the LiLi methodology, one has to identify the threads in the environment that will be able to make progress, and show how this progress is bringing us closer to acquiring lock `l1`. Consider the case when there are $n_1 > 0$ threads ahead of us in the queue for `l1`. Assume thread t_1 is the head of the queue for `l1`. It can make progress in three ways:

- if `l1` is unlocked it can acquire it;
- if `l1` is locked it can unlock it;
- if `l1` is locked it can request `l2`.

How do these actions represent progress for us? The first case makes progress by moving to the second or third case. The second case removes t_1 from the queue of `l1` bringing us closer to the front of the queue. The third case complicates matters: in this case t_1 is enqueued in the queue of `l2` with a non-deterministic number n_2 of threads ahead of it. The thread t_1 is now blocked, and to track progress we need to consider the head of the queue for `l2`, which can only make progress by acquiring the lock when unlocked, or releasing the lock when locked. What progress had been made towards us acquiring `l1`? The measure of progress needs to consider the contents of the queues for both threads: the measure before t_1 requests `l2` needs to be (n_1, ω) (ordered lexicographically) so that we can lower the measure to (n_1, n_2) once t_1 joined the queue of `l2`. Whenever t_1 reaches, finally, the head of the queue of `l2`, the measure of progress would become $(n_1, 0)$, and the only option for t_1 is to release `l2`. Now thread t_1 is back to the three options as above. This is a problem because nothing would prevent t_1 from requesting `l2` again. This could repeat ad libitum, leaving us to starve on `l1`. To rule this out, the argument needs to place a bound b on the number of times `l2` can be acquired while holding `l1`; in our example this bound can be 1. By mixing this bound in the measure (n_1, b, n_2) , the action of t_1 releasing `l2` brings real progress by taking b from 1 to 0. When that happens, the only option for t_1 is to release the lock. This brings down n_1 , the number of threads ahead of us; at the same time we want to reset n_2 to ω and b to 1 to allow the new head of the queue of `l1` to request `l2`.

This substantiates our claim that LiLi’s rely/guarantee reasoning lacks in scalability; the key reason for this is that the progress argument is forced to walk through all the possible ways the environment could be implementing progress. This in turn requires to expose the internal state of both locks (their queues) to be used in the client’s proof. In other words, the abstraction of the environment is not abstract enough. By comparison, TaDA Live’s atomic specifications allow for the termination of the lock calls in the double blocking counter to be reasoned about individually, without direct reference to the termination of the other, nor to internal state, using layers to prevent circular reasoning. The appeal to obligation κ_1 being live to justify why the call to `lock(l1)` terminates, abstracts away *how* the environment may be keeping it live. The layers capture the essential information: the only thing that is important is that to keep κ_2 live, the environment does not assume κ_1 live.

5.5 Lock-Coupling Set

To conclude this series of examples, we present a challenging fine-grained lock-based implementation of a linearizable finite set. A lock-coupling set implements a set by maintaining an ordered

linked list of the elements with fair locks (here CLH locks) guarding each individual element. The module exposes an `add` and `remove` operation to add and remove elements from the abstract set it represents. To make modifications to the nodes of the linked list, the operations traverse the list using a lock-coupling pattern. In this pattern, all threads start the traversal at the head of the list. To be at position i a thread must acquire the lock at that position. To move to position $i + 1$ the thread would first acquire the lock at $i + 1$, and then release the lock at position i . This way, the threads cannot overtake each other, and owning a lock allows the owner to safely perform modifications at that position. We sketch here the main points of interest of our proof, the full details can be found in Appendix C.

This example is challenging because it makes use of a dynamically changing list of locks with non-trivial liveness dependencies between them. In particular, the termination of the acquisition of each lock depends on the usage of the locks further down the list. Although these dependencies are acyclic, they change over time as the list grows or shrinks. At first sight, it is unclear how the seemingly static layer structure of TaDA Live, and the fixed layers decorating the specifications of lock operations can cope with this complexity, without breaking modularity.

The TaDA Live proof of this example relies on solving two key challenges:

- How can we modularly coordinate the choice of layers needed for the proof of a module and the ones needed for the proofs of its clients?
- How can we dynamically reassign layers to resources?

We solve the first challenge by introducing a style of specification that allows the client to “remap” the layers of the implementation into a larger layer structure, and the implementation to prove correctness with respect to a “local” layer structure which is opaque to the client. The key observation is that a TaDA Live derivation’s validity is preserved by transformations of the layer structure that preserve the strict order between layers. This leads to the following proof style. Given two partial orders $(\mathcal{L}_1, \leq_1, \top_1, \perp_1)$ and $(\mathcal{L}_2, \leq_2, \top_2, \perp_2)$, a function $\eta: \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is *strictly monotone* if $\forall m, n \in \mathcal{L}_1. m <_1 n \Rightarrow \eta(m) <_2 \eta(n)$. A *layer map* $\eta: \mathcal{L}_1 \rightarrow_{\text{lay}} \mathcal{L}_2$ is a strictly monotone function between the two partial orders. Using this notion, we generalise the client-facing CLH lock specifications as follows:

$$\begin{aligned} & \exists (\mathcal{L}_{\text{clh}}, \leq_{\text{clh}}, \top_{\text{clh}}, \perp_{\text{clh}}). \forall \eta: \mathcal{L}_{\text{clh}} \rightarrow_{\text{lay}} \mathcal{L}. \\ & \eta(\top_{\text{clh}}) \vdash \mathbf{W}l \in \{0, 1\} \Rightarrow_{\eta(\perp_{\text{clh}})} \{0\}. \langle L_\eta(s, x, l) \rangle \text{lock}(x) \langle L_\eta(s, x, 1) \wedge l = 0 \rangle \\ & \eta(\perp_{\text{clh}}) \vdash \langle L_\eta(s, x, 1) \rangle \text{unlock}(x) \langle L_\eta(s, x, 0) \rangle \end{aligned}$$

From the perspective of the implementation, a proof of correctness would start by defining the partial order of the “internal” layers. In the case of CLH, as we have seen in Section 5.2, we would let $\mathcal{L}_{\text{clh}} = \mathbb{N} \cup \{1, 0\}$ with $\top_{\text{clh}} = 1$ and $\perp_{\text{clh}} = 0$. Then, to be able to prove the triples with the layers remapped by the arbitrary layer map η we would reproduce the derivation presented in Section 5.2 but with η applied to every occurrence of an internal layer. For example, the `lclh` region type would also be parametrised by the layer map, `lclhr(η, x, h, l, o, t)`, so that its associated obligations and their layers can depend on η , e.g. `lay($p_\eta(t)$) = $\eta(t)$` . Since the map preserves the strict order of \mathcal{L}_{clh} , the proof goes through exactly as in the un-parametrised case.

From the perspective of the client, to use these specifications one would first obtain the arbitrary \mathcal{L}_{clh} from the existential quantification. Then the client would be able to choose a layer map from \mathcal{L}_{clh} to \mathcal{L} . Here \mathcal{L} could be the global layer structure, in the case of a closed proof, or itself being the internal layer structure of a module using the lock module internally. Note that the client needs to define η parametrically on \mathcal{L}_{clh} , since it has no control on the inner structure of \mathcal{L}_{clh} . For example, in the case of a client with a static list of locks, one would use as \mathcal{L} the lexicographically ordered set of pairs from $(\mathbb{N} \cup \{\top, \perp\}) \times \mathcal{L}_{\text{clh}}$ where the first component corresponds to the position

of the lock from the end of the list. Then, for the lock at position $i \in \mathbb{N}$, the client would instantiate the specifications choosing $\eta_i(k) \triangleq (i, k)$.

The second challenge is also solved by a slight generalisation of the lock specifications, following a proof pattern that, if adopted, always increases the generality of module specifications: adding some fractional permissions to control the update of ghost parameters of the resource. The idea is that the layer map is ghost state, and as such we should be able to update it using a viewshift. To do this without invalidating the other thread’s information about the region we are updating, we add standard fractional permissions to the lock specifications. We introduce the abstract predicate $P(s, \pi)$ representing ownership of the fraction $0 \leq \pi \leq 1$ of permissions for a lock at abstract location s . To split permissions, the predicate satisfies, for $0 \leq \pi_1 + \pi_2$, $\pi_1, \pi_2 \leq 1$, $P(s, \pi_1 + \pi_2) \Leftrightarrow P(s, \pi_1) * P(s, \pi_2)$. The generalised lock specifications would then be:

$$\begin{aligned} & \exists(\mathcal{L}_{\text{clh}}, \leq_{\text{clh}}, \top_{\text{clh}}, \perp_{\text{clh}}). \forall \eta: \mathcal{L}_{\text{clh}} \rightarrow_{\text{lay}} \mathcal{L}. \\ & \quad \eta(\perp_{\text{clh}}) \vdash \{\text{emp}\} \text{makeLock}() \{\exists s. L_\eta(s, x, 0) * P(s, 1)\} \\ & \quad \forall \pi > 0. \eta(\top_{\text{clh}}) \vdash \forall l \in \{0, 1\} \rightarrow_{\eta(\perp_{\text{clh}})} \{0\}. \langle P(s, \pi) \mid L_\eta(s, x, l) \rangle \text{lock}(x) \langle P(s, \pi) \mid L_\eta(s, x, 1) \wedge l = 0 \rangle \\ & \quad \eta(\perp_{\text{clh}}) \vdash \langle L_\eta(s, x, 1) \rangle \text{unlock}(x) \langle L_\eta(s, x, 0) \rangle \end{aligned}$$

When creating a new lock, one gets a local resource representing an unlocked lock and full permissions. Typically then permissions are distributed to the threads by splitting the full permission into smaller fractions. A non-trivial fraction of permission is now needed to perform the lock operation. We can then provide the viewshift $L_\eta(s, l) * P(s, 1) \Rightarrow L_{\eta'}(s, l) * P(s, 1)$ which allows to change the layer map without invalidating the knowledge about it in any other thread: if we own $P(s, 1)$ then no other thread can race on the lock. Adapting the proof of CLH to support permissions and the viewshift above follows standard (safety) proof patterns which we explain in Appendix C.

Let us briefly explain how we can use this viewshift in the lock-coupling set example. Conceptually, we want to organise the layers of the lock-coupling set module as for a static list of locks: they go in decreasing order from the head of the list to the tail. A thread holding a lock at position i will be able to eventually acquire the lock at position $i + 1$ because the release of such lock is associated with an obligation of strictly lower layer than the one associated with the lock at i . Each operation of the module inserts at most one element to the set per traversal of the list. We therefore arrange the proof invariants so that each thread traversing the list will shift up the layer of the lock at the thread’s current position by one. This way, when the thread finally finds the position where the new element has to be inserted, there is already a gap of 1 between the layers associated with the positions being altered by the thread. The layer sitting at the gap will be the one we associate with the lock of the new element. The layer-map-altering viewshift we explained above is used at each step of the traversal, to shift up the layer of the current lock. This is possible without breaking the information owned by other threads because when the current thread holds the lock at position i and the lock at $i + 1$ finally becomes available, the current thread is the only thread with access to the reference (and the associated resources) of the lock at $i + 1$. Formally, this means that when we obtain the lock at $i + 1$ we are able to obtain full permissions for it until we unlock the lock at i . With the full permissions we can apply the viewshift and effectively shift up the layers associated with the lock at $i + 1$.

The only exception to this scheme is the lock at the head of the queue: this is the only lock which does not need a remapping of layers as its associated layer can be $(\top, \perp_{\text{clh}})$ which is always bigger than any layer ever associated with the locks at the other positions.

It is worth noting that the LiLi proof of the same example does not use the specifications of the fair locks modularly, but instead inlines the code of the lock operations, allowing for a non-modular handling of the internal state.

Interestingly, the same lock-coupling set specifications can be implemented by using spin locks instead of CLH locks, for each element except the one at the head. In fact, the locks in the tail of the list do not experience any impedance. At first sight, it seems impossible to represent this fact using our specifications for spin lock: the lock operation needs to consume non-trivial budget, but there is no bound on the number of calls to it. The TaDA Live way of expressing the absence of impedance in this example uses a viewshift similar to the one we introduced above, which allows us to reset the budget (and the layer map) when we own full permissions. The proof in LiLi of this variant of the lock-coupling set again inlines the lock code, with the effect of being able to redefine which internal steps are susceptible of impedance and which do not, breaking modularity.

5.6 Limitations

Non-local linearization points. As with other total program logics, TaDA Live does not support helping/speculation. Such patterns are challenging for the identification of the linearization point, which is entirely a safety property. Extensions to TaDA that could support such patterns are discussed in [6]. Such extensions are orthogonal to the termination argument. We therefore choose, in line with the related literature, to explore termination in a simpler logic.

Non-structural thread creation. TaDA Live currently supports only structural parallel composition. We believe the support of non-structural fork/join would not require substantial new ideas. For comparison, LiLi does not support parallel nor fork/join.

Scheduling non-determinism. A more interesting limitation comes from our approach to specifying impedance. For non-blocking programs, the ordinal-based approach is complete. It is not complete for blocking programs. Consider $\mathbb{C}_2 \triangleq (\mathbb{C}_1 \parallel [\text{done}] := \text{true})$ where \mathbb{C}_1 is the distinguishing client with a spin lock. Scheduler fairness guarantees the right-hand thread of \mathbb{C}_2 will be eventually executed. The specification of spin lock, however, states that every call to lock needs to consume budget, forcing the client to provide an upper bound for the total number of calls to initialise the budget. Unfortunately, \mathbb{C}_2 will call lock an arbitrary unbounded number of times, determined only by the choices of the scheduler. It is, thus, not possible to provide the initial budget, and TaDA Live cannot prove that the program terminates. The impedance on the lock is only relevant when the client is unblocked (i.e. done is true) but the specifications do not allow for the distinction. To accommodate this behaviour, we could introduce $\alpha(\mathbf{d})$ to represent a prophecy of the number of steps it will take to fulfil *live* obligation \mathbf{d} . This would solve the problem for \mathbb{C}_2 , because $\alpha(\mathbf{d}) + 1$ (where \mathbf{d} is fulfilled by setting done to true) would be the required budget. How to introduce this extension soundly is future work. To the best of our knowledge, none of the approaches in the literature can handle this example.

Loop body specifications. Consider a loop invariant asserting the possession of obligation \mathbf{k} . We cannot distinguish, by only looking at the specification of the loop body, the case where \mathbf{k} is continuously held throughout the execution of the body, from the case where \mathbf{k} is fulfilled and then reacquired before the end of an iteration. The current **WHILE** rule conservatively rules out the use of assumptions with layer higher than or equal to $\text{lay}(\mathbf{k})$; doing otherwise would be unsound in the case when \mathbf{k} is held continuously. A solution would be to introduce an assertion $\text{live}(\mathbf{k})$, certifying that an obligation is fulfilled at some point in a block of code. It would allow the **WHILE** rule to only forbid layers which may depend on obligations one holds in the loop invariant and for which it was not possible to prove $\text{live}(\mathbf{k})$.

More Expressive Layers. Advanced examples like the lock-coupling set of Section 5.5 need powerful parametric specifications in order to work around the fact that the lay function is statically specified. We are not aware of any example that cannot be proved using static layers and critically requires more expressive layers. Even for current proofs, however, being able to constrain layers through assertions and allowing them to change as result of interference would allow for more concise and intuitive proofs. The lay function could in principle be encoded as “regular” ghost state and the crucial relative order between layers be enforced through invariants. It is however not clear how to ensure soundness if interference on layers is allowed. We leave this exploration as future work.

6 RELATED WORK

Primitive Blocking. There has been work on termination and deadlock-freedom of concurrent programs with primitive blocking constructs. Starting from the seminal work of [27], the idea of tracking dependencies between blocking actions and ensuring their acyclicity has been used to prove deadlock-freedom of shared-memory concurrent programs using primitive locks and (synchronous) channels [3, 28]. Similar techniques have been used in [16] to prove global deadlock-freedom (a *safety* property requiring that at least some thread can take a step), and [22] to prove termination. This entire line of work assumes the invocation of lock/channel *primitives* as the only source of blocking. As a consequence, this methodology provides no insight on the issue of understanding abstract blocking patterns arising from busy waiting and shared memory interference. Moreover, the specifications for blocking built-ins (hardcoded in the logic as ad hoc axioms) impose a usage protocol in the client, instead of just capturing the abstract effect of the operation: for instance, a call to `lock(x)` always entails an obligation to unlock the lock, regardless of how the client intends to use the lock. This has had the side effect of requiring ad hoc extensions of the reasoning principles to increase the expressivity of this hard-coded protocol, to allow, for example, for delegation of obligations [17]. Our solution uniformly handles programs that mix blocking primitives and ad-hoc synchronisation patterns, and is not imposing any specific protocol on the client.

The notion of “obligations” found in [3, 16, 22, 28] is only superficially related to our obligations. First, obligations found in the literature represent primitive blocking events (like the acquisition of a lock). They are also typically equipped with a structure to represent causal dependencies between these events, to detect deadlocks. Our layered obligations are associated with arbitrary *abstract* state changes, removing the need for ad-hoc treatment of primitives, and supporting abstraction and abstract atomicity. Moreover, our layers do not represent causal dependencies between events, but rather dependencies between liveness assumptions in a termination argument. This reflects in our specifications, e.g. a lock operation does not return an obligation in its post-condition. Whether there is a need for an obligation linked to that lock is entirely dependent on how the client will decide to use the lock. Nevertheless, the specification precisely captures the termination guarantees of lock operations. Finally, obligations in the literature have a purely safety semantics, from which one can only derive safety properties as non-blocking or deadlock-freedom. Our obligations explain how to express proper liveness invariants, how to blend them with the layers, and how to use them for proving termination.

Temporal Logics. There is substantial literature on using temporal logics to prove liveness and termination of concurrent programs, e.g. [37]. By working directly at the level of traces with liveness properties stated as temporal logic formulas, this approach is very general. It does however provide less guidance on how to prove programs, and does not tackle the problem of abstract interfaces and proof reuse. Our adoption of concurrent separation logic as the basis of our reasoning achieves superior compositionality of the reasoning including proof reuse.

History-based methods. The CertiKOS project [15, 25] developed mechanised techniques for the specification and verification of fine-grained low-level code with explicit support for abstract atomicity and progress verification. The approach is based on *histories*: the abstract state is a log of the abstract events of a trace; and the specification of an atomic operation inserts exactly one event in the log. Local reasoning is achieved by rely/guarantee through complex automata product constructions. The framework is very expressive,

with the downside that specifications are more complex and difficult to read, and verification requires manipulation of abstract traces/interleavings. Our work is similar in aim and scope, but our strategy is different. We try to specify/verify programs using the minimal machinery possible, keeping the specifications as close to the developer’s intuition as we can. As a result, our specifications are more readable (compare our fair-lock specification with the corresponding 30-line specification from Fig. 7 in [25]), and our reasoning is simpler (the layered obligation system leads to a more intuitive proof compared to the proof of MCS locks in [25]).¹⁸

Contextual refinement. Another approach to specify and prove progress of concurrent systems is to prove refinement between the implementation and simpler, abstract code acting as a specification [30, 31, 39]. By making sure the refinement preserves progress properties, one can represent the salient termination properties of the implementation by the termination properties of the specification code. The Iris implementation of this idea [39] uses a non-contextual refinement, which means that the refinement is proven between the closed-world behaviour of implementation and specification code, and does not necessarily carry over contexts. This severely hinders proof reuse. The only refinement-based work that is able to modularly verify blocking code is the LiLi logic discussed below.

There has been work on extending linearizability, characterised as a contextual refinement, to support reasoning about progress properties, e.g. [14]. This work only supports non-blocking operations. Liang et al. [33] studies the exact relationship between common progress properties of fine-grained operations and contextual refinement. The study of the contextual refinement induced by our triple semantics is future work.

LiLi. The work closest to ours is LiLi [30, 31]. LiLi was the first concurrent separation logic to prove progress specifications for linearizable concurrent objects with internal blocking [30], and it was then extended to handle external blocking [31]. Although we share most of our goals with LiLi, our approach differs in two important ways.

First, LiLi’s goal is to prove a progress-preserving contextual refinement between the implementation of a module and its specification. Termination properties of implementation code are not represented directly, but in terms of the termination properties of the specification code. Although proof of clients of the module have to be done outside of the LiLi logic (there is no rule for parallel, nor for calling a module’s operation) such proofs would need to reprove the relevant termination properties of the specification code so that the properties themselves become available in the proof. Moreover, as we outlined in Section 2 for CLH lock, the specification code for blocking operations may be non-atomic even in the case of linearizable operations. Instead we aim at specifications that directly represent termination properties as a logical statement that can be readily used in a client proof, and in the proof of the implementation. Our specification format obtains a crucial advantage: it achieves abstraction and can represent atomicity for blocking operations, enabling more scalable and reusable reasoning.

Second, LiLi’s rely/guarantee incorporates a form of liveness invariants through so-called *definite actions*. Definite actions require the identification of a logical global “queue” of threads where the

¹⁸The proof is a variation of the one for CLH.

thread at the front is always able to execute its action and that action implies global progress. This queue is maintained as shared auxiliary state manipulated through ghost code. It is due to this global view that definite actions can side-step the issue of circular reasoning. Our layered subjective obligations push the idea much further, obtaining sound liveness invariants that can be represented thread-locally and without the need for ghost code, improving proof scalability. The design choice of making both rely/guarantee and specification represent blocking via liveness assumptions is the key to making the blocking specifications directly usable in the proof system.

7 CONCLUSIONS AND FUTURE WORK

We have introduced TaDA Live, a concurrent separation logic for reasoning compositionally about the termination of fine-grained blocking concurrent programs, and proved a substantial soundness result. Our key contribution is our approach to abstract atomic blocking as the reliance of termination on the liveness properties of the environment. By wholly embracing this point of view, we have designed a rely/guarantee principle that incorporates liveness invariants using layered subjective obligations, a new form of local ghost state, and have extended TaDA’s abstract atomic specifications to provide total specifications for blocking programs using environment liveness assumptions. Through several case studies, we have illustrated how our formalisation of abstract blocking allows for the right level of abstraction in specification, and strong thread-locality of the proofs. The result is a verification system with scalable and reusable proofs.

The work presented in this paper opens a number of immediate directions for future work on concurrent separation logics. A first direction is to extend TaDA Live to prove general liveness properties beyond termination. A possible way to achieve this is to wrap a refinement calculus around TaDA Live’s atomic specifications, as was done in the safety case in TaDA Refine [35]. Specifications would be able to sequentially compose atomic triples and take fixpoints, thus being able to specify linear-time temporal properties of infinite traces. A second direction is to study general fork/join concurrency and provide a generalisation of the liveness rely/guarantee necessary to accommodate patterns typical of distributed/reactive systems, where long-lived maintenance threads interact with an environment to realise an operation’s effect. A third direction is to transfer ideas from TaDA Live to the Iris framework [23], to provide a Coq-mechanised environment for reasoning about the termination of concurrent programs. More widely, we hope that our emphasis on environment liveness invariants for proving termination will transfer to other forms of reasoning about blocking concurrent programs.

ACKNOWLEDGMENTS

We would like to thank Hongjin Liang, Xinyu Feng, Martin Bodin, Shale Xiong and Petar Maksimovic, for the helpful discussions and comments. We also thank the anonymous reviewers for their thorough critical reading of the paper and insightful feedback. This research was supported by: the EPSRC Programme Grant “REMS: Rigorous Engineering for Mainstream Systems” (EP/K008528/1); by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie project “VeSPA”, grant agreement no. 795218; by a Department of Computing PhD Scholarship from Imperial; by the UKRI Established Fellowship “VeTSpec: Verified Trustworthy Software Specification” (EP/R034567/1); and in the final stages by the ERC Consolidator Grant for the project “RustBelt”, also funded under EU Horizon 2020, grant agreement no. 683289.

REFERENCES

- [1] Ales Bizjak, Daniel Gratzer, Robbert Krebbers, and Lars Birkedal. 2019. Iron: managing obligations in higher-order concurrent separation logic. *Proc. ACM Program. Lang.* 3, POPL (2019), 65:1–65:30.

- [2] Richard Bornat, Cristiano Calcagno, and Hongseok Yang. 2005. Variables as Resource in Separation Logic. In *MFPS (Electronic Notes in Theoretical Computer Science)*, Vol. 155. Elsevier, 247–276.
- [3] Pontus Boström and Peter Müller. 2015. Modular Verification of Finite Blocking in Non-terminating Programs. In *ECOOP*. 639–663.
- [4] Stephen D. Brookes. 2004. A Semantics for Concurrent Separation Logic. In *CONCUR (Lecture Notes in Computer Science)*, Vol. 3170. Springer, 16–34.
- [5] Byron Cook, Andreas Podelski, and Andrey Rybalchenko. 2007. Proving thread termination. In *PLDI*. ACM, 320–330.
- [6] Pedro da Rocha Pinto. 2016. *Reasoning with Time and Data Abstractions*. Ph.D. Dissertation. Imperial College London.
- [7] Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *ECOOP 2014–Object-Oriented Programming*. Springer Berlin Heidelberg, 207–231.
- [8] Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. 2016. Modular Termination Verification for Non-blocking Concurrency. In *ESOP (Lecture Notes in Computer Science)*, Vol. 9632. Springer, 176–201.
- [9] Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: compositional reasoning for concurrent programs. In *POPL*. ACM, 287–300.
- [10] Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. 2010. Concurrent Abstract Predicates. In *ECOOP (Lecture Notes in Computer Science)*, Vol. 6183. Springer, 504–528.
- [11] Mike Dodds, Xinyu Feng, Matthew J. Parkinson, and Viktor Vafeiadis. 2009. Deny-Guarantee Reasoning. In *ESOP (Lecture Notes in Computer Science)*, Vol. 5502. Springer, 363–377.
- [12] Emanuele D’Osualdo, Azadeh Farzan, Philippa Gardner, and Julian Sutherland. 2021. TaDA Live: Compositional Reasoning for Termination of Fine-grained Concurrent Programs. *CoRR* abs/1901.05750 (2021).
- [13] Alexey Gotsman, Byron Cook, Matthew J. Parkinson, and Viktor Vafeiadis. 2009. Proving that non-blocking algorithms don’t block. In *POPL*. ACM, 16–28.
- [14] Alexey Gotsman and Hongseok Yang. 2011. Liveness-Preserving Atomicity Abstraction. In *ICALP*. 453–465.
- [15] Ronghui Gu, Zhong Shao, Jieung Kim, Xiongnan (Newman) Wu, Jérémie Koenig, Vilhelm Sjöberg, Hao Chen, David Costanzo, and Tahina Ramananandro. 2018. Certified concurrent abstraction layers. In *PLDI*. ACM, 646–661.
- [16] Jafar Hamin and Bart Jacobs. 2018. Deadlock-Free Monitors. In *ESOP (Lecture Notes in Computer Science)*, Vol. 10801. Springer, 415–441.
- [17] Jafar Hamin and Bart Jacobs. 2019. Transferring Obligations Through Synchronizations. In *ECOOP (LIPIcs)*, Vol. 134. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 19:1–19:58.
- [18] Maurice Herlihy and Nir Shavit. 2008. *The art of multiprocessor programming*. Morgan Kaufmann.
- [19] Maurice Herlihy and Nir Shavit. 2011. On the Nature of Progress. In *Principles of Distributed Systems - 15th International Conference, OPODIS 2011, Toulouse, France, December 13–16, 2011. Proceedings*. 313–328.
- [20] Maurice Herlihy and Jeanette M. Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (1990), 463–492.
- [21] Jan Hoffmann, Michael Marmor, and Zhong Shao. 2013. Quantitative Reasoning for Proving Lock-Freedom. In *LICS*. IEEE Computer Society, 124–133.
- [22] Bart Jacobs, Dragan Bosnacki, and Ruurd Kuiper. 2018. Modular Termination Verification of Single-Threaded and Multithreaded Programs. *ACM Trans. Program. Lang. Syst.* 40, 3 (2018), 12:1–12:59.
- [23] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20.
- [24] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. ACM, 637–650.
- [25] Jieung Kim, Vilhelm Sjöberg, Ronghui Gu, and Zhong Shao. 2017. Safety and Liveness of MCS Lock – Layer by Layer. In *APLAS (Lecture Notes in Computer Science)*, Vol. 10695. Springer, 273–297.
- [26] Naoki Kobayashi. 2000. Type Systems for Concurrent Processes: From Deadlock-Freedom to Livelock-Freedom, Time-Boundedness. In *IFIP TCS (Lecture Notes in Computer Science)*, Vol. 1872. Springer, 365–389.
- [27] Naoki Kobayashi. 2006. A New Type System for Deadlock-Free Processes. In *CONCUR (Lecture Notes in Computer Science)*, Vol. 4137. Springer, 233–247.
- [28] K. Rustan M. Leino, Peter Müller, and Jan Smans. 2010. Deadlock-Free Channels and Locks. In *ESOP (Lecture Notes in Computer Science)*, Vol. 6012. Springer, 407–426.
- [29] Ruy Ley-Wild and Aleksandar Nanevski. 2013. Subjective auxiliary state for coarse-grained concurrency. In *POPL*. ACM, 561–574.
- [30] Hongjin Liang and Xinyu Feng. 2016. A program logic for concurrent objects under fair scheduling. In *POPL*. 385–399.
- [31] Hongjin Liang and Xinyu Feng. 2018. Progress of concurrent objects with partial methods. *PACMPL* 2, POPL (2018), 20:1–20:31.
- [32] Hongjin Liang, Xinyu Feng, and Zhong Shao. 2014. Compositional verification of termination-preserving refinement of concurrent programs. In *CSL-LICS*. 65:1–65:10.

- [33] Hongjin Liang, Jan Hoffmann, Xinyu Feng, and Zhong Shao. 2013. Characterizing Progress Properties of Concurrent Objects via Contextual Refinements. In *CONCUR (Lecture Notes in Computer Science)*, Vol. 8052. Springer, 227–241.
- [34] Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating State Transition Systems for Fine-Grained Concurrent Resources. In *ESOP*. 290–310.
- [35] Gian Ntzik, Pedro da Rocha Pinto, Julian Sutherland, and Philippa Gardner. 2018. A Concurrent Specification of POSIX File Systems. In *ECOOP (LIPICs)*, Vol. 109. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 4:1–4:28.
- [36] Peter W. O’Hearn. 2004. Resources, Concurrency and Local Reasoning. In *CONCUR (Lecture Notes in Computer Science)*, Vol. 3170. Springer, 49–67.
- [37] Susan S. Owicki and Leslie Lamport. 1982. Proving Liveness Properties of Concurrent Programs. *ACM Trans. Program. Lang. Syst.* 4, 3 (1982), 455–495.
- [38] Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2015. Specifying and Verifying Concurrent Algorithms with Histories and Subjectivity. In *ESOP (Lecture Notes in Computer Science)*, Vol. 9032. Springer, 333–358.
- [39] Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *ESOP (Lecture Notes in Computer Science)*, Vol. 10201. Springer, 909–936.
- [40] Moshe Y. Vardi. 1995. Alternating Automata and Program Verification. In *Computer Science Today. Lecture Notes in Computer Science*, Vol. 1000. Springer, 471–485.

Appendix

We present here omitted definitions and details of proofs. An extended version of this paper is also available at <https://arxiv.org/abs/1901.05750> [12].

A SOME PROOFS CONVENTIONS

A.1 Specification abbreviations

Here is a summary of all the abbreviations we use in writing specifications. The full hybrid specification format is

$$m; \lambda; \mathcal{A} \vDash \mathbb{W}x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle$$

The $\exists y$ quantification is a normal existential quantification but its scope extends over both the Hoare and the atomic postconditions. We omit it when y does not occur in the triple.

$$\mathbb{W}x \triangleq \mathbb{W}x \in \text{Val}$$

$$\mathbb{W}x \in X \triangleq \mathbb{W}x \in X \rightarrow_{\perp} X$$

$$\mathbb{W}x_1 \in X_1 \rightarrow_k X'_1, x_2 \in X_2 \rightarrow_k X'_2. \triangleq \mathbb{W}(x_1, x_2) \in (X_1 \times X_2) \rightarrow_k (X'_1 \times X'_2).$$

An omitted pseudo-quantifier is to be understood as the trivial pseudo-quantifier $\mathbb{W}x \in \text{AVal} \rightarrow_{\perp} \text{AVal}$, for an unused x .

The triples

$$m, \lambda, \mathcal{A} \vdash \{P\} \mathbb{C} \{Q\}$$

$$m, \lambda, \mathcal{A} \vdash \mathbb{W}x \in X \rightarrow_k X'. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle$$

are abbreviated with

$$m; \lambda; \mathcal{A} \vdash \langle P \mid \text{emp} \rangle \mathbb{C} \langle Q \mid \text{emp} \rangle$$

$$\forall \vec{v}_0. m; \lambda; \mathcal{A} \vdash \mathbb{W}x \in X \rightarrow_k X'. \langle \vec{v}_0 \doteq \vec{v}_0 \mid P'(x) \rangle \mathbb{C} \exists \vec{v}_1. \langle \vec{v}_0 \doteq \vec{v}_0 \wedge \vec{v}_1 \doteq \vec{v}_1 \mid Q'(x) \rangle$$

respectively, where $\vec{v}_0 = \text{pv}(P(x))$, $\vec{v}_1 = \text{pv}(Q(x)) \setminus \vec{v}_0$, $P'(x) = P(x)[\vec{v}_0/\vec{v}_0]$ and $Q'(x) = Q(x)[\vec{v}_0/\vec{v}_0, \vec{v}_1/\vec{v}_1]$ (for technical reasons the atomic pre/post-conditions in the general triples cannot contain program variables). In other words, the program variables mentioned in the atomic pre/post-conditions refer to the value stored in them *at the beginning* of the execution of the command. Most commonly, the program variables used this way are actually not modified by the command.

A.2 Guard and Obligation Algebras

Defining a guard algebra can be tedious. In program proofs, we will define guard algebras by generating them from some *guard constructors* and some axioms defining the guard operation.

Consider two common guard patterns in TaDA Live: the use of an *exclusive guard* and the **U, L, K** *pattern* used to represent possession of a lock in ghost state.

An exclusive guard, **E**, is very commonly used to express some exclusive permission on some shared resource, which cannot be composed with itself: i.e. $\mathbf{E} \bullet \mathbf{E} = \perp$. Local ownership of **E** is exclusive in that no other thread can at the same time assert ownership of **E**. A ubiquitous use of this guard is in representing the resource offered by a module.

The **U, L, K** pattern is commonly used to represent ownership of a lock guarding a resource. The thread records its ownership of a lock by holding the ghost state **K**, which cannot be composed with the guard **U**, recording the lock is unlocked: $\mathbf{U} \bullet \mathbf{K} = \perp$ The region holds the associated guard

\mathbf{L} , which can be recombined with the guard \mathbf{K} once the thread releases the lock to form the guard \mathbf{U} : $\mathbf{U} = \mathbf{L} \bullet \mathbf{K}$.

We explain the construction of a guard or obligations algebra from these axioms by introducing some unsurprising auxiliary definitions.

Given a set X , the set $\mathcal{M}(X) \triangleq X \rightarrow \mathbb{N}$ is the set of *multisets* over X ; \emptyset is the empty multiset (i.e. the function mapping every element to 0) and $\oplus: \mathcal{M}(X) \times \mathcal{M}(X) \rightarrow \mathcal{M}(X)$ is multiset union (i.e. the pointwise lifting of $+$). The expression $\langle x_1, \dots, x_n \rangle$ denotes the multiset containing the elements x_1, \dots, x_n . Given a set X , the *free commutative monoid* over X is the monoid $(\mathcal{M}(X), \oplus, \emptyset)$. Given a commutative monoid $(X, \bullet, \mathbf{0})$ and a congruence relation $\cong \subseteq X \times X$, the *quotient* $(X/\cong, \bullet/\cong, [\mathbf{0}]_\cong)$ is a commutative monoid. Given a commutative monoid $(X, \bullet, \mathbf{0})$ and a set $U \subseteq X$ with $\mathbf{0} \notin U$, the *PCM over X induced by U* is $(X|_U, \bullet_U, \mathbf{0})$ where

$$X|_U \triangleq \{x \in X \mid \forall u \in U. \nexists y \in X. x = u \bullet y\}$$

and for $x, y \in X|_U$, $x \bullet_U y = x \bullet y$ if $x \bullet y \in X|_U$, otherwise undefined.

For each guard algebra to be defined, we will introduce a number of symbols $\mathbf{G}_1, \dots, \mathbf{G}_n$, called *guard constructors* each with some *guard domain* $\text{dom}(\mathbf{G}_i) \subseteq \text{AVal}^{k_i}$ for some $k_i \in \mathbb{N}$. They induce the set of *guard terms* $\text{GT} \triangleq \bigcup_{i=1}^n \{\mathbf{G}_i(\vec{a}) \mid \vec{a} \in \text{dom}(\mathbf{G}_i)\}$. By specifying some guard constructors, a congruence $\cong \subseteq \mathcal{M}(\text{GT}) \times \mathcal{M}(\text{GT})$ and a set $U \subseteq \mathcal{M}(\text{GT})/\cong$ one obtains the guard algebra $((\mathcal{M}(\text{GT})/\cong)|_U, (\oplus/\cong)_U, [\emptyset]_\cong)$.

The guard constructors are specified by listing their domains, writing $\mathbf{G}_i: D_i$ to mean $\text{dom}(\mathbf{G}_i) = D_i \subseteq \text{AVal}^{k_i}$, as, in certain cases, we may want to further restrict the domain of the guard constructors to simplify the reasoning.

The congruence \cong is specified as the smallest congruence satisfying given axioms of the form

$$\langle \mathbf{G}_{i_1}(\vec{a}_{i_1}), \dots, \mathbf{G}_{i_k}(\vec{a}_{i_k}) \rangle \cong \langle \mathbf{G}_{j_1}(\vec{a}_{j_1}), \dots, \mathbf{G}_{j_{k'}}(\vec{a}_{j_{k'}}) \rangle$$

which we write using the syntax

$$\mathbf{G}_{i_1}(\vec{a}_{i_1}) \bullet \dots \bullet \mathbf{G}_{i_k}(\vec{a}_{i_k}) = \mathbf{G}_{j_1}(\vec{a}_{j_1}) \bullet \dots \bullet \mathbf{G}_{j_{k'}}(\vec{a}_{j_{k'}})$$

The set U is specified as the smallest set satisfying given axioms of the form

$$\langle \langle \mathbf{G}_{i_1}(\vec{a}_{i_1}), \dots, \mathbf{G}_{i_k}(\vec{a}_{i_k}) \rangle \rangle \cong \in U$$

which we write using the syntax

$$\mathbf{G}_{i_1}(\vec{a}_{i_1}) \bullet \dots \bullet \mathbf{G}_{i_k}(\vec{a}_{i_k}) = \perp$$

Example A.1. The guard algebra used in Example 4.1, is expressed by using two guard constructors with empty domain, \mathbf{K} and \mathbf{D} , and axioms: $\mathbf{K} \bullet \mathbf{K} = \perp$, $\mathbf{D} \bullet \mathbf{D} = \perp$. Note that with no congruence axioms, the induced congruence relation is equality. These induce the guard algebra with elements $\{\emptyset, \langle \mathbf{K} \rangle, \langle \mathbf{D} \rangle, \langle \mathbf{K}, \mathbf{D} \rangle\}$.

A.3 Levels

Region levels are used to remove the possibility of unsound duplication of resources by opening regions. The presentation of the program proofs omits the level annotations to ease readability. The levels can be unambiguously derived from the sequence of application of rules **UPDREG** and **LIFTA**.

To see the problem consider a generic region $\mathbf{t}_r^\lambda(a)$; we have $\mathbf{t}_r^\lambda(a) \equiv \mathbf{t}_r^\lambda(a) * \mathbf{t}_r^\lambda(a)$: this is the essence of what it means for a region to be a shared resource. When we open a region however, we obtain ownership of the contents of its interpretation $\mathcal{I}(\mathbf{t}_r^\lambda(a))$; the interpretation can contain resources that are not shared, for example heap assertions, in which case we have $\mathcal{I}(\mathbf{t}_r^\lambda(a)) \neq \mathcal{I}(\mathbf{t}_r^\lambda(a)) * \mathcal{I}(\mathbf{t}_r^\lambda(a)) \equiv \text{False}$. Without constraining levels, one could start with $\mathbf{t}_r^\lambda(a)$, produce the equivalent $\mathbf{t}_r^\lambda(a) * \mathbf{t}_r^\lambda(a)$, open the first region assertion with **UPDREG** or **LIFTA**, then open the

second region assertion and end up with `False`. Levels are a mean to avoid unsound derivations that use the above chain of implications. A level λ in the context of a judgement records that all the regions of level λ or higher might have been already opened and should not be opened again. The rules that do open regions (rules `UPDREG` and `LIFTA`) can only open a region of level λ if the level in the context is $\lambda + 1$, and they record the operation by setting the context level to λ , so that the region cannot be opened again.

A.4 Region type specifications

– *Abstract state domain.* It can be tedious (and detrimental to readability) to always explicitly write the domains of quantified variables in the assertions of program proofs, especially when they can be easily inferred from context. Consider the case of regions. Some of the rules, for example `MkATOM`, need the precise domain of the abstract state ($\exists x \in X$) because it needs to match the pseudo-quantifier’s domain ($\forall x \in X$). To improve readability, we adopt the following strategy. Suppose the region type t has abstract state in the domain A . We can define the interpretation function so that it constrains the domain of the abstract state accordingly: $\mathcal{I}(t_r^\lambda(a)) = a \in A \wedge \dots$. Then we trivially have that $\lambda'; \mathcal{A} \vDash \exists a. t_r^\lambda(a) \Rightarrow \exists a \in A. t_r^\lambda(a)$. We thus can omit the domains from existential quantification and implicitly apply rule `CONS` whenever the domain information is needed in the proof.

To further ease the specification of region types, when defining a new region type we will introduce the domain of the corresponding abstract state, and omit the obvious constraint from the interpretation definition.

– *Fixed parameters.* It is very common to have a product domain as abstract state of regions, as one needs to assemble in an abstract state many bits of information that characterise region’s state. Typically, the abstract state domain A can be seen as the product of two domains $F \times S$, the domain of the *fixed parameters* F and the domain of the *non-fixed parameters* S . (Both F and S can be themselves products of simpler domains.) The fixed parameters are set at the point of creation of the region, and can never be updated; they typically define the “interface” of the region. For example, if the address of a lock module instance x is the fixed parameter of a hypothetical region `lockr(x, l)` and $l \in \{0, 1\}$ the non-fixed parameter representing the state of the lock. When introducing a new region type we will specify which parameters are fixed, and they will be omitted from the region interference specification, as they are left untouched by every transition. For example, for the region `lockr(x, l)` above, we may write $G : (0, \mathbf{0}) \rightsquigarrow (1, \mathbf{k})$ and $G : (1, \mathbf{k}) \rightsquigarrow (0, \mathbf{0})$ to denote $G : ((x, 0), \mathbf{0}) \rightsquigarrow ((x, 1), \mathbf{k})$ and $G : ((x, 1), \mathbf{k}) \rightsquigarrow ((x, 0), \mathbf{0})$.

– *Interference protocols and atomicity contexts.* Definition 3.13 requires \mathcal{T}_t to be monotone in the guards, reflexive and closed under obligation frames. Since writing the whole function can be tedious and redundant, we will only write a number of expressions of the form

$$G : (a_1, O_1) \rightsquigarrow (a_2, O_2) \tag{25}$$

which will set $\mathcal{T}_t(G) \ni \{(a_1, O_1), (a_2, O_2)\}$, and implicitly complete the function by closing \mathcal{T}_t under the properties above.

Similarly, atomicity contexts associate to some region identifier records $\mathcal{A}(r) = (X, k, X', R)$ that have (unguarded) transition relations as their last component R . We therefore borrow the syntax from (25), and write $R = (a_1, O_1) \rightsquigarrow (a_2, O_2)$ to specify R as the minimal relation that include such relations and is closed under obligation frames.

A.5 Proof patterns

There are some recurring patterns in TaDA Live proofs, which we summarise here to help the reader navigate the examples.

– *The exclusive guard.* Take for example a concurrent counter module. Abstractly we have a (fixed) location x for the module instance and an abstract state $n \in \mathbb{N}$ representing the current value of the counter. Since this is a concurrent counter it uses internally shared resources. We therefore have a region $\mathbf{cnt}_r(x, n)$ encapsulating the shared internal resources of the counter. From the perspective of the client, however, at the moment of creation of the counter with, say, an operation $\mathbf{makeCounter}()$, the counter is exclusively owned by the client. This, for example, is reflected in the fact that, until the client shares the counter or invokes operations on it, the value of the counter will be stably 0. To represent this fact, one typically defines an exclusive guard \mathbf{E} guarding each transition of the region interference: e.g. $\mathbf{E} : (n, O_1) \rightsquigarrow (m, O_2)$. Then the $\mathbf{makeCounter}()$ operation can be given the specification

$$\vdash \{\mathbf{emp}\} x := \mathbf{makeCounter}() \{ \exists r. \mathbf{cnt}_r(x, 0) * \lceil \mathbf{E} \rceil_r \}$$

which gives to the client the stable assertion $\mathbf{cnt}_r(x, 0) * \lceil \mathbf{E} \rceil_r$. (Note how $\mathbf{cnt}_r(x, 0)$ is not stable.) To re-share the counter, the client will create its own region encoding the invariants governing the interaction over the counter (and the other resources of the client) the interpretation of which will contain $\mathbf{cnt}_r(x, 0) * \lceil \mathbf{E} \rceil_r$.

Note that the assertion $\mathbf{cnt}_r(x, 0) * \lceil \mathbf{E} \rceil_r$ has a very different meaning if occurring in the *atomic* precondition of a triple, as opposed to the *Hoare* precondition: the resources in the atomic precondition are not owned by the local thread, but only acquired instantaneously at the linearisation point. For example, in the triple

$$\vdash \forall n \in \mathbb{N}. \langle \mathbf{cnt}_r(x, n) * \lceil \mathbf{E} \rceil_r \rangle \text{incr}(x) \langle \mathbf{cnt}_r(x, n+1) * \lceil \mathbf{E} \rceil_r \rangle$$

the exclusivity of \mathbf{E} is only granted *instantaneously* to the thread acting on it atomically, i.e. either the environment during the interference phase as allowed by the pseudo-quantifier, or the local thread at the linearisation point.

Since this pattern is ubiquitous, we reserve the \mathbf{E} guard constructor for this use, and will omit the $\mathbf{E} \bullet \mathbf{E} = \perp$ axiom when specifying guard algebras.

A.6 Modules

TaDA Live is a logics that emphasizes modularity of the proofs. One aspect of this is that when a program is naturally structured as a collection of modules, one would want the proof of correctness to be decomposed into independent proofs of each module exporting some specifications for the externally accessible operations, and a proof that the client of these modules is correct, which depends only on these abstract module specifications.

In our model, a module is nothing but a conceptually related set of operations f_1, \dots, f_n that are defined in a **let** statement: **let** $f_1(\vec{x}_1) = \mathbb{C}_1$ **in** ... **let** $f_n(\vec{x}_n) = \mathbb{C}_n$ **in** \mathbb{C} . Here \mathbb{C} is what we call “client” of a module offering operations f_1, \dots, f_n . The operation deals with let statements by populating a function φ associating each function name f_i to its formal parameters \vec{x}_i and its implementation \mathbb{C}_i .

Similarly, the proof of correctness of \mathbb{C} , will need to fetch the abstract specifications of the functions (which appear as free names in \mathbb{C}) from some mapping Φ from function names to their specifications. The fact that the implementation of each operation satisfies its specification is checked in the proof derivation for the let statement (rule **LET**) but then the proof of the client and of the module are done separately.

For this reason, we present proofs of just a module against its abstract specifications, which can be used as if they were axioms in the proof of any client using them. To talk about modules independently of their clients we introduce the notation $\mathbf{def} \ f(x) \ \{\mathbb{C}\}$ which can be understood as populating an entry of φ for f . We will then prove some specification for f which will populate an entry of Φ for f .

In the proof of some client, we will recall the module specifications that are assumed in Φ , and use rule **CALL** to handle the calls to the operations of the module. We will omit from the proof outlines Φ and the applications of rule **CALL** for readability.

A.7 Proof outlines

In program proof outlines, we adopt a number of notational conventions. First, unless it involves a viewshift or we want to highlight it, we will apply rule **CONS** without mentioning it. Similarly, we omit the obvious applications of rules **VAR**, **CALL** and **SUBPQ** and the axioms (i.e. the rules associated with primitive commands).

Next, in outline such as

$$\begin{array}{c}
 m; \lambda; \mathcal{A} \vdash \\
 \forall x \in X \rightarrow X'. \\
 \langle P(x) \rangle \\
 \left. \begin{array}{c} \langle P'(x) \rangle \\ \vdots \\ \langle Q'(x) \rangle \\ \langle Q(x) \rangle \end{array} \right\} \begin{array}{l} \text{OUTER} \\ \text{INNER} \end{array}
 \end{array}
 \quad \frac{\vdots}{\frac{m; \lambda; \mathcal{A} \vdash \forall x \in X \rightarrow X'. \langle P'(x) \rangle \mathbb{C} \langle Q'(x) \rangle}{m; \lambda; \mathcal{A} \vdash \forall x \in X \rightarrow X'. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}}
 \begin{array}{l} \text{INNER} \\ \text{OUTER} \end{array}$$

the specification of the inner step inherits the context and the pseudo-quantifier of the specification of the outer step, as in the derivation on the right.

B THE TADA LIVE PROOF SYSTEM

In this section, we present the full proof system of TaDA Live.

For brevity we use the metavariable \vec{X} to range over expressions of the form $X_1 \rightarrow_k X_2$ and is used in rules when the pseudo-quantification is simply copied verbatim from premise to conclusion.

In the rules we use the following abbreviation:

$$\begin{aligned}
 \vdash_{\mathcal{A}} P \triangleright k &\triangleq \forall r \in \text{Rld}. \vdash_{\mathcal{A}} P \Rightarrow r \triangleright k \\
 k \triangleright n &\triangleq \forall k' > k. k' \triangleright n
 \end{aligned}$$

The λ -safety condition is defined in Appendix B.2.1 and can be typically proven by using Lemma 4.2.

B.1 Liveness rules

For reference we reproduce the liveness-related rules.

$$\frac{\forall x \in X. \vdash_{\lambda; \mathcal{A}} P_a(x) * T \Rightarrow x \in X' \quad \begin{array}{c} n; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T \quad m \triangleright n \quad k \triangleright n \quad \text{pv}(L) \cap \text{mod}(\mathbb{C}) = \emptyset \\ m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x) \rangle \end{array}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X. \langle P_h * L \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * L \mid Q_a(x) \rangle} \text{LIVECG}$$

$$\frac{\begin{array}{c} \forall \beta \leq \beta_0. m(\beta); \lambda; \mathcal{A} \vdash L \xrightarrow{M} T(\beta) \quad \forall \beta \leq \beta_0. \vdash_{\mathcal{A}} P(\beta) \triangleright m(\beta) \leq m \\ \forall \alpha. \mathcal{A} \models \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \quad \text{pv}(T, L, M) \cap \text{mod}(\mathbb{C}) = \emptyset \\ \forall \beta \leq \beta_0. \forall b \in \text{Bool}. m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta) * (b \dot{\Rightarrow} T(\beta)) \wedge \mathbb{B}\} \mathbb{C} \{\exists \gamma. P(\gamma) \wedge \gamma \leq \beta * (b \dot{\Rightarrow} \gamma < \beta)\} \end{array}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(\beta_0) * L\} \mathbf{while}(\mathbb{B})\{\mathbb{C}\} \{\exists \gamma. P(\gamma) * L \wedge \neg \mathbb{B} \wedge \gamma \leq \beta_0\}} \text{WHILE}$$

$$\frac{\begin{array}{l} m_1; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad \vdash_{\mathcal{A}} Q_1 \triangleright m_2 \leq m \\ m_2; \lambda; \mathcal{A} \vdash_{\Phi} \{P_2\} \mathbb{C}_2 \{Q_2\} \quad \vdash_{\mathcal{A}} Q_2 \triangleright m_1 \leq m \end{array}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}} \text{PAR}$$

B.1.1— The Environment liveness rules. The Environment liveness rules use the $\text{impr}_{\mathcal{A}}$ condition (Definition 4.3) recalled here for convenience:

Definition B.1 ($\text{impr}_{\mathcal{A}}$). Given assertions $L(\alpha)$, $L'(\alpha)$ and T , the condition $\text{impr}_{\mathcal{A}}(L', L, T)$ holds if and only if, for arbitrary $\sigma \in \text{Store}$, letting

$$l(\alpha) = \mathcal{W}[[L(\alpha)]]_{\mathcal{A}}^{\sigma} \quad l'(\alpha) = \mathcal{W}[[L'(\alpha)]]_{\mathcal{A}}^{\sigma} \quad t = \mathcal{W}[[T * \text{True}]]_{\mathcal{A}}^{\sigma}$$

the following holds:

$$\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(l'(\alpha_1)) \cap l(\alpha_2) \subseteq l'(\alpha_1) \cup t$$

We reproduce below for completeness the rules to prove the environment liveness condition.

$$\frac{\begin{array}{l} \lambda; \mathcal{A} \vDash L \text{ stable} \quad \vdash_{\lambda; \mathcal{A}} L \Rightarrow L * \exists \alpha. M(\alpha) \\ m; \lambda; \mathcal{A} \vdash L * M(\alpha) : L * M(\alpha) \rightarrow T \end{array}}{m; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T} \text{ENVLIVE}$$

$$\frac{\begin{array}{l} m; \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \rightarrow T \\ m; \lambda; \mathcal{A} \vdash L(\alpha) : L_2(\alpha) \rightarrow T \end{array}}{m; \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \vee L_2(\alpha) \rightarrow T} \text{ECASE} \quad \frac{\forall x \in X. m; \lambda; \mathcal{A} \vdash L(\alpha) : L(x, \alpha) \rightarrow T}{m; \lambda; \mathcal{A} \vdash L(\alpha) : \exists x \in X. L(x, \alpha) \rightarrow T} \text{EQUANT}$$

$$\frac{\forall \alpha. \vdash_{\mathcal{A}} T'(\alpha) \Rightarrow T}{m; \lambda; \mathcal{A} \vdash L(\alpha) : T'(\alpha) \rightarrow T} \text{LIVET}$$

$$\frac{\begin{array}{l} \text{impr}_{\mathcal{A}}(L', L, T) \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright \text{lay}(O(x)) \\ \lambda < \lambda' \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x. \mathbf{t}_r^{\lambda}(x) * \lfloor O(x) \rfloor_r^E * \text{True} \wedge m > \text{lay}(O(x)) \end{array}}{m; \lambda'; \mathcal{A} \vdash L(\alpha) : L'(\alpha) \rightarrow T} \text{LIVEO}$$

$$\frac{\begin{array}{l} \text{impr}_{\mathcal{A}}(L', L, T) \quad m > k \quad \forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright k \\ (X \rightarrow_k X') = \text{live}(\mathcal{A}, r) \quad \lambda < \lambda' \quad \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x \in X \setminus X'. \mathbf{t}_r^{\lambda}(x) * r \mapsto \diamond * \text{True} \end{array}}{m; \lambda'; \mathcal{A} \vdash L(\alpha) : L'(\alpha) \rightarrow T} \text{LIVEA}$$

B.2 Atomicity rules

We first give the formal definition of λ -safety and prove its properties, and then give the general forms of rules **LIFTA**, **MkATOM** and **UPDREG**, which are the ones dealing with proving atomicity.

B.2.1— The λ -safety condition. The rules of TaDA Live dealing with opening and closing regions (rules **UPDREG** and **LIFTA**) require the λ -safety side condition for the postcondition. While the definition of λ -safety is technical, its intuition is simple: those are the assertions that preserve their meaning when interpreted at level λ or at level $\lambda + 1$. The only possible contradictions arising by increasing the level come from assertions about the state and environment obligations of regions that are open at λ but not at $\lambda + 1$.

Definition B.2 (*Havoc*). Let $\lambda \in \text{Lvl}$. The set $\text{closed}_{\lambda_1}^{\lambda_2}(\rho) \triangleq \{r \mid \rho(r) = (_, \lambda, _), \lambda_1 \leq \lambda < \lambda_2\}$ is the set of region ids of ρ that are closed at level λ_2 but not at level λ_1 . We define the function on

worlds:

$$\text{havoc}_\lambda(h, \rho, \gamma, \chi, \theta, \xi) \triangleq \left\{ (h, \rho', \gamma, \chi, \theta, \xi') \left| \begin{array}{l} \text{closed}_\lambda^{\lambda+1}(\rho) = \{r_1, \dots, r_n\}, \\ \rho(r_i) = (\mathbf{t}_i, _ , _), b_i \in \text{AVal}, w_i \in \mathcal{I}_{\mathbf{t}_i} \llbracket r_i, \lambda, b_i \rrbracket, \\ \rho' = \rho[r_1 \mapsto (\mathbf{t}_1, \lambda, b_1), \dots, r_n \mapsto (\mathbf{t}_n, \lambda, b_n)], \\ O'_i \bullet \theta_{w_i}(r_i) = \xi(r_i), \xi' \sqsupseteq \xi[r_1 \mapsto O'_1, \dots, r_n \mapsto O'_n] \end{array} \right. \right\}$$

We extend it to a function on sets of worlds in the obvious way: $\text{havoc}_\lambda(p) \triangleq \bigcup_{w \in p} \text{havoc}_\lambda(w)$.

Definition B.3 (λ -safety). A set $p \in \text{World}_{\mathcal{A}}^\dagger$ is λ -safe if $p = \text{havoc}_\lambda(p)$. An assertion P is λ -safe, written $\mathcal{A} \vDash P$ λ -safe if, for all ς , $\mathcal{W} \llbracket P \rrbracket_{\mathcal{A}}^\varsigma$ is λ -safe.

Since proving λ -safety in general involves meddling with the semantics of assertions, we provide the following lemma that can be used to immediately prove all the λ -safety side conditions involved in our program proofs.

LEMMA B.4. *The properties below hold, for arbitrary $\lambda \in \text{Lvl}$:*

- (1) $\text{emp}, \mathbb{E}_1 \mapsto \mathbb{E}_2$ and \mathbb{B} are λ -safe.
- (2) $\llbracket G \rrbracket_r$ and $\llbracket O \rrbracket_r^\dagger$ are both λ -safe.
- (3) If $\lambda' < \lambda$, then $\mathbf{t}_r^{\lambda'}(a) * \llbracket O \rrbracket_r^\dagger$ is λ -safe.
- (4) If P, Q are both λ -safe, then so are $P \wedge Q, P \vee Q$, and $P * Q$.
- (5) If $P(v)$ is λ -safe for all $v \in \text{AVal}$, then $\exists x. P(x)$ is λ -safe.

B.2.2— Generalised atomicity rules. The following rules are the general forms of rules **LIFTA**, **MkATOM** and **UPDREG** of Fig. 9.

$$\frac{\begin{array}{c} \lambda < \lambda' \quad r \notin \text{dom}(\mathcal{A}) \\ \mathcal{A}' = \mathcal{A}[r \mapsto (X, k, X', T)] \quad T \subseteq \mathcal{T}_\dagger(\mathbb{G}) \quad R = \text{io}(T) \quad \forall x \in X. \mathcal{A} \vDash \mathbf{t}_r^\lambda(x) * \llbracket G \rrbracket_r \text{ stable} \\ m; \lambda'; \mathcal{A}' \vdash_\Phi \{ P_h * \exists x \in X. \mathbf{t}_r^\lambda(x) * r \mapsto \blacklozenge \} \mathbb{C} \{ \exists x, y. R(x, y) * Q_h(x, y) * r \mapsto (x, y) \} \end{array}}{m; \lambda'; \mathcal{A} \vdash_\Phi \forall x \in X \rightarrow_k X'. \langle P_h \mid \mathbf{t}_r^\lambda(x) * \llbracket G \rrbracket_r \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid \mathbf{t}_r^\lambda(y) * \llbracket G \rrbracket_r * R(x, y) \rangle} \text{MkATOMG}$$

$$\frac{\begin{array}{c} r \in \text{dom}(\mathcal{A}) \quad \mathcal{A}' = \mathcal{A}[r \mapsto \perp] \\ \vdash_{\mathcal{A}} P_h \Rightarrow \text{emp}_{\text{Ob}}^r \quad \vdash_{\mathcal{A}} P_a(x) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} \\ \vdash_{\mathcal{A}} Q_h(x, y) \Rightarrow \text{emp}_{\text{Ob}}^r \quad \vdash_{\mathcal{A}} Q_i(x, y, z) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} \\ \mathcal{A} \vDash P_h \text{ } \lambda\text{-safe} \quad \mathcal{A} \vDash P_a(x) \text{ } \lambda\text{-safe} \end{array}}{\mathcal{A} \vDash Q_h(x, y) \text{ } \lambda\text{-safe} \quad \mathcal{A} \vDash \left(\frac{(R(x, z) \wedge Q_1(x, y, z))}{\vee (x = z \wedge Q_2(x, y))} \right) \lambda\text{-safe}}$$

$$\frac{\begin{array}{c} \{ ((x, O_0), (z, O_1(x, y))) \mid x \in X \wedge (R(x, z) \vee x = z) \wedge y \in Y(x) \} \subseteq \text{tr}(\mathcal{A}, r) \\ m; \lambda; \mathcal{A}' \vdash_\Phi \forall x \in X. \left\langle P_h \mid \frac{I(\mathbf{t}_r^\lambda(x))}{* P_a(x) * \llbracket O_0 \rrbracket_r^\dagger} \right\rangle \mathbb{C} \exists y. \left\langle Q_h(x, y) \wedge y \in Y(x) \right. \\ \left. \mid \exists z. I(\mathbf{t}_r^\lambda(z)) * \llbracket O_1(x, y) \rrbracket_r^\dagger * \left(\frac{(R(x, z) \wedge Q_1(x, y, z))}{\vee (x = z \wedge Q_2(x, y))} \right) \right\rangle \end{array}}{m; \lambda+1; \mathcal{A} \vdash_\Phi \forall x \in X. \left\langle P_h * \llbracket O_0 \rrbracket_r^\dagger \mid \mathbf{t}_r^\lambda(x) * P_a(x) * r \mapsto \blacklozenge \right\rangle \mathbb{C} \exists y. \left\langle Q_h(x, y) * \llbracket O_1(x, y) \rrbracket_r^\dagger \wedge y \in Y(x) \right. \\ \left. \mid \exists z. \mathbf{t}_r^\lambda(z) * \left(\frac{(R(x, z) \wedge Q_1(x, y, z) * r \mapsto (x, z))}{\vee (x = z \wedge Q_2(x, y) * r \mapsto \blacklozenge)} \right) \right\rangle} \text{UPDREGG}$$

$$\begin{array}{c}
\vdash_{\mathcal{A}} P_h \Rightarrow \text{emp}_{\text{Ob}}^r \quad \vdash_{\mathcal{A}} Q_h(x, y) \Rightarrow \text{emp}_{\text{Ob}}^r \\
\vdash_{\mathcal{A}} P_a(x) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} \quad \vdash_{\mathcal{A}} Q_a(x, y, z) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} \\
\mathcal{A} \models P_h \text{ } \lambda\text{-safe} \quad \mathcal{A} \models P_a(x) \text{ } \lambda\text{-safe} \\
\mathcal{A} \models Q_h(x, y) \text{ } \lambda\text{-safe} \quad \mathcal{A} \models Q_a(x, y, z) \wedge R(x, z) \text{ } \lambda\text{-safe} \\
r \in \text{dom}(\mathcal{A}) \Rightarrow R = \text{id} \quad \{((x, \text{O}_1), (z, \text{O}_2(x, y))) \mid x \in X \wedge R(x, z) \wedge y \in Y(x)\} \subseteq \mathcal{T}(\mathbb{G}) \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \left\langle P_h \left| \begin{array}{l} \mathcal{I}(t_r^\lambda(x)) * P_a(x) * \lceil G \rceil_r * \lfloor \text{O}_1 \rfloor_r^{\downarrow} \\ \mathbb{C} \exists y. \left\langle \begin{array}{l} Q_h(x, y) \wedge y \in Y(x) \\ \mid \exists z. \mathcal{I}(t_r^\lambda(z)) * Q_a(x, y, z) \\ \mid * \lfloor \text{O}_2(x, y) \rfloor_r^{\downarrow} \wedge R(x, z) \end{array} \right\rangle \right. \right\rangle \\
\hline
m; \lambda+1; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \left\langle P_h * \lfloor \text{O}_1 \rfloor_r^{\downarrow} \left| \begin{array}{l} t_r^\lambda(x) * P_a(x) * \lceil G \rceil_r \\ \mathbb{C} \exists y. \left\langle \begin{array}{l} Q_h(x, y) * \lfloor \text{O}_2(x, y) \rfloor_r^{\downarrow} \wedge y \in Y(x) \\ \mid \exists z. t_r^\lambda(z) * Q_a(x, y, z) \wedge R(x, z) \end{array} \right\rangle \right. \right\rangle
\end{array}
\right.
\end{array}$$

LIFTAG

B.3 General forms

The following rules are the general forms of some of the rules in Fig. 9.

$$\begin{array}{c}
\forall x \in X. \vdash_{\mathcal{A}} R_h * R_a(x) \triangleright m \\
\text{pv}(R_h, R_a(x)) \cap \text{mod}(\mathbb{C}) = \emptyset \quad \mathcal{A} \models R_h \text{ stable} \quad \forall x \in X. \mathcal{A} \models R_a(x) \text{ stable} \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \left\langle P_h \left| \begin{array}{l} P_a(x) \\ \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle \right. \right\rangle \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h * R_h \mid P_a(x) * R_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * R_h \mid Q_a(x, y) * R_a(x) \rangle \\
\hline
\mathcal{A} \models P_h * P \text{ stable} \quad \forall x \in X, y. \mathcal{A} \models Q(x, y) \text{ stable} \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h \mid P * P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q(x, y) * Q_a(x, y) \rangle \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h * P \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) * Q(x, y) \mid Q_a(x, y) \rangle \\
\hline
\mathcal{A} \models P_h * P \text{ stable} \quad \forall x \in X, z. \langle P_h \mid P_a(x, z) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y, z) \rangle \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h \mid \exists z \in Z. P_a(x, z) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid \exists z \in Z. Q_a(x, y, z) \rangle \\
\hline
k_1; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle \quad k_1 \leq k_2 \\
\hline
k_2; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle
\end{array}$$

FRAME

ATOMWG

A \exists ELIMG

LAYWG

B.4 Logical manipulation rules

The rules below allow for basic logical manipulation.

$$\begin{array}{c}
\mathcal{A} \models P_h \text{ stable} \quad \forall x \in X, y. \mathcal{A} \models Q_h(x, y) \text{ stable} \\
\lambda; \mathcal{A} \models P_h \Rightarrow P'_h \quad \forall x \in X, y. \lambda; \mathcal{A} \models Q'_h(x, y) \Rightarrow Q_h(x, y) \\
\forall x \in X. \lambda; \mathcal{A} \models P_a(x) \Leftrightarrow P'_a(x) \quad \forall x \in X, y. \lambda; \mathcal{A} \models Q'_a(x, y) \Rightarrow Q_a(x, y) \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P'_h \mid P'_a(x) \rangle \mathbb{C} \exists y. \langle Q'_h(x, y) \mid Q'_a(x, y) \rangle \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle \\
\hline
\forall v \in X. m; \lambda; \mathcal{A} \vdash_{\Phi} \{P(v)\} \mathbb{C} \{Q\} \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \{\exists x \in X. P(x)\} \mathbb{C} \{Q\} \\
\hline
\forall k \leq m. k; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h(k) \wedge k \leq m \mid P_a(k, x) \rangle \mathbb{C} \exists y. \langle Q_h(k, x, y) \mid Q_a(k, x, y) \rangle \\
\hline
\forall k \leq m. m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X}. \langle P_h(k) \mid P_a(k, x) \rangle \mathbb{C} \exists y. \langle Q_h(k, x, y) \mid Q_a(k, x, y) \rangle \\
\hline
f: X \rightarrow Y \quad Y' = f(X') \quad \forall x \in X. \vdash_{\mathcal{A}} P'_a(x) \Leftrightarrow P_a(f(x)) \\
\forall x \in X, z. \vdash_{\mathcal{A}} Q_h(f(x), z) \Rightarrow Q'_h(x, z) \quad \forall x \in X, z. \vdash_{\mathcal{A}} Q_a(f(x), z) \Rightarrow Q'_a(x, z) \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall y \in Y \rightarrow_k Y'. \langle P_h \mid P_a(y) \rangle \mathbb{C} \exists z. \langle Q_h(y, z) \mid Q_a(y, z) \rangle \\
\hline
m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X'. \langle P_h \mid P'_a(x) \rangle \mathbb{C} \exists z. \langle Q'_h(x, z) \mid Q'_a(x, z) \rangle
\end{array}$$

CONS

 \exists ELIM

QL

SUBPQ

$$\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X'' . \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y . \langle Q_h(x, y) \mid Q_a(x, y) \rangle \quad X' \subseteq X'' \subseteq X}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in X \rightarrow_k X' . \langle P_h \mid P_a(x) \rangle \mathbb{C} \exists y . \langle Q_h(x, y) \mid Q_a(x, y) \rangle} \text{LIVEW}$$

B.5 Axioms

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ \mathbb{E} \geq 0 \} \quad x := \mathbf{alloc}(\mathbb{E}) \left\{ \ast_{i=0}^{\mathbb{E}-1} x + i \mapsto _ \right\}} \text{ALLOC}$$

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ \mathbb{E} \mapsto _ \} \quad \mathbf{dealloc}(\mathbb{E}) \quad \{ \mathbf{emp} \}} \text{DEALLOC}$$

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall v . \langle \mathbb{E} \mapsto v \rangle \quad x := [\mathbb{E}] \langle \mathbb{E} \mapsto v \wedge x = v \rangle} \text{READ}$$

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall v . \langle \mathbb{E}_1 \mapsto v \rangle \quad [\mathbb{E}_1] := \mathbb{E}_2 \langle \mathbb{E}_1 \mapsto \mathbb{E}_2 \rangle} \text{MUTATE}$$

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall v . \langle \mathbb{E}_1 \mapsto v \rangle \quad x := \mathbf{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3) \left(\begin{array}{l} (x = 1 \wedge \mathbb{E}_1 \mapsto \mathbb{E}_3 \wedge v = \mathbb{E}_2) \vee \\ (x = 0 \wedge \mathbb{E}_1 \mapsto v \wedge v \neq \mathbb{E}_2) \end{array} \right)} \text{CAS}$$

$$\frac{}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall v . \langle \mathbb{E}_1 \mapsto v \rangle \quad x := \mathbf{FAS}(\mathbb{E}_1, \mathbb{E}_2) \langle \mathbb{E}_1 \mapsto \mathbb{E}_2 \wedge x = v \rangle} \text{FAS}$$

B.6 Standard Hoare rules

$$\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ P \} \mathbb{C}_1 \{ R \} \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \{ R \} \mathbb{C}_2 \{ Q \}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ P \} \mathbb{C}_1 ; \mathbb{C}_2 \{ Q \}} \text{SEQ}$$

$$\frac{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ P \wedge \mathbb{B} \} \mathbb{C}_1 \{ Q \} \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \{ P \wedge \neg \mathbb{B} \} \mathbb{C}_2 \{ Q \}}{m; \lambda; \mathcal{A} \vdash_{\Phi} \{ P \} \mathbf{if}(\mathbb{B}) \{ \mathbb{C}_1 \} \mathbf{else} \{ \mathbb{C}_2 \} \{ Q \}} \text{IF}$$

$$\frac{x \notin \text{fv}(P_h) \cup \text{fv}(Q_h) \cup \text{fv}(\mathbb{E}) \quad m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X} . \langle P_h \wedge x = \mathbb{E} \mid P_a(x) \rangle \mathbb{C} \langle Q_h(x, y) \mid Q_a(x, y) \rangle}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X} . \langle P_h \mid P_a(x) \rangle \mathbf{var} \ x = \mathbb{E} \mathbf{in} \mathbb{C} \langle Q_h(x, y) \mid Q_a(x, y) \rangle} \text{VAR}$$

$$\frac{(\vec{x}, \forall x \in \vec{X} . \langle P_h \mid P_a(x) \rangle \cdot \exists y . \langle Q_h(x, y, \text{ret}) \mid Q_a(x, y) \rangle)_{m; \lambda; \mathcal{A}} \in \Phi(\mathbf{f})}{m; \lambda; \mathcal{A} \vdash_{\Phi} \forall x \in \vec{X} . \langle P_h \mid \vec{\mathbb{E}} / \vec{x} \mid P_a(x) \rangle \quad z := \mathbf{f}(\vec{\mathbb{E}}) \exists y . \langle Q_h(x, y, z) \mid Q_a(x, y) \rangle} \text{CALL}$$

$$\frac{\text{pv}(\mathbb{S}) \subseteq \vec{x} \cup \{ \text{ret} \} \quad \mathbf{f} \notin \text{dom}(\Phi) \quad \Phi' = \Phi[\mathbf{f} \mapsto (\vec{x}, \mathbb{S})] \quad \vdash_{\Phi} \mathbb{C}_1 : \mathbb{S}_1 \quad \vdash_{\Phi'} \mathbb{C}_2 : \mathbb{S}_2}{\vdash_{\Phi} \mathbf{let} \ \mathbf{f}(\vec{x}) = \mathbb{C}_1 \ \mathbf{in} \ \mathbb{C}_2 : \mathbb{S}_2} \text{LET}$$

$$\frac{P, Q \in \text{SL} \quad \forall x \in X . \perp; 0; \emptyset \vdash_{\Phi} \{ P(x) \} \mathbb{C} \{ Q(x) \}}{\perp; 0; \emptyset \vdash_{\Phi} \forall x \in X . \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle} \text{PRAT}$$

B.7 On Stability Checks

A triple is well-defined, according to Definition 3.24, if the Hoare pre- and post-conditions are both stable assertions. The rules all assume the triples in the premises are well-defined and ensure that the triple in the conclusion is well-defined as well. The only exceptions are rules **MkATOMG**, **SUBPQ**, and **ΞELIM**, where the Hoare pre-/post-conditions should be checked for stability to ensure the conclusion is a well-defined triple. We omitted these stability checks from these rules to improve readability.

In practice, however, this way of handling stability has a drawback: if one starts with a goal that has unstable pre-/post-conditions, one would only see the mistake much further up in the proof, typically at applications of **ATOMW** or **FRAME** (which requires stability of the frames) just before

applications of the axioms. Therefore, in practice, to make the proof fail early in case of mistakes, one would want to additionally check stability at the top-level goal, and applications of `PAR`.

C CASE STUDY: LOCK-COUPLING SET

We develop the proof of a lock-coupling set module, which represents a set of integer numbers using an ordered linked list. The module interface presents three operations, `add`, `remove` and `member` for adding and removing elements from the abstract set representing the module’s state, as well as checking membership of an integer in this set.

Each cell of the linked list contains either a value from this set or $\pm\infty$ (representing dummy beginning and end nodes respectively), a pointer to a lock and a pointer to the next cell of the linked list (null for the final cell, with value ∞). The values of the cells in the linked list are sorted in strictly increasing order.

The value and lock associated with a cell in the linked list are immutable, however, the module’s protocol allows a thread holding the lock associated with a cell to change the value of the pointer to the next cell, allowing cells to be added and removed from the linked list.

The internal operation `locate` performs a traversal of the linked list using hand-over-hand locking so as to, given some value v , find and lock the two adjacent cells with values v' and v'' such that $v' < v \leq v''$. All the operations would use `locate` to obtain ownership of the nodes that they need to modify.

To perform this hand-over-hand locking, the `locate` operation must hold the lock associated with a cell while locking the lock associated with the next, therefore the layers of the locks associated with each cell of the linked list must strictly decrease as the list is traversed.

As we explained in Section 5.5, the example is challenging for the handling of layers. Intuitively, we want to associate layers with each lock in the list, in strictly decreasing order. This represent the dependencies between the locks introduced by the order of the traversal: the release of lock at position i from the head depends on the liveness of the lock at position $i + 1$. This introduces two challenges: we need to associate different layers to each instance of a lock while the lock specifications mention fixed layers; and we need to dynamically reassign layers to locks as the list grows. As we already anticipated, we can solve both challenges by a suitable generalisation of the lock specifications. Let us first introduce this generalisation formally, and then use it for the proof of the lock-coupling set.

C.1 Interlude: a Generalisation of Fair Lock Specifications

We generalise the fair lock specifications we used for the CLH lock in three ways:

- (1) we parametrise the specifications with client-definable layer maps;
- (2) we provide a viewshift to the client with which it is possible to reassign layers;
- (3) we add the `deleteLock` operation since the lock-coupling set’s `remove` operation disposes of the removed cells; we omit its implementation and proof as it is standard.

First let us recall the definition of a layer map. Given two partial orders $(\mathcal{L}_1, \leq_1, \top_1, \perp_1)$ and $(\mathcal{L}_2, \leq_2, \top_2, \perp_2)$, a function $\eta: \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is *strictly monotone* if $\forall m, n \in \mathcal{L}_1. m <_1 n \Rightarrow \eta(m) <_2 \eta(n)$. A *layer map* $\eta: \mathcal{L}_1 \rightarrow_{\text{lay}} \mathcal{L}_2$ is a strictly monotone function between the two partial orders.

Let $\lambda_{\text{clh}} - 2$ be the level of the **lclh** region used in the proof of the CLH lock. We generalise the client-facing CLH lock specifications as follows:

$\exists(\mathcal{L}_{\text{clh}}, \leq_{\text{clh}}, \top_{\text{clh}}, \perp_{\text{clh}}) \cdot \forall \eta: \mathcal{L}_{\text{clh}} \rightarrow_{\text{lay}} \mathcal{L}$.

$\eta(\top_{\text{clh}}); \lambda_{\text{clh}} \vdash \forall l \in \{0, 1\} \rightarrow_{\eta(\perp_{\text{clh}})} \{0\} \cdot \langle P(s, \pi) \mid L_{\eta}(s, x, l) \rangle \text{lock}(x) \langle P(s, \pi) \mid L_{\eta}(s, x, 1) \wedge l = 0 \rangle$

$\eta(\perp_{\text{clh}}); \lambda_{\text{clh}} \vdash \langle L_{\eta}(s, x, 1) \rangle \text{unlock}(x) \langle L_{\eta}(s, x, 0) \rangle$

$\eta(\perp_{\text{clh}}); \lambda_{\text{clh}} \vdash \{\text{emp}\} \text{makeLock}() \{\exists s. L_{\eta}(s, \text{ret}, 0) * P(s, 1)\}$

$\eta(\perp_{\text{clh}}); \lambda_{\text{clh}} \vdash \{L_{\eta}(s, x, _) * P(s, 1)\} \text{deleteLock}(x) \{\text{emp}\}$

In particular, the abstract predicate $L_{\eta}(s, x, l)$ represents a lock resource with abstract identifier $s \in \mathbb{S}_{\text{clh}}$ (i.e. a pair of region identifiers; the client will treat this type opaquely), concrete address $x \in \text{Addr}$, and abstract state $l \in \{0, 1\}$.

Moreover, the specifications would export to the client the following viewshifts, for every $\lambda \geq \lambda_{\text{clh}}$, and every $\eta, \eta': \mathcal{L}_{\text{clh}} \rightarrow_{\text{lay}} \mathcal{L}$:

$$\lambda \vDash L_{\eta}(s, x, l) * L_{\eta'}(s, x', l') \Rightarrow \text{False} \quad (26)$$

$$\lambda \vDash L_{\eta}(s, x, l) * P(s, 1) \Rightarrow L_{\eta'}(s, x, l) * P(s, 1) \quad (27)$$

Note that the naming choice here suggests CLH as the implementation to keep the discussion grounded, but the specification would be the same for any other fair lock implementation.

We now sketch the modifications needed to adapt the proof of CLH presented in Section 5.2 to prove the generalised specification.

First, we pick, just as in Section 5.2, $\mathcal{L}_{\text{clh}} = \mathbb{N} \uplus \{\perp_{\text{clh}}, \top_{\text{clh}}\}$. We then need to parametrise the two regions with a layer map $\eta: \mathcal{L}_{\text{clh}} \rightarrow_{\text{lay}} \mathcal{L}$, for an arbitrary \mathcal{L} . We include it in the regions abstract state: $\mathbf{clh}_r(r', x, \eta, h, l, o)$ and $\mathbf{lclh}_{r'}(x, \eta, h, l, o, t)$. The abstract predicate for the lock can then be defined as:

$$L_{\eta}(s, x, l) \triangleq \exists r, r'. s = (r, r') \wedge \exists o, h. \mathbf{clh}_r(r', x, \eta, h, l, o) * [\mathbf{E}]_{r'}$$

We similarly parametrise every obligation with a layer map as well, obtaining obligations $\mathbf{o}_{\eta}(o, t)$ and $\mathbf{p}_{\eta}(t)$ with layers $\text{lay}(\mathbf{o}_{\eta}(o, t)) = \eta(\perp_{\text{clh}})$ $\text{lay}(\mathbf{p}_{\eta}(t)) = \eta(t)$.

The protocol of the regions is extended by having each transition preserve the layer map. Before extending the protocol with a transition that can update the layer map, we motivate the need for fractional permissions by showing what goes wrong without them. Suppose we just provide a transition, guarded by \mathbf{E} , to update the current layer map to an arbitrary new one, and define $P(s, \pi) = \text{emp}$. With this protocol it would be impossible to prove the layer-map-altering viewshift (27). The reason lies in the definition of the interpretation of \mathbf{clh} :

$$I(\mathbf{clh}_r(r', x, \eta, h, l, o)) \triangleq \exists t \in \mathbb{N}. \mathbf{lclh}_{r'}(x, \eta, h, l, o, t) * [\mathbf{E}]_{r'} * \bigstar_{i=o+1}^{t-1} [\mathbf{p}_{\eta}(i)]_{r'}^{\mathbf{E}}$$

In the case where $t \neq o$, which represents the case where there are threads enqueued waiting to acquire the lock, the interpretation ensures that the environment will contain obligations $\mathbf{p}_{\eta}(i)$ for each issued ticket i . When we try to prove the viewshift, we need to obtain $[\mathbf{p}_{\eta'}(i)]_{r'}^{\mathbf{E}}$ with the new layer map, which can be obtained only by creating out-of-thin-air the corresponding $[\mathbf{p}_{\eta'}(i)]_{r'}^{\perp}$ resources. These would be created in the local state, leaving us with $\mathbf{clh}_r(r', x, \eta, h, l, o) * [\mathbf{E}]_{r'} * \bigstar_{i=o+1}^{t-1} [\mathbf{p}_{\eta'}(i)]_{r'}^{\perp}$, which cannot be viewshifted to the desired $L_{\eta'}(s, x, l) * P(s, 1)$ since there is no way to get rid of the local obligations. Conceptually this encodes the following fact: if we were to remap the layers of the lock when other threads are queued, the obligations held by those thread would become unfulfillable, and we would inherit copies of them with the new mapping, which we also would not be able to fulfil on behalf of the other threads.

To resolve this impasse, we need to allow the layer map to be update only when there is no thread queued to acquire the lock. This way we would have $t = o$ and so no environment obligation laying around. We cannot achieve this by exposing the queue in the abstract state of the lock, however, without loosing the atomicity of the lock specifications. With the introduction of fractional permissions, giving the right to enqueue to the lock, we can encode the emptiness of the queue by asserting we are the only one with that right.

To achieve this technically, we start by encoding fractional permissions as a guard algebra. We introduce guards F_π with the axioms $F_0 = \mathbf{0}$ and $F_{\pi_1 + \pi_2} = F_{\pi_1} \bullet F_{\pi_2}$. We then define the abstract predicate $P(s, \pi) = (\exists r, r'. s = (r, r') \wedge [F_\pi]_{r'})$.

For technical reasons explained later we introduce guards G_π with exactly the same axioms as the F guards. To encode the fact that full permissions imply empty queue, we adapt the interpretation of **lclh** as follows:

$$\begin{aligned} \mathcal{I}(\mathbf{lclh}_{r'}(x, \eta, h, l, o, t)) \triangleq & \exists ns. x \mapsto h, \text{last}(ns) * h \mapsto l * \text{ones}(ns) * [Q(ns, o)]_{r'} * [O_\eta(o, t)]_{r'}^\perp * \\ & \exists \pi. [F_\pi]_{r'} * [R_{1-\pi}]_{r'} * (\pi = 0 \Rightarrow t = o) \wedge t - o = |ns| \wedge ns(0) = h \end{aligned}$$

From $\mathbf{lclh}_{r'}(x, \eta, h, l, o, t) * [F_1]_{r'}$ we can deduce that $\pi = 0$ inside the region interpretation, and hence $t = o$.

Finally, we add to the protocol of **lclh** the possibility of updating the layer map when owning full permissions:

$$F_1 : ((\eta, h, l, o, t), \mathbf{0}) \rightsquigarrow ((\eta', h, l, o, t), \mathbf{0})$$

The reason for including the R_π is as follows. When a thread enqueues on the lock, it gives up a non-trivial fraction of the F_π permission it owns to be able to make $t \neq o$. When it dequeues, it should get back that fraction; the R_π guards are obtained as “leftovers” when putting F_π in the region’s interpretation. Those are proof that the region interpretation has at least F_π in it when we want to get it back.

Adapting the proof of Section 5.2 to use these generalised definitions is a routine application of standard TaDA patterns. The satisfiability of the layer constraints is preserved by strict monotonicity of layer maps.

C.2 Correctness of the Lock-Coupling Set

— *Code*. The implementation of the module’s operations is in Fig. 27 with the implementation of the constructor `makeSet` in Fig. 28. We write `dealloc(c, 3)` for the deallocation of the 3 contiguous cells from address `c`. The auxiliary operation `locate` (also in Fig. 28) is meant to only be used internally. The code uses a “record” syntax for readability, desugared as follows:

$$x.\text{lock} \triangleq [x] \qquad x.\text{val} \triangleq [x + 1] \qquad x.\text{next} \triangleq [x + 2]$$

— *Specifications*. The abstract predicate $\text{LCSet}(s, x, S)$ represents a lock-coupling set at address x abstractly representing the set S .

$$\begin{aligned} \perp_{lc} & \vdash \{\text{emp}\} \text{makeSet}() \{ \exists s. \text{LCSet}(s, \text{ret}, \emptyset) \} \\ T_{lc} & \vdash \forall S \in \wp(\mathbb{Z}). \langle \text{LCSet}(s, x, S) \wedge e \in \mathbb{Z} \rangle \text{add}(x, e) \langle \text{LCSet}(s, x, S \cup \{e\}) \rangle \\ T_{lc} & \vdash \forall S \in \wp(\mathbb{Z}). \langle \text{LCSet}(s, x, S) \wedge e \in \mathbb{Z} \rangle \text{remove}(x, e) \langle \text{LCSet}(s, x, S \setminus \{e\}) \rangle \\ T_{lc} & \vdash \forall S \in \wp(\mathbb{Z}). \langle \text{LCSet}(s, x, S) \wedge e \in \mathbb{Z} \rangle \text{member}(x, e) \langle \text{LCSet}(s, x, S) \wedge \text{ret} = (e \in S) \rangle \end{aligned}$$


```

1  def add(x, e) {
2    var p, c, v, n,
3      nl, pl, cl in
4    p := locate(x, e);
5    c := p.next;
6    v := c.val;
7
8    if(v ≠ e) {
9      n := alloc(3);
10     nl := makeLock();
11     n.lock := nl;
12     n.val := e;
13     n.next := c;
14     p.next := n;
15   }
16
17   pl := p.lock;
18   cl := c.lock;
19   unlock(cl);
20   unlock(pl)
21 }

```

```

1  def remove(x, e) {
2    var p, c, v,
3      n, pl, cl in
4    p := locate(x, e);
5    c := p.next;
6    v := c.val;
7    pl := p.lock;
8    cl := c.lock;
9
10   if(v = e) {
11     p.next = c.next;
12     deleteLock(cl);
13     dealloc(c, 3);
14   } else {
15     unlock(cl);
16   }
17   unlock(pl);
18 }

```

```

1  def member(x, e) {
2    var p, c, v,
3      n, pl, cl in
4    p := locate(x, e);
5    c := p.next;
6    v := c.val;
7
8    pl := p.lock;
9    cl := c.lock;
10   unlock(cl);
11   unlock(pl);
12
13   ret := (v = e)
14 }

```

Fig. 27. Implementation of the lock-coupling set operations.

```

1  def makeSet() {
2    var x, y in
3    y := alloc(3);
4    y.next := null;
5    y.val := ∞;
6    y.lock := makeLock();
7    x := alloc(3)
8    x.next := y;
9    x.val := -∞;
10   x.lock := makeLock();
11   ret := x
12 }

```

```

1  def locate(x, e) {
2    var p, c, c', v,
3      pl, cl, cl' in
4    p := x;
5    pl := p.lock;
6    lock(pl);
7    c := p.next;
8    cl := c.lock;
9    lock(cl);
10   v := c.value;
11   // continues...

```

```

12  // ...locate
13  while(v < e) {
14    pl := p.lock;
15    c' := c.next;
16    cl' := c'.lock;
17    lock(cl');
18    v := c'.val;
19    unlock(pl);
20    p := c;
21    c := c';
22  }
23  ret := p;
24 }

```

Fig. 28. Implementation of makeSet and the internal locate operation.

— *Region Types*. This proof will utilise two region types: $\text{lcset}_r(r', x, hl, shl, S)$ and $\text{lclist}_r(x, hl, shl, ls)$ where $r' \in \text{RId}$, $x, hl \in \text{Addr}$, $shl \in \mathbb{S}_{\text{clh}}$, $S \in \mathcal{P}(\mathbb{Z})$, $ls \in ((\mathbb{Z} \cup \{\infty, -\infty\}) \times \{0, 1\} \times (\mathbb{N} \cup \{1\}))^*$. Here r' , x , hl and shl are fixed parameters of both regions. The lock-coupling set resource is abstractly represented by the predicate

$$\text{LCSet}(s, x, S) \triangleq \exists r, r', hl, shl. s = (r, r', hl, shl) \wedge \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{E}]_r$$

– *Guards*. We introduce a number of guards that are used to represent ownership of information regarding nodes of the linked list. To ease readability, we will adopt a record notation for tuples (i.e. tuples with named positions). In particular, we will make a record n with the following information for each node: an address ($n.addr \in \text{Addr}$), a lock address ($n.lck \in \text{Addr}$), a lock abstract identifier ($n.lid \in \mathbb{S}_{\text{clh}}$), a value ($n.val \in \mathbb{Z} \cup \{\infty, -\infty\}$), and a layer ($n.lay \in \mathbb{N} \cup \{1\}$). A guard $\mathbf{c}(vs)$ records the list vs of values represented by the linked list. An unlocked node is represented by the guard $\mathbf{U}(n)$ where n is a record of the value, lock address/id, layer associated with the node of the list at address $n.addr$. So in particular, the cell at $n.addr$ would store the tuple $(n.lck, n.val, n.nxt)$ and we will have the resource $L_{\eta_{n.lay}}(n.lck, n.lid, l)$ associated with its lock (we will explain the layer map $\eta_{n.lay}$ when introducing the region interpretations). A locked node is represented by two guards $\mathbf{L}(n, a)$ and $\mathbf{K}(n, a)$, following the usual pattern for locks. These guards additionally store the address a of the next node which is stable if we hold the lock at $n.lck$. Moreover, assuming m is the node following n , if we hold the lock at $n.lck$, we know that all the information in m is stable (i.e. everything but the address of the node following m). To represent this we make use of a guard $\mathbf{W}(m)$.

The following axioms reflect the operations we desire to perform on the nodes. For locking/unlocking a non-terminal node, when $vs' \neq []$:

$$\begin{aligned} \mathbf{c}(vs \oplus [n.val, m.val] \oplus vs') \bullet \mathbf{U}(n) \bullet \mathbf{L}(m, a') &= \\ \mathbf{c}(vs \oplus [n.val, m.val] \oplus vs') \bullet \mathbf{L}(n, m.addr) \bullet \mathbf{K}(n, m.addr) \bullet \mathbf{L}(m, a') \bullet \mathbf{W}(m) & \end{aligned}$$

For locking/unlocking the last node:

$$\mathbf{c}(vs \oplus v) \bullet \mathbf{U}(n) = \mathbf{c}(vs \oplus v) \bullet \mathbf{L}(n, \text{null}) \bullet \mathbf{K}(n, \text{null})$$

For inserting a node m between n_1 and n_2 :

$$\begin{aligned} \mathbf{c}(vs \oplus [n_1.val, n_2.val] \oplus vs') \bullet \mathbf{L}(n_1, n_2.addr) \bullet \mathbf{K}(n_1, n_2.addr) \bullet \mathbf{W}(n_2) &= \\ \mathbf{c}(vs \oplus [n_1.val, m.val, n_2.val] \oplus vs') \bullet \mathbf{L}(n_1, m.addr) \bullet \mathbf{K}(n_1, m.addr) \bullet \mathbf{U}(m) \bullet \mathbf{W}(m) & \end{aligned}$$

For deleting a node m :

$$\begin{aligned} \mathbf{c}(vs \oplus [n.val, m.val] \oplus vs') \bullet \mathbf{L}(n, m.addr) \bullet \mathbf{K}(n, m.addr) \bullet \mathbf{L}(m, a) \bullet \mathbf{K}(m, a) &= \\ \mathbf{c}(vs \oplus [n.val] \oplus vs') \bullet \mathbf{L}(n, a) \bullet \mathbf{K}(n, a) & \end{aligned}$$

Then the following axioms keep the guard’s information for the nodes consistent:

$$\begin{aligned} n.val \notin vs &\Rightarrow \mathbf{c}(vs) \bullet \mathbf{K}(n) = \perp \\ n.val = n'.val &\Rightarrow \mathbf{K}(n, _) \bullet \mathbf{U}(n') = \perp \\ n.val = n'.val &\Rightarrow \mathbf{K}(n, _) \bullet \mathbf{K}(n', _) = \perp \\ (n.val = n'.val \wedge (a \neq a' \vee n \neq n')) &\Rightarrow \mathbf{K}(n, a) \bullet \mathbf{L}(n', a') = \perp \\ (n.addr = n'.addr \wedge n \neq n') &\Rightarrow \mathbf{K}(_, n.addr) \bullet \mathbf{W}(n) \bullet \mathbf{L}(n', _) = \perp \end{aligned}$$

– *Layers and Obligations*. We use the layer structure $\mathcal{L}_{\text{lc}} \triangleq (\mathbb{N} \cup \{1, 0\}) \times \mathcal{L}_{\text{clh}}$ (where $\forall n \in \mathbb{N}. 1 > n > 0$), ordered by the lexicographic ordering \leq and with $\mathbf{1}_{\text{lc}} \triangleq (1, \mathbf{1}_{\text{clh}})$ and $\mathbf{0}_{\text{lc}} \triangleq (0, \mathbf{0}_{\text{clh}})$. Roughly, take a non-initial node n which is at position ℓ from the end of the list; we will associate with it the layer $(\ell, \mathbf{1}_{\text{clh}})$, which is guaranteed to be strictly greater than any layer associated with the nodes following n in the list. Intuitively, no matter what \mathcal{L}_{clh} has been chosen for the proof of the implementation of locks, there are enough layers between $(\ell, \mathbf{1}_{\text{clh}})$ and $(\ell + 1, \mathbf{1}_{\text{clh}})$ to allow the proof of the lock of n not to conflict with the lock of the node ahead.

We construct obligations out of the atoms $\mathbf{k}(\ell)$ (representing the “key” of the lock associated with the layer ℓ) and $\mathbf{F}(\ell)$ (representing a “free” spot at layer ℓ) for $\ell \in \mathbb{N} \cup \{1\}$. We set $\text{lay}(\mathbf{k}(\ell)) \triangleq (\ell, \perp_{\text{clh}})$ and $\text{lay}(\mathbf{F}(\ell)) = \top_{\text{c}}$. We also define an obligation acting as a “reservoir” of atoms:

$$\mathbf{R}(\bar{\ell}) \triangleq \{\mathbf{k}(\ell) \mid \bar{\ell} \leq \ell \in \mathbb{N}\} \cup \{\mathbf{F}(\ell) \mid \bar{\ell} \leq \ell \in \mathbb{N}\} \quad \text{lay}(\mathbf{R}(\bar{\ell})) = (\bar{\ell}, \perp_{\text{clh}})$$

We can always split a pair of \mathbf{F} and \mathbf{k} atoms from the reservoir: $\mathbf{R}(\bar{\ell}) = \mathbf{R}(\bar{\ell} + 1) \bullet \mathbf{k}(\bar{\ell}) \bullet \mathbf{F}(\bar{\ell})$.

– *Interference protocol.* The guard-labelled transition system of the region $\text{lcset}_r(r', x, hl, shl, S)$ is:

$$\begin{aligned} \mathbf{E} &: \forall v. (S, \mathbf{0}) \rightsquigarrow (S \cup \{v\}, \mathbf{0}) \\ \mathbf{E} &: \forall v. (S, \mathbf{0}) \rightsquigarrow (S \setminus \{v\}, \mathbf{0}) \end{aligned}$$

and the guard-labelled transition system of the region $\text{lclist}_r(x, hl, shl, ls)$ is:

$$\begin{aligned} \mathbf{E} &: ((-\infty, 0, 1) \oplus ls, \mathbf{0}) \\ &\rightsquigarrow ((-\infty, 1, 1) \oplus ls, \mathbf{k}(1) \bullet \mathbf{F}(\ell)) \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 0, \ell') \oplus ls', \mathbf{0}) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell')) \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{F}(\ell^\dagger) \bullet \mathbf{k}(\ell')) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 0, \ell) \oplus (v', 1, \ell^\dagger) \oplus ls', \mathbf{k}(\ell^\dagger) \bullet \mathbf{F}(\ell')) \quad \ell > \ell^\dagger > \ell' \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{F}(\ell^\dagger) \bullet \mathbf{k}(\ell')) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 1, \ell) \oplus (v^\dagger, 0, \ell^\dagger) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{k}(\ell')) \quad \ell > \ell^\dagger > \ell', v < v^\dagger < v' \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{F}(\ell^\dagger) \bullet \mathbf{k}(\ell')) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 1, \ell) \oplus ls', \mathbf{k}(\ell)) \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{F}(\ell^\dagger) \bullet \mathbf{k}(\ell')) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 1, \ell) \oplus (v', 1, \ell') \oplus ls', \mathbf{k}(\ell) \bullet \mathbf{k}(\ell')) \quad \ell > \ell^\dagger > \ell' \\ \mathbf{k}(n, _) &: (ls \oplus (n.\text{val}, 1, \ell) \oplus ls', \mathbf{k}(\ell)) \\ &\rightsquigarrow (ls \oplus (n.\text{val}, 0, \ell) \oplus ls', \mathbf{0}) \end{aligned}$$

They represent, in order: the acquisition of the first lock, obtaining both the key for that lock and a “free” layer spot; the acquisition of a next lock, obtaining its key; the release of the previous lock, swapping layer of the next with the free one; the insertion of a node which gets assigned the free layer between the two adjacent locks held (used by add); the deletion of a node that also drops the non-needed free layer spot (used by remove); the drop of a non-needed free layer spot (used by member); the release of a lock.

– *Region interpretation.* The lock-coupling set internally represents the elements of the set with a lock-coupling linked list. To represent these in ghost state, we will use a list of quadruples of each cells value, the state of the associated lock, as well as its layer and region identifier. We introduce the predicate *ord* which verifies that the value in the list are in strictly increasing order, while the layers of the associated locks are in strictly decreasing order:

$$\text{ord}(ls) \triangleq \begin{cases} \text{True} & \text{if } ls = [] \\ v < v' \wedge \ell > \ell' \wedge \text{ord}((v', \ell', \ell') : ls') & \text{if } ls = (v, _, \ell) : (v', \ell', \ell') : ls' \end{cases}$$

We also introduce a function that allows us to extract the associated set of values from such a list, *vals*, and a function that similarly allows us to extract a list of just the values, retaining their order,

$lvals$:

$$\begin{aligned} vals(ls) &\triangleq \begin{cases} \emptyset & \text{if } ls = [] \\ \{v\} \uplus vals(ls') & \text{if } ls = (v, _, _) \oplus ls' \end{cases} \\ lvals(ls) &\triangleq \begin{cases} [] & \text{if } ls = [] \\ v \oplus vals(ls') & \text{if } ls = (v, _, _) \oplus ls' \end{cases} \end{aligned}$$

The interpretation of the outer region is a straightforward wrapper around the inner one.

$$\begin{aligned} \mathcal{I}(\text{lcset}_r(r', x, hl, shl, S)) &\triangleq \exists ls. \text{Iclist}_r(x, hl, shl, ls) * [\mathbf{E}]_r * \text{envK}_r(ls) \wedge \\ &S \uplus \{-\infty, \infty\} = vals(ls) \wedge \text{ord}(ls) \wedge ls = ((-\infty, _, _) \oplus _) \\ \text{envK}_r(ls) &\triangleq \begin{cases} \text{emp} & \text{if } ls = [] \\ (l = 1 \Rightarrow [\mathbf{K}(\ell)]_r^{\mathbf{E}}) * \text{envK}_r(ls') & \text{if } ls = (v, l, \ell) \oplus ls' \end{cases} \end{aligned}$$

As usual, the outer region has two purposes: hiding internal state so that the operations can be seen as abstractly atomic, and keeping track of the obligations held by threads.

The interpretation of the inner region encapsulates the concrete heap cells and the lock-related guards and obligations:

$$\begin{aligned} \mathcal{I}(\text{Iclist}_r(x, hl, shl, ls)) &\triangleq \exists n_0, l_0, \dots, n_{k+1}, l_{k+1}. \exists \bar{\ell} \in \mathbb{N}. \\ &[\mathbf{C}(lvals(ls))]_r * ls = [(n_0.\text{val}, l_0, n_0.\text{lay}), \dots, (n_{k+1}.\text{val}, l_{k+1}, n_{k+1}.\text{lay})] \wedge \\ &n_0.\text{val} = -\infty \wedge n_{k+1}.\text{val} = \infty \wedge n_0.\text{lay} = 1 \wedge n_{k+1}.\text{lay} = 0 \wedge \\ &n_0.\text{addr} = x \wedge n_0.\text{lck} = hl \wedge n_0.\text{lid} = shl \wedge \\ &[\mathbf{R}(\bar{\ell})]_r^{\mathbf{L}} \wedge \bar{\ell} > n_1.\text{lay} * \text{Node}_r^0(n_0, l_0, n_1.\text{addr}) * \\ &\text{Nodes}_r(\bar{\ell}, l_0, [(n_1, l_1), \dots, (n_{k+1}, l_{k+1})], \bar{\ell}) \end{aligned}$$

where the resources associated with each node are described by the following auxiliary predicates:

$$\begin{aligned} \text{Node}_r^0(n, l, a) &\triangleq (n.\text{addr} \mapsto n.\text{lck}, -\infty, a) * \\ &\mathbf{L}_{\eta_1}(n.\text{lid}, n.\text{lck}, l) * \exists \pi > 0. \text{P}(n.\text{lid}, \pi) * [\mathbf{F}(1)]_r^{\mathbf{L}} * \\ &((l = 0 \wedge [\mathbf{U}(n)]_r * [\mathbf{K}(1)]_r^{\mathbf{L}}) \vee (l = 1 \wedge [\mathbf{L}(n, a)]_r)) \\ \text{Nodes}_r(\ell_p, l_p, [(n, l)]) &\triangleq (n.\text{addr} \mapsto n.\text{lck}, n.\text{val}, \text{null}) * \text{Gaps}_r(\ell_p, n.\text{lay}) * \\ &\mathbf{L}_{\eta_{n.\text{lay}}}(n.\text{lid}, n.\text{lck}, l) * \text{P}(n.\text{lid}, (l_p=1 ? \frac{1}{2} : 1)) * [\mathbf{F}(n.\text{lay})]_r^{\mathbf{L}} * \\ &((l = 0 \wedge [\mathbf{U}(n)]_r * [\mathbf{K}(n.\text{lay})]_r^{\mathbf{L}}) \vee (l = 1 \wedge [\mathbf{L}(n, \text{null})]_r)) \\ \text{Nodes}_r(\ell_p, l_p, [(n, l), (n', l')]) \oplus ns &\triangleq (n.\text{addr} \mapsto n.\text{lck}, n.\text{val}, n'.\text{addr}) * \text{Gaps}_r(\ell_p, n.\text{lay}) * \\ &\mathbf{L}_{\eta_{n.\text{lay}}}(n.\text{lid}, n.\text{lck}, l) * \text{P}(n.\text{lid}, (l_p=1 ? \frac{1}{2} : 1)) * [\mathbf{F}(n.\text{lay})]_r^{\mathbf{L}} * \\ &((l = 0 \wedge [\mathbf{U}(n)]_r * [\mathbf{K}(n.\text{lay})]_r^{\mathbf{L}}) \vee (l = 1 \wedge [\mathbf{L}(n, n'.\text{addr})]_r)) * \\ &\text{Nodes}_r(n.\text{lay}, l, (n', l') \oplus ns) \\ \text{Gaps}_r(\ell_1, \ell_2) &\triangleq \bigstar_{\ell=\ell_1+1}^{\ell_2-1} ([\mathbf{K}(\ell)]_r^{\mathbf{L}} \vee [\mathbf{K}(\ell) \bullet \mathbf{F}(\ell)]_r^{\mathbf{L}}) \end{aligned}$$

The layer map η_ℓ maps the layers of \mathcal{L}_{clh} to the ones of \mathcal{L}_{lc} as follows:

$$\eta_\ell(k) = (\ell, k)$$

– *Proof of locate.* We use the following specification for the internal operation `locate`:

$$\mathbf{T}_{lc} \vdash \{ \exists S. \text{lcset}_r(r', x, hl, shl, S) \} \text{locate}(x, e) \{ \exists S. \text{lcset}_r(r', x, hl, shl, S) * \text{Loc}(r', x, e, \text{ret}) \}$$

where $\text{Loc}(r', x, e, p)$ represents the ownership of two adjacent list nodes representing value v and v' with $v < e \leq v'$ (where e is the value we wanted to locate in the list):

$$\begin{aligned} \text{Loc}(r', x, e, p) \triangleq & \exists n_1, n_2, n_3, \ell. n_1.\text{addr} = p \wedge \\ & n_1.\text{val} < e \leq n_2.\text{val} \wedge n_1.\text{lay} > \ell > n_2.\text{lay} \wedge \\ & [\mathbf{k}(n_1, n_2.\text{addr})]_{r'} * [\mathbf{k}(n_1.\text{lay})]_{r'}^L * [\mathbf{F}(\ell)]_{r'}^L * \\ & [\mathbf{k}(n_2, n_3.\text{addr})]_{r'} * [\mathbf{k}(n_2.\text{lay})]_{r'}^L * [\mathbf{w}(n_2)]_{r'} * \text{P}(n_2.\text{lid}, \frac{1}{2}) * \\ & (n_3.\text{addr} \neq \text{null} \Rightarrow ([\mathbf{w}(n_3)]_{r'} * \text{P}(n_3.\text{lid}, \frac{1}{2}))) \end{aligned}$$

The proof of `locate` is shown in Fig. 31. In the outlines, we expand the record notation to tuples, e.g. $\mathbf{k}(n.\text{addr}, n.\text{lck}, n.\text{lid}, n.\text{val}, n.\text{lay}, a)$. We detail here the application of `LIVEC`. The associated environment liveness condition is proved by:

$$\frac{\frac{\forall \alpha. \vdash_{\mathcal{A}} T'(\alpha) \Rightarrow T \quad (28)}{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : T'(\alpha) \rightarrow T} \text{LIVET} \quad \frac{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \rightarrow T}{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : L(\alpha) \rightarrow T} \text{EQUNT}}{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : L(\alpha) \rightarrow T} \text{ECASE} \quad \text{ENVLIVE}$$

where $L(\alpha) \triangleq L * M(\alpha)$ and

$$\begin{aligned} M(\alpha) & \triangleq \exists l. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', l, _) \oplus _) \wedge \alpha = l \\ L & \triangleq \exists l, \ell''. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', l, \ell'') \oplus _) * [\mathbf{w}(c', c1', \text{snl}, v'', \ell'')]_{r'} \\ & * l = 1 \Rightarrow [\mathbf{k}(\ell'')]_{r'}^E \wedge \ell' > \ell'' \\ L_1(\alpha) & \triangleq \exists \ell''. L'_{\ell''}(\alpha) \\ L'_{\ell''}(\alpha) & \triangleq \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', 1, \ell'') \oplus _) * [\mathbf{w}(c', c1', \text{snl}, v'', \ell'')]_{r'} \\ & * [\mathbf{k}(\ell'')]_{r'}^E \wedge \ell' > \ell'' \wedge \alpha = 1 \\ T'(\alpha) & \triangleq \exists \ell''. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', 0, \ell'') \oplus _) * [\mathbf{w}(c', c1', \text{snl}, v'', \ell'')]_{r'} \wedge \\ & \ell' > \ell'' \wedge \alpha = 0 \end{aligned}$$

$$\frac{\frac{\text{impr}_{\mathcal{A}}(L'_{\ell''}, L, T) \quad \forall \alpha. \vdash_{\mathcal{A}} L'_{\ell''}(\alpha) \triangleright \text{lay}(\mathbf{k}(\ell''))}{\forall \alpha. \vdash_{\mathcal{A}} L'_{\ell''}(\alpha) \Rightarrow \text{lclist}_{r'}^{\lambda}(x, hl, shl, _ \oplus (v'', 1, \ell'') \oplus _) * [\mathbf{k}(\ell'')]_{r'}^E * \text{True} \wedge (\ell', \mathbf{T}_{clh}) > \text{lay}(\mathbf{k}(\ell''))} \text{LIVEO}}{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : L'_{\ell''}(\alpha) \rightarrow T} \text{EQUNT}}{(\ell', \mathbf{T}_{clh}); \lambda; \mathcal{A} \vdash L(\alpha) : L_1(\alpha) \rightarrow T} \text{EQUNT} \quad (28)$$

– *Proof of add.* The proof of the `add` operation builds on the specification of `locate`. We show its outline in Fig. 32 with a more detailed derivation showing how the first unlock operation is handled in Fig. 33.

– *Proof of makeSet, member and remove.* We omit the proofs of the `makeSet`, `member` and `remove` operations as they do not add much to the presentation. `makeSet` can be proved as standard by keeping track of the nodes created locally and with a final viewshift to create the two nested regions

PROOF OF locate(x, e):

$\mathbb{T}_{IC}; \emptyset \vdash$
 $\{ \exists S. \text{lcset}_r(r', x, hl, shl, S) \}$
 $p := x;$
 $\{ \exists S. \text{lcset}_r(r', x, hl, shl, S) \wedge p = x \}$
 $pl := p. \text{lock};$
 $\{ \exists S. \text{lcset}_r(r', x, hl, shl, S) * \exists \pi > 0. P(shl, \pi) * \}$
 $\{ \exists l. \text{lclist}_r(x, hl, shl, (-\infty, l, 1) \oplus _) * l = 1 \Rightarrow \lfloor \mathbf{K}(1) \rfloor_r^E, \wedge p = x \wedge pl = hl \}$
 $\text{lock}(pl);$
 $\{ \exists \ell^\dagger, \ell', S, c, scl. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, hl, shl, -\infty, 1, c) \rfloor_{r'} * \lfloor \mathbf{K}(1) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, scl, _) \rfloor_{r'} * P(scl, \frac{1}{2}) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge 1 > \ell^\dagger > \ell' \}$
 $c := p. \text{next};$
 $\{ \exists \ell^\dagger, \ell', S, scl. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, hl, shl, -\infty, 1, c) \rfloor_{r'} * \lfloor \mathbf{K}(1) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, scl, _) \rfloor_{r'} * P(scl, \frac{1}{2}) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge 1 > \ell^\dagger > \ell' \}$
 $c1 := c. \text{lock};$
 $\{ \exists \ell^\dagger, \ell', S, scl, v'. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, hl, shl, -\infty, 1, c) \rfloor_{r'} * \lfloor \mathbf{K}(1) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, c1, scl, v', \ell') \rfloor_{r'} * P(scl, \frac{1}{2}) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge 1 > \ell^\dagger > \ell' * \}$
 $\{ \exists l. \text{lclist}_r(x, hl, shl, _ \oplus (v', l, \ell') \oplus _) * l = 1 \Rightarrow \lfloor \mathbf{K}(\ell') \rfloor_r^E \}$
 $\text{lock}(c1);$
 $\{ \exists \ell^\dagger, \ell', S, scl, snl, n, v'. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, hl, shl, -\infty, 1, c) \rfloor_{r'} * \lfloor \mathbf{K}(1) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, _) \rfloor_{r'} * \lfloor \mathbf{K}(c, c1, scl, v', \ell', n) \rfloor_{r'} * \lfloor \mathbf{K}(\ell') \rfloor_r^L * P(scl, \frac{1}{2}) * \}$
 $\{ (n \neq \text{null} \Rightarrow (\lfloor \mathbf{W}(n, _, snl, _) \rfloor_{r'} * P(snl, \frac{1}{2}))) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge 1 > \ell^\dagger > \ell' \}$
 $v := c. \text{value};$
 $\{ \exists \ell^\dagger, \ell', S, scl, snl, n. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, hl, shl, -\infty, 1, c) \rfloor_{r'} * \lfloor \mathbf{K}(1) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, _) \rfloor_{r'} * \lfloor \mathbf{K}(c, c1, scl, v, \ell', n) \rfloor_{r'} * \lfloor \mathbf{K}(\ell') \rfloor_r^L * P(scl, \frac{1}{2}) * \}$
 $\{ (n \neq \text{null} \Rightarrow (\lfloor \mathbf{W}(n, _, snl, _) \rfloor_{r'} * P(snl, \frac{1}{2}))) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge 1 > \ell^\dagger > \ell' \}$
 $\{ \exists \ell, \ell^\dagger, \ell', S, scl, snl, n, v. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, _, _, v, \ell, c) \rfloor_{r'} * \lfloor \mathbf{K}(\ell) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, _) \rfloor_{r'} * \lfloor \mathbf{K}(c, _, scl, v, \ell', n) \rfloor_{r'} * \lfloor \mathbf{K}(\ell') \rfloor_r^L * P(scl, \frac{1}{2}) * \}$
 $\{ (n \neq \text{null} \Rightarrow (\lfloor \mathbf{W}(n, _, snl, _) \rfloor_{r'} * P(snl, \frac{1}{2}))) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge \ell > \ell^\dagger > \ell' \wedge v < e \wedge v < v' \}$
 $\text{while}(v < e) \{$
 $pl := p. \text{lock};$
 $c' := c. \text{next};$
 $c1' := c'. \text{lock};$
 $\text{lock}(c1');$
 $v := c'. \text{val};$
 $\text{unlock}(pl);$
 $p := c;$
 $c := c';$
 $\}$
 $\{ \exists \ell, \ell^\dagger, \ell', S, scl, snl, c, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(p, _, _, v, \ell, c) \rfloor_{r'} * \lfloor \mathbf{K}(\ell) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, _) \rfloor_{r'} * \lfloor \mathbf{K}(c, _, scl, v', \ell', n) \rfloor_{r'} * \lfloor \mathbf{K}(\ell') \rfloor_r^L * P(scl, \frac{1}{2}) * \}$
 $\{ (n \neq \text{null} \Rightarrow (\lfloor \mathbf{W}(n, _, snl, _) \rfloor_{r'} * P(snl, \frac{1}{2}))) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \}$
 $\text{ret} := p;$
 $\{ \exists \ell, \ell^\dagger, \ell', S, scl, snl, c, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * \lfloor \mathbf{K}(\text{ret}, _, _, v, \ell, c) \rfloor_{r'} * \lfloor \mathbf{K}(\ell) \rfloor_r^L * \}$
 $\{ \lfloor \mathbf{W}(c, _, _) \rfloor_{r'} * \lfloor \mathbf{K}(c, _, scl, v', \ell', n) \rfloor_{r'} * \lfloor \mathbf{K}(\ell') \rfloor_r^L * P(scl, \frac{1}{2}) * \}$
 $\{ (n \neq \text{null} \Rightarrow (\lfloor \mathbf{W}(n, _, snl, _) \rfloor_{r'} * P(snl, \frac{1}{2}))) * \lfloor \mathbf{F}(\ell^\dagger) \rfloor_r^L, \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \}$

Fig. 29. Proof outline of locate.

$$\begin{array}{l}
\mathbf{T}_c: \emptyset \vdash \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', S, scl, snl, n, v. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v, \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * \\
(n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) * [\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \ell > \ell^\dagger > \ell' \wedge v < e \wedge v < v
\end{array} \right\} \\
\text{while}(v < e) \{ \\
\forall \beta. \mathbf{T}_c: \emptyset \vdash \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', S, scl, snl, n, v. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v, \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * P(snl, \frac{1}{2}) * \\
[\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \beta \geq \ell > \ell^\dagger > \ell' \wedge v < v < e
\end{array} \right\} \\
p1 := p. \text{lock}; \\
c' := c. \text{next}; \\
c1' := c'. \text{lock}; \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, v, v'. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, p1, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v, \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * [\mathbf{W}(c', c1', snl, v'', \ell'')]_{r'} * P(snl, \frac{1}{2}) * \\
[\mathbf{F}(\ell^\dagger)]_{r'}^\perp * \exists l. \text{Icllist}_r(x, hl, shl, _ \oplus (v'', l, \ell'') \oplus _) * l = 1 \Rightarrow [\mathbf{K}(\ell'')]_{r'}^\perp \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v < e \wedge v < v''
\end{array} \right\} \\
\text{lock}(c1'); \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, p1, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v, \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * [\mathbf{W}(c', c1', snl, v'', \ell'')]_{r'} * \\
[\mathbf{K}(c', _ _ _ snl, v'', \ell'', n)]_{r'} * [\mathbf{K}(\ell'')]_{r'}^\perp * P(snl, \frac{1}{2}) * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl', \frac{1}{2})) * [\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v < e \wedge v < v''
\end{array} \right\} \\
v := c'. \text{val}; \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, p1, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v', \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * [\mathbf{W}(c', c1', snl, v, \ell'')]_{r'} * \\
[\mathbf{K}(c', _ _ _ snl, v, \ell'', n)]_{r'} * [\mathbf{K}(\ell'')]_{r'}^\perp * P(snl, \frac{1}{2}) * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl', \frac{1}{2})) * [\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v' < e \wedge v' < v
\end{array} \right\} \\
\text{unlock}(p1); \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * \\
[\mathbf{K}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v', \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * [\mathbf{W}(c', c1', snl, v'', \ell'')]_{r'} * \\
[\mathbf{K}(c', _ _ _ snl, v, \ell'', n)]_{r'} * [\mathbf{K}(\ell'')]_{r'}^\perp * P(snl, \frac{1}{2}) * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl', \frac{1}{2})) * [\mathbf{F}(\ell')]_{r'}^\perp \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v' < e \wedge v' < v
\end{array} \right\} \\
p := c; \\
c := c'; \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', S, scl, snl, n, v. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v, \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * \\
(n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) * [\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \ell > \ell^\dagger > \ell' \wedge v < e \wedge v < v
\end{array} \right\} \\
\} \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', S, scl, snl, c, n, v, v'. \text{lcset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^\perp * \\
[\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'}^\perp * P(scl, \frac{1}{2}) * \\
(n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) * [\mathbf{F}(\ell^\dagger)]_{r'}^\perp \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v'
\end{array} \right\}
\end{array}$$

Fig. 30. Details of while loop in locate.

$$\begin{array}{l}
\mathbf{T}_c; \emptyset \vdash \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, v, v''. \text{lset}_r(r', x, hl, shl, S) * [\mathbf{K}(p, pl, _ , v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^L * \\
[\mathbf{W}(c, _ , _ , _)]_{r'} * [\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^L * P(scl, \frac{1}{2}) * [\mathbf{W}(c', cl', snl, v'', \ell'')]_{r'} * P(snl, \frac{1}{2}) * \\
[\mathbf{F}(\ell^\dagger)]_{r'}^L * \exists l. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', l, \ell'') \oplus _) * l = 1 \Rightarrow [\mathbf{K}(\ell'')]_{r'}^E \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v < e \wedge v < v''
\end{array} \right\} \\
(\ell', \mathbf{T}_{clh}); \emptyset \vdash \\
\mathbf{W}ls, ls' \in \mathbb{Z}^*, \ell', l \in \{0, 1\}. \\
\left\{ \begin{array}{l}
\exists l, \ell''. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', l, \ell'') \oplus _) * P(snl, \frac{1}{2}) * |\text{lclist}_{r'}(x, hl, shl, ls \oplus ((v'', l, \ell'') \oplus ls')) * [\mathbf{E}]_{r'} * \\
[\mathbf{W}(c', cl', snl, v'', \ell'')]_{r'} * l = 1 \Rightarrow [\mathbf{K}(\ell'')]_{r'}^E \wedge \ell' > \ell'' \mid [\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} \wedge \ell' > \ell''
\end{array} \right\} \\
(\ell', \mathbf{T}_{clh}); \emptyset \vdash \\
\mathbf{W}ls, ls' \in \mathbb{Z}^*, l \in \{0, 1\} \Rightarrow (\ell', \perp_{clh}) \{0\}. \\
\left\{ \begin{array}{l}
P(snl, \frac{1}{2}) \mid \text{lclist}_{r'}(x, hl, shl, ls \oplus ((v'', l, \ell'') \oplus ls')) * [\mathbf{E}]_{r'} * \\
[\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} * [\mathbf{W}(c', cl', snl, v'', \ell'')]_{r'} \wedge \ell' > \ell''
\end{array} \right\} \\
\text{STEP 8} \quad \text{LIVEC} \quad (\ell'', \mathbf{T}_{clh}); \emptyset \vdash \\
\mathbf{W}l \in \{0, 1\} \Rightarrow (\ell'', \perp_{clh}) \{0\}. \\
\left\{ \begin{array}{l}
P(snl, \frac{1}{2}) \mid L_{\eta_{\ell''}}(snl, cl', l) \\
\text{lock}(cl'); \\
P(snl, \frac{1}{2}) \mid L_{\eta_{\ell''}}(snl, cl', 1) \wedge l = 0
\end{array} \right\} \\
\left\{ \begin{array}{l}
P(snl, \frac{1}{2}) \mid \left\{ \begin{array}{l}
\exists n. \text{lclist}_{r'}(x, hl, shl, ls \oplus ((v'', 1, \ell'') \oplus ls')) * [\mathbf{E}]_{r'} * \\
[\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} * [\mathbf{K}(c', _ , snl, v'', n)]_{r'} * [\mathbf{K}(\ell'')]_{r'}^L * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ , snl', _)]_{r'} * P(snl', \frac{1}{2}))
\end{array} \right\} \\
\left\{ \begin{array}{l}
\exists l, \ell''. \text{lclist}_{r'}(x, hl, shl, _ \oplus (v'', l, \ell'') \oplus _) * \left\{ \begin{array}{l}
\exists n. \text{lclist}_{r'}(x, hl, shl, ls \oplus ((v'', 1, \ell'') \oplus ls')) * [\mathbf{E}]_{r'} * \\
P(snl, \frac{1}{2}) * [\mathbf{W}(c', cl', snl, v'', \ell'')]_{r'} * \left\{ \begin{array}{l}
[\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} * [\mathbf{K}(c', _ , snl, v'', n)]_{r'} * [\mathbf{K}(\ell'')]_{r'}^L * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ , snl', _)]_{r'} * P(snl', \frac{1}{2}))
\end{array} \right\}
\end{array} \right\} \\
l = 1 \Rightarrow [\mathbf{K}(\ell'')]_{r'}^E \wedge \ell' > \ell'' \mid n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ , snl', _)]_{r'} * P(snl', \frac{1}{2}))
\end{array} \right\} \\
\left\{ \begin{array}{l}
\exists \ell, \ell^\dagger, \ell', \ell'', S, scl, snl, n, v, v''. \text{lset}_r(r', x, hl, shl, S) * \\
[\mathbf{K}(p, pl, _ , v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'}^L * [\mathbf{W}(c, _ , _ , _)]_{r'} * \\
[\mathbf{K}(c, _ , scl, v, \ell', c')]_{r'} * [\mathbf{K}(\ell')]_{r'}^L * P(scl, \frac{1}{2}) * [\mathbf{W}(c', _ , _ , _)]_{r'} * \\
[\mathbf{K}(c', _ , snl, v'', \ell'')]_{r'} * [\mathbf{K}(\ell'')]_{r'}^L * P(snl, \frac{1}{2}) * \\
n \neq \text{null} \Rightarrow (\exists snl'. [\mathbf{W}(n, _ , snl', _)]_{r'} * P(snl', \frac{1}{2})) * [\mathbf{F}(\ell^\dagger)]_{r'}^L \wedge \\
\beta \geq \ell > \ell^\dagger > \ell' > \ell'' \wedge v < e \wedge v < v''
\end{array} \right\}
\end{array}$$

Fig. 31. Details of lock in while loop of locate.

STEP 8 is $\exists\text{ELIM}$, ATOMW , $\text{A}\exists\text{ELIM}$, LIFTA , QL , FRAME .STEP 9 is LIFTA , QL , CONS , FRAME .

representing an empty set. The hard part of the proofs of member and remove is the call to locate which has been already presented in detail. The rest is handled analogously to add.

 PROOF OF $\text{add}(x, e)$:

 $\top_{\mathbb{C}}; \emptyset \vdash \mathbf{WS} \in \mathcal{P}(\mathbb{Z})$.

 $\langle \text{LCSet}(s, x, S) \rangle$
 $\langle \text{lcsset}_r(r', x, hl, S) * [\mathbb{E}]_r \rangle$
 $\top_{\mathbb{C}}; \mathcal{A} = [r \mapsto (\wp(\mathbb{Z}), \perp, \wp(\mathbb{Z}), \{(S, \mathbf{0}), (S \cup \{e\}, \mathbf{0}) \mid S \subseteq \mathbb{Z}\})] \vdash$
 $\{ \exists S. \text{lcsset}_r(r', x, hl, S) * r \Rightarrow \diamond \}$
 $p := \text{locate}(x, e);$

$$\left\{ \begin{array}{l} \{ \exists \ell, \ell', S, scl, snl, c, n, v, v'. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * \\ [\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * \\ (n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) * [\mathbf{F}(\ell^\dagger)]_{r'} \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \end{array} \right\}$$

$$\left\{ \begin{array}{l} \{ \exists S. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * \\ [\mathbf{W}(c, _ _ _ _)]_{r'} * [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * \\ (n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) * [\mathbf{F}(\ell^\dagger)]_{r'} \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \end{array} \right\}$$

$$\left\{ \begin{array}{l} \{ \exists S. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(c, _ _ _ _)]_{r'} * \\ [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * [\mathbf{F}(\ell^\dagger)]_{r'} \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \end{array} \right\}$$
 $c := p.\text{next};$
 $v := c.\text{val};$

$$\left\{ \begin{array}{l} \{ \exists S. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(c, _ _ _ _)]_{r'} * \\ [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * [\mathbf{F}(\ell^\dagger)]_{r'} \wedge \ell > \ell^\dagger > \ell' \wedge v < e \leq v' \wedge v = v' \end{array} \right\}$$
 $\text{if}(v \neq e) \{$

$$\left\{ \begin{array}{l} \{ \exists S. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(c, _ _ _ _)]_{r'} * \\ [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * [\mathbf{F}(\ell^\dagger)]_{r'} \wedge \ell > \ell^\dagger > \ell' \wedge v < e < v' \end{array} \right\}$$
 $n := \text{alloc}(3);$
 $n1 := \text{makeLock}();$
 $n.\text{lock} := n1;$
 $n.\text{val} := e;$
 $n.\text{next} := c;$

$$\left\{ \begin{array}{l} \{ \exists S. \text{lcsset}_r(r', x, hl, shl, S) * r \Rightarrow \diamond * [\mathbf{K}(p, _ _ _ v, \ell, c)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(c, _ _ _ _)]_{r'} * \\ [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * P(scl, \frac{1}{2}) * [\mathbf{F}(\ell^\dagger)]_{r'} * \\ \{ \exists s. n \mapsto n1, e, c * L_{\eta_{\ell^\dagger}}(s, n1, 0) * P(s, 1) \wedge \ell > \ell^\dagger > \ell' \wedge v < e < v' \end{array} \right\}$$
 $p.\text{next} := n;$

$$\left\{ \begin{array}{l} \{ \exists S, S'. snl. \text{lcsset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * [\mathbf{K}(p, _ _ _ v, \ell, n)]_{r'} * [\mathbf{K}(\ell)]_{r'} * \\ [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}) * [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} \wedge \ell > \ell' \end{array} \right\}$$
 $\}$

$$\left\{ \begin{array}{l} \{ \exists S, S'. snl. \text{lcsset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * [\mathbf{K}(p, _ _ _ v, \ell, n)]_{r'} * [\mathbf{K}(\ell)]_{r'} * \\ [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}) * [\mathbf{K}(c, _ _ _ scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} \wedge \ell > \ell' \end{array} \right\}$$
 $p1 := p.\text{lock};$
 $c1 := c.\text{lock};$

$$\left\{ \begin{array}{l} \{ \exists S, S'. snl. \text{lcsset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * [\mathbf{K}(p, p1, _ _ _ v, \ell, n)]_{r'} * [\mathbf{K}(\ell)]_{r'} * \\ [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}) * [\mathbf{K}(c, c1, scl, v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} \wedge \ell > \ell' \end{array} \right\}$$

$$\left\{ \begin{array}{l} \{ \exists S, S, snl'. \text{lcsset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * \\ [\mathbf{K}(p, p1, _ _ _ v, \ell, n)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl', \frac{1}{2}) * \\ [\mathbf{K}(c, c1, _ _ _ v', \ell', n)]_{r'} * [\mathbf{K}(\ell')]_{r'} * \\ (n \neq \text{null} \Rightarrow ([\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl, \frac{1}{2}))) \wedge \ell > \ell' \end{array} \right\}$$
 $\text{unlock}(c1);$

$$\left\{ \begin{array}{l} \{ \exists S, S', snl'. \text{lcsset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * \\ [\mathbf{K}(p, p1, _ _ _ v, \ell, n)]_{r'} * [\mathbf{K}(\ell)]_{r'} * [\mathbf{W}(n, _ _ _ _)]_{r'} * P(snl', \frac{1}{2}) \end{array} \right\}$$
 $\text{unlock}(p1);$
 $\{ \exists S. r \Rightarrow (S, S \cup \{e\}) \}$
 $\{ \exists S. r \Rightarrow (S, S \cup \{e\}) \}$
 $\langle \text{lcsset}_r(r', hl, x, S \cup \{e\}) * [\mathbb{E}]_r \rangle$
 $\langle \text{LCSet}(s, x, S \cup \{e\}) \rangle$

 $\text{Cons}; \text{Sub } s = (r, r', hl)$
 MikATOM
 $\exists \text{ELIM}$
 FRAME

Fig. 32. Proof outline of add operation.

$$\begin{array}{l}
\text{Ti}_c; \mathcal{A} \vdash \\
\left\{ \begin{array}{l} \exists S, S', snl'. \text{lcset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * \\ \llbracket \mathbf{k}(p, pl, _, v, \ell, n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell) \rrbracket_{r'}^L * \llbracket \mathbf{w}(n, _, snl', _) \rrbracket_{r'} * P(snl', \frac{1}{2}) * \\ \llbracket \mathbf{k}(c, cl, _, v', \ell', n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell') \rrbracket_{r'}^L * \\ (n \neq \text{null} \Rightarrow (\llbracket \mathbf{w}(n, _, snl, _) \rrbracket_{r'} * P(snl, \frac{1}{2}))) \wedge \ell > \ell' \end{array} \right\} \\
\left(\ell', \perp_{\text{clh}} \right); \mathcal{A} \vdash \\
\left\{ \begin{array}{l} \exists S. \text{lcset}_r(r', x, hl, shl, S) * \llbracket \mathbf{k}(c, cl, _, v', \ell', n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell') \rrbracket_{r'}^L * \\ (n \neq \text{null} \Rightarrow (\llbracket \mathbf{w}(n, _, snl, _) \rrbracket_{r'} * P(snl, \frac{1}{2}))) \end{array} \right\} \\
\left(\ell', \perp_{\text{clh}} \right); \mathcal{A} \vdash \\
\mathbb{W}S \in \mathcal{P}(\mathbb{Z}). \\
\left\{ \begin{array}{l} \llbracket \text{lcset}_r(r', x, hl, shl, S) * \llbracket \mathbf{k}(c, cl, _, v', \ell', n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell') \rrbracket_{r'}^L * \\ (n \neq \text{null} \Rightarrow (\llbracket \mathbf{w}(n, _, snl, _) \rrbracket_{r'} * P(snl, \frac{1}{2}))) \end{array} \right\} \\
\left(\ell', \perp_{\text{clh}} \right); \mathcal{A} \vdash \\
\mathbb{W}ls, ls' \in ((\mathbb{Z} \uplus \{-\infty, \infty\}) \times \{0, 1\} \times \mathbb{N})^*. \\
\left\{ \begin{array}{l} \llbracket \text{lclist}_{r'}(x, hl, shl, ls \oplus (v', 1, \ell') \oplus ls') * \llbracket \mathbf{k}(c, cl, _, v', \ell', n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell') \rrbracket_{r'}^L * \\ (n \neq \text{null} \Rightarrow (\llbracket \mathbf{w}(n, _, snl, _) \rrbracket_{r'} * P(snl, \frac{1}{2}))) \end{array} \right\} \\
\left(\ell', \perp_{\text{clh}} \right); \mathcal{A} \vdash \\
\langle L_{\eta_{\ell'}}(s, la, 1) \rangle \\
\text{unlock}(cl); \\
\langle L_{\eta_{\ell'}}(s, la, 0) \rangle \\
\langle \text{lclist}_{r'}(x, hl, shl, ls \oplus (v', 0, \ell') \oplus ls') \rangle \\
\langle \text{lcset}_r(r', x, hl, shl, S) \rangle \\
\left\{ \exists S. \text{lcset}_r(r', x, hl, shl, S) \right\} \\
\left\{ \begin{array}{l} \exists S, S', snl'. \text{lcset}_r(r', x, hl, shl, S') * r \Rightarrow (S, S \cup \{e\}) * \\ \llbracket \mathbf{k}(p, pl, _, v, \ell, n) \rrbracket_{r'} * \llbracket \mathbf{k}(\ell) \rrbracket_{r'}^L * \llbracket \mathbf{w}(n, _, snl', _) \rrbracket_{r'} * P(snl', \frac{1}{2}) \end{array} \right\}
\end{array}$$

Fig. 33. Details of unlock(cl) in add. STEP 10 is LIFTA, CONS, FRAME.

D PROGRAMMING LANGUAGE DEFINITION

We will make regular use of partial functions. We write $X \rightarrow Y$ for the set of partial function from X to Y , and $X \rightarrow_f Y$ for the set of finite partial function. Given $f: X \rightarrow Y$, we write $f(x) = \perp$ if f is undefined on x , and $\text{dom}(f) \triangleq \{x \mid f(x) \neq \perp\}$. We will use the notation $[x_1 \mapsto y_1, \dots, x_n \mapsto y_n]$ for the finite function that maps each of the x_i to y_i and is undefined on any other input. Given elements $x \in X$ and $y \in Y$, and functions $f: X \rightarrow Y$ and $g: X' \rightarrow Y'$, we define the functions $f[x \mapsto y]$ and $f \uplus g$ by:

$$(f[x \mapsto y])(z) \triangleq \begin{cases} y & \text{if } z = x \\ f(z) & \text{otherwise} \end{cases}$$

$$(f \uplus g)(x) \triangleq \begin{cases} f(x) & \text{if } x \in \text{dom}(f) \\ g(x) & \text{if } x \in \text{dom}(g) \end{cases} \quad \text{if } \text{dom}(f) \cap \text{dom}(g) = \emptyset$$

We write $f[x \mapsto \perp]$ for the partial function that is undefined on x but otherwise behaves like f . The union of two partial function $f \cup g$ is a well-defined partial function as long as $f(x) = g(x)$ where their domains overlap.

We use the *set of Booleans*, $\text{Bool} \triangleq \{\text{True}, \text{False}\} \ni b, b_1, b_2$, a *set of values*, $\text{Val} \triangleq \mathbb{Z} \cup \text{Bool} \ni v, v_1, v_2, \dots$, a *set of program variables*, $\text{PVar} \ni x, y, \dots$, and a *set of function names*, $\text{FName} \ni f, g, \dots$. The set PVar contains a special element, ret , that holds a function's return value. Heap addresses are represented by natural numbers, $\text{Addr} \triangleq \mathbb{N}$. The natural numbers in Val represent both numeric values and heap addresses.

Definition D.1 (Numeric and Boolean Expressions). Let Vars be an arbitrary set of variables, and Values and arbitrary set of values. The *set of numerical expressions*, $\text{Exp}(\text{Vars}, \text{Values}) \ni \mathbb{E}, \mathbb{E}_1, \mathbb{E}_2, \dots$, and the *set of boolean expressions*, $\text{BExp}(\text{Vars}, \text{Values}) \ni \mathbb{B}, \mathbb{B}_1, \mathbb{B}_2, \dots$, are defined by the grammars:

$$\begin{aligned} \mathbb{E} &::= v \mid x \mid \mathbb{E} + \mathbb{E} \mid \mathbb{E} - \mathbb{E} \mid \mathbb{E} * \mathbb{E} \mid \dots & \text{where } v \in \text{Values}, x \in \text{Vars} \\ \mathbb{B} &::= b \mid x \mid \neg \mathbb{B} \mid \mathbb{B} \wedge \mathbb{B} \mid \mathbb{E} = \mathbb{E} \mid \mathbb{E} < \mathbb{E} \mid \dots & \text{where } b \in \text{Bool}, x \in \text{Vars} \end{aligned}$$

The numeric and Boolean program expressions are defined by the sets $\text{Exp}(\text{PVar}, \text{Val})$ and $\text{BExp}(\text{PVar}, \text{Val})$ respectively. In Section 3.3, we also work with logical expressions built from both program and logical variables and values, hence the reason for the expression definition defined over an arbitrary variable and value sets.

The functions $\text{fv}_{\mathcal{E}}$ and $\text{fv}_{\mathcal{B}}$ provide the sets of free variables for the numeric and Boolean expressions respectively. They are defined inductively on the structure of expressions by:

$$\begin{aligned} \text{fv}_{\mathcal{E}}(v) &= \emptyset & v \in \text{Values} & & \text{fv}_{\mathcal{B}}(b) &= \emptyset & b \in \{\text{True}, \text{False}\} \\ \text{fv}_{\mathcal{E}}(x) &= \{x\} & x \in \text{Vars} & & \text{fv}_{\mathcal{B}}(x) &= \{x\} & x \in \text{Vars} \\ \text{fv}_{\mathcal{E}}(\mathbb{E}_1 + \mathbb{E}_2) &= \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) & & & \text{fv}_{\mathcal{B}}(\neg \mathbb{B}) &= \text{fv}_{\mathcal{B}}(\mathbb{B}) \\ \text{fv}_{\mathcal{E}}(\mathbb{E}_1 - \mathbb{E}_2) &= \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) & & & \text{fv}_{\mathcal{B}}(\mathbb{B}_1 \wedge \mathbb{B}_2) &= \text{fv}_{\mathcal{B}}(\mathbb{B}_1) \cup \text{fv}_{\mathcal{B}}(\mathbb{B}_2) \\ \text{fv}_{\mathcal{E}}(\mathbb{E}_1 * \mathbb{E}_2) &= \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) & & & \text{fv}_{\mathcal{B}}(\mathbb{E}_1 = \mathbb{E}_2) &= \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) \\ \dots & & & & \text{fv}_{\mathcal{B}}(\mathbb{E}_1 < \mathbb{E}_2) &= \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) \\ & & & & \dots & \end{aligned}$$

Definition D.2 (Commands). The *set of commands*, $\text{Cmd} \ni \mathbb{C}$, is defined by the grammar in Fig. 34 where $\mathbb{E} \in \text{Exp}(\text{PVar}, \text{Val})$, $\mathbb{B} \in \text{BExp}(\text{PVar}, \text{Val})$, $x \in \text{PVar}$, $\vec{x} \in \text{PVar}^*$ is a list of pairwise distinct variables, and $f \in \text{FName}$.

We use $[\mathbb{E}]$ to denote the value of the heap cell with address given by \mathbb{E} . In Fig. 35, we define operators fv and mods , which identify the variables that a command can access and the variables

$\mathbb{C} ::= \text{skip}$	(skip)
$x := \mathbb{E}$	(assignment)
$x := [\mathbb{E}]$	(read)
$[\mathbb{E}] := \mathbb{E}$	(write)
$x := \text{CAS}(\mathbb{E}, \mathbb{E}, \mathbb{E})$	(compare-and-swap)
$x := \text{FAS}(\mathbb{E}, \mathbb{E}, \mathbb{E})$	(fetch-and-set)
$x := \text{alloc}(\mathbb{E})$	(allocate)
$\text{dealloc}(\mathbb{E})$	(deallocate)
$\mathbb{C}; \mathbb{C}$	(sequential composition)
$\mathbb{C} \parallel \mathbb{C}$	(parallel composition)
$\text{let } f(\vec{x}) = \mathbb{C} \text{ in } \mathbb{C}$	(function definition)
$\text{var } x = \mathbb{E} \text{ in } \mathbb{C}$	(local variable binding)
$\text{if}(\mathbb{B})\{\mathbb{C}\}\text{else}\{\mathbb{C}\}$	(if)
$\text{while}(\mathbb{B})\{\mathbb{C}\}$	(while loop)
$x := f(\vec{\mathbb{E}})$	(function call)
$\langle \mathbb{C} \rangle$	(primitive atomic block)

Fig. 34. Syntax of commands

$\text{pv}(\text{skip}) = \emptyset$	$\text{mods}(\text{skip}) = \emptyset$
$\text{pv}(x := \mathbb{E}) = \{x\} \cup \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(x := \mathbb{E}) = \{x\}$
$\text{pv}(x := [\mathbb{E}]) = \{x\} \cup \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(x := [\mathbb{E}]) = \{x\}$
$\text{pv}([\mathbb{E}_1] := \mathbb{E}_2) = \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2)$	$\text{mods}([\mathbb{E}_1] := \mathbb{E}_2) = \emptyset$
$\text{pv}(x := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3)) =$ $\{x\} \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_1) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_2) \cup \text{fv}_{\mathcal{E}}(\mathbb{E}_3)$	$\text{mods}(x := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3)) = \{x\}$
$\text{pv}(x := \text{alloc}(\mathbb{E})) = \{x\} \cup \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(x := \text{alloc}(\mathbb{E})) = \{x\}$
$\text{pv}(\text{dealloc}(\mathbb{E})) = \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(\text{dealloc}(\mathbb{E})) = \emptyset$
$\text{pv}(\text{let } f(\vec{x}) = \mathbb{C}_f \text{ in } \mathbb{C}) = \text{pv}(\mathbb{C})$	$\text{mods}(\text{let } f(\vec{x}) = \mathbb{C}_f \text{ in } \mathbb{C}) = \text{mods}(\mathbb{C})$
$\text{pv}(\text{var } x = \mathbb{E} \text{ in } \mathbb{C}) = (\text{pv}(\mathbb{C}) \setminus \{x\}) \cup \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(\text{var } x = \mathbb{E} \text{ in } \mathbb{C}) = \text{mods}(\mathbb{C}) \setminus \{x\}$
$\text{pv}(\text{if}(\mathbb{B})\{\mathbb{C}_1\}\text{else}\{\mathbb{C}_2\}) = \text{fv}_{\mathcal{B}}(\mathbb{B}) \cup \text{pv}(\mathbb{C}_1) \cup \text{pv}(\mathbb{C}_2)$	$\text{mods}(\text{if}(\mathbb{B})\{\mathbb{C}_1\}\text{else}\{\mathbb{C}_2\}) = \text{mods}(\mathbb{C}_1) \cup \text{mods}(\mathbb{C}_2)$
$\text{pv}(\text{while}(\mathbb{B})\{\mathbb{C}\}) = \text{fv}_{\mathcal{B}}(\mathbb{B}) \cup \text{pv}(\mathbb{C})$	$\text{mods}(\text{while}(\mathbb{B})\{\mathbb{C}\}) = \text{mods}(\mathbb{C})$
$\text{pv}(x := f(\vec{\mathbb{E}})) = \{x\} \cup \text{fv}_{\mathcal{E}}(\mathbb{E})$	$\text{mods}(x := f(\vec{\mathbb{E}})) = \{x\}$
$\text{pv}(\mathbb{C}_1; \mathbb{C}_2) = \text{pv}(\mathbb{C}_1) \cup \text{pv}(\mathbb{C}_2)$	$\text{mods}(\mathbb{C}_1; \mathbb{C}_2) = \text{mods}(\mathbb{C}_1) \cup \text{mods}(\mathbb{C}_2)$
$\text{pv}(\mathbb{C}_1 \parallel \mathbb{C}_2) = \text{pv}(\mathbb{C}_1) \cup \text{pv}(\mathbb{C}_2)$	$\text{mods}(\mathbb{C}_1 \parallel \mathbb{C}_2) = \text{mods}(\mathbb{C}_1) \cup \text{mods}(\mathbb{C}_2)$

Fig. 35. The sets of free and modified program variables

that are potentially modified by a command, respectively. In a command $\mathbb{C}_1 \parallel \mathbb{C}_2$, we apply a strong syntactic restriction that $\text{mods}(\mathbb{C}_1) = \text{mods}(\mathbb{C}_2) = \emptyset$. Each individual thread is still able to modify variables that are created locally and to modify shared heap cells, but are not allowed to modify the free variables.¹⁹ In a function definition $\text{let } f(x_1, \dots, x_n) = \mathbb{C}_1 \text{ in } \mathbb{C}_2$, we use the natural restriction $\text{fv}(\mathbb{C}_1) \subseteq \{x_1, \dots, x_n, \text{ret}\}$. Also for simplicity, we assume each function name is given a definition at most once. The function $\text{fn}: \text{Cmd} \rightarrow \mathcal{P}(\text{FName})$ returns the function names occurring in Cmd that are not bound by a **let**.

Definition D.3 (Variable Store). A program variable store, $\sigma \in \text{Store} \triangleq \text{PVar} \rightarrow \text{Val}$, is a finite partial function from program variables to values. The *right-biased union* of variable stores, $\sigma_1 \triangleleft \sigma_2$,

¹⁹To lift this restriction, one could use standard techniques, such as “variables as resources” [2]. Our restriction minimises the noise generated by handling local state in the formalisation of the model and the assertions. Note that expressivity is not really limited by our restriction: any local variable in the scope common to both threads, that needs to be modified, can be instead implemented by using a shared memory cell.

is defined by:

$$(\sigma_1 \triangleleft \sigma_2)(x) = \begin{cases} \sigma_2(x) & \text{if } x \in \text{dom}(\sigma_2) \\ \sigma_1(x) & \text{otherwise} \end{cases}$$

Definition D.4 (Expression evaluation). Let $\zeta : \text{Vars} \rightarrow_f \text{Values}$ be an arbitrary function from an arbitrary set of variables to values. The *numeric expression evaluation function*, $\mathcal{E}[\cdot]_\zeta : \text{Exp}(\text{Vars}, \text{Values}) \rightarrow \text{Values}$, and the *Boolean expression evaluation function*, $\mathcal{B}[\cdot]_\zeta : \text{BExp}(\text{Vars}, \text{Values}) \rightarrow \text{Bool}$, are defined by:

$$\begin{array}{ll} \mathcal{E}[v]_\zeta = v & \mathcal{B}[b]_\zeta = b \\ \mathcal{E}[x]_\zeta = \zeta(x) & \mathcal{B}[\neg B]_\zeta = \neg \mathcal{B}[B]_\zeta \\ \mathcal{E}[\mathbb{E}_1 + \mathbb{E}_2]_\zeta = \mathcal{E}[\mathbb{E}_1]_\zeta + \mathcal{E}[\mathbb{E}_2]_\zeta & \mathcal{B}[\mathbb{B}_1 \wedge \mathbb{B}_2]_\zeta = \mathcal{B}[\mathbb{B}_1]_\zeta \wedge \mathcal{B}[\mathbb{B}_2]_\zeta \\ \mathcal{E}[\mathbb{E}_1 - \mathbb{E}_2]_\zeta = \mathcal{E}[\mathbb{E}_1]_\zeta - \mathcal{E}[\mathbb{E}_2]_\zeta & \mathcal{B}[\mathbb{E}_1 = \mathbb{E}_2]_\zeta = (\mathcal{E}[\mathbb{E}_1]_\zeta = \mathcal{E}[\mathbb{E}_2]_\zeta) \\ \mathcal{E}[\mathbb{E}_1 \cdot \mathbb{E}_2]_\zeta = \mathcal{E}[\mathbb{E}_1]_\zeta \cdot \mathcal{E}[\mathbb{E}_2]_\zeta & \mathcal{B}[\mathbb{E}_1 < \mathbb{E}_2]_\zeta = (\mathcal{E}[\mathbb{E}_1]_\zeta < \mathcal{E}[\mathbb{E}_2]_\zeta) \\ \dots & \dots \end{array}$$

The program expressions are evaluated using program store $\sigma \in \text{Store}$. In Section 3.3, we also work with logical expressions which are evaluated over both program and logical variables and values. The right-biased union of stores is used to describe how, when nesting scopes, a variable occurrence is bound by the innermost binder surrounding it. The notation $\mathbf{var} \ x_1 = \theta \ \mathbf{in} \ \mathbf{var} \ x_2 = \theta \ \mathbf{in} \ \dots \ \mathbf{var} \ x_n = \theta \ \mathbf{in} \ \mathbb{C}$ denotes $\mathbf{var} \ x_1 = \theta \ \mathbf{in} \ \mathbf{var} \ x_2 = \theta \ \mathbf{in} \ \dots \ \mathbf{var} \ x_n = \theta \ \mathbf{in} \ \mathbb{C}$.

Definition D.5 (Heap). A *heap*, $h \in \text{Heap} \triangleq \text{Addr} \rightarrow_f \text{Val}$, is a finite partial function from addresses to values. The *set of heaps*, Heap , forms a PCM $(\text{Heap}, \uplus, \{\emptyset\})$ with $h_1 \uplus h_2$ defined only if $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$.

Definition D.6 (Function Implementation Context). A *function implementation context*, $\varphi \in \text{Flmpl} \triangleq \text{FName} \rightarrow (\text{PVar}^*, \text{Cmd})$, is a finite partial function from function names to pairs comprising a finite list of distinct variables and a command.

We write $\varphi(f) = (\vec{x}, \mathbb{C})$, where variable list \vec{x} represents the function arguments and \mathbb{C} represents the function body. We use the notation φ_{var} and φ_{cmd} to refer to the arguments and function body of f respectively.

In order to describe the behaviour of local variable binding and function calls, we define program states which extends commands with variable stores. For example, the program state (σ, \mathbb{C}) indicates that the command \mathbb{C} is evaluated in the current store updated with the variables in σ .

Definition D.7 (Program States). The *set of program states*, $\text{PState} \ni C, C_1, C_2, \dots$ is defined by the grammar:

$$C ::= \checkmark \mid (\sigma, C) \mid C; C \mid \mathbf{let} \ f(\vec{x}) = \mathbb{C} \ \mathbf{in} \ C \mid C \parallel C \mid \mathbb{C}$$

The \checkmark indicates a terminated program. It is a technical device so that every $\mathbb{C} \in \text{Cmd}$, including **skip**, takes at least one step.

In the operational semantics, we need to keep track of which thread is originating each step to be able to define later concepts of fairness of the scheduling. We do this tracking using *thread identifiers* $t \in \text{TId} \triangleq \{L, R\}^*$ which are strings of letters L (for the left thread) and R (for the right thread). ϵ will be used to denote the thread identifier which is an empty sequence. Intuitively, such a string identifies a single thread as the path in the syntax tree of parallel compositions at which the thread is found.

Definition D.8 (Command Semantics). A scheduler annotation t is an element of the set

$$\text{Sched} \triangleq \{\text{loc}_t \mid t \in \text{Tid}\} \uplus \{\text{env}\}.$$

A *program configuration* c is an element of the set $\text{PConf} \triangleq (\text{Store} \times \text{Heap} \times \text{PState}) \uplus \{\zeta\}$. Let $\varphi \in \text{Flmpl}$. The *operational semantics of the commands* is given by the labelled relation, $\longrightarrow_\varphi \subseteq \text{PConf} \times \text{Sched} \times \text{PConf}$, defined in Fig. 36 and Fig. 37. We write $a \xrightarrow{t}_\varphi b$ for $(a, t, b) \in \longrightarrow_\varphi$. We also define $\xrightarrow{\text{loc}}_\varphi^* \triangleq (\bigcup_{t \in \text{Tid}} \xrightarrow{\text{loc}_t}_\varphi)^*$.

To simplify the development, in our programming language the initial state’s store assigns arbitrary values to the free variables of a program. With such assumption, every reference to a local variable will be in the domain of the current store. This ensures that in every application of the rules in Fig. 36 and Fig. 37 to construct a trace, the evaluations of (boolean) expressions are well-defined.

Definition D.9 (Threads). Given a program state $c \in \text{PConf}$, the set $\text{threads}(c)$ is the set of threads of c that can take a step. The function $\text{threads}: \text{PConf} \rightarrow \wp(\text{Tid})$ is defined as follows:

$$\begin{aligned} \text{threads}(\zeta) &\triangleq \emptyset \\ \text{threads}(c) &\triangleq \{t \in \text{TID} \mid c \xrightarrow{\text{loc}_t}_\varphi _ \} \end{aligned}$$

Definition D.10 (Program Traces and Fairness). We call *program traces*, the infinite sequences of the form $c_0 \pi_0 c_1 \pi_1 \dots$ where, for all $i \in \mathbb{N}$, $c_i \in \text{PConf}$, $\pi_i \in \text{Sched}$. We use τ for ranging over infinite suffixes of program traces and PTrace for the set of all program traces. For a program trace $\tau = c_0 \pi_0 c_1 \pi_1 \dots$, we define $\tau(i) \triangleq (c_i, \pi_i)$, and $\tau_{/i} \triangleq c_i \pi_i c_{i+1} \pi_{i+1} \dots$. We define the *set of φ -program traces*

$$\text{PTrace}_\varphi \triangleq \{c_0 \pi_0 c_1 \pi_1 \dots \mid \forall i \in \mathbb{N}. c_i \xrightarrow{\pi_i}_\varphi c_{i+1}\}.$$

A program trace $(c_0 \pi_0 c_1 \pi_1 \dots) \in \text{PTrace}_\varphi$ is (*weakly*) *fair* if and only if:

$$\forall i \in \mathbb{N}. \forall t \in \text{threads}(c_i). \exists j \geq i. (\pi_j = \text{loc}_t \vee c_j = \zeta) \quad (29)$$

$$\forall i \in \mathbb{N}. \exists j \geq i. \pi_j = \text{env} \quad (30)$$

That is: a trace is fair if, at any point in time, every thread that can take a step (and the environment) will eventually be scheduled.

The open-world program semantics defines the behaviour of a command when run concurrently with an arbitrary environment. This semantics interleaves steps from two “players”: the local thread given by the loc relation; and its environment given by the env relation, respectively.

Definition D.11 (Open World Semantics). We call *traces* the infinite sequences $c_0 \pi_0 c_1 \pi_1 \dots$ where, for all $i \in \mathbb{N}$, $c_i \in \text{Conf} \triangleq (\text{Store} \times \text{Heap}) \cup \{\zeta\}$, $\pi_i \in \{\text{loc}, \text{env}\}$. We use τ for ranging over infinite suffixes of traces and Trace for the set of all traces. For a trace $\tau = c_0 \pi_0 c_1 \pi_1 \dots$, we define $\tau(i) \triangleq (c_i, \pi_i)$, and $\tau_{/i} \triangleq c_i \pi_i c_{i+1} \pi_{i+1} \dots$. The function $[\cdot]: \text{PTrace} \rightarrow \text{Trace}$ is defined by $[c_0 \pi_0 c_1 \pi_1 \dots] \triangleq c_0 \pi_0 c_1 \pi_1 \dots$ where

$$c_i \triangleq \begin{cases} (\sigma, h) & \text{if } c_i = (\sigma, h, _ _) \\ \zeta & \text{if } c_i = \zeta \end{cases} \quad \pi_i \triangleq \begin{cases} \text{loc} & \text{if } \pi_i \in \text{Sched} \setminus \{\text{env}\} \\ \text{env} & \text{if } \pi_i = \text{env} \end{cases}$$

The *open-world program semantics function*, $[\cdot]_\varphi: \text{Cmd} \rightarrow \wp(\text{Trace})$ is the function such that

$$[\mathbb{C}]_\varphi \triangleq \{ [c_0 \tau] \mid (c_0 \tau) \in \text{PTrace}_\varphi, \text{fv}(\mathbb{C}) \subseteq \text{dom}(\sigma_0), c_0 = (\sigma_0, _ _, \mathbb{C}), c_0 \tau \text{ is fair} \}$$

The notation $[\mathbb{C}]$ is syntactic sugar for $[\mathbb{C}]_\emptyset$.

$$\begin{array}{c}
\frac{}{\sigma, h, \text{skip} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \checkmark} \qquad \frac{}{\sigma, h, x := \mathbb{B} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto \mathcal{E}[\mathbb{B}]_\sigma], h, \checkmark} \\
\frac{\mathcal{E}[\mathbb{B}]_\sigma \in \text{dom}(h)}{\sigma, h, x := [\mathbb{B}] \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto h(\mathcal{E}[\mathbb{B}]_\sigma)], h, \checkmark} \qquad \frac{\mathcal{E}[\mathbb{B}_1]_\sigma \in \text{dom}(h)}{\sigma, h, [\mathbb{B}_1] := \mathbb{B}_2 \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h[\mathcal{E}[\mathbb{B}_1]_\sigma \mapsto \mathcal{E}[\mathbb{B}_2]_\sigma], \checkmark} \\
\frac{\mathcal{E}[\mathbb{B}_1]_\sigma \in \text{dom}(h) \quad h(\mathcal{E}[\mathbb{B}_1]_\sigma) = \mathcal{E}[\mathbb{B}_2]_\sigma}{\sigma, h, x := \text{CAS}(\mathbb{B}_1, \mathbb{B}_2, \mathbb{B}_3) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto 1], h[\mathcal{E}[\mathbb{B}_1]_\sigma \mapsto \mathcal{E}[\mathbb{B}_3]_\sigma], \checkmark} \\
\frac{\mathcal{E}[\mathbb{B}_1]_\sigma \in \text{dom}(h) \quad h(\mathcal{E}[\mathbb{B}_1]_\sigma) \neq \mathcal{E}[\mathbb{B}_2]_\sigma}{\sigma, h, x := \text{CAS}(\mathbb{B}_1, \mathbb{B}_2, \mathbb{B}_3) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto 0], h, \checkmark} \qquad \frac{a = \mathcal{E}[\mathbb{B}_1]_\sigma \in \text{dom}(h) \quad v = \mathcal{E}[\mathbb{B}_2]_\sigma}{\sigma, h, x := \text{FAS}(\mathbb{B}_1, \mathbb{B}_2) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto h(a)], h[a \mapsto v], \checkmark} \\
\frac{l = \mathcal{E}[\mathbb{B}]_\sigma \quad l > 0 \quad \{r, r+1, \dots, r+l-1\} \cap \text{dom}(h) = \emptyset \quad v_0, v_1, \dots, v_{l-1} \in \text{Val}}{\sigma, h, x := \text{alloc}(\mathbb{B}) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma[x \mapsto r], h[r \mapsto v_0, r+1 \mapsto v_1, \dots, r+l-1 \mapsto v_{l-1}], \checkmark} \\
\frac{\mathcal{E}[\mathbb{B}]_\sigma \in \text{dom}(h)}{\sigma, h, \text{dealloc}(\mathbb{B}) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h[\mathcal{E}[\mathbb{B}]_\sigma \mapsto \perp], \checkmark} \\
\frac{\sigma, h, C \xrightarrow{\text{loc}_t}_\varphi \sigma', h', C' \quad \varphi' = \varphi[f \mapsto (\vec{x}, C_f)]}{\sigma, h, \text{let } f(\vec{x}) = C_f \text{ in } C \xrightarrow{\text{loc}_t}_\varphi \sigma', h', \text{let } f(\vec{x}) = C_f \text{ in } C'} \qquad \frac{}{\sigma, h, \text{let } f(\vec{x}) = C_f \text{ in } \checkmark \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \checkmark} \\
\frac{}{\sigma, h, \text{var } x = \mathbb{B} \text{ in } C \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, ([x \mapsto \mathcal{E}[\mathbb{B}]_\sigma], C)} \qquad \frac{}{\sigma, h, (\sigma', \checkmark) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \checkmark} \\
\frac{\sigma \triangleleft \sigma_1, h, C \xrightarrow{\text{loc}_t}_\varphi \sigma' \triangleleft \sigma'_1, h', C' \quad \text{dom}(\sigma) = \text{dom}(\sigma') \quad \text{dom}(\sigma_1) = \text{dom}(\sigma'_1)}{\sigma, h, (\sigma_1, C) \xrightarrow{\text{loc}_t}_\varphi \sigma', h', (\sigma'_1, C')} \\
\frac{\mathcal{B}[\mathbb{B}]_\sigma}{\sigma, h, \text{if}(\mathbb{B})\{C_1\}\text{else}\{C_2\} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, C_1} \qquad \frac{\neg \mathcal{B}[\mathbb{B}]_\sigma}{\sigma, h, \text{if}(\mathbb{B})\{C_1\}\text{else}\{C_2\} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, C_2} \\
\frac{\mathcal{B}[\mathbb{B}]_\sigma}{\sigma, h, \text{while}(\mathbb{B})\{C\} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, C; \text{while}(\mathbb{B})\{C\}} \qquad \frac{\neg \mathcal{B}[\mathbb{B}]_\sigma}{\sigma, h, \text{while}(\mathbb{B})\{C\} \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \checkmark} \\
\frac{\varphi(f) = (\vec{x}, C)}{\sigma, h, y := f(\vec{\mathbb{B}}) \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \text{var } \text{ret} = \emptyset \text{ in } (\text{var } \vec{x} = \vec{\mathbb{B}} \text{ in } C); y := \text{ret}} \qquad \frac{\sigma, h, C_1 \xrightarrow{\text{loc}_t}_\varphi \sigma', h', C'_1}{\sigma, h, C_1; C_2 \xrightarrow{\text{loc}_t}_\varphi \sigma', h', C'_1; C'_2} \\
\frac{}{\sigma, h, \checkmark; C \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, C} \qquad \frac{\sigma, h, C_1 \xrightarrow{\text{loc}_t}_\varphi \sigma', h', C'_1}{\sigma, h, C_1 \parallel C_2 \xrightarrow{\text{loc}_{tL}}_\varphi \sigma', h', C'_1 \parallel C_2} \qquad \frac{\sigma, h, C_2 \xrightarrow{\text{loc}_t}_\varphi \sigma', h', C'_2}{\sigma, h, C_1 \parallel C_2 \xrightarrow{\text{loc}_{rL}}_\varphi \sigma', h', C_1 \parallel C'_2} \\
\frac{}{\sigma, h, \checkmark \parallel \checkmark \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma, h, \checkmark} \qquad \frac{\sigma, h, C \xrightarrow{\text{loc}^*}_\varphi \sigma', h', \checkmark}{\sigma, h, \langle C \rangle \xrightarrow{\text{loc}_\epsilon}_\varphi \sigma', h', \checkmark} \qquad \frac{h' \in \text{Heap}}{\sigma, h, C \xrightarrow{\text{env}}_\varphi \sigma, h', C}
\end{array}$$

Fig. 36. The small-step operational semantics

$$\begin{array}{c}
\frac{\mathcal{E}[\mathbb{E}]_{\sigma} \notin \text{dom}(h)}{\sigma, h, x := [\mathbb{E}] \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \quad \frac{\mathcal{E}[\mathbb{E}_1]_{\sigma} \notin \text{dom}(h)}{\sigma, h, [\mathbb{E}_1] := \mathbb{E}_2 \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \quad \frac{\mathcal{E}[\mathbb{E}_1]_{\sigma} \notin \text{dom}(h)}{\sigma, h, x := \text{CAS}(\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3) \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \\
\\
\frac{\mathcal{E}[\mathbb{E}_1]_{\sigma} \notin \text{dom}(h)}{\sigma, h, x := \text{FAS}(\mathbb{E}_1, \mathbb{E}_2) \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \quad \frac{\mathcal{E}[\mathbb{E}]_{\sigma} \notin \text{dom}(h)}{\sigma, h, \text{dealloc}(\mathbb{E}) \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \\
\\
\frac{\sigma, h, C \xrightarrow{\text{loc}_t}_{\varphi'} \zeta \quad \varphi' = \varphi[f \mapsto (\vec{x}, C_f)]}{\sigma, h, \text{let } f(\vec{x}) = C_f \text{ in } C \xrightarrow{\text{loc}_t}_{\varphi} \zeta} \quad \frac{\sigma \triangleleft \sigma', h, C \xrightarrow{\text{loc}_t}_{\varphi} \zeta}{\sigma, h, (\sigma', C) \xrightarrow{\text{loc}_t}_{\varphi} \zeta} \\
\\
\frac{f \notin \text{dom}(\varphi)}{\sigma, h, y := f(\vec{\mathbb{E}}) \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \quad \frac{\sigma, h, C_1 \xrightarrow{\text{loc}_t}_{\varphi} \zeta}{\sigma, h, C_1; C_2 \xrightarrow{\text{loc}_t}_{\varphi} \zeta} \quad \frac{\sigma, h, C_1 \xrightarrow{\text{loc}_t}_{\varphi} \zeta}{\sigma, h, C_1 \parallel C_2 \xrightarrow{\text{loc}_{lt}}_{\varphi} \zeta} \\
\\
\frac{\sigma, h, C \xrightarrow{\text{loc}}_{\varphi}^* \zeta}{\sigma, h, \langle C \rangle \xrightarrow{\text{loc}_{\epsilon}}_{\varphi} \zeta} \quad \frac{\sigma, h, C_2 \xrightarrow{\text{loc}_t}_{\varphi} \zeta}{\sigma, h, C_1 \parallel C_2 \xrightarrow{\text{loc}_{rt}}_{\varphi} \zeta} \quad \frac{c \in \text{PConf}}{c \xrightarrow{\text{env}}_{\varphi} \zeta}
\end{array}$$

Fig. 37. The small-step operational semantics, failure cases

Definition D.12. A trace $\tau \in \text{Trace}$ is *locally terminating*, written $\text{lterm}(\tau)$, if it contains finitely many occurrences of loc .

Remark 3 (Design of semantics). We made some design choices in crafting this semantics, with the motivation of making manipulation easier in the proofs. The first choice is to model environmental steps explicitly. These steps drive the argument about progress in the presence of blocking, where the local thread is not able to make progress in isolation but is relying on the environment actively performing some state changes that would lead to local progress.

The second choice we highlight is that the semantics of a program only contains infinite traces. This might seem odd when the goal is proving termination. Traces that locally terminate simply have an infinite tail of environment steps. To simulate a closed system one can select for the traces where the environment steps preserve the heaps. More importantly, we strip the information about threads and program state, which means that information about when the local thread terminated (in the form of \checkmark or end_t) has been erased. However, by construction, traces obtained from fair program traces can only contain finitely many local steps if the program terminated, justifying our definition of local termination.

Example D.13. The traces in $\llbracket [x] := y \rrbracket$ can be characterised as follows. They all start from some configuration (σ, h_0) with $x, y \in \text{dom}(\sigma)$. A (possibly zero) finite number of environment steps follow; these steps preserve the store, but arbitrarily alter the heap, or they lead to a fault, terminating the trace with an infinite tail of $\zeta \text{ env } \zeta \text{ env } \dots$ steps. If no fault happened, a local step is taken from some configuration (σ, h) for an arbitrary $h \in \text{Heap}$. If $\sigma(x) \notin \text{dom}(h)$ then the local step leads to a fault, leading again to a $\zeta \text{ env } \zeta \text{ env } \dots$ tail. Otherwise, it leads to the configuration $(\sigma, h[\sigma(x) \mapsto \sigma(y)])$. After that there is an infinite number of environment steps, which again preserve the store but arbitrarily mutate the heap, or lead to an infinite fault tail.

E SOUNDNESS

In this section, we provide the details of the soundness of three rules: **LIVEC**, **PAR**, **WHILE**, **FRAME**, **LIVEO** and **LIVEA**. These are the only proof rules in TaDA-Live that bring in non-trivial liveness information. All other proof rules follow in the same way as for TaDA, with the liveness constraints

on the traces being identical between the antecedent and consequent of such rules or being trivial in the case of command axioms. We will focus particularly on the liveness argument for these rules.

We start by giving some technical definitions omitted from the main text, and then move to the soundness argument.

E.1 Atomic World Rely

Recall that the *atomic world rely relation*, $\mathbf{R}_{\mathcal{A}}^a$, coincides with the smallest reflexive relation closed under the rules of the world rely (Fig. 6), with the restriction that rules \mathbf{WR}_1 and \mathbf{WR}_2 can be applied at most once per region identifier.

Definition E.1 (Atomic World Rely Relation). The atomic world rely relation, $\mathbf{R}_{\mathcal{A}}^a$, is defined as $\mathbf{R}_{\mathcal{A}}^a = \mathbf{R}_{\mathcal{A}}^{\emptyset}$, where $\mathbf{R}_{\mathcal{A}}^R$, taking $R \subseteq \text{Rld}$, is defined to be the smallest reflexive relation closed under:

$$\frac{\gamma(r) \# G \quad ((a_1, O_1), (a_2, O_2)) \in \mathcal{T}_i(G) \quad \chi(r) \in \{\diamond, \diamond\} \Rightarrow a_2 \in \text{safe}(\mathcal{A}, r) \quad O_2 \# \theta(r) \quad r \notin R \quad (h, \rho[r \mapsto (t, \lambda, a_2)], \gamma, \chi, \theta, \xi[r \mapsto O_2]) \mathbf{R}_{\mathcal{A}}^{R \uplus \{r\}} w'}{(h, \rho[r \mapsto (t, \lambda, a_1)], \gamma, \chi, \theta, \xi[r \mapsto O_1]) \mathbf{R}_{\mathcal{A}}^R w'} \quad \mathbf{WR}_1$$

$$\frac{O_2 \# \theta(r) \quad r \notin R \quad ((a_1, O_1), (a_2, O_2)) \in \text{tr}(\mathcal{A}, r) \quad (h, \rho[r \mapsto (t, \lambda, a_2)], \gamma, \chi[r \mapsto (a_1, a_2)], \theta, \xi[r \mapsto O_2]) \mathbf{R}_{\mathcal{A}}^{R \uplus \{r\}} w'}{(h, \rho[r \mapsto (t, \lambda, a_1)], \gamma, \chi[r \mapsto \diamond], \theta, \xi[r \mapsto O_1]) \mathbf{R}_{\mathcal{A}}^R w'} \quad \mathbf{WR}_2$$

E.2 Environment Liveness Judgement Semantics

We give semantics to the judgements defined in Fig. 10.

Definition E.2 (Auxiliary Environment Liveness Judgement Semantics). Let $m \in \mathcal{L}, \lambda \in \text{Lvl}, \mathcal{A} \in \text{ACtxt}, L, L' \in \mathbb{O} \rightarrow \text{Assrt}, T \in \text{Assrt}$ such that

- $\lambda; \mathcal{A} \vDash \exists \alpha. L(\alpha)$ stable.
- $\forall \alpha. \lambda; \mathcal{A} \vDash L'(\alpha) \Rightarrow L(\alpha)$.

and let

$$t_\sigma = \mathcal{W}[\![T * \text{True}]\!]_{\mathcal{A}}^\sigma \quad l_\sigma(\alpha) = \mathcal{W}[\![L(\alpha)]\!]_{\mathcal{A}}^\sigma \quad l'_\sigma(\alpha) = \mathcal{W}[\![L'(\alpha)]\!]_{\mathcal{A}}^\sigma$$

Then, the auxiliary semantic environmental liveness judgement $m; \lambda; \mathcal{A} \vDash L : L' \rightarrow T$ holds when, for arbitrary $\sigma \in \text{Store}$, there exist $P \subseteq \mathcal{P}(\mathbb{O} \rightarrow \text{World}_{\mathcal{A}}^\uparrow)$ such that $l'_\sigma(\alpha) = \bigcup_{lf \in P} lf(\alpha)$ and for all

$lf \in P$, either $\forall \alpha. lf(\alpha) \subseteq t$ or there exists some $r \in \text{Rld}$ and

$$\widehat{\mathcal{O}} \in \text{AOB}_{< m} \uplus \{ \text{live}(\mathcal{A}, r) \mid \text{lay}(\text{live}(\mathcal{A}, r)) < m \}$$

such that

- $\forall \alpha \in \mathbb{O}, w \in lf(\alpha). \text{active}_{r; \lambda}(w, \widehat{\mathcal{O}})$
- $\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(lf(\alpha_1)) \cap l_\sigma(\alpha_2) \subseteq lf(\alpha_1) \cup t$

hold, where:

$$\text{active}_{r; \lambda}(w, \widehat{\mathcal{O}}) \triangleq \begin{cases} \text{dep}_{r; \lambda}(w, \widehat{\mathcal{O}}) \wedge \xi_w(r) \sqsupseteq \widehat{\mathcal{O}} & \widehat{\mathcal{O}} \in \text{AOB} \\ \text{dep}_{r; \lambda}(w, \widehat{\mathcal{O}}) \wedge \chi_w(r) \# \diamond \wedge \text{ast}_w(r) \in X \setminus X' & \widehat{\mathcal{O}} = X \rightarrow_k X' \end{cases}$$

and

$$\text{dep}_{r; \lambda}(w, \widehat{\mathcal{O}}) \triangleq \forall r' \in \text{dom}(\theta_w). \text{lay}(\theta_w(r')) > \text{lay}(\widehat{\mathcal{O}}) \wedge \text{lvl}_w(r) < \lambda$$

Definition E.3 (Environment Liveness Judgement Semantics). The semantic environmental liveness judgement:

$$m; \lambda; \mathcal{A} \vDash L \xrightarrow{M} T$$

where $m \in \mathcal{L}$, $\lambda \in \text{Lvl}$, $\mathcal{A} \in \text{Actxt}$, $L \in \text{Assrt}$, $M \in \mathbb{O} \rightarrow \text{Assrt}$, $T \in \text{Assrt}$, holds when

$$\begin{aligned} \lambda; \mathcal{A} \vDash L \text{ stable} \\ \vdash_{\lambda; \mathcal{A}} L \Rightarrow \exists \alpha. L * M(\alpha) \\ m; \lambda; \mathcal{A} \vDash L * M(\alpha) : L * M(\alpha) \rightarrow T \end{aligned}$$

THEOREM E.4. For arbitrary $m \in \mathcal{L}$, $\lambda \in \text{Lvl}$, \mathcal{A} an atomicity context, $L, L' \in \mathbb{O} \rightarrow \text{Assrt}$, $T \in \text{Assrt}$ such that

$$\lambda; \mathcal{A} \vDash \exists \alpha. L(\alpha) \text{ stable} \tag{31}$$

$$\forall \alpha. \lambda; \mathcal{A} \vDash L'(\alpha) \Rightarrow L(\alpha) \tag{32}$$

if $m; \lambda; \mathcal{A} \vdash L : L' \rightarrow T$, then $m; \lambda; \mathcal{A} \vDash L : L' \rightarrow T$.

PROOF. Assuming $m; \lambda; \mathcal{A} \vdash L : L' \rightarrow T$ and taking $\sigma \in \text{Store}$ arbitrary, the proof proceeds by induction on the structure of derivation trees of the auxiliary environmental liveness condition. We start of with the bases cases: **LIVEO**, **LIVEA** and **LIVET**.

— case **LIVEO**. In this case, for some $r \in \text{Rld}$, $t \in \text{RType}$, $\lambda' \in \text{Lvl}$, $\mathbb{O} \in \text{AOB}$,

$$\text{impr}_{\mathcal{A}}(L', L, T) \tag{33}$$

$$m > \text{lay}(\mathbb{O}) \tag{34}$$

$$\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright \text{lay}(\mathbb{O}) \tag{35}$$

$$\lambda' < \lambda \tag{36}$$

$$\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x. t_r^\lambda(x) * [\mathbb{O}]_r^E * \text{True} \tag{37}$$

hold. From this, we need to show $m; \lambda; \mathcal{A} \vDash L(\alpha) : L'(\alpha) \rightarrow T$.

Let $P = \{l'_\sigma(\alpha)\}$, clearly the union of the elements of this set equals $l'_\sigma(\alpha)$ as required. Assuming $l'_\sigma(\alpha) \not\subseteq t_\sigma$ and setting $\widehat{\mathbb{O}} = \mathbb{O}$, which is in $\text{AOB}_{< m}$ given (34), it suffices to show

$$\forall \alpha \in \mathbb{O}, w \in l'_\sigma(\alpha). \text{active}_{r; \lambda}(w, \mathbb{O}) \tag{38}$$

$$\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(l'_\sigma(\alpha_1)) \cap l_\sigma(\alpha_2) \subseteq l'_\sigma(\alpha_1) \cup t \tag{39}$$

hold to complete the proof.

We start off by showing that (38) holds. Taking $\alpha \in \mathbb{O}$ and $w \in l'_\sigma(\alpha)$ arbitrary, given (35), it is clear that $\text{lay}(\theta_w(r)) \geq \text{lay}(\mathbb{O})$ holds and, given (36) and (37), $\text{lvl}_w(r) < \lambda$ holds. From these two conclusions, we can infer that $\text{dep}_{r; \lambda}(w, \mathbb{O})$ holds. Then, from (37), it is clear that $\xi_w(r) \supseteq \mathbb{O}$ holds, and therefore, $\text{active}_{r; \lambda}(w, \mathbb{O})$.

Finally, (39) follows immediately from (33) and the definition of $\text{impr}_{\mathcal{A}}$.

– case *LIVEA*. In this case, for some $r \in \text{Rld}$, $t \in \text{RType}$, $\lambda' \in \text{Lvl}$,

$$\text{impr}_{\mathcal{A}}(L', L, T) \quad (40)$$

$$m > k \quad (41)$$

$$\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \triangleright k \quad (42)$$

$$\text{live}(\mathcal{A}, r) = X \twoheadrightarrow_k X' \quad (43)$$

$$\lambda' < \lambda \quad (44)$$

$$\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow \exists x. t_r^\lambda(x) * r \Rightarrow \diamond \wedge x \in X \setminus X' * \text{True} \quad (45)$$

hold. From this, we need to show $m; \lambda; \mathcal{A} \vDash L(\alpha) : L'(\alpha) \twoheadrightarrow T$.

Let $P = \{l'_\sigma(\alpha)\}$, clearly the union of the elements of this set equals $l'_\sigma(\alpha)$ as required. Assuming $l'_\sigma(\alpha) \not\subseteq t_\sigma$ and setting $\widehat{O} = X \twoheadrightarrow_k X'$, which is in $\{\text{live}(\mathcal{A}, r) \mid \text{lay}(\text{live}(\mathcal{A}, r)) < m\}$ given (41) and (43), it suffices to show

$$\forall \alpha \in \mathbb{O}, w \in l'_\sigma(\alpha). \text{active}_{r;\lambda}(w, X \twoheadrightarrow_k X') \quad (46)$$

$$\forall \alpha_1, \alpha_2 \geq \alpha_1. \mathbf{R}_{\mathcal{A}}^a(l'_\sigma(\alpha_1)) \cap l_\sigma(\alpha_2) \subseteq l'_\sigma(\alpha_1) \cup t \quad (47)$$

to complete the proof.

We start off by showing that (46) holds. Taking $\alpha \in \mathbb{O}$ and $w \in l'_\sigma(\alpha)$ arbitrary, given (42), it is clear that $\text{lay}(\theta_w(r)) \geq \text{lay}(X \twoheadrightarrow_k X')$ holds and, given (44) and (45), $\text{lvl}_w(r) < \lambda$ holds. From these two conclusions, we can infer that $\text{dep}_{r;\lambda}(w, X \twoheadrightarrow_k X')$ holds. Then, from (45), it is clear that $\chi_w(r) \# \diamond$ holds, and therefore, $\text{active}_{r;\lambda}(w, X \twoheadrightarrow_k X')$.

Finally, (47) follows immediately from (40) and the definition of $\text{impr}_{\mathcal{A}}$.

– case *LIVET*. In this case $\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow T$ holds. From this, we need to show $m; \lambda; \mathcal{A} \vDash L(\alpha) : L'(\alpha) \twoheadrightarrow T$. Let $P = \{l'_\sigma(\alpha)\}$, clearly the union of the elements of this set equals $l'_\sigma(\alpha)$ as required. From $\forall \alpha. \vdash_{\mathcal{A}} L'(\alpha) \Rightarrow T$, clearly $l'_\sigma(\alpha) \subseteq t_\sigma$, therefore $m; \lambda; \mathcal{A} \vDash L(\alpha) : L'(\alpha) \twoheadrightarrow T$ holds, as required.

Finally, we complete this theorem's proof with a proof of the soundness of the one inductive case, *EQUANT*. Note that *ECASE* can be derived directly from *EQUANT*.

– case *EQUANT*. In this case, $L'(\alpha) = \exists x \in X. L''(x, \alpha)$ for some $L'' \in X \times \mathbb{O} \rightarrow \text{Asrt}$ and

$$\forall x \in X. m; \lambda; \mathcal{A} \vdash L(\alpha) : L''(x, \alpha) \twoheadrightarrow T \quad (48)$$

hold. Letting

$$l'_{x,\sigma}(\alpha) = \mathcal{W}[\![L''(x, \alpha)]\!]_{\mathcal{A}}^\sigma$$

From (48), for any $x \in X$ there exists $P_x \subseteq \mathcal{P}(\text{World}_{\mathcal{A}})$ such that $l'_{x,\sigma}(\alpha) = \bigcup P_x$ with the appropriate conditions holding for each $l_f \in \bigcup_{x \in X} P_x$.

Setting $P = \bigcup_{x \in X} P_x$, given the definition of $L'(\alpha)$, clearly $l'_\sigma(\alpha) = \bigcup P$ and for each $l \in P$, there exists some $x \in X$ such that $l \in P_x$ and therefore, the appropriate properties hold due to (48) as required.

By induction on the structure of derivation trees of the auxiliary environmental liveness condition, $m; \lambda; \mathcal{A} \vDash L : L' \twoheadrightarrow T$ holds, as required. □

THEOREM E.5. *If $m; \lambda; \mathcal{A} \vdash L \xrightarrow{M} T$ then $m; \lambda; \mathcal{A} \vDash L \xrightarrow{M} T$.*

PROOF. This theorem follows trivially from theorem E.4. □

E.3 Soundness of FRAME

For the rest of the section, we let

$$\begin{aligned} \mathbb{S} &= \forall x \in \vec{X}. \langle P_h \mid P_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle_{m; \lambda; \mathcal{A}} \\ \mathbb{S}' &= \forall x \in \vec{X}. \langle P_h * R_h \mid P_a(x) * R_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) * R_h \mid Q_a(x, y) * R_a(x) \rangle_{m; \lambda; \mathcal{A}} \end{aligned}$$

such that

$$\begin{aligned} \mathcal{A} &\vDash R_h \text{ stable} \\ \forall x \in X. \mathcal{A} &\vDash R_a(x) \text{ stable} \end{aligned}$$

LEMMA E.6. For arbitrary $\lambda \in \text{Lvl}$, \mathcal{A} and atomicity context, $h_0, h_1 \in \text{Heap}$, $p, q \in \text{World}_{\mathcal{A}}^{\uparrow}$ and $r \in \text{View}_{\mathcal{A}}$:

$$(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p \rightarrow q \Rightarrow (h_0, h_1) \vDash_{\lambda; \mathcal{A}} p * r \rightarrow q * r$$

PROOF. Assume $(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p \rightarrow q$, which is equivalent to:

$$\forall f \in \text{World}_{\mathcal{A}}^{\uparrow}. h_1 \in \llbracket p * f \rrbracket_{\lambda} \Rightarrow h_2 \in \llbracket p_2 * \mathbf{R}_{\mathcal{A}}^a(f) \rrbracket_{\lambda}$$

Substituting $f = r * f'$, this is equivalent to:

$$\forall f' \in \text{World}_{\mathcal{A}}^{\uparrow}. h_1 \in \llbracket p_1 * r * f' \rrbracket_{\lambda} \Rightarrow h_2 \in \llbracket p_2 * \mathbf{R}_{\mathcal{A}}^a(r * f') \rrbracket_{\lambda}$$

As $r \in \text{View}_{\mathcal{A}}$, $\mathbf{R}_{\mathcal{A}}^a(r * f') = r * \mathbf{R}_{\mathcal{A}}^a(f')$ holds, and therefore, as required:

$$\forall f' \in \text{World}_{\mathcal{A}}^{\uparrow}. h_1 \in \llbracket p_1 * r * f' \rrbracket_{\lambda} \Rightarrow h_2 \in \llbracket p_2 * r * \mathbf{R}_{\mathcal{A}}^a(f') \rrbracket_{\lambda} \quad \square$$

LEMMA E.7. For arbitrary $\lambda \in \text{Lvl}$, \mathcal{A} and atomicity context, $h_0, h_1 \in \text{Heap}$, $p, q \in \text{World}_{\mathcal{A}}^{\uparrow}$ and $f \in \text{View}_{\mathcal{A}}$:

$$h_0 \in \llbracket p * f \rrbracket \wedge (h_0, h_1) \vDash_{\lambda; \mathcal{A}} p \rightarrow q \Rightarrow h_1 \in \llbracket q * f \rrbracket$$

PROOF. To start off, assume $h_0 \in \llbracket p * f \rrbracket$ and $(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p \rightarrow q$. Clearly, this second assumption entails $(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p \rightarrow^* q$, which is equivalent to:

$$\forall f \in \text{View}_{\mathcal{A}}. h_0 \in \llbracket p * f \rrbracket \Rightarrow h_1 \in \llbracket q * f \rrbracket$$

Choosing the initial f and applying the first assumption yields $h_1 \in \llbracket q * f \rrbracket$ as required. \square

Definition E.8. For arbitrary $V \subseteq \text{PVar}$ and $\tau \in \text{Trace}$, we define the predicate $\text{noMods}_V(\tau)$, identifying traces that only modify the program variables in V :

$$\text{noMods}_V((\sigma_0, h_0) \pi (\sigma_1, h_1) \tau') \triangleq \text{noMods}_V((\sigma_1, h_1) \tau') \wedge \forall v \in V. \sigma_0(v) = \sigma_1(v)$$

Definition E.9. For $V \subseteq \text{PVar}$: $\text{Trace}_V \triangleq \{\tau \in \text{Trace} \mid \text{noMods}_V(\tau)\}$.

LEMMA E.10. Given $\mathbb{C} \in \text{Cmd}$, $\varphi \in \text{FImpl}$ and $V \subseteq \text{PVar}$ arbitrary such that $V \cap \text{mods}(\mathbb{C}) = \emptyset$, then: $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \text{Trace}_V$.

PROOF. Easy coinduction on the small-step operational semantics of commands. \square

Definition E.11. We define an auxiliary operation that takes a Hoare frame $r_h \in \text{View}_{\mathcal{A}}$ and an atomic frame $r_a \in \text{World}_{\mathcal{A}}^{\uparrow}$ and applies the frames at each position of a specification trace, if the heaps at each position are compatible with said frames (and returns the empty set otherwise).

$$((\sigma, h, p_h, p_a, v) \pi \hat{\tau}) \otimes (r_h, r_a) \triangleq \left\{ (\sigma, h, p_h * r_h, p_a * r_a, v) \pi \hat{\tau}' \mid \begin{array}{l} \hat{\tau}' \in (\hat{\tau} \otimes (r_h, r_a)) \wedge \\ h \in \llbracket p_h * r_h * p_a(v) * r_a(v) * \text{True} \rrbracket_{\lambda} \end{array} \right\}$$

This can be lifted to sets of specification traces, $\mathbb{T} \subseteq \text{STrace}$:

$$\mathbb{T} \otimes (r_h, r_a) \triangleq \bigcup_{\hat{\tau} \in \mathbb{T}} \hat{\tau} \otimes (r_h, r_a)$$

LEMMA E.12. For arbitrary $(\sigma_0, h_0)\tau \in \text{Trace}_{\text{fv}(R_h)}$, $p_h \in \text{View}_{\mathcal{A}}$, $v_0 \in \text{AVal}'$ and $\mathbb{T} \in \mathcal{P}(\text{STrace})$, then

$$h_0 \in \llbracket p_h * r_h * p_a(v_0) * r_a(v_0) * \text{True} \rrbracket \wedge (\sigma_0, h_0) \tau \vDash_{\mathbb{S}} p_h, p_a, v_0 : \mathbb{T} \Rightarrow (\sigma_0, h_0) \tau \vDash_{\mathbb{S}'} p_h * r_h, p_a * r_a, v_0 : \mathbb{T} \otimes (r_h, r_a)$$

holds, where

$$\begin{aligned} r_h &= \mathcal{W} \llbracket R_h \rrbracket_{\mathcal{A}}^{\sigma_0} \\ p_a(v) &= \begin{cases} \mathcal{W} \llbracket P_a(v) \wedge v \in X \rrbracket_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \\ r_a(v) &= \begin{cases} \mathcal{W} \llbracket R_a(v) \wedge v \in X \rrbracket_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \end{aligned}$$

PROOF. Taking $(\sigma_0, h_0)\tau \in \text{Trace}_{\text{fv}(R_h) \cap \text{PVar}}$, $p_h \in \text{View}_{\mathcal{A}}$, $v_0 \in \text{AVal}'$ and $\mathbb{T} \in \mathcal{P}(\text{STrace})$ arbitrary such that:

$$h_0 \in \llbracket p_h * r_h * p_a(v_0) * r_a(v_0) * \text{True} \rrbracket \quad (49)$$

$$(\sigma_0, h_0) \tau \vDash_{\mathbb{S}} p_h, p_a, v_0 : \mathbb{T} \quad (50)$$

The proof proceeds by coinduction on the structure of τ . We consider the rules can apply from the trace safety judgement: **STUTTER**, **LINPT**, **ENV**, **ENV'** and **ENV₂**.

– Case **STUTTER**. In this case, $(\sigma_0, h_0)\tau = (\sigma_0, h_0) \text{loc } (\sigma_1, h_1)\tau'$ and $\mathbb{T} = (\sigma_0, h_0, p_h, p_a, v) \text{loc } \mathbb{T}'$. From (50), for some $p'_h \in \text{View}_{\mathcal{A}}$, the following hold:

$$(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p_h * p_a(v_0) \rightarrow p'_h * p_a(v_0) \quad (51)$$

$$(\sigma_1, h_1) \tau' \vDash_{\mathbb{S}} p'_h, p_a, v_0 : \mathbb{T}' \quad (52)$$

$$\text{term}(\tau') \Rightarrow \exists v_1, v_2. v = \langle v_1, v_2 \rangle \wedge p'_h = \mathcal{W} \llbracket Q_h(v_1, v_2) \rrbracket_{\mathcal{A}}^{\sigma_1} \quad (53)$$

Given that $r_h, r_a(v_0) \in \text{View}_{\mathcal{A}}$, using Lemma E.6, (51) implies

$$(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p_h * r_h * p_a(v_0) * r_a(v_0) \rightarrow p'_h * r_h * p_a(v_0) * r_a(v_0) \quad (54)$$

By Lemma E.7, (49) and (51) imply:

$$h_1 \in \llbracket p'_h * r_h * p_a(v_0) * r_a(v_0) * \text{True} \rrbracket \quad (55)$$

Given that $(\sigma_0, h_0)\tau \in \text{Trace}_{\text{fv}(R_h) \cap \text{PVar}}$, $\forall v \in \text{fv}(R_h) \cap \text{PVar}. \sigma_0(v) = \sigma_1(v)$ holds, and therefore:

$$r_h = \mathcal{W} \llbracket R_h \rrbracket_{\mathcal{A}}^{\sigma_1} \quad (56)$$

From this, given (56), (55) and (52) and using the inductive assumption, we derive:

$$(\sigma_1, h_1) \tau' \vDash_{\mathbb{S}'} p'_h * r_h, p_a * r_a, v_0 : \mathbb{T}' \otimes (r_h, r_a) \quad (57)$$

Finally, assuming $\text{term}(\tau')$, given (53), we know $\exists v_1, v_2. v = \langle v_1, v_2 \rangle \wedge p'_h = \mathcal{W}[[Q_h(v_1, v_2)]]_{\mathcal{A}}^{\sigma_1}$. From this and (56), we infer that $p'_h * r_h = \mathcal{W}[[Q_h(v, v') * R_h]]_{\mathcal{A}}^{\sigma_1}$ and therefore,

$$\text{term}(\tau') \Rightarrow \exists v_1, v_2. v = \langle v_1, v_2 \rangle \wedge p'_h * r_h = \mathcal{W}[[Q_h(v, v') * R_h]]_{\mathcal{A}}^{\sigma_1} \quad (58)$$

From (54), (57) and (58) by **STUTTER**, $(\sigma_0, h_0) \tau \vDash_{\mathbb{S}'} p_h * r_h, p_a * r_a, v_0 : \mathbb{T} \otimes (r_h, r_a)$ holds as required.

– *Case LINPT*. In this case, $(\sigma_0, h_0)\tau = (\sigma_0, h_0) \text{loc } (\sigma_1, h_1)\tau'$ and $\mathbb{T} = (\sigma_0, h_0, p_h, p_a, v) \text{loc } \mathbb{T}'$. From (50), the following hold for some $v' \in \text{AVal}$:

$$(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p_h * p_a(v_0) \rightarrow q'_h * \mathcal{W}[[Q_a(v_0, v')]]_{\mathcal{A}} \quad (59)$$

$$(\sigma_1, h_1) \tau' \vDash_{\mathbb{S}} q'_h, \text{emp}, \langle v_0, v' \rangle : \mathbb{T}' \quad (60)$$

$$\text{term}(\tau') \Rightarrow q'_h = \mathcal{W}[[Q_h(v_0, v')]]_{\mathcal{A}}^{\sigma_1} \quad (61)$$

Given that $r_h, r_a(v_0) \in \text{View}_{\mathcal{A}}$, using Lemma E.6, (59) implies

$$(h_0, h_1) \vDash_{\lambda; \mathcal{A}} p_h * r_h * p_a(v_0) * r_a(v_0) \rightarrow q'_h * r_h * \mathcal{W}[[Q_a(v, v')]]_{\mathcal{A}} * r_a(v_0) \quad (62)$$

By Lemma E.7, (49) and (59) imply:

$$h_1 \in [[q'_h * r_h * \mathcal{W}[[Q_a(v_0, v')]]_{\mathcal{A}} * r_a(v_0) * \text{True}]] \quad (63)$$

Given that $(\sigma_0, h_0)\tau \in \text{Trace}_{\text{fv}(R_h) \cap \text{PVar}}$, $\forall v \in \text{fv}(R_h) \cap \text{PVar}. \sigma_0(v) = \sigma_1(v)$ holds, and therefore:

$$r_h = \mathcal{W}[[R_h]]_{\mathcal{A}}^{\sigma_1} \quad (64)$$

From this, given (64), (63) and (60) and using the inductive assumption, we derive:

$$(\sigma_1, h_1) \tau' \vDash_{\mathbb{S}'} q'_h * r_h, p_a * r_a, \langle v_0, v' \rangle : \mathbb{T}' \otimes (r_h, r_a) \quad (65)$$

Finally, assuming $\text{term}(\tau')$, given (61), we know $q'_h = \mathcal{W}[[Q_h(v_0, v')]]_{\mathcal{A}}^{\sigma_1}$. From this and (64), we infer that $q'_h * r_h = \mathcal{W}[[Q_h(v_0, v') * R_h]]_{\mathcal{A}}^{\sigma_1}$ and therefore,

$$\text{term}(\tau') \Rightarrow q'_h * r_h = \mathcal{W}[[Q_h(v_0, v') * R_h]]_{\mathcal{A}}^{\sigma_1} \quad (66)$$

From (62), (65) and (66) by **LINPT**, $(\sigma_0, h_0) \tau \vDash_{\mathbb{S}'} p_h * r_h, p_a * r_a, v_0 : \mathbb{T} \otimes (r_h, r_a)$ holds as required.

– *case ENV*. In this case, $(\sigma_0, h_0)\tau = (\sigma_0, h_0) \text{env } (\sigma_1, h_1)\tau'$ and

$$\mathbb{T} = \bigcup \{ (\sigma, h_1, p_h, p_a, v_0) \text{env } \mathbb{T}'_{v'} \mid v' \in X, E(v') \}.$$

From (50) we have that $\forall v' \in X. E(v') \Rightarrow (\sigma, h_2) \tau \vDash_{\mathbb{S}} p_h, p_a, v' : \mathbb{T}_{v'}$. Taking $v' \in X$ arbitrary and, assuming $E(v')$ given some $\overline{p_e}, \overline{p'_e}$, for the goal specification:

$$\begin{aligned} h_0 &\in [[p_h * r_h * p_a(v_0) * r_a(v_0) * \overline{p_e}]]_{\lambda} \\ (h_0, h_1) &\vDash_{\lambda; \mathcal{A}} p_a(v_0) * r_a(v_0) * \overline{p_e} \rightarrow p_a(v') * r_a(v') * \overline{p'_e} \end{aligned}$$

It then suffices to show that $(\sigma_1, h_1) \tau' \vDash_{\mathbb{S}'} p_h * r_h, p_a * r_a, v' : \mathbb{T}_{v'}$ holds. This follows from Lemma E.6 by using $p_e = \overline{p_e} * r_h * r_a(v_0)$ and $p'_e = \overline{p'_e} * r_h * r_a(v')$, yielding:

$$h_1 \in [[p_h * p_a(v) * p_e]]_{\lambda} \quad (h_1, h_2) \vDash_{\lambda; \mathcal{A}} p_a(v) * p_e \rightarrow p_a(v') * p'_e$$

as required.

– *Case ENV'*. This case follows similarly to the **ENV** case.

– Case ENV_{\neq} . This case is trivially true. \square

LEMMA E.13. *Letting*

$$r_h = \mathcal{W} \llbracket R_h \rrbracket_{\mathcal{A}}^{\sigma_0}$$

$$r_a(v) = \begin{cases} \mathcal{W} \llbracket R_a(v) \wedge v \in X \rrbracket_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases}$$

and assuming $\vdash_{\mathcal{A}} R_h * R_a(x) \triangleright m$, then, given $\mathbb{T} \subseteq \text{STrace}$:

$$\forall \bar{\tau}' \in \llbracket \mathbb{T} \otimes (r_h, r_a) \rrbracket. \text{liveEnv}_{\mathbb{S}'}(\bar{\tau}') \Rightarrow \exists \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S}}(\bar{\tau})$$

PROOF. Taking $\bar{\tau}' \in \llbracket \mathbb{T} \otimes (r_h, r_a) \rrbracket$ arbitrary such that $\text{liveEnv}_{\mathbb{S}'}(\bar{\tau}')$. This implies that:

$$\forall \widehat{O} \in \text{POb}_{< m}^{\mathbb{S}'}. \text{if } \forall r, O \in \text{AOB}_{\leq \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_{\lambda}(r, O, \bar{\tau}'(j))$$

$$\text{then } \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(\widehat{O}, \bar{\tau}'(j))$$

As $\bar{\tau}' \in \llbracket \mathbb{T} \otimes (r_h, r_a) \rrbracket$, there must be some $\hat{\tau} \in \mathbb{T}$ such that $\bar{\tau}' \in \llbracket \hat{\tau} \otimes (r_h, r_a) \rrbracket$. Taking $\bar{\tau} \in \llbracket \hat{\tau} \rrbracket$ arbitrary, to show $\text{liveEnv}_{\mathbb{S}}(\bar{\tau})$, take $\widehat{O} \in \text{POb}_{< m}^{\mathbb{S}}$ arbitrary such that

$$\forall r, O \in \text{AOB}_{\leq \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_{\lambda}(r, O, \bar{\tau}(j))$$

Given $\vdash_{\mathcal{A}} R_h * R_a(x) \triangleright m$ and the definition of \otimes , the following holds:

$$\forall r, O \in \text{AOB}_{\leq \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_{\lambda}(r, O, \bar{\tau}'(j))$$

Now, from $\text{liveEnv}_{\mathbb{S}'}(\bar{\tau}')$:

$$\forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(\widehat{O}, \bar{\tau}'(j))$$

From this, as required:

$$\forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(\widehat{O}, \bar{\tau}(j))$$

\square

THEOREM E.14 (SOUNDNESS OF **FRAME**). *Assuming*

$$\forall x \in X. \vdash_{\mathcal{A}} R_h * R_a(x) \triangleright m \quad (67)$$

and given arbitrary $\mathbb{C} \in \text{Cmd}$ such that

$$\text{pv}(R_h) \cap \text{mod}(\mathbb{C}) = \emptyset \quad (68)$$

and arbitrary $\Phi \in \text{FSpec}$ such that

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S}$$

then

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S}'$$

PROOF. To start off, as $\mathcal{A} \vDash R_h$ stable, clearly $P_h * R_h \in \text{Stable}_{\mathcal{A}}$ and therefore, $\mathbb{S}' \in \text{Spec}$.

Taking $\mathbb{C} \in \text{Cmd}$ arbitrary such that (68) holds, $\Phi \in \text{FSpec}$ arbitrary such that $\vDash_{\Phi} \mathbb{C} : \mathbb{S}$ holds and arbitrary $\varphi \in \text{FImpL}$ such that $\vDash \varphi : \Phi$ holds, then $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \llbracket \mathbb{S} \rrbracket$. From Lemma E.10 and (68), we can also infer that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \text{Trace}_{\text{pv}(R_h)}$ and therefore, it is clear that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \text{Trace}_{\text{pv}(R_h)} \cap \llbracket \mathbb{S} \rrbracket$. From this, we know that it is sufficient to show that $\text{Trace}_{\text{pv}(R_h)} \cap \llbracket \mathbb{S} \rrbracket \subseteq \llbracket \mathbb{S}' \rrbracket$, to show that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \llbracket \mathbb{S}' \rrbracket$ holds, and therefore, $\vDash_{\Phi} \mathbb{C} : \mathbb{S}'$, as required.

Therefore, taking $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket \cap \text{Trace}_{\text{pv}(R_h)}$ arbitrary, it is sufficient to show $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S}' \rrbracket$.

Let

$$\begin{aligned} p_h &= \mathcal{W}[[P_h]]_{\mathcal{A}}^{\sigma_0} \\ r_h &= \mathcal{W}[[R_h]]_{\mathcal{A}}^{\sigma_0} \\ p_a(v) &= \begin{cases} \mathcal{W}[[P_a(v) \wedge v \in X]]_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \\ r_a(v) &= \begin{cases} \mathcal{W}[[R_a(v) \wedge v \in X]]_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \end{aligned}$$

To show $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S}' \rrbracket$, for some arbitrary $v_0 \in X$, assume $h_0 \in \llbracket p_h * r_h * p_a(v_0) * r_a(v_0) * \text{True} \rrbracket_{\lambda}$, from which it follows that $h_0 \in \llbracket p_h * p_a(v_0) * \text{True} \rrbracket_{\lambda}$. Then, as $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket$, for some $\mathbb{T} \subseteq \text{STrace}$:

$$(\sigma_0, h_0) \tau \vDash_{\mathbb{S}} p_h, p_a, v_0 : \mathbb{T} \quad (69)$$

$$\forall \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau}) \quad (70)$$

From Lemma E.12 and (69), $(\sigma_0, h_0) \tau \vDash_{\mathbb{S}} p_h * r_h, p_a * r_a, v_0 : \mathbb{T} \otimes (r_h, r_a)$. To reach the goal now, it suffices to show that for some arbitrary $\bar{\tau}' \in \llbracket \mathbb{T} \otimes (r_h, r_a) \rrbracket$:

$$\text{liveEnv}_{\mathbb{S}'}(\bar{\tau}') \Rightarrow \text{lterm}(\bar{\tau}')$$

This holds trivially from Lemma E.13, (67) and (70). □

Note that Theorem E.14 has the side condition $\text{pv}(R_h) \cap \text{mod}(\mathbb{C}) = \emptyset$ rather than $\forall x \in X. \text{pv}(R_h, R_a(x)) \cap \text{mod}(\mathbb{C}) = \emptyset$ as in **FRAME**. This is because this theorem applies to TaDA Live specifications without the syntactic sugar that permits program variables to be directly referenced in the atomic precondition and postcondition of a TaDA Live hybrid triple. The side condition present in the **FRAME** rule permits it to be applied directly to sugared hybrid specifications, as it guarantees that the necessary side condition for the corresponding desugared specification holds.

E.4 Soundness of **LIVEC**

Let

$$\begin{aligned} \mathbb{S}_{\rightarrow} &= \forall x \in X \rightarrow_k X'. \langle P_h \mid P_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle_{m;\lambda;\mathcal{A}} \\ \mathbb{S} &= \forall x \in X. \langle P_h * L \mid P_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) * L \mid Q_a(x, y) \rangle_{m;\lambda;\mathcal{A}} \end{aligned}$$

where $L \in \text{Stable}_{\mathcal{A}}$.

Definition E.15. For atomicity context \mathcal{A} and layer m from the context of \mathbb{S} and sets X and X' as well as the layer k from the pseudo-quantifier of \mathbb{S} , let

$$\text{POb}^{\mathbb{S}} \triangleq \{ (r, O) \mid r \in \text{RId}, O \in \text{AOB} \} \uplus \{ (r, \text{live}(\mathcal{A}, r)) \mid r \in \text{dom}(\mathcal{A}) \} \uplus \{ X \rightarrow_k X' \}$$

Then $\text{liveEnv}_{\mathbb{S}}(\hat{\tau})$ predicate checks whether the environment is satisfying the liveness assumptions of the specification:

$$\begin{aligned} \text{liveEnv}_{\mathbb{S}}(\hat{\tau}) &\triangleq \forall \widehat{O} \in \text{POb}_{< m}^{\mathbb{S}}. \text{if } \forall r, O \in \text{AOB}_{\leq \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_{\lambda}(r, O, \hat{\tau}(j)) \\ &\quad \text{then } \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(\widehat{O}, \hat{\tau}(j)) \end{aligned}$$

LEMMA E.16. Given $M \in \mathbb{O} \rightarrow \text{World}_{\mathcal{A}}^{\uparrow}, T \in \text{Assrt}, n \leq m, k$ such that

$$n; \lambda; \mathcal{A} \vDash L \xrightarrow{M} T \quad (71)$$

$$\forall x \in X. \vdash_{\lambda; \mathcal{A}} P(x) * T \Rightarrow x \in X' \quad (72)$$

hold. Take $(\sigma_0, h_0)\tau \in \text{Trace}$ and let

$$p_a(v) = \begin{cases} \mathcal{W}[[P_a(v) \wedge v \in X]]_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases}$$

$$l = \mathcal{W}[[L]]_{\mathcal{A}}^{\sigma_0}$$

$$l(\alpha) = \mathcal{W}[[L * M(\alpha)]]_{\mathcal{A}}^{\sigma_0}$$

$$t = \mathcal{W}[[T * \text{True}]]_{\mathcal{A}}^{\sigma_0}$$

Taking arbitrary $p'_h \in \text{View}_{\mathcal{A}}, \mathbb{T} \subseteq \text{STrace}$ and $v_0 \in X$ such that

$$h_0 \in [[p'_h * l * p_a(v_0) * \text{True}]]_{\lambda} \quad (73)$$

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p'_h, p_a, v_0 : \mathbb{T} \quad (74)$$

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p'_h * l, p_a, v_0 : \mathbb{T} \otimes (l, \text{emp}) \quad (75)$$

for arbitrary $\bar{\tau}' \in [[\mathbb{T} \otimes (l, \text{emp})]]_{\lambda; \mathcal{A}}$ such that

$$\text{liveEnv}_{\mathbb{S}}(\bar{\tau}') \quad (76)$$

there exists $\bar{\tau} \in [[\mathbb{T}]]_{\lambda; \mathcal{A}}$, such that, if

$$\forall \bar{O}' \in \text{RId} \times \text{AOb}_{\leq k}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{loheld}_{\lambda}(\bar{O}', \bar{\tau}(j)) \quad (77)$$

then $\forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}_{\lambda}(X \rightarrow_k X', \bar{\tau}(j))$.

PROOF. Taking $(\sigma_0, h_0)\tau \in \text{Trace}$, $p'_h \in \text{View}_{\mathcal{A}}, \mathbb{T} \subseteq \text{STrace}$ and $v_0 \in X$ arbitrary such that (73), (74) and (75) hold and $\bar{\tau}' \in [[\mathbb{T} \otimes (l, \text{emp})]]_{\lambda; \mathcal{A}}$ satisfying (76). From this, we know that, there exists $\bar{\tau} \in [[\mathbb{T}]]_{\lambda; \mathcal{A}}$ and $ls \in l^{\omega}$, such that, for all $i \in \mathbb{N}$:

$$ls(i) \mathbf{R}_{\mathcal{A}}^a ls(i+1) \quad (78)$$

$$\exists h \in \text{Heap}, w_h, w_a, w_e \in \text{View}_{\mathcal{A}}, v \in \text{Val}'. \quad \bar{\tau}(i) = ((h, w_h, w_a, w_e \odot ls(i), v), _) \wedge$$

$$\bar{\tau}'(i) = ((h, w_h \odot ls(i), w_a, w_e, v), _) \quad (79)$$

hold. If (77) does not hold, our proof is complete, otherwise, from (71), there exists some $P \subseteq \wp(\mathbb{O} \rightarrow \text{World}_{\mathcal{A}}^{\uparrow})$ such that $\forall \alpha. l(\alpha) = \bigcup_{lf \in P} lf(\alpha)$.

We now show by transfinite induction on $\alpha \in \mathbb{O}$, that

$$\forall \alpha \in \mathbb{O}, i \in \mathbb{N}. ls(i) \in l(\alpha) \Rightarrow \exists j \geq i. \neg \text{envheld}_{\lambda}(X \rightarrow_k X', \bar{\tau}(j))$$

Base case ($\alpha = 0$): Take $i \in \mathbb{N}$ and assume $ls(i) \in l(0)$. Since $l(0) = \bigcup_{lf \in P} lf(0)$, for some $lf \in P$ we have $ls(i) \in lf(0)$. We now assume, towards a contradiction, that $\forall j \geq i. \text{envheld}_{\lambda}(X \rightarrow_k X', \bar{\tau}(j))$ and therefore $\forall j \geq i. \exists v \in X \setminus X'. \bar{\tau}(j) = ((_, _, _, v), _)$. We now demonstrate, that under this assumption, by induction on $j \geq i$ that

$$\forall j \geq i. ls(j) \in lf(0) \quad (80)$$

Assume that for $j \geq i$, $ls(j) \in lf(0)$ holds. From (78), $ls(j+1) \in \mathbf{R}_{\mathcal{A}}^a(lf(0))$ holds and from (71), setting $\alpha_1 = 0$, $\mathbf{R}_{\mathcal{A}}^a(lf(0)) \subseteq lf(0) \cup t$, therefore, either $ls(j+1) \in lf(0)$ or $ls(j+1) \in t$ hold. In the latter case, from (72), $\exists v \in X'. \bar{\tau}(j+1) = ((_, _, _, v), _)$, a contradiction. Therefore, $ls(j+1) \in lf(0)$ holds, proving (80).

From (71), there exists some $r \in \text{RId}$ and $\bar{O} \in \text{AOb}_{< n} \uplus \{\text{live}(\mathcal{A}, r) \mid \text{lay}(\text{live}(\mathcal{A}, r)) < n\}$ such that:

$$\forall w \in lf(0). \text{active}_{r, \lambda}(w, \bar{O}) \quad (81)$$

Taking $r' \in \text{RId}$ and $\mathbf{O} \in \text{AOb}_{\leq \text{lay}(\bar{O})}$ arbitrary, since $\text{lay}(\bar{O}) < n$, $\text{lay}(\mathbf{O}) < n$ and therefore $\text{lay}(\mathbf{O}) < k$ and $\text{lay}(\mathbf{O}) < m$ hold. As $\text{lay}(\mathbf{O}) < k$ holds, given (77), for some $j' \geq i$,

$\neg \text{locheld}_\lambda((r', \mathbf{O}), \bar{\tau}(j'))$ holds. Given (79), we know $\bar{\tau}(j') = (h, w_h, w_a, w_e \odot ls(j'), v)$ and $\bar{\tau}'(j') = (h, w_h \odot ls(j'), w_a, w_e, v)$ for some $h \in \text{Heap}$, $w_h, w_a, w_e \in \text{World}_{\mathcal{A}}$, $v \in X \setminus X'$, and therefore, $\theta_{w_h}(r') \not\sqsubseteq \mathbf{O}$. Given (80), $ls(j') \in lf(0)$ holds, and, therefore, given (81), we know that:

$$\text{lay}(\theta_{ls(j')}(r')) > \text{lay}(\widehat{\mathbf{O}}) \quad (82)$$

$$\text{lvl}_{ls(j')}(r) < \lambda \quad (83)$$

Given (82), $\theta_{ls(j')}(r') \not\sqsubseteq \mathbf{O}$, as otherwise, $\text{lay}(\theta_{ls(j')}(r')) \leq \text{lay}(\mathbf{O})$ and therefore, as $\text{lay}(\mathbf{O}) < \text{lay}(\widehat{\mathbf{O}})$, we reach a contradiction. Then, from $\theta_{w_h}(r) \not\sqsubseteq \mathbf{O}$ and $\theta_{ls(j')}(r) \not\sqsubseteq \mathbf{O}$, we know $\theta_{w_h \odot ls(j')}(r) \not\sqsubseteq \mathbf{O}$, as $\mathbf{O} \in \text{AOb}$, and therefore, $\neg \text{locheld}_\lambda((r', \mathbf{O}), \hat{\tau}'(j'))$. From this, it follows that

$$\forall \widehat{\mathbf{O}}' \in \text{Rld} \times \text{AOb}_{\leq \text{lay}(\widehat{\mathbf{O}})}, i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}_\lambda(\widehat{\mathbf{O}}', \bar{\tau}'(j)) \quad (84)$$

holds.

Letting $\bar{\tau}'(i) = ((h^i, w_h^i \odot ls(i), w_a^i, w_e^i, v^i), _)$, first, consider the case $\widehat{\mathbf{O}} \in \text{AOb}_{< n}$. As an obligation's composition with itself within the obligation algebra of the region type of the shared region r must be undefined, one of

$$\begin{aligned} & \text{locheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \wedge \neg \text{envheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \\ & \neg \text{locheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \wedge \text{envheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \\ & \neg \text{locheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \wedge \neg \text{envheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \end{aligned}$$

holds, as if both the environment and local worlds hold $\widehat{\mathbf{O}}$, their composition would be undefined.

If $\neg \text{locheld}_\lambda(\widehat{\mathbf{O}}, \hat{\tau}'(i)) \wedge \neg \text{envheld}_\lambda(\widehat{\mathbf{O}}, \hat{\tau}'(i))$ holds, then $\theta_{w_h^i \odot ls(i)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$ and $\xi_{w_h^i \odot ls(i)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$ hold. From (80), we know that $ls(i) \in lf(0)$, which, given (81) and $\widehat{\mathbf{O}} \in \text{AOb}_{< n}$ implies that $\xi_{ls(i)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$. Given $\xi_{w_h^i \odot ls(i)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$ and $\xi_{ls(i)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$ and the definition of \odot , we know that $\theta_{w_h^i}(r) \sqsupseteq \widehat{\mathbf{O}}$, which in turn implies $\theta_{w_h^i \odot ls(i)}(r) \sqsupseteq \widehat{\mathbf{O}}$, a contradiction, as required.

Otherwise, if $\neg \text{locheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i)) \wedge \text{envheld}_\lambda(\widehat{\mathbf{O}}, \bar{\tau}'(i))$ holds, from (76)

$$\begin{aligned} & \text{if } \forall \widehat{\mathbf{O}}' \in \text{Rld} \times \text{AOb}_{\leq \text{lay}(\widehat{\mathbf{O}})}, i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(\widehat{\mathbf{O}}', \bar{\tau}'(j)) \\ & \text{then } \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{envheld}(\widehat{\mathbf{O}}, \bar{\tau}'(j)) \end{aligned}$$

holds, and therefore, given (84), there exists some minimal $j \geq i$, such that $\neg \text{envheld}(\widehat{\mathbf{O}}, \bar{\tau}'(j+1))$ holds. Since j is minimal, $\forall k \in \mathbb{N}. i \leq k \leq j \Rightarrow \text{envheld}(\widehat{\mathbf{O}}, \bar{\tau}'(k))$ also holds. From this, letting

$$\begin{aligned} \bar{\tau}'(j) &= ((h^j, w_h^j \odot ls(j), w_a^j, w_e^j, v^j), \pi) \\ \bar{\tau}'(j+1) &= ((h^{j+1}, w_h^{j+1} \odot ls(j+1), w_a^{j+1}, w_e^{j+1}, v^{j+1}), _) \end{aligned}$$

we know that $\xi_{w_h^j \odot ls(j)}(r) \sqsupseteq \widehat{\mathbf{O}}$ and, given (81), $\xi_{w_h^{j+1} \odot ls(j+1)}(r) \not\sqsubseteq \widehat{\mathbf{O}}$. Then, given (83) and the invariant on atomic resources, we know $\theta_{w_a^j}(r) = \mathbf{0}$ and $\theta_{w_a^{j+1}}(r) = \mathbf{0}$, and therefore, by the definition of \odot , $\xi_{w_h^j \odot ls(j) \odot w_a^j}(r) \sqsupseteq \widehat{\mathbf{O}}$ and $\xi_{w_h^{j+1} \odot ls(j+1) \odot w_a^{j+1}}(r) \not\sqsubseteq \widehat{\mathbf{O}}$.

If $\pi = \text{loc}$, by construction of $\bar{\tau}'$, $w_e^j \mathbf{R}_{\mathcal{A}}^a w_e^{j+1}$ holds, and therefore, from the definition of $\mathbf{R}_{\mathcal{A}}^a$, $\theta_{w_e^j}(r) = \theta_{w_e^{j+1}}(r)$. By construction of $\bar{\tau}'$, we know $h^j \in \llbracket w_h^j \odot ls(j) \odot w_a^j \odot w_e^j \rrbracket_\lambda$ and therefore, $\xi_{w_h^j \odot ls(j) \odot w_a^j \odot w_e^j}(r) = \mathbf{0}$, given the definition of reification. Similarly, $\xi_{w_h^{j+1} \odot ls(j+1) \odot w_a^{j+1} \odot w_e^{j+1}}(r) = \mathbf{0}$. From this, we can infer $\theta_{w_e^j}(r) \sqsupseteq \widehat{\mathbf{O}}$ and $\theta_{w_e^{j+1}}(r) \not\sqsubseteq \widehat{\mathbf{O}}$, a contradiction.

Otherwise, if $\pi = \text{env}$, by construction of $\bar{\tau}'$, $w_h^j \odot ls(j) \mathbf{R}_{\mathcal{A}}^a w_h^{j+1} \odot ls(j+1)$ holds, and therefore, from the definition of $\mathbf{R}_{\mathcal{A}}^a$, $\theta_{w_h^j \odot ls(j)}(r) = \theta_{w_h^{j+1} \odot ls(j+1)}(r)$. Given (78), $\theta_{ls(j)}(r) = \theta_{ls(j+1)}(r)$, and therefore, from the definition of \odot , $\theta_{w_h^j}(r) = \theta_{w_h^{j+1}}(r)$. From $\xi_{w_h^j \odot ls(j)}(r) \sqsupseteq \widehat{\mathcal{O}}$, we know that $\theta_{w_h^j}(r) \not\sqsupseteq \widehat{\mathcal{O}}$. Given (80) and (81), $\xi_{ls(j+1)}(r) \sqsupseteq \widehat{\mathcal{O}}$ holds and therefore $\xi_{w_h^{j+1} \odot ls(j+1)}(r) \not\sqsupseteq \widehat{\mathcal{O}}$ implies $\theta_{w_h^{j+1}}(r) \sqsupseteq \widehat{\mathcal{O}}$, a contradiction.

Finally, the case $\text{loheld}_\lambda(\widehat{\mathcal{O}}, \bar{\tau}'(i)) \wedge \neg \text{envheld}_\lambda(\widehat{\mathcal{O}}, \bar{\tau}'(i))$ reaches a contradiction similarly. To finish the base case, it now suffices to consider the case $\widehat{\mathcal{O}} = X \rightarrow_k X'$. Given (76) and (84), we know that for some $j \geq i$, $\neg \text{envheld}_\lambda(\widehat{\mathcal{O}}, \bar{\tau}'(j))$ holds. Letting $\bar{\tau}'(j) = ((_, w_h^j \odot ls(j), _, _, _))$, given (81), $\text{lvl}_{ls(j)}(r) < \lambda$ and $\text{ast}_{ls(j)}(r) \in X \setminus X'$ hold. Given $\text{lvl}_{ls(j)}(r) < \lambda$ and $\neg \text{envheld}_\lambda(\widehat{\mathcal{O}}, \bar{\tau}'(j))$, $\text{ast}_{ls(j)}(r) \in X'$ holds, a contradiction.

By contradiction, the base case holds as required.

Inductive case: Take $\alpha \in \mathbb{O}$, $i \in \mathbb{N}$ and assume $ls(i) \in l(\alpha)$. Since $l(\alpha) = \bigcup_{lf \in P} lf(\alpha)$ holds for some $lf \in P$, we have $ls(i) \in lf(\alpha)$. Now assume, towards a contradiction, that $\forall j \geq i$. $\text{envheld}_\lambda(X \rightarrow_k X', \bar{\tau}(j))$ and therefore $\forall j \geq i$. $\exists v \in X \setminus X'$. $\bar{\tau}(j) = ((_, _, _, v), _)$. We now demonstrate, that under this assumption the following holds:

$$(\forall j \geq i. ls(j) \in lf(\alpha)) \vee (\exists j > i, \beta < \alpha. ls(j) \in l(\beta)) \quad (85)$$

To show this, it is sufficient to show that $\neg(\exists j > i, \beta < \alpha. ls(j) \in l(\beta))$ implies $\forall j \geq i. ls(j) \in lf(\alpha)$. We proceed to prove $\forall j \geq i. ls(j) \in lf(\alpha)$ by induction on $j \geq i$. The base case holds by our assumptions. Now for the inductive case, assume that for $j \geq i$, $ls(j) \in lf(\alpha)$ holds. From (78), $ls(j+1) \in \mathbf{R}_{\mathcal{A}}^a(lf(\alpha))$ holds and from (71), setting $\alpha_1 = \alpha$, $\mathbf{R}_{\mathcal{A}}^a(lf(\alpha)) \subseteq lf(\alpha) \cup \bigcup_{\beta < \alpha} l(\beta) \cup t$, therefore, either $ls(j+1) \in lf(\alpha)$, $ls(j+1) \in \bigcup_{\beta < \alpha} l(\beta)$ or $ls(j+1) \in t$ hold. In the case where $ls(j+1) \in t$ holds, from (72), $\exists v \in X'$. $\bar{\tau}(j+1) = ((_, _, _, v), _)$, a contradiction and in the case where $ls(j+1) \in \bigcup_{\beta < \alpha} l(\beta)$ holds, we reach a contradiction with $\neg(\exists j > i, \beta < \alpha. ls(j) \in l(\beta))$ which implies $\forall j > i, \beta < \alpha. ls(j) \notin l(\beta)$. Therefore, $ls(j+1) \in lf(\alpha)$ holds, as required, completing the proof by induction.

The inductive case then follows from (85). The goal follows similarly to the base case for the first disjunct and by inductive assumption in the second. \square

LEMMA E.17. Given $M \in \mathbb{O} \rightarrow \text{World}_{\mathcal{A}}^\dagger$, $T \in \text{Assrt}$, $n \leq m, k$ such that

$$n; \lambda; \mathcal{A} \vDash L \xrightarrow{M} T \quad (86)$$

$$\forall x \in X. \vdash_{\lambda; \mathcal{A}} P(x) * T \Rightarrow x \in X' \quad (87)$$

hold. Take $(\sigma_0, h_0)\tau \in \text{Trace}$ and let

$$\begin{aligned} \forall v \in \text{AVal}. p_a(v) &= \mathcal{W} \llbracket v \in X \wedge P_a(v) \rrbracket_{\mathcal{A}} \\ l &= \mathcal{W} \llbracket L \rrbracket_{\mathcal{A}}^{\sigma_0} \\ t &= \mathcal{W} \llbracket T * \text{True} \rrbracket_{\mathcal{A}}^{\sigma_0} \end{aligned}$$

Taking arbitrary $p'_h, p_e \in \text{View}_{\mathcal{A}}$, $\mathbb{T} \subseteq \text{STrace}$ and $v_0 \in X$ such that

$$h_0 \in \llbracket p'_h * l * p_a(v_0) * \text{True} * p_e \rrbracket_{\lambda} \quad (88)$$

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S} \rightarrow} p'_h, p_a, v_0 : \mathbb{T} \quad (89)$$

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p'_h * l, p_a, v_0 : \mathbb{T} \otimes (l, \text{emp}) \quad (90)$$

and arbitrary $\bar{\tau}' \in \llbracket \mathbb{T} \otimes (l, \text{emp}) \rrbracket$ such that

$$\text{liveEnv}_{\mathbb{S}}(\bar{\tau}') \quad (91)$$

holds, then, there exists $\bar{\tau} \in \llbracket \mathbb{T} \rrbracket$, such that:

$$\text{liveEnv}_{\mathbb{S} \rightarrow}(\bar{\tau})$$

PROOF. This lemma follows straightforwardly from lemma E.16. \square

THEOREM E.18. Taking $n \in \mathcal{L}, T \in \text{Assrt}$ and $M \in \mathbb{O} \rightarrow \text{Assrt}$ such that $m \geq n, k \geq n$ and

$$n; \lambda; \mathcal{A} \vDash L \xrightarrow{M} T \quad (92)$$

$$\forall x \in X. \vdash_{\lambda; \mathcal{A}} P_a(x) * T \Rightarrow x \in X' \quad (93)$$

Then, for any $\Phi \in \text{FSpec}$ and $\mathbb{C} \in \text{Cmd}$, if

$$\text{pv}(L) \cap \text{mod}(\mathbb{C}) = \emptyset \quad (94)$$

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S} \rightarrow \quad (95)$$

then

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S}$$

PROOF. Taking $n \in \mathcal{L}, T \in \text{Assrt}$ and $M \in \mathbb{O} \rightarrow \text{Assrt}$ arbitrary such that $m \geq n, k \geq n$, (92) and (93) hold. Then, to start off, given (92), $\mathcal{A} \vDash L$ stable holds, and therefore, $P_h * L \in \text{Stable}_{\mathcal{A}}$. From this we can infer, $\mathbb{S} \in \text{Spec}$.

Then, taking $\mathbb{C} \in \text{Cmd}$ arbitrary such that (68) holds, $\Phi \in \text{FSpec}$ arbitrary such that $\vDash_{\Phi} \mathbb{C} : \mathbb{S} \rightarrow$ holds and arbitrary $\varphi \in \text{FImpl}$ such that $\vDash_{\varphi} : \Phi$ holds, then $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \llbracket \mathbb{S} \rightarrow \rrbracket$. From Lemma E.10 and (94), we can also infer that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \text{Trace}_{\text{pv}(L)}$ and therefore, it is clear that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \text{Trace}_{\text{pv}(L)} \cap \llbracket \mathbb{S} \rightarrow \rrbracket$. From this, we know that it is sufficient to show that $\text{Trace}_{\text{pv}(R_h)} \cap \llbracket \mathbb{S} \rightarrow \rrbracket \subseteq \llbracket \mathbb{S} \rrbracket$, to show that $\llbracket \mathbb{C} \rrbracket_{\varphi} \subseteq \llbracket \mathbb{S} \rrbracket$ holds, and therefore, $\vDash_{\Phi} \mathbb{C} : \mathbb{S}$, as required.

Therefore, taking $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rightarrow \rrbracket \cap \text{Trace}_{\text{pv}(L)}$ arbitrary, it is sufficient to show $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket$. Let

$$\begin{aligned} p_h &= \mathcal{W} \llbracket P_h \rrbracket_{\mathcal{A}}^{\sigma_0} & p_a(v) &= \begin{cases} \mathcal{W} \llbracket P_a(v) \wedge v \in X \rrbracket_{\mathcal{A}} & \text{if } x \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \\ m(\alpha) &= \mathcal{W} \llbracket M(\alpha) \rrbracket_{\mathcal{A}}^{\sigma_0} & l(\alpha) &= \mathcal{W} \llbracket L * M(\alpha) \rrbracket_{\mathcal{A}}^{\sigma_0} \\ l &= \mathcal{W} \llbracket L \rrbracket_{\mathcal{A}}^{\sigma_0} & t &= \mathcal{W} \llbracket T \rrbracket_{\mathcal{A}}^{\sigma_0} \end{aligned}$$

To show $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket$, assume for some $v_0 \in X, h_0 \in \llbracket p_h * p_a(v_0) * l * \text{True} \rrbracket_{\lambda}$. Then, given $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rightarrow \rrbracket$, for some $\mathbb{T} \in \wp(\text{STrace})$:

$$(\sigma_0, h_0) \tau \vDash_{\mathbb{S} \rightarrow} p_h, p_a, v_0 : \mathbb{T} \wedge \forall \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S} \rightarrow}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$$

Given Lemma E.12 and that the definition of the trace safety judgement does not depend on the good states of a specification, X' , clearly, $(\sigma_0, h_0) \tau \vDash_{\mathbb{S}} p_h * l, p_a, v_0 : \mathbb{T} \otimes (L, \text{emp})$ holds. To complete the proof, it suffices show that

$$\forall \bar{\tau} \in \llbracket \mathbb{T} \otimes (L, \text{emp}) \rrbracket. \text{liveEnv}_{\mathbb{S} \rightarrow}(\bar{\tau}) \Rightarrow \text{liveEnv}_{\mathbb{S}}(\bar{\tau})$$

This follows straightforwardly from Lemma E.17. \square

E.5 Soundness of WHILE

Definition E.19 (Concrete trace sequence operator).

$$\tau = \tau_1 \wp \tau_2 \Leftrightarrow \left(\begin{array}{l} (\neg \text{lterm}(\tau_1) \wedge \tau = \tau_1) \vee \\ \left(\begin{array}{l} \exists \sigma \in \text{Store}, h \in \text{Heap}, \tau'_1 \text{ loc } (\sigma, h) \tau''_1, (\sigma, h) \tau'_2 \in \text{Trace}. \\ \tau_1 = \tau'_1 \text{ loc } (\sigma, h) \tau''_1 \wedge \tau_2 = (\sigma, h) \tau'_2 \wedge \text{term}((\sigma, h) \tau''_1) \wedge \tau = \tau'_1 \text{ loc } (\sigma, h) \text{ loc } (\sigma, h) \tau'_2 \end{array} \right) \end{array} \right)$$

A similarly defined overloading of this operator exists for specification traces, $\hat{\tau}_1 \mathbin{\text{;}} \hat{\tau}_2$ and the obvious lifting to sets $\mathbb{T}_1 \mathbin{\text{;}} \mathbb{T}_2$.

LEMMA E.20. For arbitrary $\varphi \in \text{Flmpl}$, $(\sigma_0, h_0)\tau \in \llbracket \text{while}(\mathbb{B})\{\mathbb{C}\} \rrbracket_\varphi$, either $\neg \mathcal{B}[\mathbb{B}]_{\sigma_0}$ and $(\sigma_0, h_0)\tau \in \llbracket \text{skip} \rrbracket_\varphi$, or $\mathcal{B}[\mathbb{B}]_{\sigma_0}$ and there exists $(\sigma_0, h_0)\tau' \in \llbracket \mathbb{C} \rrbracket_\varphi$ and $\tau'' \in \llbracket \text{while}(\mathbb{B})\{\mathbb{C}\} \rrbracket_\varphi$, such that $(\sigma_0, h_0)\tau = (\sigma_0, h_0)\text{loc}(\sigma_0, h_0)\tau' \mathbin{\text{;}} \tau''$.

PROOF. Straightforward by induction on \rightarrow_φ . \square

LEMMA E.21. Given an arbitrary specification

$$\mathbb{S} = \mathbb{W}x \in X \rightarrow X'. \langle P_h \mid P_a(x) \rangle \cdot \exists y. \langle Q_h(x, y) \mid Q_a(x, y) \rangle_{\lambda; \mathcal{A}}$$

for an arbitrary trace $(\sigma_0, h_0)\tau \in \text{Trace}$, let

$$p_h = \mathcal{W}[\llbracket P_h \rrbracket_{\mathcal{A}}^{\sigma_0}] \quad p_a(v) = \mathcal{W}[\llbracket P_a(v) \rrbracket_{\mathcal{A}}^{\sigma_0}]$$

If for some $v \in X$ and $\mathbb{T} \in \mathcal{P}(\text{STrace})$, $h_0 \in \llbracket p_h * p_a(v) * \text{True} \rrbracket$ and $(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p_h, p_a, v : \mathbb{T}$, then:

$\forall \hat{\tau} \in \mathbb{T}. \forall i \in \mathbb{N}. \text{term}(\hat{\tau}_{/i}) \Rightarrow \exists h \in \text{Heap}, \sigma \in \text{Store}, p_h \in \text{View}_{\mathcal{A}}, v \in X, v' \in \text{AVal}$.

$$\hat{\tau}(i) = (\sigma, h, p_h, \text{emp}, \langle v, v' \rangle) \wedge p_h = \mathcal{W}[\llbracket Q_h(v, v') \rrbracket_{\mathcal{A}}^{\sigma}] \wedge h \in \llbracket p_h * \text{True} \rrbracket_{\lambda} \quad (96)$$

PROOF. Straightforward by induction on the specification semantics rules. \square

For the rest of the section, let

$$\begin{aligned} \mathbb{S}'(\beta, b) &= \{P(\beta) * (b \Rightarrow T(\beta)) \wedge \mathbb{B}\} \cdot \{\exists \gamma. P(\gamma) \wedge \gamma \leq \beta * (b \Rightarrow \gamma < \beta)\}_{m; \lambda; \mathcal{A}} \\ \mathbb{S}(\beta_0) &= \{P(\beta_0) * L\} \cdot \{\exists \beta. P(\beta) * L \wedge \neg \mathbb{B} \wedge \beta_0 \geq \beta\}_{m; \lambda; \mathcal{A}} \end{aligned}$$

LEMMA E.22. Take $\varphi \in \text{Flmpl}$ and $\beta_0 \in \mathbb{O}$ arbitrary and take $(\sigma_0, h_0)\tau \in \llbracket \text{while}(\mathbb{B})\{\mathbb{C}\} \rrbracket_\varphi$ such that $\mathcal{B}[\mathbb{B}]_{\sigma_0}$. Let

$$p'(\beta, b) = \mathcal{W}[\llbracket P(\beta) * (b \Rightarrow T) \rrbracket_{\mathcal{A}}^{\sigma_0}] \quad l = \mathcal{W}[\llbracket L \rrbracket_{\mathcal{A}}^{\sigma_0}]$$

As $\mathcal{B}[\mathbb{B}]_{\sigma_0}$, by lemma E.20, there exists $(\sigma_0, h_0)\tau' \in \llbracket \mathbb{C} \rrbracket_\varphi$ and $(\sigma_1, h_1)\tau'' \in \llbracket \text{while}(\mathbb{B})\{\mathbb{C}\} \rrbracket_\varphi$, such that $(\sigma_0, h_0)\tau = (\sigma_0, h_0)\tau' \mathbin{\text{;}} (\sigma_1, h_1)\tau''$. If, for arbitrary $\beta' \leq \beta \leq \beta_0$ and $\mathbb{T}', \mathbb{T}'' \in \mathcal{P}(\text{STrace})$, there exists $b \in \text{Bool}$, such that:

$$\begin{aligned} h_0 &\in \llbracket p'(\beta, b) * l * \text{True} \rrbracket_{\lambda} \\ (\sigma_0, h_0)\tau' &\vDash_{\mathbb{S}'(\beta, b)} p'(\beta, b), \text{emp}, 1 : \mathbb{T}' \\ (\sigma_1, h_1)\tau'' &\vDash_{\mathbb{S}(\beta')} p'(\beta', \text{False}) * l, \text{emp}, 1 : \mathbb{T}'' \end{aligned}$$

then:

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}(\beta)} p'(\beta, \text{False}) * l, \text{emp}, 1 : \mathbb{T}' \mathbin{\text{;}} \mathbb{T}''$$

and one of the following hold:

$$\begin{aligned} &\text{lterm}((\sigma_0, h_0)\tau) \\ &\forall \bar{\tau} \in \llbracket \mathbb{T}' \mathbin{\text{;}} \mathbb{T}'' \rrbracket. \neg \text{liveEnv}_{\mathbb{S}(\beta)}(\bar{\tau}) \\ &\forall \hat{\tau} \in \mathbb{T}' \mathbin{\text{;}} \mathbb{T}''. \forall i \in \mathbb{N}. \exists j \geq i, \beta. \hat{\tau}(j) = (\sigma, h, p'(\beta, \text{False}), \text{emp}, 1) \end{aligned}$$

PROOF. This lemma is proven by coinduction on the structure of $(\sigma_0, h_0)\tau$. First, assume:

$$h_0 \in \llbracket p'(\beta, b) * l * \text{True} \rrbracket_{\lambda} \quad (97)$$

$$(\sigma_0, h_0)\tau' \vDash_{\mathbb{S}'(\beta, b)} p'(\beta, b), \text{emp}, 1 : \mathbb{T}' \quad (98)$$

$$(\sigma_1, h_1)\tau'' \vDash_{\mathbb{S}(\beta')} p'(\beta', \text{False}) * l, \text{emp}, 1 : \mathbb{T}'' \quad (99)$$

As, clearly, $\forall \beta. p'(\beta, \text{True}) \subseteq p'(\beta, \text{False})$, using (97), (98), (99), Lemma E.12 and Lemma E.21, by coinduction, we can derive:

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}(\beta)} p'(\beta, \text{False}) * l, \text{emp}, 1 : \mathbb{T}' \ ; \ \mathbb{T}''$$

Now, split on $\text{lterm}((\sigma_0, h_0)\tau)$. If $\text{lterm}((\sigma_0, h_0)\tau)$, then the goal holds, otherwise, split again on $\text{lterm}((\sigma_0, h_0)\tau')$. If $\neg \text{lterm}((\sigma_0, h_0)\tau')$, then $\mathbb{T}' \ ; \ \mathbb{T}'' = \mathbb{T}'$, so from (98), $\forall \bar{\tau} \in \llbracket \mathbb{T}' \ ; \ \mathbb{T}'' \rrbracket. \text{liveEnv}_{\mathbb{S}'(\beta, b)}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$, from this we infer that $\forall \bar{\tau} \in \llbracket \mathbb{T}' \ ; \ \mathbb{T}'' \rrbracket. \neg \text{liveEnv}_{\mathbb{S}'(\beta, b)}(\bar{\tau})$. Given that the definition of liveEnv only reference the pseudo-quantifier, context layer and atomicity context of the parametrising specification, this clearly implies our goal, $\forall \bar{\tau} \in \llbracket \mathbb{T}' \ ; \ \mathbb{T}'' \rrbracket. \neg \text{liveEnv}_{\mathbb{S}(\beta)}(\bar{\tau})$, as required. Otherwise, $\neg \text{lterm}((\sigma_1, h_1)\tau'')$. To not terminate, the while loop must iterate at least one more time, as $(\sigma_1, h_1)\tau''$ is a fair trace, therefore $\mathcal{B}[\mathbb{B}]_{\sigma_1}$ holds. We can then use lemma E.20 and our coinductive assumption to obtain $h_1 \in \llbracket p'(\beta, b) * l * \text{True} \rrbracket_{\lambda}$ and that one of the following holds:

$$\forall \bar{\tau} \in \llbracket \mathbb{T}'' \rrbracket. \neg \text{liveEnv}_{\mathbb{S}(\beta')}(\bar{\tau})$$

$$\forall \hat{\tau} \in \mathbb{T}'' . \forall i \in \mathbb{N}. \exists j \geq i, \beta. \hat{\tau}(j) = (\sigma, h, p'(\beta, \text{False}), \text{emp}, 1)$$

If the first holds, then $\forall \hat{\tau} \in \mathbb{T}' \ ; \ \mathbb{T}'' . \text{liveEnv}_{\mathbb{S}(\beta)}(\hat{\tau}) \Rightarrow \text{lterm}(\hat{\tau})$, so the goal is proven; if the second holds, then from $h_1 \in \llbracket p'(\beta, b) * l * \text{True} \rrbracket_{\lambda}$:

$$\forall \hat{\tau} \in \mathbb{T}' \ ; \ \mathbb{T}'' . \forall i \in \mathbb{N}. \exists j \geq i, \beta. \hat{\tau}(j) = (\sigma, h, p'(\beta, \text{False}), \text{emp}, 1) \quad \square$$

THEOREM E.23. *Given*

$$\forall \beta \leq \beta_0. \forall b \in \{0, 1\}. \vDash_{\Phi} \mathbb{C} : \mathbb{S}'(\beta, b) \quad (100)$$

$$\forall \beta \leq \beta_0. m(\beta); \lambda; \mathcal{A} \vDash L \xrightarrow{M} T(\beta) \quad (101)$$

$$\forall \alpha. \mathcal{A} \vDash \exists \alpha'. L * M(\alpha') \wedge \alpha' \leq \alpha \text{ stable} \quad (102)$$

$$\mathcal{A} \vDash L \text{ stable} \quad (103)$$

$$\forall \beta \leq \beta_0. \vdash_{\mathcal{A}} P(\beta) \triangleright m(\beta) \leq m \quad (104)$$

$$\text{pv}(T, L, M) \cap \text{mod}(\mathbb{C}) = \emptyset \quad (105)$$

then:

$$\vDash_{\Phi} \text{while}(\mathbb{B})\{\mathbb{C}\} : \mathbb{S}(\beta_0)$$

PROOF. Taking $\varphi \in \text{Flmpl}$ arbitrary such that $\vDash \varphi : \Phi$ and $(\sigma_0, h_0)\tau \in \llbracket \text{while}(\mathbb{B})\{\mathbb{C}\} \rrbracket_{\varphi}$ arbitrary. We need to show $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S}(\beta_0) \rrbracket$. Let

$$p'(\beta, b) = \mathcal{W}[\llbracket P(\beta) * (b \dot{\Rightarrow} T) \rrbracket_{\mathcal{A}}^{\sigma_0}] \quad l = \mathcal{W}[\llbracket L \rrbracket_{\mathcal{A}}^{\sigma_0}]$$

To reach the goal, assume $h_0 \in \llbracket p'(\beta_0, \text{False}) * l * \text{True} \rrbracket_{\lambda}$. By lemma E.22, in the case that $\mathcal{B}[\mathbb{B}]_{\sigma_0}$, and our assumptions, there exists $\mathbb{T} \subseteq \text{STrace}$:

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} \mathcal{W}[\llbracket P(\beta_0) * L \rrbracket_{\mathcal{A}}^{\sigma_0}], \text{emp}, 1 : \mathbb{T}$$

and one of the following holds:

$$\text{lterm}((\sigma_0, h_0)\tau)$$

$$\forall \hat{\tau} \in \mathbb{T}. \neg \text{liveEnv}_{\mathbb{S}}(\hat{\tau})$$

$$\forall \hat{\tau} \in \mathbb{T}. \forall i \in \mathbb{N}. \exists j \geq i, \beta. \hat{\tau}(j) = (\sigma, h, p'(\beta, \text{False}), \text{emp}, 1)$$

In the first case, $\forall \hat{\tau} \in \mathbb{T}. \text{lterm}(\hat{\tau})$, therefore, $\forall \hat{\tau} \in \mathbb{T}. \text{liveEnv}_{\mathbb{S}(\beta_0)}(\hat{\tau}) \Rightarrow \text{lterm}(\hat{\tau})$, as required. In the second, $\forall \hat{\tau} \in \mathbb{T}. \text{liveEnv}_{\mathbb{S}(\beta_0)}(\hat{\tau}) \Rightarrow \text{lterm}(\hat{\tau})$ clearly also holds. Finally, we consider the third case. Take $\hat{\tau} \in \mathbb{T}$ arbitrary and assume $\text{liveEnv}_{\mathbb{S}(\beta_0)}(\hat{\tau})$. Now, for a contradiction, assume $\neg \text{lterm}(\hat{\tau})$. In this case, due to (101), with an argument similar to that in the soundness of (LIVEC), at every

point, every $\hat{\tau} \in \mathbb{T}$ eventually reaches a state satisfying $T(\beta_0)$. This must eventually be stable due to the metric stably decreasing due to assumption (102), holding till the next iteration, at which point, the loop variant decreases due to (100) with $b = \text{True}$. By repeating this argument with the continuation, by well-foundedness of ordinals, the while loop must eventually terminate if $\text{liveEnv}(\hat{\tau})$ holds, leading to a contradiction. Therefore, in all cases, $\forall \hat{\tau} \in \mathbb{T}. \text{liveEnv}_{\mathbb{S}(\beta_0)}(\hat{\tau}) \Rightarrow \text{lterm}(\hat{\tau})$ holds, as required, concluding the proof. \square

E.6 Soundness of PAR

Definition E.24 (Bowtie operator). The bowtie operator, \bowtie , which interleaves the subjective traces of two commands executed in parallel into a command from their combined perspective:

$$\begin{aligned} (\sigma, h) \text{ env } \tau'_1 \bowtie (\sigma, h) \text{ env } \tau'_2 &= (\sigma, h) \text{ env } (\tau'_1 \bowtie \tau'_2) \\ (\sigma, h) \text{ env } \tau'_1 \bowtie (\sigma, h) \text{ loc } \tau'_2 &= (\sigma, h) \text{ loc } (\tau'_1 \bowtie \tau'_2) \\ (\sigma, h) \text{ loc } \tau'_1 \bowtie (\sigma, h) \text{ env } \tau'_2 &= (\sigma, h) \text{ loc } (\tau'_1 \bowtie \tau'_2) \end{aligned}$$

All other cases are undefined.

Definition E.25 (Specification Bowtie operator). The specification bowtie operator, $\overset{\mathbb{S}}{\bowtie}$, which interleaves the subjective specification traces of two commands executed in parallel into a command from their combined perspective:

$$\begin{aligned} (\sigma, h, p_1, \text{emp}, 1) \text{ env } \tau'_1 \overset{\mathbb{S}}{\bowtie} (\sigma, h, p_2, \text{emp}, 1) \text{ env } \tau'_2 &= (\sigma, h, p_1 * p_2, \text{emp}, 1) \text{ env } (\tau'_1 \overset{\mathbb{S}}{\bowtie} \tau'_2) \\ (\sigma, h, p_1, \text{emp}, 1) \text{ env } \tau'_1 \overset{\mathbb{S}}{\bowtie} (\sigma, h, p_2, \text{emp}, 1) \text{ loc } \tau'_2 &= (\sigma, h, p_1 * p_2, \text{emp}, 1) \text{ loc } (\tau'_1 \overset{\mathbb{S}}{\bowtie} \tau'_2) \\ (\sigma, h, p_1, \text{emp}, 1) \text{ loc } \tau'_1 \overset{\mathbb{S}}{\bowtie} (\sigma, h, p_2, \text{emp}, 1) \text{ env } \tau'_2 &= (\sigma, h, p_1 * p_2, \text{emp}, 1) \text{ loc } (\tau'_1 \overset{\mathbb{S}}{\bowtie} \tau'_2) \end{aligned}$$

All other cases are undefined.

LEMMA E.26. For any $\varphi \in \text{Flmpl}$:

$$\forall \tau \in \llbracket \mathbb{C}_1 \parallel \mathbb{C}_2 \rrbracket_{\varphi}. \exists \tau_1 \in \llbracket \mathbb{C}_1 \rrbracket_{\varphi}, \tau_2 \in \llbracket \mathbb{C}_2 \rrbracket_{\varphi}. \tau = \tau_1 \bowtie \tau_2$$

PROOF. Straightforward by induction on \Rightarrow_{φ} . \square

LEMMA E.27. For any trace $(\sigma_0, h_0) \tau (\sigma_1, h_1) \tau' \in \llbracket \mathbb{C} \rrbracket_{\varphi}$, we have $\forall x \in \text{PVar} \setminus \text{mods}(\mathbb{C}). \sigma_0(x) = \sigma_1(x)$.

PROOF. Straightforward by induction on the length of the trace. \square

For the rest of the section, we name the specifications involved in the PAR rule as follows:

$$\mathbb{S}_1 = \{P_1\} \cdot \{Q_1\}_{m_1; \lambda; \mathcal{A}} \quad \mathbb{S}_2 = \{P_2\} \cdot \{Q_2\}_{m_2; \lambda; \mathcal{A}} \quad \mathbb{S} = \{P_1 * P_2\} \cdot \{Q_1 * Q_2\}_{m; \lambda; \mathcal{A}}$$

LEMMA E.28. For arbitrary $(\sigma_0, h_0) \tau, (\sigma_0, h_0) \tau_1, (\sigma_0, h_0) \tau_2 \in \text{Trace}$, $\mathbb{T}_1, \mathbb{T}_2 \in \mathcal{P}(\text{STrace})$, $v_1, v_2 \in \{1, \langle 1, 1 \rangle\}$, and $p'_1, p'_2 \in \text{View}_{\mathcal{A}}$, then:

$$\left. \begin{array}{l} (\sigma_0, h_0) \tau = (\sigma_0, h_0) \tau_1 \bowtie (\sigma_0, h_0) \tau_2 \\ (\sigma_0, h_0) \tau_1 \vDash_{\mathbb{S}_1} p'_1, \text{emp}, v_1 : \mathbb{T}_1 \\ (\sigma_0, h_0) \tau_2 \vDash_{\mathbb{S}_2} p'_2, \text{emp}, v_2 : \mathbb{T}_2 \\ h_0 \in \llbracket p'_1 * p'_2 * \text{True} \rrbracket_{\lambda} \\ \text{term}((\sigma_0, h_0) \tau_1) \Rightarrow p'_1 = \mathcal{W} \llbracket Q_1 \rrbracket_{\mathcal{A}}^{\sigma} \\ \text{term}((\sigma_0, h_0) \tau_2) \Rightarrow p'_2 = \mathcal{W} \llbracket Q_2 \rrbracket_{\mathcal{A}}^{\sigma} \end{array} \right\} \Rightarrow \begin{array}{l} \exists \mathbb{T} \in \mathcal{P}(\text{STrace}), v \in \{1, \langle 1, 1 \rangle\}. \\ (\sigma, h) \tau \vDash_{\mathbb{S}} p'_1 * p'_2, \text{emp}, v : \mathbb{T} \wedge \\ \forall \hat{\tau} \in \mathbb{T}. \exists \hat{\tau}_1 \in \mathbb{T}_1, \hat{\tau}_2 \in \mathbb{T}_2. \hat{\tau} = \hat{\tau}_1 \overset{\mathbb{S}}{\bowtie} \hat{\tau}_2 \wedge \\ (v_1 = \langle 1, 1 \rangle \wedge v_2 = \langle 1, 1 \rangle) \Leftrightarrow v = \langle 1, 1 \rangle \end{array}$$

PROOF. This lemma is proven by coinduction on the structure of $(\sigma_0, h_0)\tau$.

The trace either starts with a local, or an environment step. We split on the two cases:

Case $(\sigma, h)\tau = (\sigma, h) \text{env } (\sigma, h')\tau'$: Take $(\sigma_0, h_0)\tau_1, (\sigma_0, h_0)\tau_2 \in \text{Trace}$, $\mathbb{T}_1, \mathbb{T}_2 \in \mathcal{P}(\text{STrace})$, $v_1, v_2 \in \{1, \langle 1, 1 \rangle\}$, and, $p'_1, p'_2 \in \text{View}_{\mathcal{A}}$ arbitrary, and assume:

$$(\sigma_0, h_0)\tau = (\sigma_0, h_0)\tau_1 \bowtie (\sigma_0, h_0)\tau_2 \quad (106)$$

$$(\sigma_0, h_0)\tau_1 \vDash_{\mathbb{S}_1} p'_1, \text{emp}, v_1 : \mathbb{T}_1 \quad (107)$$

$$(\sigma_0, h_0)\tau_2 \vDash_{\mathbb{S}_2} p'_2, \text{emp}, v_2 : \mathbb{T}_2 \quad (108)$$

$$h_0 \in \llbracket p'_1 * p'_2 * \text{True} \rrbracket_{\lambda} \quad (109)$$

$$\text{term}((\sigma_0, h_0)\tau_1) \Rightarrow p'_1 = \mathcal{W}[\llbracket Q_1 \rrbracket_{\mathcal{A}}]_{\sigma}^{\sigma} \quad (110)$$

$$\text{term}((\sigma_0, h_0)\tau_2) \Rightarrow p'_2 = \mathcal{W}[\llbracket Q_2 \rrbracket_{\mathcal{A}}]_{\sigma}^{\sigma} \quad (111)$$

Given (106) and the definition of \bowtie :

$$(\sigma_0, h_0)\tau_1 = (\sigma_0, h_0) \text{env } (\sigma_0, h')\tau'_1$$

$$(\sigma_0, h_0)\tau_2 = (\sigma_0, h_0) \text{env } (\sigma_0, h')\tau'_2$$

$$(\sigma_0, h')\tau' = (\sigma_0, h')\tau'_1 \bowtie (\sigma_0, h')\tau'_2$$

Now to prove the goal, consider the case $v_1 = \langle 1, 1 \rangle$ and $v_2 = \langle 1, 1 \rangle$. In this case, take $v = \langle 1, 1 \rangle$, so **ENV'** must hold for the goal as well as (107) and (108). Note that this choice of v immediately satisfies the third conjunct of the goal. To show **ENV'** holds for the goal, given some $p_e, p'_e \in \text{View}_{\mathcal{A}}$, assume:

$$h_0 \in \llbracket p'_1 * p'_2 * p_e \rrbracket \wedge (h_0, h') \vDash_{\lambda, \mathcal{A}} p'_1 * p'_2 * p_e \rightarrow p'_1 * p'_2 * p'_e$$

By substitution, this implies both:

$$\exists p_e, p'_e. h_0 \in \llbracket p'_1 * p_e \rrbracket \wedge (h, h') \vDash_{\lambda, \mathcal{A}} p'_1 * p_e \rightarrow p'_1 * p'_e$$

$$\exists p_e, p'_e. h_0 \in \llbracket p'_2 * p_e \rrbracket \wedge (h, h') \vDash_{\lambda, \mathcal{A}} p'_2 * p_e \rightarrow p'_2 * p'_e$$

Given 107 and 108, these imply:

$$(\sigma_0, h')\tau'_1 \vDash_{\mathbb{S}_1} p'_1, \text{emp}, v_1 : \mathbb{T}'_1 \quad \mathbb{T}_1 = (\sigma_0, h_0, p_1, \text{emp}, \langle 1, 1 \rangle) \text{env } \mathbb{T}'_1$$

$$(\sigma_0, h')\tau'_2 \vDash_{\mathbb{S}_2} p'_2, \text{emp}, v_2 : \mathbb{T}'_2 \quad \mathbb{T}_2 = (\sigma_0, h_0, p_2, \text{emp}, \langle 1, 1 \rangle) \text{env } \mathbb{T}'_2$$

Assumption (109) and $(h_0, h') \vDash_{\lambda, \mathcal{A}} p'_1 * p'_2 * p_e \rightarrow p'_1 * p'_2 * p'_e$ yield:

$$h' \in \llbracket p'_1 * p'_2 * \text{True} \rrbracket_{\lambda} \quad (112)$$

Now, by using the inductive assumption, as (110) and (111) clearly imply the same assertions for $(\sigma_0, h')\tau'_1$ and $(\sigma_0, h')\tau'_2$ respectively, for some $\mathbb{T}' \in \mathcal{P}(\text{STrace})$:

$$(\sigma_0, h')\tau' \vDash_{\mathbb{S}} p'_1 * p'_2, \text{emp}, v : \mathbb{T}' \quad (113)$$

$$\forall \hat{\tau} \in \mathbb{T}'. \exists \hat{\tau}_1 \in \mathbb{T}'_1, \hat{\tau}_2 \in \mathbb{T}'_2. \hat{\tau} = \hat{\tau}_1 \overset{\mathbb{S}}{\bowtie} \hat{\tau}_2 \quad (114)$$

From this first consequence:

$$(\sigma_0, h)\tau \vDash_{\mathbb{S}} p'_1 * p'_2, \text{emp}, v : \mathbb{T}$$

holds, where $\mathbb{T} = (\sigma_0, h, p'_1 * p'_2, \text{emp}, v) \text{env } \mathbb{T}'$. This is the first conjunct of the goal.

Finally, taking $\hat{\tau} \in \mathbb{T}$ arbitrary, there exists $\hat{\tau}' \in \mathbb{T}'$ such that $\hat{\tau} = (\sigma_0, h, p'_1 * p'_2, \text{emp}, v) \text{env } \hat{\tau}'$. From the second consequence of our inductive assumption, it follows that there exist $\hat{\tau}'_1 \in \mathbb{T}'_1$ and $\hat{\tau}'_2 \in \mathbb{T}'_2$ such that $\hat{\tau}' = \hat{\tau}'_1 \overset{\mathbb{S}}{\bowtie} \hat{\tau}'_2$. Then, from the definitions of \mathbb{T}_1 and \mathbb{T}_2 ,

$(\sigma_0, h_0, p_1, \text{emp}, \langle 1, 1 \rangle) \text{ env } \hat{\tau}'_1 \in \mathbb{T}_1$ and $(\sigma_0, h_0, p_2, \text{emp}, \langle 1, 1 \rangle) \text{ env } \hat{\tau}'_2 \in \mathbb{T}_2$ hold, and $\hat{\tau} = (\sigma_0, h_0, p_1, \text{emp}, \langle 1, 1 \rangle) \text{ env } \hat{\tau}'_1 \stackrel{\text{S}}{\bowtie} (\sigma_0, h_0, p_2, \text{emp}, \langle 1, 1 \rangle) \text{ env } \hat{\tau}'_2 \in \mathbb{T}_2$ holds as required.

Other cases for v_1, v_2 follow similarly.

Case $(\sigma, h)\tau = (\sigma, h) \text{ loc } (\sigma, h')\tau'$: Here the variable store does not change as $\text{mods}(\mathbb{C}_1 || \mathbb{C}_2) = \emptyset$, due to Lemma E.27 and the syntactic restriction on parallel commands, requiring both threads to not modify the value of any variable. To prove the goal, take $(\sigma_0, h_0)\tau_1, (\sigma_0, h_0)\tau_2 \in \text{Trace}$, $\mathbb{T}_1, \mathbb{T}_2 \in \mathcal{P}(\text{STrace})$, $v_1, v_2 \in \{1, \langle 1, 1 \rangle\}$, and, $p'_1, p'_2 \in \text{View}_{\mathcal{A}}$ arbitrary, and assume:

$$(\sigma_0, h_0)\tau = (\sigma_0, h_0)\tau_1 \bowtie (\sigma_0, h_0)\tau_2 \quad (115)$$

$$(\sigma_0, h_0)\tau_1 \vDash_{\mathbb{S}_1} p'_1, \text{emp}, v_1 : \mathbb{T}_1 \quad (116)$$

$$(\sigma_0, h_0)\tau_2 \vDash_{\mathbb{S}_2} p'_2, \text{emp}, v_2 : \mathbb{T}_2 \quad (117)$$

$$h_0 \in \llbracket p'_1 * p'_2 * \text{True} \rrbracket_{\lambda} \quad (118)$$

$$\text{term}((\sigma_0, h_0)\tau_1) \Rightarrow p'_1 = \mathcal{W} \llbracket Q_1 \rrbracket_{\mathcal{A}}^{\sigma} \quad (119)$$

$$\text{term}((\sigma_0, h_0)\tau_2) \Rightarrow p'_2 = \mathcal{W} \llbracket Q_2 \rrbracket_{\mathcal{A}}^{\sigma} \quad (120)$$

Given (115) and the definition of \bowtie , either:

$$(\sigma_0, h_0)\tau_1 = (\sigma_0, h_0) \text{ loc } (\sigma_0, h')\tau'_1$$

$$(\sigma_0, h_0)\tau_2 = (\sigma_0, h_0) \text{ env } (\sigma_0, h')\tau'_2$$

or:

$$(\sigma_0, h_0)\tau_1 = (\sigma_0, h_0) \text{ env } (\sigma_0, h')\tau'_1$$

$$(\sigma_0, h_0)\tau_2 = (\sigma_0, h_0) \text{ loc } (\sigma_0, h')\tau'_2$$

and in both cases:

$$(\sigma_0, h')\tau' = (\sigma_0, h')\tau'_1 \bowtie (\sigma_0, h')\tau'_2$$

Consider the first case, the second will follow symmetrically. Assume that the **STUTTER** rule holds for $(\sigma, h) \text{ loc } (\sigma, h')\tau_1 \vDash_{\mathbb{S}_1} p'_1, \text{emp}, v_1 : \mathbb{T}_1$, then, for some $p''_1 \in \text{View}_{\mathcal{A}}$:

$$(h_0, h') \vDash_{\lambda; \mathcal{A}} p'_1 \rightarrow p''_1$$

$$(\sigma_0, h') \tau'_1 \vDash_{\mathbb{S}_1} p''_1, \text{emp}, v_1 : \mathbb{T}'_1$$

$$\text{term}((\sigma_0, h')\tau'_1) \Rightarrow v_1 = \langle 1, 1 \rangle \wedge p''_1 = \mathcal{W} \llbracket Q_1 \rrbracket_{\mathcal{A}}^{\sigma_0}$$

where $\mathbb{T}'_1 = (\sigma_0, h_0, p_1, \text{emp}, v_1) \text{ loc } \mathbb{T}'_1$. Given 118 and $(h_0, h') \vDash_{\lambda; \mathcal{A}} p'_1 \rightarrow p''_1$, $h' \in \llbracket p''_1 * p'_2 * \text{True} \rrbracket_{\lambda}$ holds. Given $(h_0, h') \vDash_{\lambda; \mathcal{A}} p'_1 \rightarrow p''_1$, $(h_0, h') \vDash_{\lambda; \mathcal{A}} p'_1 * p_2 \rightarrow p''_1 * p_2$, also holds. Using this and **ENV** or **ENV'**:

$$(\sigma, h') \tau'_2 \vDash_{\mathbb{S}_2} p'_2, \text{emp}, v_2 : \mathbb{T}'_2$$

where $\mathbb{T}'_2 = (\sigma_0, h_0, p_2, \text{emp}, v_2) \text{ loc } \mathbb{T}'_2$. Now using the inductive assumption, as, once again, (119) and (120) clearly imply the same assertions for $(\sigma_0, h')\tau'_1$ and $(\sigma_0, h')\tau'_2$ respectively, for some $\mathbb{T}' \in \mathcal{P}(\text{STrace})$:

$$(\sigma_0, h')\tau' \vDash_{\mathbb{S}} p''_1 * p'_2, \text{emp}, v : \mathbb{T}' \wedge \quad (121)$$

$$\forall \hat{\tau} \in \mathbb{T}'. \exists \hat{\tau}_1 \in \mathbb{T}'_1, \hat{\tau}_2 \in \mathbb{T}'_2. \hat{\tau} = \hat{\tau}_1 \stackrel{\text{S}}{\bowtie} \hat{\tau}_2 \wedge \quad (122)$$

$$(v_1 = \langle 1, 1 \rangle \wedge v_2 = \langle 1, 1 \rangle) \Leftrightarrow v = \langle 1, 1 \rangle \quad (123)$$

The second and third consequents imply the equivalent conjuncts of the goal with the same method as in the env case and directly respectively. As we have shown $(h, h') \vDash_{\lambda; \mathcal{A}} p'_1 * p_2 \rightarrow$

$p_1'' * p_2$ holds, using the **STUTTER** rule, to show that $(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p_1' * p_2', \text{emp}, v : \mathbb{T}$ holds, where $\mathbb{T} = (\sigma_0, h_0, p_1' * p_2', \text{emp}, v) \text{ loc } \mathbb{T}'$, it suffices to show:

$$\text{term}((\sigma_0, h')\tau') \Rightarrow v = \langle 1, 1 \rangle \wedge p_1'' * p_2' = \mathcal{W}[[Q_1 * Q_2]]_{\mathcal{A}}^{\sigma_0}$$

Assuming $\text{term}((\sigma_0, h')\tau')$ holds, then $\text{term}((\sigma_0, h')\tau_1')$ and $\text{term}((\sigma_0, h')\tau_2')$ hold. From this it follows that $v_1, v_2 = \langle 1, 1 \rangle$, so, due to (123), $v = \langle 1, 1 \rangle$.

Finally, due to $\text{term}((\sigma_0, h')\tau_1')$ and $\text{term}((\sigma_0, h')\tau_2')$, $p_1'' = \mathcal{W}[[Q_1]]_{\mathcal{A}}^{\sigma_0}$ and $p_2' = \mathcal{W}[[Q_2]]_{\mathcal{A}}^{\sigma}$ hold respectively, yielding $p_1'' * p_2' = \mathcal{W}[[Q_1 * Q_2]]_{\mathcal{A}}^{\sigma_0}$, as required.

The **LINPT** rule follows similarly. \square

THEOREM E.29. *Given*

$$m_1; \lambda; \mathcal{A} \vDash_{\Phi} \{P_1\} \mathbb{C}_1 \{Q_1\} \quad (124)$$

$$m_2; \lambda; \mathcal{A} \vDash_{\Phi} \{P_2\} \mathbb{C}_2 \{Q_2\} \quad (125)$$

$$\lambda; \mathcal{A} \vdash Q_1 \triangleright m_2 \leq m \quad (126)$$

$$\lambda; \mathcal{A} \vdash Q_2 \triangleright m_1 \leq m \quad (127)$$

then:

$$m; \lambda; \mathcal{A} \vDash_{\Phi} \{P_1 * P_2\} \mathbb{C}_1 || \mathbb{C}_2 \{Q_1 * Q_2\}$$

PROOF. Taking $\varphi \in \text{Flmpl}$ arbitrary such that $\vDash_{\Phi} \varphi$, from (124) and (125), $[[\mathbb{C}_1]]_{\varphi} \subseteq [[\mathbb{S}_1]]$ and $[[\mathbb{C}_2]]_{\varphi} \subseteq [[\mathbb{S}_2]]$ hold. Given an arbitrary $(\sigma_0, h_0)\tau \in [[\mathbb{C}_1 || \mathbb{C}_2]]_{\varphi}$, need to show $(\sigma_0, h_0)\tau \in [[\mathbb{S}]]$. Let

$$p_1 = \mathcal{W}[[P_1]]_{\mathcal{A}}^{\sigma_0} \quad p_2 = \mathcal{W}[[P_2]]_{\mathcal{A}}^{\sigma_0}$$

To reach the goal, assume $h_0 \in [[p_1 * p_2 * \text{True}]]_{\lambda}$. Then $h_0 \in [[p_1 * \text{True}]]_{\lambda}$ and $h_0 \in [[p_2 * \text{True}]]_{\lambda}$ hold. From E.26 and the definition of \bowtie , there exists $(\sigma_0, h_0)\tau_1 \in [[\mathbb{C}_1]]_{\varphi}$ and $(\sigma_0, h_0)\tau_2 \in [[\mathbb{C}_2]]_{\varphi}$ such that $(\sigma_0, h_0)\tau = (\sigma_0, h_0)\tau_1 \bowtie (\sigma_0, h_0)\tau_2$. As $[[\mathbb{C}_1]]_{\varphi} \subseteq [[\mathbb{S}_1]]$ and $[[\mathbb{C}_2]]_{\varphi} \subseteq [[\mathbb{S}_2]]$, $(\sigma_0, h_0)\tau_1 \in [[\mathbb{S}_1]]$ and $(\sigma_0, h_0)\tau_2 \in [[\mathbb{S}_2]]$ hold. Now, as $h_0 \in [[p_1 * \text{True}]]_{\lambda}$ and $h_0 \in [[p_2 * \text{True}]]_{\lambda}$, then for some $\mathbb{T}_1, \mathbb{T}_2 \in \mathcal{P}(\text{STrace})$:

$$\begin{aligned} (\sigma_0, h_0)\tau_1 \vDash_{\mathbb{S}} p_1, \text{emp}, 1 : \mathbb{T}_1 & \quad \forall \hat{\tau}_1 \in [[\mathbb{T}_1]]. \text{liveEnv}_{\mathbb{S}}(\hat{\tau}_1) \Rightarrow \text{lterm}(\hat{\tau}_1) \\ (\sigma_0, h_0)\tau_2 \vDash_{\mathbb{S}} p_2, \text{emp}, 1 : \mathbb{T}_2 & \quad \forall \hat{\tau}_2 \in [[\mathbb{T}_2]]. \text{liveEnv}_{\mathbb{S}}(\hat{\tau}_2) \Rightarrow \text{lterm}(\hat{\tau}_2) \end{aligned}$$

As all commands must take at least one step, $\neg \text{term}((\sigma_0, h_0)\tau_1)$ and $\neg \text{term}((\sigma_0, h_0)\tau_2)$ hold, therefore:

$$\begin{aligned} \text{term}((\sigma_0, h_0)\tau_1) \Rightarrow p_1 = \mathcal{W}[[Q_1]]_{\mathcal{A}}^{\sigma} \\ \text{term}((\sigma_0, h_0)\tau_2) \Rightarrow p_2 = \mathcal{W}[[Q_2]]_{\mathcal{A}}^{\sigma} \end{aligned}$$

hold. Now, using lemma E.28, there exists $\mathbb{T} \in \mathcal{P}(\text{STrace})$ such that:

$$(\sigma_0, h_0)\tau \vDash_{\mathbb{S}} p_1 * p_2, \text{emp}, 1 : \mathbb{T}$$

and for any $\hat{\tau} \in \mathbb{T}$, there exist $\hat{\tau}_1 \in \mathbb{T}_1$ and $\hat{\tau}_2 \in \mathbb{T}_2$, such that $\hat{\tau} = \hat{\tau}_1 \bowtie \hat{\tau}_2$. It now suffices to show that $\forall \bar{\tau} \in [[\mathbb{T}]]. \text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$. Take $\hat{\tau} \in \mathbb{T}$ arbitrary and $\hat{\tau}_1 \in \mathbb{T}_1$ and $\hat{\tau}_2 \in \mathbb{T}_2$ such that $\hat{\tau} = \hat{\tau}_1 \bowtie \hat{\tau}_2$ and $\bar{\tau} \in [[\hat{\tau}]], \bar{\tau}_1 \in [[\hat{\tau}_1]], \bar{\tau}_2 \in [[\hat{\tau}_2]]$. From above:

$$\text{liveEnv}_{\mathbb{S}_1}(\bar{\tau}_1) \Rightarrow \text{lterm}(\bar{\tau}_1) \quad (128)$$

$$\text{liveEnv}_{\mathbb{S}_2}(\bar{\tau}_2) \Rightarrow \text{lterm}(\bar{\tau}_2) \quad (129)$$

holds. To reach the goal, split on $\text{lterm}(\bar{\tau}_1)$ and $\text{lterm}(\bar{\tau}_2)$.

Case $\text{lterm}(\bar{\tau}_1) \wedge \text{lterm}(\bar{\tau}_2)$: In this case, clearly $\text{lterm}(\bar{\tau})$ holds, therefore $\text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$ holds trivially, as required.

Case $\text{lterm}(\bar{\tau}_1) \wedge \neg \text{lterm}(\bar{\tau}_2)$: From $\neg \text{lterm}(\bar{\tau}_2)$, by (129), $\neg \text{liveEnv}_{\mathbb{S}_2}(\bar{\tau}_2)$ holds:

$$\begin{aligned} \exists \widehat{O} \in \text{POb}_{< m_2}^{\mathbb{S}_2}. (\forall O \in \text{AOB}_{< \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}_2(j))) \wedge \\ (\exists i \in \mathbb{N}. \forall j \geq i. \text{envheld}(\widehat{O}, \bar{\tau}_2(j))) \end{aligned}$$

As $\text{lterm}(\bar{\tau}_1)$, by Lemma E.21, there exists some $i_1 \in \mathbb{N}$, an index after which the trace $\bar{\tau}_1$ only performs env steps, in particular, for any $j \geq i_1$, $\bar{\tau}_1(j) = (\sigma, h, w_h^1, w_a^1, \langle 1, 1 \rangle)$, where $w_h^1 \in \mathcal{W}[[Q_1]]_{\sigma}^{\mathcal{A}}$ and $w_a^1 \in \text{Emp}_{\mathcal{A}}$. Therefore $\bar{\tau}(j) = (\sigma, h, w_h^1 \odot w_h^2, w_a^1 \odot w_a^2, \langle 1, 1 \rangle)$, where $\bar{\tau}_2(j) = (\sigma, h, w_h^2, w_a^2, \langle 1, 1 \rangle)$, such that $w_a^2 \in \text{Emp}_{\mathcal{A}}$. Given $\lambda; \mathcal{A} \vdash Q_1 \triangleright m_2$, it is clear that:

$$\begin{aligned} (\forall O \in \text{AOB}_{< \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}_2(j))) \Rightarrow \\ (\forall O \in \text{AOB}_{< \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}(j))) \end{aligned}$$

and similarly as $\widehat{O} \in \text{POb}_{< m_2}^{\mathbb{S}_2}$:

$$(\exists i \in \mathbb{N}. \forall j \geq i. \text{envheld}(\widehat{O}, \bar{\tau}_2(j))) \Rightarrow (\exists i \in \mathbb{N}. \forall j \geq i. \text{envheld}(\widehat{O}, \bar{\tau}(j)))$$

As $m_2 \leq m$, $\widehat{O} \in \text{POb}_{< m}^{\mathbb{S}}$. Finally, from this $\neg \text{liveEnv}_{\mathbb{S}}(\bar{\tau})$ holds, implying $\text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$, as required.

Case $\neg \text{lterm}(\hat{\tau}_1) \wedge \text{lterm}(\hat{\tau}_2)$: Similarly to the previous case.

Case $\neg \text{lterm}(\bar{\tau}_1) \wedge \neg \text{lterm}(\bar{\tau}_2)$: Given (128) and (129), we can infer $\neg \text{liveEnv}_{\mathbb{S}_1}(\bar{\tau}_1)$ and $\neg \text{liveEnv}_{\mathbb{S}_2}(\bar{\tau}_2)$.

Assume $\text{liveEnv}_{\mathbb{S}}(\bar{\tau})$ for a contradiction. From $\neg \text{liveEnv}_{\mathbb{S}_1}(\bar{\tau}_1)$, for some $\widehat{O} \in \text{POb}_{< m_1}^{\mathbb{S}}$:

$$(\forall O \in \text{AOB}_{< \text{lay}(\widehat{O})}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}_1(j))) \wedge (\exists i \in \mathbb{N}. \forall j \geq i. \text{envheld}(\widehat{O}, \bar{\tau}_1(j)))$$

From this and $\text{liveEnv}_{\mathbb{S}}(\bar{\tau})$, there is some $i \in \mathbb{N}$ such that:

$$\forall j \geq i. \text{locheld}(\widehat{O}, \bar{\tau}_2(j))$$

From $\neg \text{liveEnv}_{\mathbb{S}_2}(\bar{\tau}_2)$, for some $\widehat{O}' \in \text{POb}_{< m_2}^{\mathbb{S}_2}$:

$$(\forall O \in \text{AOB}_{< \text{lay}(\widehat{O}')}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}_2(j))) \wedge (\exists i \in \mathbb{N}. \forall j \geq i. \text{envheld}(\widehat{O}', \bar{\tau}_2(j)))$$

Given that $\forall j \geq i. \text{locheld}(\widehat{O}, \bar{\tau}_2(j))$, for $\forall O \in \text{AOB}_{< \text{lay}(\widehat{O}')}. \forall i \in \mathbb{N}. \exists j \geq i. \neg \text{locheld}(O, \bar{\tau}_2(j))$ to hold, it must be the case that $\text{lay}(\widehat{O}) > \text{lay}(\widehat{O}')$. This argument can be repeated ad-infinitum, which, by the well-foundedness of layers, leads to a contradiction, and therefore $\neg \text{liveEnv}_{\mathbb{S}}(\bar{\tau})$ holds. This implies $\text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$.

From these cases, we deduce that $\forall \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}(\bar{\tau})$.

From this, we can infer $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket$ and consequently, $\llbracket \mathbb{C}_1 \parallel \mathbb{C}_2 \rrbracket_{\varphi} \subseteq \llbracket \mathbb{S} \rrbracket$, as required. \square

E.7 Soundness of LIFTAG

Recall the triples of the premise and conclusion of rule LIFTAG:

$$\begin{aligned} \mathbb{S} &= \forall x \in \vec{X}. \left\langle P_h \left| \mathcal{I}(t_r^\lambda(x)) * P_a(x) * \lceil G \rceil_r * \lfloor O_1 \rfloor_r^\perp \right. \right\rangle \cdot \exists y. \left\langle \begin{array}{l} Q_h(x, y) \wedge y \in Y(x) \\ \exists z. \mathcal{I}(t_r^\lambda(z)) * Q_a(x, y, z) \\ * \lfloor O_2(x, y) \rfloor_r^\perp \wedge R(x, z) \end{array} \right\rangle_{m; \lambda; \mathcal{A}} \\ \mathbb{S}' &= \forall x \in \vec{X}. \left\langle P_h * \lfloor O_1 \rfloor_r^\perp \left| t_r^\lambda(x) * P_a(x) * \lceil G \rceil_r \right. \right\rangle \cdot \exists y. \left\langle \begin{array}{l} Q_h(x, y) * \lfloor O_2(x, y) \rfloor_r^\perp \wedge y \in Y(x) \\ \exists z. t_r^\lambda(z) * Q_a(x, y, z) \wedge R(x, z) \end{array} \right\rangle_{m; \lambda+1; \mathcal{A}} \end{aligned}$$

and let us name the semantic counterparts of their atomic pre/postconditions as follows:

$$\begin{aligned}\bar{p}_a(v) &= \begin{cases} \mathcal{W}[\![I(\mathbf{t}_r^\lambda(v)) * P_a(v) * [G]_r * [O_1]_r^\perp \wedge v \in X]\!]_{\mathcal{A}} & \text{if } v \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \\ p_a(v) &= \begin{cases} \mathcal{W}[\![\mathbf{t}_r^\lambda(v) * P_a(v) * [G]_r \wedge v \in X]\!]_{\mathcal{A}} & \text{if } v \in \text{AVal} \\ \text{Emp}_{\mathcal{A}} & \text{otherwise} \end{cases} \\ \bar{q}_a(v, v', z) &= \mathcal{W}[\![I(\mathbf{t}_r^\lambda(z)) * Q_a(v, v', z) * [O_2(v, v')]_r^\perp \wedge R(v, z)]\!]_{\mathcal{A}}^0 \\ q_a(v, v', z) &= \mathcal{W}[\![\mathbf{t}_r^\lambda(z) * Q'_a(v, v', z) \wedge R(v, z)]\!]_{\mathcal{A}}^0\end{aligned}$$

Definition E.30 (lift).

$$\text{lift}((\sigma, h, p_h, p_a, v) \pi \hat{\tau}) \triangleq \left\{ (\sigma, h, p_h * [O_1]_r^\perp, p'_a, v) \pi \hat{\tau}' \mid \begin{array}{l} \hat{\tau}' \in \text{lift}(\hat{\tau}), \\ \forall v. p'_a(v) * [O_1]_r^\perp * \mathcal{I}_t[r, \lambda, v] = \\ \text{havoc}_\lambda(p_a(v)) * \mathbf{t}_r^\lambda(v) \end{array} \right\}$$

$$\text{lift}((\sigma, h, p_h, p_a, \langle v, v' \rangle) \pi \hat{\tau}) \triangleq \{ (\sigma, h, p_h * [O_2(v, v')]_r^\perp, p_a, \langle v, v' \rangle) \pi \hat{\tau}' \mid \hat{\tau}' \in \text{lift}(\hat{\tau}) \}$$

This can be lifted to sets of specification traces, $\mathbb{T} \subseteq \text{STrace}$:

$$\text{lift}(\mathbb{T}) \triangleq \bigcup_{\hat{\tau} \in \mathbb{T}} \text{lift}(\hat{\tau})$$

As a technical tool of our proofs, we use the function $\text{obliv}_\lambda(w)$ which removes the information about states of regions that are open at level λ from w .

Definition E.31 (obliv). Let $\lambda \in \text{Lvl}$, we then define the function on worlds:

$$\text{obliv}_\lambda(h, \rho, \gamma, \chi, \theta, \xi) \triangleq \left\{ (h, \rho', \gamma, \chi, \theta, \xi) \mid \begin{array}{l} \text{closed}_\lambda^{\lambda+1}(\rho) = \{r_1, \dots, r_n\}, \\ \rho(r_i) = (\mathbf{t}_i, _ , _), b_i \in \text{AVal}, \\ \rho' = \rho[r_1 \mapsto (\mathbf{t}_1, \lambda, b_1), \dots, r_1 \mapsto (\mathbf{t}_n, \lambda, b_n)] \end{array} \right\}$$

We extend it to a function on sets of worlds in the obvious way: $\text{obliv}_\lambda(p) \triangleq \bigcup_{w \in p} \text{obliv}_\lambda(w)$.

LEMMA E.32. For arbitrary $p, f \in \text{World}_{\mathcal{A}}^\dagger, v \in \text{AVal}$:

$$\llbracket p * \mathbf{t}_r^\lambda(v) * f \rrbracket_{\lambda+1} \subseteq \llbracket \text{havoc}_\lambda(p) * \mathbf{t}_r^\lambda(v) * f \rrbracket_{\lambda+1}$$

LEMMA E.33. For arbitrary $h, h' \in \text{Heap}, \bar{p}, \bar{p}' \in \text{View}_{\mathcal{A}}, v \in \text{AVal}$ such that

$$(h, h') \vDash_{\lambda; \mathcal{A}} \bar{p} * \bar{p}_a(v) \rightarrow \bar{p}' * \bar{p}_a(v) \quad (130)$$

then $(h, h') \vDash_{\lambda+1; \mathcal{A}} \text{havoc}_\lambda(\bar{p} * [O_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v) \rightarrow \text{havoc}_\lambda(\bar{p}' * [O_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v)$.

PROOF. Given arbitrary $f \in \text{World}_{\mathcal{A}}^\dagger$, take $h \in \llbracket \text{havoc}_\lambda(\bar{p} * [O_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v) * f \rrbracket_{\lambda+1}$ arbitrary. Then, there exists $w_l \in \text{havoc}_\lambda(\bar{p} * [O_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v)$ and $w_f \in f$ such that $h \in \llbracket w_l \odot w_f \rrbracket_{\lambda+1}^{\mathcal{A}}$. Given that $w_l \# w_f$:

$$w_l = (h_l, \rho_l, \gamma_l, \chi_l, \theta_l, \xi_l) \quad (131)$$

$$w_f = (h_f, \rho_f, \gamma_f, \chi_f, \theta_f, \xi_f) \quad (132)$$

$$\forall r \in \text{dom}(\rho_l). h_l \# h_f \wedge \gamma_l(r) \# \gamma_f(r) \wedge \chi_l(r) \# \chi_f(r) \wedge \theta_l(r) \# \theta_f(r) \quad (133)$$

We also know that $\rho_l(r) = (\mathbf{t}, \lambda, v)$. Now, letting $\text{closed}_\lambda^{\lambda+1}(\rho_l) = \{r, r_1, \dots, r_n\}$ and $\rho_l(r_i) = (\mathbf{t}_i, \lambda, a_i)$, from the definition of $\llbracket _ \rrbracket_{\lambda+1}^{\mathcal{A}}$ and \odot , we also know:

$$\forall r \in \text{dom}(\rho_l). \theta_l(r) \supseteq \xi_f(r) \wedge \theta_f(r) \supseteq \xi_l(r) \quad (134)$$

$$h \in \llbracket w_l \odot w_r \odot w_1 \odot \dots \odot w_n \odot w_f \rrbracket_\lambda^{\mathcal{A}} \quad (135)$$

for some $w_r \in \mathcal{I}_t[[r, \lambda, v]]$ and $w_i \in \mathcal{I}_t[[r_i, \lambda, a_i]]$.

Given that $w_l \in \text{havoc}_\lambda(\bar{p} * [\mathbf{O}_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v)$, there exists, $\bar{w}_h \in \bar{p}$, $\bar{w}_o \in [\mathbf{O}_1]_r^\perp$, $\bar{w}_a \in p_a(v)$, \bar{w}_l such that $w_l \in \text{havoc}_\lambda(\bar{w}_l)$ and $\bar{w}_l = \bar{w}_h \odot \bar{w}_o \odot \bar{w}_a$. From the definition of havoc_λ , we know that $\bar{w}_l = (h_l, \bar{\rho}_l, \gamma_l, \chi_l, \theta_l, \xi_l)$, where $\bar{\rho}_l$ and ξ_l are such that $\text{dom}(\bar{\rho}_l) = \text{dom}(\rho_l)$ and

$$\forall r \in \text{dom}(\rho_l) \setminus \text{closed}_\lambda^{\lambda+1}(\rho_l). \rho_l(r) = \bar{\rho}_l(r) \quad (136)$$

$$\forall r \in \text{closed}_\lambda^{\lambda+1}(\rho_l). \text{rty}_{w_l}(r) = \text{rty}_{\bar{w}_l}(r) \wedge \text{lvl}_{\bar{w}_l}(r) = \lambda \quad (137)$$

$$\forall r \in \text{closed}_\lambda^{\lambda+1}(\rho_l). \exists w_I \in \mathcal{I}_{\text{rty}_{w_l}(r)}[[r, \lambda, \text{ast}_{w_l}(r)]], \mathbf{O}, \bar{\xi}_l(r) = \mathbf{O} \bullet \theta_{w_I}(r) \wedge \mathbf{O} \sqsubseteq \bar{\xi}_l(r) \quad (138)$$

Let $\bar{f} = \text{obliv}_\lambda(f)$. From the definition of obliv_λ , and from (136) and (137), we know that $(h_f, \bar{\rho}_l, \gamma_f, \chi_f, \theta_f, \xi_f) \in \bar{f}$.

Then, since all our the region interpretations at level λ are λ -safe, there exists $\bar{w}_i \in \mathcal{I}_t[[r_i, \lambda, a_i]]$, such that $h \in [\bar{w}_h \odot (\bar{w}_o \odot \bar{w}_a \odot \bar{w}_r) \odot \bar{w}_1 \odot \dots \odot \bar{w}_n \odot \bar{w}_f]_{\lambda}^{\mathcal{A}}$. As $\bar{w}_o \odot \bar{w}_a \odot \bar{w}_r \in \bar{p}_a(v)$, $h \in [\bar{p} * \bar{p}_a(v) * \mathcal{I}_t[[r, \lambda, v]] * \bigstar_{i=1}^n \mathcal{I}_t[[r_i, \lambda, a_i]] * \bar{f}]_{\lambda}$. By (130), this implies that $h' \in [\bar{p}' * \bar{p}_a(v) * \mathbf{R}_{\mathcal{A}}^a(\mathcal{I}_{\mathcal{A}, \lambda+1}^\lambda * \bar{f})]_{\lambda}$. As $\mathcal{I}_{\mathcal{A}, \lambda+1}^\lambda \in \text{View}_{\mathcal{A}}$, this is equivalent to $h' \in [\bar{p}' * \bar{p}_a(v) * \mathcal{I}_{\mathcal{A}, \lambda+1}^\lambda * \mathbf{R}_{\mathcal{A}}^a(\bar{f})]_{\lambda}$.

From this we can infer there exists $\bar{w}_h' \in \bar{p}'$, $\bar{w}_o' \in [\mathbf{O}_1]_r^\perp$, $\bar{w}_a' \in p_a(v)$, $\bar{w}_r' \in \mathcal{I}_t[[r, \lambda, v]]$, $\bar{w}_i' \in \mathcal{I}_t[[r_i, \lambda, a_i]]$ and $\bar{w}_f' \in \mathbf{R}_{\mathcal{A}}^a(\bar{f})$ such that:

$$h' \in [\bar{w}_h' \odot \bar{w}_o' \odot \bar{w}_a' \odot \bar{w}_r' \odot \bar{w}_1' \odot \dots \odot \bar{w}_n' \odot \bar{w}_f']_{\lambda}^{\mathcal{A}}$$

There exists some $w'_l \in \text{havoc}_\lambda(\bar{w}_h' \odot \bar{w}_o' \odot \bar{w}_a')$ and $w'_f \in \mathbf{R}_{\mathcal{A}}^a(\bar{f})$, such that, from the definition of reification and the fact that $\bar{w}_r' \in \mathcal{I}_t[[r, \lambda, v]]$, this implies:

$$h' \in [w'_l \odot w'_f]_{\lambda+1}^{\mathcal{A}}$$

where $\rho_{w'_l}(r) = (t, \lambda, v)$ and therefore:

$$h' \in [\text{havoc}_\lambda(\bar{p}' * [\mathbf{O}_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v)]_{\lambda+1}$$

as required. \square

LEMMA E.34. For arbitrary $h, h' \in \text{Heap}$, $\bar{q}, \bar{q}' \in \text{AVal} \times \text{AVal} \rightarrow \text{View}_{\mathcal{A}}, v, v' \in \text{AVal}$ such that

$$(h, h') \vDash_{\lambda; \mathcal{A}} \bar{q}(v, v') \rightarrow \bar{q}'(v, v') \quad (139)$$

then $(h, h') \vDash_{\lambda+1; \mathcal{A}} \text{havoc}_\lambda(\bar{q}(v, v') * [\mathbf{O}_2(v, v')]_r^\perp) \rightarrow \text{havoc}_\lambda(\bar{q}'(v, v') * [\mathbf{O}_2(v, v')]_r^\perp)$.

PROOF. Proof follows similarly to Lemma E.33. \square

LEMMA E.35. For arbitrary $h, h' \in \text{Heap}$, $\bar{q} \in \text{View}_{\mathcal{A}}, \bar{q}' \in \text{AVal} \times \text{AVal} \rightarrow \text{View}_{\mathcal{A}}, v, v', z \in \text{AVal}$ such that

$$(h, h') \vDash_{\lambda; \mathcal{A}} \bar{q} * \bar{p}_a(v) \rightarrow \bar{q}'(v, v') * \bar{q}_a(v, v', z) \quad (140)$$

then

$$(h, h') \vDash_{\lambda+1; \mathcal{A}} \text{havoc}_\lambda(\bar{q} * [\mathbf{O}_1]_r^\perp * p_a(v)) * \mathbf{t}_r^\lambda(v) \rightarrow \text{havoc}_\lambda(\bar{q}'(v, v') * [\mathbf{O}_2(v, v')]_r^\perp * q_a(v, v', z)) * \mathbf{t}_r^\lambda(z)$$

PROOF. Proof follows similarly to Lemma E.33. \square

LEMMA E.36. For arbitrary $h, h' \in \text{Heap}$, $p_e, p'_e \in \text{View}_{\mathcal{A}}, v, v' \in \text{AVal}$ such that

$$(h, h') \vDash_{\lambda+1; \mathcal{A}} p_e * \text{havoc}_\lambda(p_a(v)) * \mathbf{t}_r^\lambda(v) \rightarrow p'_e * \text{havoc}_\lambda(p_a(v')) * \mathbf{t}_r^\lambda(v') \quad (141)$$

then there exists $\bar{p}_e, \bar{p}'_e \in \text{View}_{\mathcal{A}}$ such that:

$$(h, h') \vDash_{\lambda; \mathcal{A}} \bar{p}_e * \bar{p}_a(v) \rightarrow \bar{p}'_e * \bar{p}_a(v')$$

LEMMA E.37. For arbitrary $h, h' \in \text{Heap}$, $p_e, p'_e \in \text{View}_{\mathcal{A}}$ such that

$$(h, h') \vDash_{\lambda+1; \mathcal{A}} p_e \rightarrow p'_e \quad (142)$$

then there exists $\bar{p}_e, \bar{p}'_e \in \text{View}_{\mathcal{A}}$ such that $(h, h') \vDash_{\lambda; \mathcal{A}} \bar{p}_e \rightarrow \bar{p}'_e$.

LEMMA E.38. For arbitrary \mathbb{T} :

$$(\exists \bar{\tau} \in \llbracket \text{lift}(\mathbb{T}) \rrbracket. \text{liveEnv}_{\mathbb{S}'}(\bar{\tau})) \Rightarrow (\exists \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S}}(\bar{\tau}))$$

THEOREM E.39 (SOUNDNESS OF RULE LIFTAG). Assuming

$$\begin{array}{ll} \vdash_{\mathcal{A}} P_h \Rightarrow \text{emp}_{\text{Ob}}^r & \vdash_{\mathcal{A}} Q_h(x, y) \Rightarrow \text{emp}_{\text{Ob}}^r \\ \vdash_{\mathcal{A}} P_a(x) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} & \vdash_{\mathcal{A}} Q_a(x, y, z) \Rightarrow \text{emp}_{\text{Ob}}^{\lambda+1} \\ \mathcal{A} \vDash Q_h(x, y) \text{ } \lambda\text{-safe} & \mathcal{A} \vDash Q_a(x, y, z) \wedge R(x, z) \text{ } \lambda\text{-safe} \\ r \in \text{dom}(\mathcal{A}) \Rightarrow R = \text{id} & \end{array}$$

for $R \subseteq \text{AVal} \times \text{AVal}$ such that

$$\{ ((x, \text{O}_1), (z, \text{O}_2(x, y))) \mid x \in X \wedge R(x, z) \wedge y \in Y(x) \} \subseteq \mathcal{T}_t(\mathbb{G})$$

then, given arbitrary $\Phi \in \text{FSpec}$ such that

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S} \quad (143)$$

then

$$\vDash_{\Phi} \mathbb{C} : \mathbb{S}'$$

PROOF. To reach the goal, it suffices to show that $\llbracket \mathbb{S} \rrbracket \subseteq \llbracket \mathbb{S}' \rrbracket$. Taking $(\sigma_0, h_0)\tau \in \llbracket \mathbb{S} \rrbracket$ arbitrary, let

$$\begin{aligned} \bar{p}_h &= \mathcal{W}[[P_h]_{\mathcal{A}}^{\sigma_0}] \\ p_h &= \mathcal{W}[[P_h * \lfloor \text{O}_1 \rfloor_r^{\perp}]_{\mathcal{A}}^{\sigma_0}] \end{aligned}$$

Then, assume that for arbitrary $v_0 \in X$, $h \in \llbracket p_h * p_a(v_0) * \text{True} \rrbracket_{\lambda+1}$ holds. Then, clearly $h \in \llbracket \bar{p}_h * \bar{p}_a(v_0) * \text{True} \rrbracket_{\lambda}$, and therefore, from (143), for some \mathbb{T} :

$$(\sigma_0, h_0) \tau \vDash_{\mathbb{S}} \bar{p}_h, \bar{p}_a, v_0 : \mathbb{T} \quad (144)$$

$$\forall \bar{\tau} \in \llbracket \mathbb{T} \rrbracket_{\lambda; \mathcal{A}}. \text{liveEnv}_{\mathbb{S}}(\bar{\tau}) \Rightarrow \text{lterm}((\sigma_0, h_0) \tau) \quad (145)$$

From (144), by coinduction over the structure of the trace safety judgement, using our assumptions and Lemmas E.32 to E.37, the following holds:

$$(\sigma_0, h_0) \tau \vDash_{\mathbb{S}'} p_h, p_a, v_0 : \text{lift}(\mathbb{T})$$

Finally, taking $\bar{\tau} \in \text{lift}(\mathbb{T})$ arbitrary such that $\text{liveEnv}_{\mathbb{S}'}(\bar{\tau})$ holds, from Lemma E.38, $\exists \bar{\tau} \in \llbracket \mathbb{T} \rrbracket. \text{liveEnv}_{\mathbb{S}}(\bar{\tau})$. By (145), the following holds, as required:

$$\forall \bar{\tau} \in \llbracket \text{lift}(\mathbb{T}) \rrbracket_{\lambda; \mathcal{A}}. \text{liveEnv}_{\mathbb{S}'}(\bar{\tau}) \Rightarrow \text{lterm}((\sigma_0, h_0) \tau)$$

□