

Communicating State Transition Systems for Fine-Grained Concurrent Resources

Extended Version

Aleksandar Nanevski¹, Ruy Ley-Wild², Ilya Sergey¹, and Germán Andrés Delbianco¹

¹ IMDEA Software Institute, Spain

{aleks.nanevski, ilya.sergey, german.delbianco}@imdea.org

² LogicBlox, USA

ruy.leywild@logicblox.com

Abstract. We present a novel model of concurrent computations with shared memory and provide a simple, yet powerful, logical framework for uniform Hoare-style reasoning about partial correctness of coarse- and fine-grained concurrent programs. The key idea is to specify arbitrary resource protocols as communicating *state transition systems* (STS) that describe valid states of a resource and the transitions the resource is allowed to make, including transfer of heap ownership. We demonstrate how reasoning in terms of communicating STS makes it easy to crystallize behavioral invariants of a resource. We also provide *entanglement* operators to build large systems from an arbitrary number of STS components, by interconnecting their lines of communication. Furthermore, we show how the classical rules from the Concurrent Separation Logic (CSL), such as scoped resource allocation, can be generalized to fine-grained resource management. This allows us to give specifications as powerful as Rely-Guarantee, in a concise, scoped way, and yet regain the compositionality of CSL-style resource management. We proved the soundness of our logic with respect to the denotational semantics of action trees (variation on Brookes’ action traces). We formalized the logic as a shallow embedding in Coq and implemented a number of examples, including a construction of coarse-grained CSL resources as a modular composition of various logical and semantic components.

1 Introduction

There are two main styles of program logics for shared-memory concurrency, customarily divided according to the supported kind of granularity of program interference. Logics for coarse-grained concurrency such as Concurrent Separation Logic (CSL) [12, 14] restrict the interference to critical sections only, but generally lead to more modular specifications and simpler proofs of program correctness. Logics for fine-grained concurrency, such as Rely-Guarantee (RG) [8] admit arbitrary interference, but their specifications have traditionally been more monolithic, as we shall illustrate. In this paper, we identify the essential ingredients required for compositional specification of concurrent programs, and combine them in a novel way to reconcile the two approaches. We present a semantic model and a logic that enables specification and reasoning about

fine-grained programs, but in the style of CSL. To describe our contribution more precisely, we first compare the relevant properties of CSL and RG.

CSL employs *shared resources* and associated *resource invariants* [13], to abstract the interference between threads. A resource r is a chunk of shared state, and a resource invariant I is a predicate over states, which holds of r whenever all threads are outside the critical section. By mutual exclusion, when a thread enters a critical section for r , it acquires ownership and hence exclusive access to r 's state. The thread may mutate the shared state and violate the invariant I , but it must restore I before releasing r and leaving the critical section, as given by the following CSL rule [2].

$$\frac{\Gamma \vdash \{p * I\} c \{q * I\}}{\Gamma, r : I \vdash \{p\} \text{ with } r \text{ do } c \{q\}} \text{CRITSECCSL}$$

Γ is a context of currently existing resources. The rule for parallel composition assumes that forked threads don't share any state beyond that of the resources in Γ , and may divide the private state of the parent thread disjointly among the children.

$$\frac{\Gamma \vdash \{p_1\} c_1 \{q_1\} \quad \Gamma \vdash \{p_2\} c_2 \{q_2\}}{\Gamma \vdash \{p_1 * p_2\} c_1 \parallel c_2 \{q_1 * q_2\}} \text{PARCSL}$$

A private heap of a thread may be promoted into a freshly allocated shared resource in a scoped manner by the following rule.

$$\frac{\Gamma, r : I \vdash \{p\} c \{q\}}{\Gamma \vdash \{p * I\} \text{ resource } r \text{ in } c \{q * I\}} \text{RESOURCECSL}$$

One may see from these rules that resources are abstractions that promote modularity. In particular, one may verify a thread wrt. the smallest resource context required. By context weakening, the introduction of new resources will not invalidate the existing verification. Thread-local resources can be hidden from the environment by the RESOURCECSL rule.

In RG, the interaction between threads is directly specified by the rule for parallel composition.³

$$\frac{R \vee G_2, G_1 \vdash \{p\} c_1 \{q_1\} \quad R \vee G_1, G_2 \vdash \{p\} c_2 \{q_2\}}{R, G_1 \vee G_2 \vdash \{p\} c_1 \parallel c_2 \{q_1 \wedge q_2\}} \text{PARRG}$$

The rely transition R and guarantee transitions G_1 and G_2 are relations on states. A rely specifies the thread's expectations of state transitions made by its environment. A guarantee specifies the state transitions made by the thread itself. The disjunctive combinations of R and G 's in the rule captures the idea we call *forking shuffle*, whereby upon forking, the thread c_1 becomes part of the environment for c_2 and vice-versa.

RG is more expressive than CSL because transitions can encode arbitrary protocols on shared state, whereas CSL is specialized to a fixed mutual exclusion protocol on critical sections. But, CSL is more compositional in manipulating resources. Where a CSL resource invariant specifies the behavior of an individual chunk of shared state, the transitions in RG treat the whole state as monolithically shared. Feng's work on Local Rely Guarantee (LRG) [5] has made first steps in improving RG in this respect.

³ In the presence of heaps, the rule is more complicated [6, 18], but we elide the issue here.

1.1 Contributions

We propose that a logic for fine-grained concurrency can be based on a notion of a *fine-grained resource*. Fine-grained resources serve as building blocks for program specification, and generalize CSL-style coarse-grained resource management. A fine-grained resource is specified by a resource invariant, as in CSL, but it also adds transitions in the form of relations between resource states. Thus, it is best viewed as a *state transition system* (STS), where the resource invariant specifies the state space. We identify a number of properties that an STS has to satisfy to specify a fine-grained resource, and refer to such STSs as *concurroids*. We refer to our generalization of CSL as Fine-grained CSL (FCSL).

There are two main ideas that we build on in FCSL, and which separate FCSL from LRG and other recent related work [4, 15, 17] (see Section 6 for details): (a) *subjectivity* and (b) *communication*. Subjectivity [10] means that each state of a concurroid STS describes not only the shared resource, but also *two* abstractions of it that represent the views of the state by the thread, and by its environment, respectively. Subjectivity will enable us to capture the idea of forking shuffle by a rule for parallel composition akin to PARCSL (but with a somewhat generalized notion of separating conjunction $*$) [10]), rather than in the monolithic style of PARRG.

To compositionally build large systems out of a number of smaller ones, we make concurroids communicate. In addition to standard for STS *internal* transitions between states, concurroids contain *external* transitions. These may be thought of as “wires” whose one end is connected to a state in the STS, but whose other end is dangling, representing either an “input” into or an “output” out of the STS. Concurroids can be *entangled*, *i.e.*, composed by interconnecting their dangling wires of opposite polarity, where the interconnections serve to transfer heap ownership between concurroids. Communication and entanglement endow FCSL with the compositionality of CSL. For example, entanglement generalizes the notion of adding a resource to the context Γ in RESOURCECSL. We also rely on entanglement to formulate a rule generalizing the scoped resource allocation of RESOURCECSL. More precisely, our contributions are:

- We identify STSs with subjectively-shaped states (concurroids) and a number of algebraic properties, as a natural model for scalable concurrency verification. We show how communication enables composing larger STSs out of smaller ones.
- We present FCSL—a simple and expressive logic for fine-grained resources that combines expressivity of RG with the compositional resource management of CSL.
- We illustrate FCSL by showing how to implement a coarse-grained resource of CSL by a fine-grained resource of FCSL in which an explicit spin lock protects the resource’s state. We also implemented examples such as ticketed locks, that go beyond coarse-grained CSL resources, and present them in the appendix.
- We implemented FCSL [11] as a shallow embedding within the type theory of the Calculus of Inductive Constructions (*i.e.*, Coq [1, 16]). Thus, FCSL naturally reconciles with features such as higher-order functions, abstract predicates, modules and functors. We formally instantiated the whole stack of abstractions: the semantic model is formalized in Coq, FCSL is built on top of the semantic model, CSL is built on top of FCSL, and then verified programs are built on top of CSL.

2 An Overview of Fine-Grained Resources

There are three different aspects along which fine-grained resources can be composed: space (*i.e.*, states), ownership, and time (*i.e.*, transitions). In this section, we describe how to represent these aspects in the assertion logic of FCSL.

Space The heap belonging to a fine-grained resource,⁴ is explicitly identified by a *resource label*. We use assertions in the “points-to” style of separation logic, to name resources and identify their respective heaps. For example, the assertion

$$\ell_1 \overset{j}{\mapsto} h_1 * \ell_2 \overset{j}{\mapsto} h_2$$

describes a state in which the heaps h_1 and h_2 are associated with the resources labeled ℓ_1 and ℓ_2 , respectively. The connective $*$ ensures that ℓ_1 and ℓ_2 are distinct labels, and that h_1 and h_2 are disjoint heaps. The superscript j indicates that the heaps are *joint* (shared), *i.e.*, can be accessed by any thread, even though they are owned by the resources ℓ_1 and ℓ_2 , respectively.

The heaps h_1 and h_2 are not described by means of points-to assertions, but are built using operators for singleton heaps $x \rightarrow v$ and disjoint union \cup . For example, the heap of the resource `lock`, which explicitly encodes a coarse-grained resource with the resource invariant I [12] may be described by the assertion

$$\text{lock} \overset{j}{\mapsto} ((lk \rightarrow b) \cup h) \wedge \text{if } b \text{ then } h = \text{empty else } I h. \quad (1)$$

The assertion exposes the fact that the heap owned by `lock` contains a boolean pointer lk encoding a lock that protects the heap h . The conditional conjunct is a *pure* (*i.e.*, label-free) assertion, which describes an aspect of the ownership transfer protocol of CSL. When the lock is not taken (*i.e.* $b = \text{false}$), the heap h satisfies the resource invariant. When the lock is taken, the heap is transferred to the private ownership of the locking thread, so h equals the empty heap, but lk remains in the ownership of `lock`.

Ownership Data in FCSL may be owned by a resource, as illustrated above, or by individual threads. The thread-owned data, however, is also associated with a resource, which it refines with *thread-relative* information. For example, the resource `lock` owns a pointer lk which operationally implements a lock. However, just knowing that the lock is taken or not is not enough for reasoning purposes; we need to know which thread has taken it, if any. Thus, we associate with each thread an extra bit of lock-related information, `Own` or `Ownπ`, which will identify the lock-owning thread as follows.

Following the idea of *subjectivity* [10], FCSL assertions are interpreted in a thread-relative way. We use *self* to name the interpreting thread, and *other* to name the combination of all other threads running concurrently with *self* (*i.e.*, the environment of *self*). We use two different assertions to describe thread-relative views: $\ell \overset{s}{\mapsto} v$ and $\ell \overset{o}{\mapsto} v$. The first is true in the *self* thread, if *self*’s view of the resource ℓ is v . The second is true in

⁴ Or just resource for short. Later on, we explicitly identify CSL resources as *coarse-grained*.

the *self* thread, if *other*'s view of the resource ℓ is v . In this sense, the $\ell \mapsto^j v$ describes the resource's view of the data. In the case of *lock*, the thread that acquired the lock will validate the assertion:

$$\text{lock} \mapsto^s \text{Own} \wedge \text{lock} \mapsto^j (lk \rightarrow \text{true}),$$

while the symmetric assertion holds in all other threads at the same moment of time:

$$\text{lock} \mapsto^j (lk \rightarrow \text{true}) \wedge \text{lock} \mapsto^o \text{Own}.$$

In general, the values of the *self* and *other* views for *any resource* are elements of some resource-specific *partial commutative monoid* (PCM) [10]. A PCM is a set with a commutative and associative operation \bullet with a unit element. \bullet combines the *self* and *other* views into a view of the parallel composition of *self* and *other* threads. The \bullet operation is commutative and associative because parallel composition of threads is commutative and associative, and the unit element models the view of the idle thread. Partiality models impossible thread combinations. For example, the elements of $O = \{\text{Own}, \text{Own}\}$ represent thread-relative views of the lock lk . O forms a PCM under the operation defined as $x \bullet \text{Own} = \text{Own} \bullet x = x$, with $\text{Own} \bullet \text{Own}$ undefined. The unit element is Own , and the undefinedness of the last combination captures that two threads can't simultaneously own the lock. Notice that heaps form a PCM under disjoint union, with the empty heap as unit. Thus, they too obey the discipline required of the general *self* and *other* components.

Anticipating lock-related examples in Section 3, we combine thread-relative views of the lock with thread-relative views of the lock-protected heap h . We parametrize the resource lock by a PCM U , which the user may choose depending on the application. Then we use assertions over *pairs*, such as $\text{lock} \mapsto^s (m_S, a_S)$ and $\text{lock} \mapsto^o (m_O, a_O)$, to express that $m_S, m_O \in O$ are views of the lock lk , and $a_S, a_O \in U$ are views of the heap h . The following assertion illustrates how the different FCSL primitives combine. It generalizes (1) and defines the valid states of the resource lock.

$$\begin{aligned} & \text{lock} \mapsto^s (m_S, a_S) \wedge \text{lock} \mapsto^o (m_O, a_O) \wedge \text{lock} \mapsto^j ((lk \rightarrow b) \cup h) \wedge \\ & \text{if } b \text{ then } h = \text{empty} \wedge m_S \bullet m_O = \text{Own} \text{ else } I (a_S \bullet a_O) h \wedge m_S \bullet m_O = \text{Own} \end{aligned} \quad (2)$$

The assertion states that if the lock is taken ($b = \text{true}$) then the heap h is given away, otherwise it satisfies the resource invariant I . In either case, the thread-relative views m_S, m_O, a_S and a_O are consistent with the resource's views of lk and h . Indeed, notice how m_S, m_O and a_S, a_O are first \bullet -joined (by the \bullet -operations of O and U , respectively) and then related to b and h ; the former implicitly by the conditional, the latter explicitly, by the resource invariant I , which is now parametrized by $a_S \bullet a_O$.

Private heaps In addition to a private view of a resource, a thread may own a private heap as well. We describe such thread-private heaps by means of the same thread-relative assertions, but with a different resource label. We consider a dedicated resource for *private heaps*, with a dedicated label *priv*. Then we can write, say, $\text{priv} \mapsto^s x \rightarrow 4$ to describe a heap consisting of a pointer x private to the *self* thread. By definition, $\text{priv} \mapsto^j \text{empty}$, *i.e.*, the *joint* heap of the *priv* resource is always empty.

Time Fine-grained reasoning requires characterization of the possible changes the threads can make to the state. We encode such a characterization as relations between states of possibly *multiple resources* (*i.e.*, using multiple labels). For example, coarse-grained resources require that upon successful acquisition, the resource’s heap is transferred into the private ownership of the acquiring thread. In our fine-grained encoding, the transition can be represented as follows:

$$\begin{aligned} \text{priv} \xrightarrow{s} h_S \quad * (\text{lock} \xrightarrow{s} (\text{Own}, a_S) \wedge \text{lock} \xrightarrow{j} ((lk \rightarrow \text{false}) \cup h)) \rightsquigarrow \\ \text{priv} \xrightarrow{s} (h_S \cup h) * (\text{lock} \xrightarrow{s} (\text{Own}, a_S) \wedge \text{lock} \xrightarrow{j} (lk \rightarrow \text{true})) \end{aligned} \quad (3)$$

This transition preserves heap *footprints*, in the sense that the domain of the combined heaps in the source of the transition equals the domain in the target of the transition. We refer to such transitions as *internal*. Footprint preservation is an essential property, as it facilitates composing and framing transitions. In particular, adding additional labels and heaps with non-overlapping footprint to a source of an internal transition is guaranteed to produce non-overlapping footprints in the target of the transition as well.

We also consider *external* transitions that *can* acquire and release heaps. We use external transitions to build internal ones. For example, the above internal transition over *priv* and *lock* resources can be obtained as an interconnection (to be defined in Section 4) of two external transitions, each operating on an individual label.

$$\begin{aligned} \text{priv} \xrightarrow{s} h_S \xrightarrow{+h} \text{priv} \xrightarrow{s} (h_S \cup h) \\ \text{lock} \xrightarrow{s} (\text{Own}, a_S) \wedge \text{lock} \xrightarrow{j} ((lk \rightarrow \text{false}) \cup h) \xrightarrow{-h} \text{lock} \xrightarrow{s} (\text{Own}, a_S) \wedge \text{lock} \xrightarrow{j} lk \rightarrow \text{true} \end{aligned} \quad (4)$$

The transition over *priv* takes a heap h as an input and attaches it to the *self* heap. The transition over *lock* gives the heap h as an output. When interconnected, the two transitions exchange the ownership of h between the *lock* and *priv*, producing (3).

A *concurroid* is an STS that formally represents a collection of resources. Each state of the STS contains a number of components, identified by the labels naming the individual resources. Each concurroid contains one internal transition, and an arbitrary number of external ones. The internal transition describes how threads specified by the concurroid may change their state in a single step. The external transitions are the “dangling wires”, which provide means for composing different concurroids by *entangling* them, *i.e.*, interconnecting (some or all of) their dually polarized external transitions, to obtain a larger concurroid.

For example, if P is the concurroid for private heaps (containing a single label *priv*), and $L_{\{\text{lock}, lk, I\}}$ is the concurroid for a lock (with a single label *lock*, lock pointer lk and protected heap described by the coarse-grained resource invariant I), we could construct the entangled concurroid $CSL_{\{\text{lock}, lk, I\}} = P \times L_{\{\text{lock}, lk, I\}}$ that captures the heap ownership-exchange protocol (3) of CSL for programs with *one coarse-grained resource*.⁵ The entanglement can be iterated, to obtain an STS for *two coarse-grained resources* $CSL_{\{\text{lock}, lk, I, \text{lock}', lk', I'\}} = CSL_{\{\text{lock}, lk, I\}} \times L_{\{\text{lock}', lk', I'\}}$, and so on. In this way, concurroids generalize the notion of resource context from the RESOURCECSL rule, with entanglement modeling the addition of new resources to the context.

⁵ The formal definition of the \times is postponed until Section 4.

Fig. 1 Semantics of selected FCSL assertions.

$w \models \top$	iff always
$w \models \ell \xrightarrow{s} v$	iff valid w , and $w. s = \ell \rightarrow v$
$w \models \ell \xrightarrow{j} h$	iff valid w , and $w. j = \ell \rightarrow h$
$w \models \ell \xrightarrow{o} v$	iff valid w , and $w. o = \ell \rightarrow v$
$w \models p \wedge q$	iff $w \models p$ and $w \models q$
$w \models p * q$	iff valid w , and $w = w_1 \cup w_2$, and $w_1 \models p$ and $w_2 \models q$
$w \models p \multimap q$	iff for every w_1 , valid $w \cup w_1$ and $w_1 \models p$ implies $w \cup w_1 \models q$
$w \models p \otimes q$	iff valid w , and $w. s = s_1 \cup s_2$, and $[s_1 \mid w. j \mid s_2 \circ w. o] \models p$ and $[s_2 \mid w. j \mid s_1 \circ w. o] \models q$
$w \models \text{this } w'$	if $w = w'$
$w \downarrow h$	iff for every valid w , $w \models p$ implies $\lfloor w \rfloor = h$

3 Reasoning with Concurroids

Auxiliary definitions A PCM-map is a finite map from labels (isomorphic to nat) to $\Sigma_{\mathbb{U}; \text{pcm}} \mathbb{U}$. It associates each label with a pair of a PCM \mathbb{U} and a value $v \in \mathbb{U}$. A heap-map is a finite map from labels to heaps. If m_1, m_2 are PCM-maps, then $m_1 \circ m_2$ is defined as $\text{empty} \circ \text{empty} = \text{empty}$, and $((\ell \rightarrow_{\mathbb{U}} v_1) \cup m'_1) \circ ((\ell \rightarrow_{\mathbb{U}} v_2) \cup m'_2) = (\ell \rightarrow_{\mathbb{U}} v_1 \bullet v_2) \cup (m'_1 \circ m'_2)$, and undefined otherwise. By overloading the notation, we define state w as a triple $[s \mid j \mid o]$, where s, o are PCM-maps, and j is a heap-map. We abbreviate $[\ell \rightarrow v_s \mid \ell \rightarrow v_j \mid \ell \rightarrow v_o]$ with $\ell \rightarrow [v_s \mid v_j \mid v_o]$. w is valid if $w. s, w. j, w. o$ have the same domain as PCM-maps, $w. s \circ w. o$ is defined, and the heaps in $w. s, w. j$ and $w. o$ are disjoint (if $w. s$ and $w. o$ contain heaps in their codomain). State flattening $\lfloor w \rfloor$ is the disjoint union of all such heaps. $w_1 \cup w_2$ is the pairwise disjoint union of component maps of w_1 and w_2 . The semantics of the main FCSL assertions is provided in Figure 1. The subjective assertions (e.g., $w \models \ell \xrightarrow{s} v$) constrain the value of one state component, assuming others to be existentially quantified over.

FCSL specifications take the form of Hoare 4-tuple $\{p\} c \{q\} @ U$ expressing that the thread c has a precondition p , postcondition q , in a state space and under transitions defined by the concurroid U , which in FCSL takes the role of a resource context from CSL. We next present the characteristic inference rules of FCSL.

Parallel composition The rule for parallel composition in FCSL is similar to PARCSL, with Γ replaced by a concurroid U , which we will define formally in Section 4.

$$\frac{\{p_1\} c_1 \{q_1\} @ U \quad \{p_2\} c_2 \{q_2\} @ U}{\{p_1 \otimes p_2\} c_1 \parallel c_2 \{q_1 \otimes q_2\} @ U} \text{PAR}$$

The PAR rule uses *subjective separating conjunction* \otimes (see [10] and Figure 1) to split the state of $c_1 \parallel c_2$ into two. The split states contain the same labels, and equal *joint* portions, but the *self* and *other* portions are recombined to match the thread-relative views of c_1 and c_2 . When the parent thread forks the children c_1 and c_2 , the PCM values in the parent's *self* components are split between the children (similarly $*$ splits heaps in CSL), while the children's *other* component are implicitly induced to preserve overall

•-total (i.e., c_1 's *other* view includes c_2 's *self* view, and vice versa). For example, in the case of one label ℓ , we have

$$\ell \mapsto^s a \bullet b \wedge \ell \mapsto^o c \implies (\ell \mapsto^s a \wedge \ell \mapsto^o c \bullet b) \otimes (\ell \mapsto^s b \wedge \ell \mapsto^o c \bullet a).$$

The implication encodes the idea of a forking shuffle from RG, but via states, rather than transitions as in RG. It allows us to use the *same* concurrroid U to specify the transitions of both c_1 and c_2 in PAR, much like PARCSL uses the same context Γ . Essentially, we rely on the recombination of views to select the transitions of U available to each of c_1 and c_2 , instead of providing distinct transitions for c_1 and c_2 as in PARRG.

We commonly encounter cases where the *other* views are existentially abstracted, hence the conjuncts $\ell \mapsto^o -$ are omitted. In those cases, we have the simplified bi-implication:

$$\ell \mapsto^s a \bullet b \iff \ell \mapsto^s a \otimes \ell \mapsto^s b \quad (5)$$

The implications generalize to \otimes -separated assertions with more than one distinct label.

We illustrate PAR and \otimes with the example of concurrent incrementation [10, 13] in a setting of a concurrroid $CSL_{\text{lock}, lk, I}$ (i.e., private state and one lock). The lock lk protects a shared integer pointer x , that is, the resource invariant is $I (a : \text{nat}) (h : \text{heap}) \hat{=} h = x \rightarrow a$. For the nat argument, we chose the PCM structure under addition; thus, an assertion $\text{lock} \mapsto^s (-, a_S)$ expresses that the *self* thread has added a_S to x , and dually for $\text{lock} \mapsto^o (-, a_O)$. Therefore, whenever the lock is not taken, x stores the sum $a_S + a_O$. This follows from interpreting \bullet with $+$ in the lock state invariant (2).

Procedure $\text{incr}(n)$ acquires the lock to ensure exclusive access to x , increments x by n , and releases the lock. In FCSL, it has the following specification:

$$\left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0) \right\} \text{incr}(n) \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, n) \right\} @_{CSL_{\text{lock}, lk, I}}$$

The specification states that incr runs in an empty private heap (hence by framing, in any larger heap), the lock is not owned by the calling thread initially, and will not be owned in the end. The addition of calling thread to x increases from 0 to n (hence by framing, from m to $m + n$). We now prove that $\text{incr}(i) \parallel \text{incr}(j)$ increments x by $i + j$.

$$\begin{aligned} & \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0) \right\} \\ & \left\{ \text{priv} \mapsto^s \text{empty} \cup \text{empty} * \text{lock} \mapsto^s (\text{Own} \bullet \text{Own}, 0 + 0) \right\} \\ & \left\{ (\text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0)) \otimes (\text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0)) \right\} \\ & \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0) \right\} \parallel \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, 0) \right\} \\ & \quad \text{incr}(i) \qquad \qquad \qquad \text{incr}(j) \\ & \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, i) \right\} \parallel \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, j) \right\} \\ & \left\{ (\text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, i)) \otimes (\text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, j)) \right\} \\ & \left\{ \text{priv} \mapsto^s \text{empty} * \text{lock} \mapsto^s (\text{Own}, i + j) \right\} \end{aligned}$$

The proof uses the bi-implication (5) to move between \otimes -separated assertions and \bullet -joined views. The proof is compositional in the sense that the same verification of incr is used as a black box in both parallel threads, with the subproofs merely instantiating the parameter n with i and j respectively.

Injection The PAR rule requires c_1 and c_2 to share the same concurroid U , which describes the totality of their resources. If the threads use different concurroids, they first must be brought into a common entanglement, via the rule INJECT.

$$\frac{\{p\} c \{q\}@U \quad r \text{ stable under } V}{\{p * r\} \text{ inject } c \{q * r\}@U \bowtie V} \text{ INJECT}$$

If c is verified wrt. concurroid U , it can be *injected* (i.e. coerced) into a larger concurroid $U \bowtie V$. In programs, we use the explicit coercion `inject` to describe the change of “type” from U to $U \bowtie V$. Reading the rule bottom-up, it says we can ignore V , as V ’s transitions and c operate on disjointly-labeled state. V may change U ’s state by communication, but the change is bounded by U ’s external transitions. Thus, we are justified in verifying c against U alone. In this sense, INJECT may be seen as generalizing the rule for resource context weakening of CSL.

The connective $*$ splits the state according to labels of U and V ; p and q describe the part labeled by U , and r describes the part labeled by V . Since r describes both the prestate and poststate, it has to be *stable* (see Appendix B) under V ; that is, determine a subset of V ’s states that remains fixed under transitions the *other* thread takes over the labels from V .

We illustrate INJECT and stability by verifying `incr`. To set the stage, we need atomic commands for reading from and writing to a pointer x . These have the following obvious specification relative to the concurroid P for private state:

$$\left\{ \begin{array}{l} \text{priv} \xrightarrow{s} x \rightarrow v \\ \text{priv} \xrightarrow{s} x \rightarrow - \end{array} \right\} \text{ read } x \quad \left\{ \begin{array}{l} \text{priv} \xrightarrow{s} x \rightarrow v \wedge \text{res} = v \\ \text{priv} \xrightarrow{s} x \rightarrow v \end{array} \right\} @P$$

The commands for acquiring and releasing lock exchange ownership of the protected pointer x . Thus, they have specifications relative to the concurroid $CSL_{\text{lock},lk,I} = P \bowtie L_{\text{lock},lk,I}$, which we have already used before.

$$\begin{array}{c} \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{lock} \xrightarrow{s} (\text{Own}, 0) \right\} \\ \text{acquire} \\ \left\{ \exists a_0. \text{priv} \xrightarrow{s} x \rightarrow a_0 * (\text{lock} \xrightarrow{s} (\text{Own}, 0) \wedge \text{lock} \xrightarrow{o} (-, a_0)) \right\} @CSL_{\text{lock},lk,I} \\ \left\{ \text{priv} \xrightarrow{s} x \rightarrow a_S + a_0 * (\text{lock} \xrightarrow{s} (\text{Own}, 0) \wedge \text{lock} \xrightarrow{o} (-, a_0)) \right\} \\ \text{release} \\ \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{lock} \xrightarrow{s} (\text{Own}, a_S) \right\} @CSL_{\text{lock},lk,I} \end{array}$$

`acquire` assumes that `lock` is not taken, and that the *self* thread so far has added 0 to x . Thus, the overall contents of x is $0 + a_0 = a_0$, where a_0 is the addition of the *other* threads. Note that `acquire` does not have to be atomic:⁶ as implemented, it just spins on `lk`, and after acquisition, x is transferred into the private heap of *self*. a_0 must be existentially quantified, because *other*’s may add to x while `acquire` is spinning.

⁶ The implementation of `acquire` and `release` relies on atomic actions (Section 5), specific for a particular concurroid, e.g. $CSL_{\text{lock},lk,I}$.

`release` assumes that `lock` is taken by `self`, and that prior to taking `lock`, `self` and `other` have added 0 and a_O to x , respectively. After acquiring x , `self` has mutated it, so that its contents is $a_S + a_O$. After releasing, x is moved from the private heap to the *joint* portion of `lock`. The postcondition does not mention x , as once in *joint*, x 's contents becomes unstable. Indeed, `other` may acquire the lock and change x after `release` terminates. However, `other` can't change the `self` view of x , which is now set to a_S .

The following proof outline presents the implementation and verification of `incr(n)`.

$$\begin{array}{l}
\{ \text{priv} \stackrel{s}{\mapsto} \text{empty} * \text{lock} \stackrel{s}{\mapsto} (\text{Own}, 0) \} \\
\text{acquire;} \\
\{ \exists a_O. \text{priv} \stackrel{s}{\mapsto} x \rightarrow a_O * (\text{lock} \stackrel{s}{\mapsto} (\text{Own}, 0) \wedge \text{lock} \stackrel{o}{\mapsto} (-, a_O)) \} \\
\text{res} \leftarrow \text{inject}(\text{read } x); \\
\{ \exists a_O. \text{priv} \stackrel{s}{\mapsto} x \rightarrow a_O \wedge \text{res} = a_O * (\text{lock} \stackrel{s}{\mapsto} (\text{Own}, 0) \wedge \text{lock} \stackrel{o}{\mapsto} (-, a_O)) \} \\
\text{inject}(\text{write } x(\text{res} + n)); \\
\{ \exists a_O. \text{priv} \stackrel{s}{\mapsto} x \rightarrow n + a_O * (\text{lock} \stackrel{s}{\mapsto} (\text{Own}, 0) \wedge \text{lock} \stackrel{o}{\mapsto} (-, a_O)) \} \\
\text{release} \\
\{ \text{priv} \stackrel{s}{\mapsto} \text{empty} * \text{lock} \stackrel{s}{\mapsto} (\text{Own}, n) \}
\end{array}$$

INJECT is used twice, to coerce `read` and `write` from the concurroid P to $CSL_{\text{lock},lk,l}$. These commands manipulate the contents of `priv`, but retain the framing predicate $\text{lock} \stackrel{s}{\mapsto} (\text{Own}, 0) \wedge \text{lock} \stackrel{o}{\mapsto} (-, a_O)$. This predicate is stable wrt. $L_{\text{lock},lk,l}$. Intuitively, because `self` owns `lock`, `other` can't acquire x and add to it. Thus, no matter what `other` does, a_O and the framing predicate remain invariant.

To simplify the proof, we have not emphasized the invariance of a_O between calls to `acquire` and `release`, even though it is the case (we could do it using the rule EXIST from Figure 2). However, this invariance is what allowed us to calculate the contribution of `self` to x as n (i.e., final contents of x minus a_O). Without tracking a_O , we would not know how much of the final contents of x is attributable to `self`, and how much to `other`.

Hiding refers to the ability to construct a concurroid V from the thread-private heap, in a scope of a thread c . The children forked by c can interfere on V 's state, respecting V 's transitions, but V is hidden from the environment of c . To the environment, V 's state changes look like changes of the private heap of c . In this sense, hiding generalizes the RESOURCECSL rule to fine-grained resources.

$$\frac{\{ \text{priv} \stackrel{s}{\mapsto} h * p \} c \{ \text{priv} \stackrel{s}{\mapsto} h' * q \} @ (P \times U) \times V \quad (\text{omitted side condition on } U \text{ and } V)}{\{ \Psi g h * (\Phi(g) \multimap p) \} \text{hide}_{\phi,g} c \{ \exists g'. \Psi g' h' * (\Phi(g') \multimap q) \} @ P \times U} \text{HIDE}$$

where $\Psi g h = \exists k: \text{heap}. \text{priv} \stackrel{s}{\mapsto} h \cup k \wedge \Phi(g) \downarrow k$

Since installing V consumes a chunk of private heap, the rule requires the overall concurroid to support private heaps, i.e., to be an entanglement $P \times U$, where P is the concurroid for private heaps, and U is arbitrary (it is also possible to generalize the rule so as to be not tied to the specific concurroid P , see [11]). The omitted side condition on U and V is essential for the existence of entanglement and will be explained in Section 5. When U is of no interest, we set it to the empty concurroid E (Section 4), for which $P \times E = P$.

In programs, we use the explicit coercion $\text{hide}_{\Phi, g}$ to indicate the change of type from $(P \times U) \times V$ to $P \times U$. The annotation $\Phi(g)$ corresponds to a set of *concrete states* of a concurroid V to be created. Its parameter g is a meaningful abstraction of such a set (e.g., (m_S, a_S) for the $L_{\{\text{lock}, lk, l\}}$ concurroid) and can be thought of as an “abstract state”. In the rule HIDE, g is the initial abstract state, i.e., upon creation, the state of V satisfies $\Phi(g)$. In the premise of the HIDE rule, the predicates $\text{priv} \mapsto -$ describe the behavior of c on the private heaps, while p and q describe the state of the labels belonging to U and V . In the conclusion, $\Psi \ g \ h$ and $\Psi \ g' \ h'$ map the abstract states g and g' into private heaps h and h' . This follows from the definition of Ψ , in which $\Phi(g) \downarrow k$ indicates that states satisfying $\Phi(g)$ *erase* to the private heap k (see Figure 1). Thus, changes that c imposes on abstract states, appear as changes to private heaps for $\text{hide}_{\Phi, g} \ c$.

In the conclusion, the assertion $\Phi(g) \multimap p$ states that attaching any state satisfying $\Phi(g)$ to the chunk of the initial state identified by the labels from U produces a state in which p holds, “compensating” for the component k in Ψ . That is, p corresponds to an abstract state g and c can be safely executed in such a state. The rule guarantees that if c terminates with a postcondition q , then q corresponds to some abstract state g' .

We illustrate the rule with a proof outline for program $\text{hide}_{\Phi, g}(\text{incr}(n))$. We show how to choose Φ and g so that the program implements the following functionality. It starts with only the concurroid P , and the private heap containing pointers lk and x . It locally installs $L_{\{\text{lock}, lk, l\}}$, which makes x a shared pointer, protected by the lock lk . It runs $\text{incr}(n)$, after which the local concurroid is disposed, and lk and x return to the private heap. We prove that if initially $x \rightarrow 0$, then in the end $x \rightarrow n$. The abstract states are pairs (m_S, a_S) , encodings of the *self* views of the concrete state of lock. Φ maps a *self* view into a predicate on the full state of lock, specifying *joint* and *other* views as well.

$$\begin{aligned} \Phi(m_S, a_S) &= \text{lock} \xrightarrow{s} (m_S, a_S) \wedge \text{lock} \xrightarrow{o} (\text{Own}, 0) \wedge \\ &\text{if } m_S = \text{Own} \text{ then } \text{lock} \xrightarrow{j} ((lk \rightarrow \text{false}) \cup (x \rightarrow a_S)) \text{ else } \text{lock} \xrightarrow{j} (lk \rightarrow \text{true}) \end{aligned}$$

We choose the initial state $g = (m_S, a_S) = (\text{Own}, 0)$: indicating that the lock is installed with lk unlocked, and x set to 0.

The proof outline uses the facts that $\Phi(\text{Own}, a_S) \downarrow lk \rightarrow \text{false} \cup x \rightarrow a_S$, and thus $\Psi(\text{Own}, a_S) \text{ empty} = \text{priv} \xrightarrow{s} lk \rightarrow \text{false} \cup x \rightarrow 0$. Also, $\Phi(m_S, a_S) \multimap \text{lock} \xrightarrow{s} (m'_S, a'_S)$ is equivalent to $(m_S, a_S) = (m'_S, a'_S)$ in the label-free state.

$$\begin{aligned} &\left\{ \text{priv} \xrightarrow{s} lk \rightarrow \text{false} \cup x \rightarrow 0 \right\} @ P \\ &\left\{ \Psi(\text{Own}, 0) \text{ empty} \right\} @ P \\ &\left\{ \Psi(\text{Own}, 0) \text{ empty} * (\Phi(\text{Own}, 0) \multimap \text{lock} \xrightarrow{s} (\text{Own}, 0)) \right\} @ P (= P \times E) \\ &\text{hide}_{\Phi, (\text{Own}, 0)} \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{lock} \xrightarrow{s} (\text{Own}, 0) \right\} @ \text{CSL}_{\{\text{lock}, lk, l\}} (= P \times E \times L_{\{\text{lock}, lk, l\}}) \\ &\quad \text{incr}(n) \\ &\quad \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{lock} \xrightarrow{s} (\text{Own}, n) \right\} @ \text{CSL}_{\{\text{lock}, lk, l\}} \\ &\left\{ \exists g_2. \Psi \ g_2 \ \text{empty} * (\Phi \ g_2 \ \multimap \ \text{lock} \xrightarrow{s} (\text{Own}, n)) \right\} @ P \\ &\left\{ \Psi(\text{Own}, n) \text{ empty} \right\} @ P \\ &\left\{ \text{priv} \xrightarrow{s} lk \rightarrow \text{false} \cup x \rightarrow n \right\} @ P \end{aligned}$$

The soundness of HIDE depends on a number of semantic properties of Φ (see Appendix C). The most important one is that states in the range of Φ have fixed *other* views for every label ℓ of V ; equivalently, that environment threads for the program $\text{hide}_{\Phi, g_1} c$ do not interfere with c on the states of V : all interference on V is *hidden* within the hide-section.

$$\text{if } w_1 \models \Phi g_1 \wedge (\ell \overset{o}{\mapsto} v_1 * \top) \text{ and } w_2 \models \Phi g_2 \wedge (\ell \overset{o}{\mapsto} v_2 * \top) \text{ then } v_1 = v_2$$

Concretely for our example, $\Phi g \wedge (\text{lock} \overset{o}{\mapsto} v)$ implies $v = (\text{Own}, 0)$, thus the above property clearly holds.

4 Concurroids Abstractly

A concurroid is a 4-tuple $V = (\mathcal{L}, \mathcal{W}, \tau, \mathcal{E})$ where: (1) \mathcal{L} is a set of labels, where a label is a nat; (2) \mathcal{W} is the *set of states*, each state $w \in \mathcal{W}$ having the structure described in Section 3; (3) τ is the *internal transition*, which is a relation on \mathcal{W} ; (4) \mathcal{E} is a set of pairs (α, ρ) , where α and ρ are *external transitions* of V . An external transition is a function, mapping a heap h into a relation on \mathcal{W} . The components must satisfy a further set of requirements, discussed next.

State properties Every state $w \in \mathcal{W}$ is valid as defined in Figure 1, and its label footprint is \mathcal{L} , *i.e.* $\text{dom}(w.s) = \text{dom}(w.j) = \text{dom}(w.o) = \mathcal{L}$. Additionally, \mathcal{W} satisfies the property:

$$\begin{aligned} \text{Fork-join closure: } \forall t: \text{PCM-map. } w \triangleleft t \in \mathcal{W} &\iff w \triangleright t \in \mathcal{W}, \\ \text{where } w \triangleleft t &= [t \circ w.s \mid w.j \mid w.o], \text{ and } w \triangleright t = [w.s \mid w.j \mid t \circ w.o] \end{aligned}$$

The property requires that \mathcal{W} is closed under the realignment of *self* and *other* components, when they exchange a PCM-map t between them. Such realignment is part of the definition of \otimes , and thus appears in proofs whenever the rule PAR is used, *i.e.* whenever threads fork or join. Fork-join closure ensures that if a parent thread forks in a state from \mathcal{W} , then the child threads are supplied with states which also are in \mathcal{W} , and dually for joining.

Transition properties A concurroid transition γ is a relation on \mathcal{W} satisfying:

$$\begin{aligned} \text{Guarantee: } (w, w') \in \gamma &\implies w.o = w'.o \\ \text{Locality: } \forall t: \text{PCM-map. } w.o = w'.o &\implies (w \triangleright t, w' \triangleright t) \in \gamma \implies (w \triangleleft t, w' \triangleleft t) \in \gamma \end{aligned}$$

Guarantee restricts γ to only modify the *self* and *joint* components. Therefore, γ describes the behavior of a viewing thread in the subjective setting, but not of the thread's environment. In the terminology of Rely-Guarantee logics [5, 6, 8, 18], γ is a *guarantee* relation. To describe the behavior of the thread's environment, *i.e.* obtain a *rely* relation, we merely *transpose* the *self* and other components of γ :

$$\gamma^\top = \{(w_1^\top, w_2^\top) \mid (w_1, w_2) \in \gamma\}, \text{ where } w^\top = [w.o \mid w.j \mid w.s].$$

In this sense, FCSL transitions always encode *both* guarantee and rely relations.

Locality ensures that if γ relates states with a certain *self* components, then γ also relates states in which the *self* components have been simultaneously *framed* by a PCM-map t , *i.e.*, enlarged according to t . It thus generalizes the notion of locality from separation logic [14], with a notable difference. In separation logic, the frame t materializes out of nowhere, whereas in FCSL, t has to be appropriated from *other*; that is, taken out from the ownership of the environment.

An *internal* transition τ is a transition which is *reflexive* and preserves heap footprints. An *acquire* transition α , and a *release* transition ρ are functions mapping heaps to transitions which extend and reduce heap footprints, respectively, as formalized below. An external transition is either an acquire or a release transition. If $(\alpha, \rho) \in \mathcal{E}$, then α is an acquire transition, and ρ is a release transition.

$$\begin{aligned} \text{Footprint preservation: } & (w, w') \in \tau \implies \text{dom } \lfloor w \rfloor = \text{dom } \lfloor w' \rfloor \\ \text{Footprint extension: } & \forall h:\text{heap}. (w, w') \in (\alpha h) \implies \text{dom } (\lfloor w \rfloor \cup h) = \text{dom } \lfloor w' \rfloor \\ \text{Footprint reduction: } & \forall h:\text{heap}. (w, w') \in (\rho h) \implies \text{dom } (\lfloor w' \rfloor \cup h) = \text{dom } \lfloor w \rfloor \end{aligned}$$

Internal transitions are reflexive so that programs specified by such transitions may be *idle* (*i.e.*, transition from a state to itself). Footprint preservation requires internal transitions to preserve the domains of heaps obtained by state flattening. Internal transitions may exchange the ownership of subheaps between the *self* and *joint* components, or change the contents of individual heap pointers, or change the values of non-heap (*i.e.*, auxiliary) state, which flattening erases. However, they cannot add new pointers to a state or remove old ones, which is the task of external transitions, as formalized by Footprint extension and reduction.

Example 1 (The concurroid for private state). $P = (\{\text{priv}\}, \mathcal{W}_P, \tau_P, \{(\alpha_P, \rho_P)\})$, with

$$\begin{aligned} \mathcal{W}_P &= \{ \text{priv} \rightarrow [h_S \mid \text{empty} \mid h_O] \mid h_S \text{ and } h_O \text{ disjoint heaps} \}, \text{ and} \\ (w, w') \in \tau_P &\iff w.s = \text{priv} \rightarrow h_S, w'.s = \text{priv} \rightarrow h'_S, \text{dom } h_S = \text{dom } h'_S, w.o = w'.o \\ (w, w') \in \alpha_P h &\iff w.s = \text{priv} \rightarrow h_S, w'.s = \text{priv} \rightarrow h_S \cup h, w.o = w'.o \\ (w, w') \in \rho_P h &\iff w.s = \text{priv} \rightarrow h_S \cup h, w'.s = \text{priv} \rightarrow h_S, w.o = w'.o \end{aligned}$$

The internal transition admits arbitrary footprint-preserving change to the private heap h_S , while the acquire and release transitions simply add and remove the heap h from h_S .

Example 2 (The concurroid for a lock). $L_{\text{lock}, lk, l} = (\{\text{lock}\}, \mathcal{W}_L, \tau_L, \{(\alpha_L, \rho_L)\})$, with $\mathcal{W}_L = \{ w \mid w \models \text{assertion (2)} \}$, and (assuming $w.o = w'.o$ everywhere):

$$\begin{aligned} (w, w') \in \tau_L &\iff w = w' \\ (w, w') \in \alpha_L h &\iff w.s = \text{lock} \rightarrow (\text{Own}, a_S), w.j = \text{lock} \rightarrow (lk \rightarrow \text{true}), \\ & \quad w'.s = \text{lock} \rightarrow (\text{Own}, a'_S), w'.j = \text{lock} \rightarrow ((lk \rightarrow \text{false}) \cup h) \\ (w, w') \in \rho_L h &\iff w.s = \text{lock} \rightarrow (\text{Own}, a_S), w.j = \text{lock} \rightarrow ((lk \rightarrow \text{false}) \cup h), \\ & \quad w'.s = \text{lock} \rightarrow (\text{Own}, a_S), w'.j = \text{lock} \rightarrow (lk \rightarrow \text{true}) \end{aligned}$$

The internal transition admits no changes to the state w . The α_L transition corresponds to unlocking, and hence to the acquisition of the heap h . It flips the ownership bit from **Own** to **Own**, the contents of the lk pointer from true to false, and adds the heap h to the resource state. The ρ_L transition corresponds to locking, and is dual to α_L . When locking, the ρ_L transition keeps the auxiliary view a_S unchanged. Thus, the resource

“remembers” the auxiliary view at the point of the last lock. Upon unlocking, the α_L transition changes this view into \mathfrak{a}'_S , where \mathfrak{a}'_S is some value that is coherent with the acquired heap h , *i.e.*, which makes the resource invariant $I(\mathfrak{a}_S \bullet \mathfrak{a}_O) h$ hold, and thus, the whole state belongs to \mathcal{W}_L .

Entanglement Let $U = (\mathcal{L}_U, \mathcal{W}_U, \tau_U, \mathcal{E}_U)$ and $V = (\mathcal{L}_V, \mathcal{W}_V, \tau_V, \mathcal{E}_V)$, be concurroids. The entanglement $U \bowtie V$ is a concurroid with the label component $\mathcal{L}_{U \bowtie V} = \mathcal{L}_U \cup \mathcal{L}_V$. The state set component combines the individual states of U and V by unioning their labels, while ensuring that the labels contain only non-overlapping heaps.

$$\mathcal{W}_{U \bowtie V} = \{w \cup w' \mid w \in \mathcal{W}_U, w' \in \mathcal{W}_V, \text{ and } [w] \text{ disjoint from } [w']\}$$

To define the transition components of $U \bowtie V$, we first need the auxiliary concept of transition interconnection. Given transitions γ_U and γ_V over \mathcal{W}_U and \mathcal{W}_V , respectively, the interconnection $\gamma_1 \bowtie \gamma_2$ is a transition on $\mathcal{W}_{U \bowtie V}$ which behaves as γ_U (resp. γ_V) on the part of the states labeled by U (resp. V).

$$\gamma_1 \bowtie \gamma_2 = \{(w_1 \cup w_2, w'_1 \cup w'_2) \mid (w_i, w'_i) \in \gamma_i, w_1 \cup w_2, w'_1 \cup w'_2 \in \mathcal{W}_{U \bowtie V}\}.$$

The internal transition of $U \bowtie V$ is defined as follows, where id_U is the diagonal of \mathcal{W}_U .

$$\tau_{U \bowtie V} = (\tau_U \bowtie \text{id}_V) \cup (\text{id}_U \bowtie \tau_V) \cup \bigcup_{h, (\alpha_U, \rho_U) \in \mathcal{E}_U, (\alpha_V, \rho_V) \in \mathcal{E}_V} (\alpha_U h \bowtie \rho_V h) \cup (\alpha_V h \bowtie \rho_U h).$$

Thus, $U \bowtie V$ steps internally whenever U steps and V stays idle, or when V steps and U stays idle, or when there exists a heap h which U and V exchange ownership over by synchronizing their external transitions.

Example 3. The transitions α_P of P and ρ_L of $L_{\text{lock},lk,l}$ have already been described in display (4) of Section 2, but using assertions, rather than semantically. The display (3) of Section 2 presents the interconnection $\alpha_P h \bowtie \rho_L h$, which moves h from $L_{\text{lock},lk,l}$ to P , and is part of the definition of $\tau_{P \bowtie L_{\text{lock},lk,l}}$. The latter further allows moving h in the opposite direction ($\alpha_L h \bowtie \rho_P h$), independent stepping of P ($\tau_P \bowtie \text{id}_L$) and of $L_{\text{lock},lk,l}$ ($\text{id}_P \bowtie \tau_L$).

The external transitions of $U \bowtie V$ are those of U , framed wrt. the labels of V .

$$\mathcal{E}_{U \bowtie V} = \{(\lambda h. (\alpha_U h) \bowtie \text{id}_V, \lambda h. (\rho_V h) \bowtie \text{id}_U) \mid (\alpha_U, \rho_U) \in \mathcal{E}_U\}$$

We note that $\mathcal{E}_{U \bowtie V}$ somewhat arbitrarily chooses to frame on the transitions of U rather than those of V . In this sense, the definition interconnects the external transitions of U and V , but it keeps those of U “open” in the entanglement, while it “shuts down” those of V . The notation $U \bowtie V$ is meant to symbolize this asymmetry. The asymmetry is important for our example of encoding CSL resources, as it enables us to iterate the (non-associative) addition of new resources as $((P \bowtie L_{\text{lock}_1, lk_1, l_1}) \bowtie L_{\text{lock}_2, lk_2, l_2}) \bowtie \dots$ while keeping the external transitions of P open to exchange heaps with new resources.

Clearly, many ways exist to interconnect transitions of two concurroids and select which transitions to keep open. In our implementation, we have identified several operators implementing common interconnection choices, and proved a number of equations

and properties about them (e.g., all of them validate an instance of the INJECT rule). We also show a version of the INJECT rule with a different operator (\bowtie) (see Appendix E). However, as none of these operators is needed for the examples in this paper, we omit them.

Lemma 1. $U \bowtie V$ is a concurroid.

We can also reorder the iterated addition of lock concurroids.

Lemma 2 (Exchange law). $(U \bowtie V) \bowtie W = (U \bowtie W) \bowtie V$.

We close the section with the definition of the concurroid E which is the right unit of the entanglement operator \bowtie . E is defined as $E = (\emptyset, \mathcal{W}_E, id, \emptyset)$, where \mathcal{W}_E contains only the empty state (i.e. the state with no labels).

5 Language and Logic

In the tradition of axiomatic program logics, the language of FCSL splits into purely functional expressions e (v when the expression is a value), and commands c with the effects of divergence, state and concurrency. We also include procedures F , for commands with arguments.

FCSL commands A command c satisfies the Hoare tuple $\{p\}c : A \{q\}@U$ if c 's effect on states respects the internal transition of the concurroid U , c is *memory-safe* when executed from a state satisfying p , and concurrently with any environment that respects the transitions (internal and external) of U . Furthermore, if c terminates, it returns a value of type A in a state satisfying q . Formally, q may use a dedicated variable res of type A to name the return result.⁷ FCSL uses a *procedure tuple*, $\forall x:B. \{p\}f(x) : A \{q\}@U$, to specify a potentially recursive higher-order procedure f taking an argument x of type B to a result of type A . The assertions p and q may depend on x . FCSL does not treat first-order looping commands, as these are special cases of recursive procedures. In the case of recursive procedures, p and q in the procedure tuple together correspond to a loop invariant, and typically are provided by the programmer.

The syntax of commands and procedures is as follows.

$$\begin{aligned} c &::= x \leftarrow c_1; c_2 \mid c_1 \parallel c_2 \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \mid F(e) \mid \text{return } v \mid \text{act } a \mid \text{inject } c \mid \text{hide}_{\phi,g} c \\ F &::= f \mid \text{fix } f. x. c \end{aligned}$$

Commands and procedures include *atomic actions* $\text{act } a$, a monadic unit $\text{return } v$ that returns v and terminates, a monadic bind (i.e. sequential composition) $x \leftarrow c_1; c_2$ that runs c_1 then substitutes its result v_1 for x to run c_2 (we write $c_1; c_2$ when $x \notin \text{FV}(c_2)$), parallel composition $c_1 \parallel c_2$, a conditional, a procedure application $F(e)$, a procedure variable f , a fixed-point construct for recursion, and injection and hiding commands.

Judgments and inference rules The FCSL judgments are *hypothetical* under a context Γ that maps *program variables* x to their type and *procedure variables* f to their specification. We allow each specification to depend on the variables declared to the left.

$$\Gamma ::= \cdot \mid \Gamma, x:A \mid \Gamma, \forall x:B. \{p\}f(x) : A \{q\}@U$$

⁷ When $A = \text{unit}$, we suppress the type and the variable res , as we did in previous sections.

Fig. 2 FCSL inference rules.

$$\begin{array}{c}
\frac{\Gamma \vdash \{p\} c_1 : B \{q\}@U \quad \Gamma, x : B \vdash \{[x/res]q\} c_2 : A \{r\}@U \quad x \notin \text{FV}(r)}{\Gamma \vdash \{p\} x \leftarrow c_1; c_2 : A \{r\}@U} \text{SEQ} \\
\\
\frac{\Gamma \vdash \{p_1\} c_1 : A_1 \{q_1\}@U \quad \Gamma \vdash \{p_2\} c_2 : A_2 \{q_2\}@U}{\Gamma \vdash \{p_1 \otimes p_2\} c_1 \parallel c_2 : A_1 \times A_2 \{[\pi_1 res/res]q_1 \otimes [\pi_2 res/res]q_2\}@U} \text{PAR} \quad \frac{\forall x:B. \{p\} f(x) : A \{q\}@U \in \Gamma}{\Gamma \vdash \forall x:B. \{p\} f(x) : A \{q\}@U} \text{HYP} \\
\\
\frac{\Gamma \vdash \{p_1\} c : A \{q_1\}@U \quad \Gamma \vdash (p_1, q_1) \sqsubseteq (p_2, q_2)}{\Gamma \vdash \{p_2\} c : A \{q_2\}@U} \text{CONSEQ} \quad \frac{\Gamma \vdash \{p\} c : A \{q\}@U \quad r \text{ stable under } U}{\Gamma \vdash \{p \otimes r\} c : A \{q \otimes r\}@U} \text{FRAME} \\
\\
\frac{\Gamma \vdash \{e = \text{true} \wedge p\} c_1 : A \{q\}@U \quad \Gamma \vdash \{e = \text{false} \wedge p\} c_2 : A \{q\}@U}{\Gamma \vdash \{p\} \text{if } e \text{ then } c_1 \text{ else } c_2 : A \{q\}@U} \text{IF} \\
\\
\frac{\Gamma \vdash \{p_1\} c : A \{q_1\}@U \quad \Gamma \vdash \{p_2\} c : A \{q_2\}@U}{\Gamma \vdash \{p_1 \wedge p_2\} c : A \{q_1 \wedge q_2\}@U} \text{CONJ} \quad \frac{\Gamma \vdash \{p\} c : A \{q\}@U \quad \alpha \notin \text{dom } \Gamma}{\Gamma \vdash \{\exists \alpha:B. p\} c : A \{\exists \alpha:B. q\}@U} \text{EXIST} \\
\\
\frac{\Gamma \vdash e : A \quad p \text{ stable under } U}{\Gamma \vdash \{p\} \text{return } e : A \{p \wedge res = e\}@U} \text{RET} \quad \frac{\Gamma, \forall x:B. \{p\} f(x) : A \{q\}@U, x:B \vdash \{p\} c : A \{q\}@U}{\Gamma \vdash \forall x:B. \{p\} (\text{fix } f. x. c)(x) : A \{q\}@U} \text{FIX} \\
\\
\frac{\Gamma \vdash \forall x:B. \{p\} F(x) : A \{q\}@U \quad \Gamma \vdash e : B}{\Gamma \vdash \{[e/x]p\} F(e) : A \{[e/x]q\}@U} \text{APP} \quad \frac{\Gamma \vdash \{p\} c : A \{q\}@U \quad r \text{ stable under } V}{\Gamma \vdash \{p * r\} \text{inject } c : A \{q * r\}@U \times V} \text{INJECT} \\
\\
\frac{\Gamma \vdash \left\{ \text{priv} \xrightarrow{s} h * p \right\} c \left\{ \text{priv} \xrightarrow{s} h' * q \right\} @ (P \times U) \times V \quad P, U \text{ and } V \text{ have disjoint sets of labels}}{\Gamma \vdash \{\Psi g h * (\Phi(g) * p)\} \text{hide}_{\Phi, g} c \{\exists g'. \Psi g' h' * (\Phi(g') * q)\} @ P \times U} \text{HIDE} \\
\\
\text{where } \Psi g h = \exists k: \text{heap. priv} \xrightarrow{s} h \cup k \wedge \Phi(g) \downarrow k \\
\\
\frac{a = (U, A, \sigma, \mu) \text{ is an action} \quad \Gamma \vdash (\sigma \wedge \text{this } w, \lambda w'. (w, w', res) \in \mu) \sqsubseteq (p, q) \quad p, q \text{ stable under } U}{\Gamma \vdash \{p\} \text{act } a : A \{q\}@U} \text{ACTION}
\end{array}$$

Γ does not bind logical variables. In first-order Hoare logics, logical variables are implicitly universally quantified with global scope. In FCSL, we limit their scope to the Hoare tuples in which they appear. This is required for specifying recursive procedures, where a logical variable may be instantiated differently in each recursive call [9]. We also assume a formation requirement on Hoare tuples $\text{FLV}(p) \supseteq \text{FLV}(q)$, *i.e.*, that all free logical variables of the postcondition also appear in the precondition.

The inference rules of the Hoare tuple judgments for commands and procedures are presented in Figure 2. We note that the assertions and the annotations in the rules (*e.g.*, Φ in the HIDE rule) may freely use the variables in Γ . To reduce clutter, we silently assume the checks that all such specification level-entities are well-typed in their respective contexts Γ .

We have already discussed PAR, INJECT and HIDE rules in their versions where the return type $A = \text{unit}$. The generalization to arbitrary A is straightforward. A side condition of HIDE ensures that the sets of labels of P , U and V don't clash, so the entanglement $(P \times U) \times V$ is defined. The rule FRAME is a special case of PAR when c_2 is taken to be the idle thread (*i.e.*, $c_2 = \text{return}()$). Just like in the rule RET, we need to prove the framing assertion r stable, to account for the interference of the *other* threads. The rule FIX requires proving a Hoare tuple for the procedure body, under a

hypothesis that the recursive calls satisfy the same tuple. The procedure APPLICATION rule uses the typing judgment for expressions $\Gamma \vdash e : A$, which is the customary one from a typed λ -calculus, so we omit its rules; in our formalization in Coq, this judgment will correspond to the CiC's typing judgment. The CONSEQ rule uses the judgment $\Gamma \vdash (p_1, q_1) \sqsubseteq (p_2, q_2)$, which generalizes the customary side conditions $p_2 \implies p_1$ for strengthening the precondition and $q_1 \implies q_2$ for weakening the postcondition, to deal with the local scope of logical variables. The other rules are standard from Hoare logic, except the ACTION rule for *atomic actions*. We devote the rest of the section to it.

Atomic actions Actions perform atomic steps from state to state, such as, *e.g.*, realigning the boundaries between, or changing the contents of *self*, *joint* and *other* state components. The actions thus serve to *synchronize* the changes to operational state (*i.e.*, heaps), with changes to the logical information required for verification (*i.e.* *auxiliary*, or *abstract*, parts of the state: $\mathbf{a}_S, \mathbf{a}_O$, *etc.*). If the logical information is erased, that is, if the states are flattened to heaps, then an action implements a single atomic memory operation such as looking up or mutating a heap pointer, CAS-ing over a heap pointer, or performing some other atomic *Read-Modify-Write* operation [7, § 5.6]. How an action manipulates the logical state is up to the user, depending on the application: we provide a formal definition of actions, and require that user's choices adhere to the definition.

An action is a 4-tuple $a = (U, A, \sigma, \mu)$ where: (1) the concurroid U whose internal transition a respects, (2) the type A of the action's return value, (3) the predicate σ on states describing the states in which the action could be executed, and (4) the relation μ relating the initial state, the ending state, and the ending result of the action. σ and μ are given in a large-footprint style, giving fully the heaps and the auxiliaries they accept.

For example, consider the action `release` used in Section 3 to release a lock and transfer the pointer x from a private heap of a thread to the ownership of the lock resource. This action is over the entangled concurroid $CSL_{\text{lock}, lk, I} = P \times L_{\text{lock}, lk, I}$ as it transfers the ownership of $(x \rightarrow -)$. Its return value type is $A = \text{unit}$. It can be executed in states in which the lock is taken by the *self* thread, and the pointer x is in the private heap. The contents of x is $\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O$, for some \mathbf{a}_S and \mathbf{a}'_S , so that once x is transferred to the ownership of the lock resource, it satisfies the resource invariant. Thus:

$$\begin{aligned}
 w \in \sigma & \iff w = \text{priv} \rightarrow [x \rightarrow (\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O) \cup h_S \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{lock} \rightarrow [(\text{Own}, \mathbf{a}'_S) \mid lk \rightarrow \text{true} \mid (\text{Own}, \mathbf{a}_O)] \\
 (w, w', \text{res}) \in \mu & \iff w = \text{priv} \rightarrow [x \rightarrow (\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O) \cup h_S \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{lock} \rightarrow [(\text{Own}, \mathbf{a}'_S) \mid lk \rightarrow \text{true} \mid (\text{Own}, \mathbf{a}_O)] \wedge \\
 & \quad w' = \text{priv} \rightarrow [h_S \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{lock} \rightarrow [(\text{Own}, \mathbf{a}_S + \mathbf{a}'_S) \mid lk \rightarrow \text{false} \cup x \rightarrow (\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O) \mid (\text{Own}, \mathbf{a}_O)]
 \end{aligned}$$

Once the states are flattened into heaps, the σ and μ components of `release` reduce to describing the behavior of a memory mutation on the pointer lk . For example, the relation $\llbracket \mu \rrbracket = \{(\llbracket w \rrbracket, \llbracket w' \rrbracket, r) \mid (w, w', r) \in \mu\}$ relates (h, h', r) iff

$$\begin{aligned}
 h & = (x \rightarrow (\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O)) \cup h_S \cup (lk \rightarrow \text{true}) \cup h_O \\
 h' & = (x \rightarrow (\mathbf{a}_S + \mathbf{a}'_S + \mathbf{a}_O)) \cup h_S \cup (lk \rightarrow \text{false}) \cup h_O
 \end{aligned}$$

Thus, operationally, `release` can be implemented as a single mutation to the lk pointer.

The inference rule ACTION takes an action $a = (U, A, \sigma, \mu)$ and checks that a satisfies that σ can be strengthened into p and μ can be weakened into q . As μ is not a

postcondition itself, but a relation taking input states, we first introduce a fresh logical variable w to name the input state using a predicate `this`. Then the predicate expressing post states for the action is computed out of μ and w , and it is this predicate that’s weakened into q . p and q must be stable wrt. U , in order to account for the possibility that an interference of the environment appears just before, or just after, the action is executed.

Soundness and Implementation We have established the soundness of FCSL by exhibiting a denotational model based on *action trees* [10, 11], which are a variation on Brookes’ action trace semantics, so we can formulate the following theorem.

Theorem 1. *FCSL is sound with respect to the denotational model of action trees.*

We developed the model in the logics of Calculus of Inductive Constructions, thus, the model is a shallow embedding in Coq, and its implementation is available on-line [11]. The implementation also defines denotational semantics for constructs and ascribes them types corresponding to rules in Figure 2. These type ascriptions require proofs, and together establish soundness of the logic, although rules/types in the implementation differ somewhat from those in Figure 2, facilitating encoding in Coq: (1) they use binary postconditions, (2) pre-/postconditions are in higher-order logic over heaps and PCMs, instead of notation from Figure 1, (3) they infer weakest-pre-/strongest-postconditions and (4) assertions are stabilized. The correspondence between the implementation and Figure 2 is straightforward, but established by hand.

6 Related Work

FCSL builds on the previous work on subjective auxiliary state and SCSL logic [10]. The SCSL logic contained the distinction between *self* and *other* views, which was essential for compositional implementation of *auxiliary state*. However, it contained exactly one coarse-grained resource, with no ability to create and dispose new resources. In contrast, FCSL can introduce any number of fine-grained resources in a scoped way.

The work on Concurrent Abstract Predicates (CAP) [4] introduces a notion of *shared region* that serves a similar purpose as concurroids, in that regions circumscribe a chunk of shared heap with a protocol governing its evolution. A *protocol* is defined by a set of atomic actions, which are RG-style transitions on private state and a region. In addition to heaps, regions may contain abstract capabilities that identify enabled actions. Thus there is a subtle mutual recursion in a protocol definition between an action and the capability to perform the action. A recurring pattern for this approach is quantification over *all* possible capabilities and placing them in a shared region, to be used up if needed in the execution of the protocol. The CAP framework could atomically change only one region; a restriction lifted in the recent work on Views [3] and HOCAP [15] that introduced *view shifts* to synchronize changes in several regions. Once allocated, CAP’s regions have dynamically-scoped lifetime, and they can be disposed by a particular thread if it collects all corresponding region’s capabilities. To the best of our knowledge, HOCAP does not allow the removal or scoped hiding of a shared region.

In contrast with CAP and their successors, FCSL does not require capabilities to perform actions, as these are naturally represented in the *self* and *other* views associated with a resource (and can also be seen as auxiliary state). Such auxiliary state is simpler

than capabilities; it is not subject to ownership transfer, and there is no need to quantify over all capabilities. In our experience, this simplicity extends to the specification of invariants and transitions, and to the proofs of stability. In FCSL, synchronizing changes over a number of *concurroids* is achieved directly at the level of transitions by means of entanglement, and at the level of programs by allowing actions to be defined over any *concurroid*, including entangled ones. Thus, no view shifts are required. The burden of stability proofs is further reduced in FCSL by formulating private heaps as a separate *concurroid* that one may, but need not, entangle with. Thus, when an action manipulates only the internal state of a resource, the attendant stability proofs can ignore private heaps, *e.g.*, the *take* action of a ticketed lock [4, 11]. Moreover, the communication in FCSL makes it possible for *concurroids* to pass heaps between each other directly, rather than going through private state. While the current paper does not present examples that exploit this ability, we have found it useful when verifying in FCSL a more advanced example of readers-writers, which we will present in future work.

CaReSL [17] uses the same notion of shared region as CAP, though it specifies the transitions in a manner closer to FCSL, namely by means of STS's. CaReSL does not directly provide subjective *self* and *other* views of a resource, but it provides a notion of *tokens*, whose ownership is exchanged between a thread and its environment. CaReSL assertions explicitly allow statements only about self-owned tokens, not *other*-owned ones. Thus, reasoning about the lack of logical changes to environment-owned data has to be encoded with a level of indirection, potentially quantifying over all tokens, similar to CAP's quantification over capabilities. A frequent side condition in CaReSL rules is that various assertions are *token-pure*, which does not have a direct correspondent in FCSL. Similar to CAP, CaReSL currently allows actions that work over only a single region, and will require an extension akin to view shifts to enable synchronized updates. CaReSL does not consider removal or scoped hiding of shared regions, although it can be emulated by introducing an empty "final" protocol state. Instead of stability checks in FCSL, in CaReSL one may stabilize assertions by composing them with environment stepping. In our experience, this does not change the proofs: the same obligations reappear in proofs out of stabilized hypotheses. On the other hand, CaReSL can reason about fine-grained data structure by means of refinement (a generalization of linearizability). FCSL supports higher-order functions by means of shallow embedding into CiC [1, 16], but we have not considered linearizability so far, which is future work.

Feng's Local Rely-Guarantee (LRG) [5] is, to the best of our knowledge, the first work that reconciled fine-grained reasoning in the style of RG with framing and hiding at the level of transitions (similar to our INJECT and HIDE). We differ from LRG in that we introduce communication and subjectivity into the mix; thus our injection and hiding rules take *self* and *other* views into account. The latter are a compositional form of auxiliary state, whereas LRG in practice has to use the classical, non-compositional form of auxiliary state [10, 13].

7 Conclusion and Future Work

We presented *concurroids*—a novel model for scalable shared-memory concurrency verification, based on communicating STS, and FCSL—a logic for *concurroids*.

In the future work, we are going to build a number of concurroids to encode common programming patterns. For example, dynamic allocation and deallocation of memory can be encoded via an allocator concurroid (without extensions of FCSL), and similarly for dynamic allocation and deallocation of locks. We hope to investigate if concurroids can be endowed with analogues of channel relabeling and restriction operators from process algebras, to provide finer control over interconnection and closure of external transitions. Finally, we plan to consider refinement which allows weakening the ascribed concurroid U of a program, to a coarser-grained concurroid V , if U can be shown to simulate V . One could then verify fine-grained concurrent ADTs against V , and afterwards hide the granularity by switching to U .

Acknowledgments We thank Anindya Banerjee, Thomas Dinsdale-Young and the ESOP 2014 anonymous reviewers for their comments. This research was partially supported by Spanish MINECO projects TIN2012-39391-C04-01 Strongsoft, TIN2010-20639 Paran10, AMAROUT grant PCOFUND-GA-2008-229599, and Ramon y Cajal grant RYC-2010-0743.

References

1. Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Springer Verlag, 2004.
2. S. Brookes. A semantics for concurrent separation logic. *Th. Comp. Sci.*, 375(1-3), 2007.
3. T. Dinsdale-Young, L. Birkedal, P. Gardner, M. J. Parkinson, and H. Yang. Views: compositional reasoning for concurrent programs. In *POPL'13*.
4. T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP 2010*.
5. X. Feng. Local rely-guarantee reasoning. In *POPL'09*.
6. X. Feng, R. Ferreira, and Z. Shao. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *ESOP 2007*.
7. M. Herlihy and N. Shavit. *The art of multiprocessor programming*. M. Kaufmann, 2008.
8. C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Trans. Prog. Lang. Syst.*, 5(4), 1983.
9. T. Kleymann. Hoare logic and auxiliary variables. *Formal Asp. Comput.*, 11(5), 1999.
10. R. Ley-Wild and A. Nanevski. Subjective auxiliary state for coarse-grained concurrency. In *POPL'13*.
11. A. Nanevski, R. Ley-Wild, I. Sergey, and G. A. Delbianco. Supporting Material. <http://software.imdea.org/~aleks/fcs1/>.
12. P. W. O'Hearn. Resources, concurrency, and local reasoning. *Th. Comp. Sci.*, 375(1-3), 2007.
13. S. S. Owicki and D. Gries. Verifying properties of parallel programs: An axiomatic approach. *Commun. ACM*, 19(5), 1976.
14. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.
15. K. Svendsen, L. Birkedal, and M. J. Parkinson. Modular reasoning about separation of concurrent data structures. In *ESOP 2013*.
16. The Coq Development Team. *The Coq Proof Assistant Reference Manual - Version V8.4*, 2012. <http://coq.inria.fr/>.
17. A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP'13*.
18. V. Vafeiadis and M. J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR 2007*.

Optional Appendices

In the optional appendices we describe in detail some technical insights behind FCSL and the model of concurroids. We also provide an extended overview of another case study, which was omitted in the paper body due to space constraints. The remainder of the paper is structured as follows. **Appendix A** describes a number of properties and requirements of the actions, presented in Section 5 of the main paper body. **Appendix B** formalizes the familiar notions of *Rely* and *Guarantee* relations, as well as a notion of *stability*, in terms of concurroids and their transitions. **Appendix C** lists the necessary properties of the abstraction functions Φ , used in the HIDE rule of FCSL. **Appendix D** defines Hoare ordering $(p_1, q_1) \sqsubseteq (p_2, q_2)$, used in the rule CONSEQ of the logic. **Appendix E** presents development of a large example: a concurroid for the ticketed lock algorithm. Finally, in **Appendix F** we present the denotational model of FCSL, based on *action trees* and *schedules*, and prove the main technical result: soundness of FCSL as a shallow embedding into the Calculus of Inductive Constructions.

A Abstract Properties of Actions

An action is a 4-tuple $a = (U, A, \sigma, \mu)$. U is a concurroid whose internal transition a is supposed to respect. A is the type of the return value of a . σ is the *safety predicate* (over states of U) describing under which conditions a may be executed. μ is the *stepping relation*; it relates the input state, output state, and the result of a .

Definition 1 (Action erasures). *Given an action a , the erasures $\lfloor \sigma \rfloor$ and $\lfloor \mu \rfloor$ of a 's safety predicate and stepping relation are relations on heaps defined as follows.*

$$\begin{aligned} \lfloor w \rfloor \in \lfloor \sigma \rfloor & \iff w \in \sigma \\ (\lfloor w \rfloor, \lfloor w' \rfloor, r) \in \lfloor \mu \rfloor & \iff (w, w', r) \in \mu \end{aligned}$$

An *atomic* is a triple $\alpha = (A, \sigma, \mu)$. It's a special kind of actions, but over concrete heaps, rather than over states. States differ from heaps in that they are decorated with additional information such as auxiliary state and partitioning between *self*, *joint* and *other*. As with actions, A is the return type, σ is the safety predicate and μ is the stepping relation, but they all range over heaps.

We consider four different (parametrized classes of) atomics, corresponding to the four (parametrized) primitive memory operations that we consider.

Definition 2 (Primitive atomic actions).

$$\begin{aligned} \text{Read}_x^A &= (A, (x \rightarrow_A -) \cup h, (x \rightarrow v) \cup h \rightsquigarrow (x \rightarrow v) \cup h \wedge \text{res} = v) \\ \text{Write}_{x\ v} &= (\text{unit}, (x \rightarrow -) \cup h, (x \rightarrow -) \cup h \rightsquigarrow (x \rightarrow v) \cup h) \\ \text{Skip} &= (\text{unit}, h, h \rightsquigarrow h) \\ \text{RMW}_{x\ f\ g}^A\ B &= (B, (x \rightarrow_A -) \cup h, (x \rightarrow v) \cup h \rightsquigarrow (x \rightarrow f(v)) \cup h \wedge \text{res} = g(v)) \end{aligned}$$

The last class $\text{RMW}_{x\ f\ g}^A\ B$ corresponds to the family of *Read-Modify-Write* operations: they all atomically replace the current register value v with $f(v)$ for some pure function f , and return the result according to the function g [7, §5.6]. One particular

representative of this family is the CAS operation, which instantiates the parameters of RMW as follows:

$$\text{CAS}_{A \ x \ v_1 \ v_2} \hat{=} \text{RMW}_{x \ f(v_1, v_2) \ g(v_1, v_2)}^{A \ \text{bool}}, \text{ where}$$

$$\begin{aligned} f(v_1, v_2)(v) &= \text{if } (v = v_1) \text{ then } v_2 \text{ else } v_1 \\ g(v_1, v_2)(v) &= (v = v_1) \end{aligned}$$

Another representative of the RMW class is the atomic *fetch-and-increment* functions, which we will employ in Appendix E, defined as follows:

$$\text{FAI}_x \hat{=} \text{RMW}_{x \ \text{inc} \ \text{id}}^{\text{nat} \ \text{nat}}, \text{ where}$$

$$\text{inc}(v) = v + 1 \quad \text{id}(v) = v$$

Definition 3 (Operational action). *An action a is operational if its erasure corresponds to one of the atomics, i.e., if there exists $b \in \{\text{Read}_x^A, \text{Write } x \ v, \text{Skip}, \text{RMW}_{x \ f \ g}^{A \ B}\}$ such that*

$$\llbracket \sigma_a \rrbracket \subseteq \sigma_b \wedge \forall h \in \llbracket \sigma_a \rrbracket \ h' \ r. (h, h', r) \in \llbracket \mu_a \rrbracket \implies (h, h', r) \in \mu_b$$

In our examples we only consider operational actions, though the inference rules and the implementation in Coq don't currently enforce this requirement (the operability of actions in the examples has been proved by hand).

Let $U = (\mathcal{L}, \mathcal{W}, \tau, \mathcal{E})$. The action $a = (U, A, \sigma, \mu)$ is required to satisfy the following properties.

$$\begin{aligned} \text{Coherence} &: w \in \sigma \implies w \in \mathcal{W} \\ \text{Safety monotonicity} &: w \triangleright t \in \sigma \implies w \triangleleft t \in \sigma \\ \text{Step safety} &: (w, w', r) \in \mu \implies w \in \sigma \\ \text{Internal stepping} &: (w, w', r) \in \mu \implies (w, w') \in \tau \\ \text{Framing} &: w \triangleright t \in \sigma \implies (w \triangleleft t, w', r) \in \mu \implies \\ &\quad \exists w''. w' = w'' \triangleleft t \wedge (w \triangleright t, w'' \triangleright t, v) \in \mu \\ \text{Erasure} &: \text{defined}(\llbracket w \rrbracket \cup h) \implies \llbracket w \rrbracket \cup h = \llbracket w' \rrbracket \cup h' \implies \\ &\quad (w, w_1, r) \in \mu \implies (w', w'_1, r') \in \mu \implies \\ &\quad r = r' \wedge \llbracket w \rrbracket_1 \cup h = \llbracket w'_1 \rrbracket \cup h' \\ \text{Totality} &: \forall w. w \in \sigma \implies \exists w' \ v. (w, w', v) \in \mu \end{aligned}$$

The properties of Coherence, Step safety and Internal stepping are straightforward. Safety monotonicity states that if the action is safe in a state with a smaller *self* component (because the other component is enlarged by t), the action is also safe if we increase the *self* component by t .

Framing property says that if a steps in a state with a large *self* component $w \triangleleft t$, but is already safe to step in a state with a smaller *self* component $w \triangleright t$, then the result state and value obtained by stepping in $w \triangleleft t$ can be obtained by stepping in $w \triangleright t$, and moving t afterwards.

The Erasure property shows that the behavior of the action on the concrete input state obtained after erasing the auxiliary fields and the logical partition, doesn't depend on the erased auxiliary fields and the logical partition. In other words, if the input state have *compatible* erasures (that is, erasures which are subheaps of a common heap), then executing the action in the two states results in equal values, and final states that also

have compatible erasures. This is a standard property proved in concurrency logics that deal with auxiliary state and code [2, 13].

The Totality property shows that an action whose safety predicate is satisfied always produces a result state and value. It doesn't loop forever, and more importantly, it doesn't crash. We will use this property of actions in the semantics of programs to establish that if the program's precondition is satisfied, then all of the approximations in the program's denotation are either done stepping, or can actually make a step (i.e., they make progress).

So far, the actions are defined in a so-called *large footprint* style, e.g., the definition of the action `release` in Section 5 mentions existentially-quantified heap h_S explicitly. To enable writing various actions in a *small footprint* style, we also enforce the property

$$\text{Locality} : w.o = w'.o \implies (w \triangleright t, w' \triangleright t, v) \in \mu \implies (w \triangleleft t, w' \triangleleft t, v) \in \mu$$

Curiously, if the default use of the logic is in a large footprint notation, then this property is not necessary as it is not used in any proofs.

B Self/Environment Stepping and Stability

Definition 4 (Self-stepping (Guarantee)).

Given a concurroid $U = (\mathcal{L}, \mathcal{W}, \tau, \mathcal{E})$, the environment stepping relation \mathcal{G}_U of U unions its internal and all external transitions of U , and then takes a transitive closure.

$$\mathcal{G}_U = (\tau_U \cup \bigcup_{h, (\alpha, \rho) \in \mathcal{E}_X} (\alpha h) \cup (\rho h))^*$$

Definition 5 (Environment stepping (Rely)).

Given a concurroid $U = (\mathcal{L}, \mathcal{W}, \tau, \mathcal{E})$, the environment stepping relation \mathcal{R}_U of U unions the transposed versions of the internal and all external transitions of U , and then takes a transitive closure.

$$\mathcal{R}_U = (\tau^T \cup \bigcup_{h, (\alpha, \rho) \in \mathcal{E}} (\alpha h)^T \cup (\rho h)^T)^*$$

Definition 6 (Stability). An assertion (i.e., a predicate over states) p is stable under concurroid U iff for every w, w' such that $(w, w') \in \mathcal{R}_U$, $w \models p \implies w' \models p$.

C Properties of Φ Functions from the Hiding Rule

The abstraction function Φ is a user-specified annotation on the hide command. It maps values $g : G$ (where G is also user specified) to assertions, that is, predicates over states (equivalently, sets of states) of a concurroid V . For the soundness of the hiding rule, Φ is required to satisfy the following properties.

$$\begin{aligned} \text{Coherence} & : w \in \Phi(g) \implies w \in \mathcal{W}_V \\ \text{Injectivity} & : w \in \Phi(g_1) \implies w \in \Phi(g_2) \implies g_1 = g_2 \\ \text{Surjectivity} & : w_1 \in \Phi(g_1) \implies w_2 \in \mathcal{W}_W \implies w_1.o = w_2.o \implies \exists g_2. w_2 \in \Phi(g_2) \\ \text{Guarantee} & : w_1 \in \Phi(g_1) \implies w_2 \in \Phi(g_2) \implies w_1.o = w_2.o \\ \text{Precision} & : w_1 \in \Phi(g) \implies w_2 \in \Phi(g) \implies \lfloor w_1 \rfloor \cup h_1 = \lfloor w_2 \rfloor \cup h_2 \implies w_1 = w_2 \end{aligned}$$

Coherence and Injectivity are obvious. Surjectivity states that for every state w_2 of the concurroid W one can find an image g , under the condition that the *other* component of w_2 is well-formed according to Φ (typically, that the *other* component is equal to the unit of the PCM-map monoid for W). Guarantee is the property stated in the paper at the end of Section 3, but through assertions, rather than semantically. This property formalizes that environment of `hide` can't interfere on V , as V is installed locally. Thus, whatever the environment does, it can't influence the *other* component of the states w described by Φ .

Precision is a technical property common to separation-style logics, though here it has a somewhat different flavor. Precision ensures that for every value g , $\Phi(g)$ precisely describes the underlying heaps of its circumscribed states; that is, each state $\Phi(g)$ is uniquely determined by its heap erasure.

D Definition of Hoare Ordering $(p_1, q_1) \sqsubseteq (p_2, q_2)$

Given preconditions p_1, p_2 and postconditions q_1, q_2 , the judgment $(p_1, q_1) \sqsubseteq (p_2, q_2)$ generalizes the usual Hoare logic side conditions on the rule of consequence about strengthening the precondition $p_2 \implies p_1$ and weakening the postcondition $q_1 \implies q_2$. The generalization is required in FCSL because of the local scope of logical variable. In first order Hoare logics, the logical variables have global scope, so the above implications over p_1, p_2 and q_1, q_2 suffice. In FCSL, the logical variables have scope locally over Hoare triples, and this scope has to be reflected in the semantic definition of \sqsubseteq by introducing quantifiers. The definition is similar to the one of Kleymann [9].

$$\begin{aligned} (p_1, q_1) \sqsubseteq (p_2, q_2) &\iff \\ &\forall w w'. (w \models \exists \bar{v}_2. p_2 \implies w \models \exists \bar{v}_1. p_1) \wedge \\ &\quad ((\forall \bar{v}_1 \text{ res}. w \models p_1 \implies w' \models q_1) \implies (\forall \bar{v}_2 \text{ res}. w \models p_2 \implies w' \models q_2)) \end{aligned}$$

where $v_i = \text{FLV}(p_i, q_i)$ are the free logical variables. The definition makes it apparent that the Hoare triple $\{p\} c \{q\}@U$ is essentially a syntactic sugar for a different kind of Hoare triple, which may be written as:

$$\{w. \exists \bar{v}. w \models p\} c \{\text{res } w w'. \forall \bar{v}. w \models p \implies w' \models q\}@U$$

where $v = \text{FLV}(p, q)$. In this alternative Hoare triple, the postconditions are predicates ranging over input and output states w and w' (they are thus called binary postconditions). The advantage of the alternative Hoare triple is that the logical variables are explicitly bound, making their scoping explicit. In our Coq implementation we use this alternative formulation of Hoare triples.

E A Concurroid for the Ticketed Lock

In this section, we instantiate the presented framework by constructing from scratch a concurroid for another famous fine-grained locking algorithm—a ticketed lock [4]. Unlike the CAS-based spin lock, which we took as a running example, this algorithm, although simple, provides fairness guarantees for competing threads. We start by giving a simplified implementation of the algorithm and building the intuition about the states and transitions of the concurroid, implementing its protocol.

E.1 Reference implementation

As a reference implementation of the ticketed lock, we use the one from the work on Concurrent Abstract Predicates [4].

```
lock = {
    n := TAKE();
    while (!TRY(n)) skip;
}
unlock = {
    DROP_TKT();
}
```

The three *atomic* actions TAKE, TRY(*n*) and DROP_TKT have the following operational meaning:

```
TAKE()      = { fetch_and_increment(next); }
TRY(n)      = { return (n == owner); }
DROP_TKT() = { fetch_and_increment(owner); }
```

As a minor difference, we assume a pointer to the lock structure being fixed and omit it from the code, referring instead to its two essential fields: *owner* and *next*. These two fields are natural numbers, describing the state of the ticketed lock system. We implemented TAKE() and DROP_TKT() as `fetch_and_increment`, which is a standard atomic operation implemented in most of the modern architectures, operationally equivalent to the following sequence of commands:

```
fetch_and_increment(ptr) = { tmp := ptr; ptr++; return tmp; }
```

The intuition behind the protocol comes from the idea of the organization of the service in a bakery: *next* corresponds to the current ticket available at the dispenser (but not yet taken!), whereas *owner* indicates a ticket number, an owner of which is called to be served, or is currently being served. Hence, the `lock` procedure corresponds to drawing a ticket from the dispenser and a subsequent repetitive attempt to be served, whereas `unlock` corresponds to throwing the ticket away and increasing the ticket counter, so the next client could be served (if there is one in a waiting line).

E.2 States of the concurroid

The intuition for the construction of a ticketed lock concurroid can be built from the idea of *splitting* the resources in the spirit of the Subjective Concurrent Separation Logic [10]. What are the primary resources that a ticketed lock deals with? Those are tickets, of course! Some of them can be taken by a current thread, some by others. If we make an assumption that each ticket is thrown away right after its owner has been served, we will see a simple invariant: all tickets in the system are from the range $[n_1, n_2)$, where n_1 and n_2 are natural numbers, corresponding to the values of *owner* (currently being served or about to be) and *next* (next ticked to be issued), hence non-inclusion on the right side of the range.

As in the case of the CAS-based spin lock (Section 2), the *self* and *other* components of the ticketed lock concurroid are pairs of two PCMs: the PCM of *finite sets* of

natural numbers, accounting for the tickets (t_S and t_O), and a generic PCM accounting in order to the heap invariant associated with this particular lock: a_S and a_O . Similarly to the CAS-based spin lock, the ticketed lock is parametrized by a resource invariant I . We will denote the whole concurroid with the label `tlock`, pointers `own` and `next` and invariant I by $T_{\{\text{tlock}, \text{own}, \text{next}, I\}}$. Let us start by defining a predicate, describing valid states of the ticketed lock concurroid.

Definition 7 ($T_{\{\text{tlock}, \text{own}, \text{next}, I\}}$ state invariant \mathcal{W}_T).

$$\mathcal{W}_T \hat{=} \text{tlock} \xrightarrow{s} (t_S, a_S) \wedge \text{tlock} \xrightarrow{o} (t_O, a_O) \wedge \text{tlock} \xrightarrow{j} (\text{own} \rightarrow \langle n_1, b \rangle) \cup (\text{next} \rightarrow n_2) \cup h \wedge$$

$$n_1 \leq n_2 \wedge (t_S \bullet t_O) = \{n \mid n_1 \leq n < n_2\} \wedge$$

$$\left(\begin{array}{l} n_1 \in (t_S \bullet t_O) \wedge b = \text{true} \wedge h = \text{empty} \vee \\ \text{if } n_1 = n_2 \text{ then } b = \text{false} \wedge I(a_S \bullet a_O) h \\ \text{else } n_1 \in (t_S \bullet t_O) \wedge b = \text{false} \wedge I(a_S \bullet a_O) h \end{array} \right)$$

Definition 7 highlights a number of important invariants of the ticketed lock:

- The first line of the conjunction specifies the shape of the concurroid states. Since the *joint* component is always a heap, we needed to introduce a “ghost” value b , such that $\text{own} \rightarrow \langle n_1, b \rangle$, which tracks if the lock is taken or not. This is a purely logical artifact, which simplifies the reasoning and it can be safely erased in the actual implementation, making this entry look just as $(\text{own} \rightarrow n_1)$. To make it clear that b is auxiliary state, we could have introduced a separate label whose *joint* part contains only b . In our implementation, *joint* components need not contain only heaps, but can contain arbitrary non-heap types, whose values are by default considered auxiliary. Thus, the addition of a new label with a boolean contents b would suffice to declare b as an auxiliary state. However, for simplicity, we omitted such generalization.
- The second line states the invariant with respect to the distributed tickets. Indeed, the number of a ticket currently being (or just about to be) served is less than or equal to the number of the ticket in the dispenser, hence, $n_1 \leq n_2$. The tickets in the queue range from n_1 to n_2 , and n_2 has not yet been assigned to any thread.
- The last three lines outline the essential “states” of the system.
 - One case is the lock being locked. It corresponds to a thread with some ticket n_1 (captured either by t_S or t_O) being served. Therefore, the ownership over the heap is being transferred to the client ($h = \text{empty}$) and the logical flag b is indicating the locking ($b = \text{true}$). We will refer to this state as *locked*.
 - Another case corresponds to the situation when there are no more threads awaiting for service or being served. This is the case $n_1 = n_2$. Therefore, the lock is not taken ($b = \text{false}$) and has the ownership over the heap with the lock invariant holding over it ($I(a_S \bullet a_O) h$). We will refer to this state as *unlocked*.
 - The least obvious case is the situation when *no* thread is being served, but one is just about to be. This is the thread having the ticket n_1 (hence $n_1 \in (t_S \bullet t_O)$), but the lock is not yet taken ($b = \text{false}$) and has the ownership over the heap ($I(a_S \bullet a_O) h$). We will refer to such state as *transit*.

It is important to notice that the awkward addition of the logical flag b to the *joint* state is a byproduct of co-existence of “locked” and “transit” states of the concurroid and our generality with respect to the invariant I . In particular, had we assumed $I(a_S \bullet a_O) h \implies h \neq \text{empty}$, we could use this fact in order to distinguish between the “locked” and “transit” states. Finally, as it will be shown, the flag b plays an important role only when verifying the body of the `lock` procedure, and can be abstracted away for convenience of the clients of `lock` and `unlock`.

E.3 Transitions of the concurroid

In contrast with the CAS-based spin lock, whose internal transitions do not allow any components of the state to change, the internal transitions of the ticketed lock concurroid account for *taking* the tickets from the dispenser and changing the corresponding counter nxt in the *joint* part of the state. Again, describing the transitions on states w and w' , we implicitly assume $w.o = w'.o$.

Definition 8 ($T_{\{\text{lock}, \text{own}, \text{nxt}, I\}}$ **internal transitions** τ_T).

$$(w, w') \in \tau_T \iff \begin{array}{l} w.s = \text{lock} \rightarrow (t_S, a_S), w.j = \text{lock} \rightarrow (nxt \rightarrow n_2 \cup h), \\ w'.s = \text{lock} \rightarrow (t'_S, a_S), w.j = \text{lock} \rightarrow (nxt \rightarrow n'_2 \cup h), \\ n_2 \leq n'_2 \end{array}$$

Notice that in Definition 8 we don't specify the rest of the *joint* heap h . What is important is that it remains unchanged. The nxt counter may only increase, and this change is reflected by changing the *self* auxiliary component t_S to t'_S , assuming that the resulting state still adheres to the state predicate \mathcal{W}_T (Definition 7).

Definition 9 ($T_{\{\text{lock}, \text{own}, \text{nxt}, I\}}$ **external transitions**, α_T and ρ_T).

$$(w, w') \in \alpha_T h \iff \begin{array}{l} w.s = \text{lock} \rightarrow (t_S \cup \{n_1\}, a_S), \\ w.j = \text{lock} \rightarrow (\text{own} \rightarrow \langle n_1, \text{true} \rangle \cup nxt \rightarrow n_2), \\ w'.s = \text{lock} \rightarrow (t_S, a'_S), \\ w'.j = \text{lock} \rightarrow (\text{own} \rightarrow \langle n_1 + 1, \text{false} \rangle \cup nxt \rightarrow n_2 \cup h) \end{array}$$

$$(w, w') \in \rho_T h \iff \begin{array}{l} w.s = \text{lock} \rightarrow (t_S \cup \{n_1\}, a_S), \\ w.j = \text{lock} \rightarrow (\text{own} \rightarrow \langle n_1, \text{false} \rangle \cup nxt \rightarrow n_2 \cup h), \\ w'.s = \text{lock} \rightarrow (t_S \cup \{n_1\}, a_S), \\ w'.j = \text{lock} \rightarrow (\text{own} \rightarrow \langle n_1, \text{true} \rangle \cup nxt \rightarrow n_2) \end{array}$$

In Definition 9, the acquire transition α_T transfers the heap h to the *joint* part of the state $w'.j$, increases the service counter, hence $n_1 + 1$, and flips the flag b from `true` to `false`. The transition is only possible when the ticket number n_1 is “registered” in the *self*-set t_S in the initial state. Similarly, the release transition ρ_T is only possible when the ticked being called is owned by the *self* set t_S . In such a case, the ownership of the heap h is given away to the calling thread.

We now have all ingredients to define the concurroid for ticketed locks.

Definition 10 (The concurroid for the ticketed lock).

$T_{\{\text{tlock}, \text{own}, \text{next}, I\}} = (\{\text{tlock}\}, \mathcal{W}_T, \tau_T, (\alpha_T, \rho_T))$, where \mathcal{W}_T is given by Definition 7 and τ_T, α_T and ρ_T are defined in Definitions 8 and 9.

E.4 Atomic actions of a ticketed lock

We now define the three atomic actions, used in the implementation of the locking/unlocking procedures in Section E.1.

The action `take` corresponds to internal transitions of the concurroid $T_{\{\text{tlock}, \text{own}, \text{next}, I\}}$ and returns the number of the drawn ticket as its result. Hence the following specifications, according to the definitions of Section 5:

Definition 11 (take action). `take` = $(T_{\{\text{tlock}, \text{own}, \text{next}, I\}}, \text{nat}, \sigma, \mu)$, where

$$\begin{aligned} w \in \sigma & \iff w = \text{tlock} \rightarrow [(\text{ts}, \text{as}) \mid \text{next} \rightarrow n_2 \cup h \mid (\text{to}, \text{ao})] \\ (w, w', \text{res}) \in \mu & \iff w = \text{tlock} \rightarrow [(\text{ts}, \text{as}) \mid \text{next} \rightarrow n_2 \cup h \mid (\text{to}, \text{ao})] \wedge \\ & w' = \text{tlock} \rightarrow [(\text{ts} \cup \{n_2\}, \text{as}) \mid \text{next} \rightarrow (n_2 + 1) \cup h \mid (\text{to}, \text{ao})] \wedge \\ & \text{res} = n_2 \end{aligned}$$

By doing *flattening*, one can see that operationally `take` corresponds to an atomic *fetch-and-increment* of the pointer `next`. Formally, the fetch-and-increment operation can be defined in Section A. Two other actions, `try` and `droptkt` account for the heap ownership transfer and, hence, correspond to internal transitions of the entanglement $P \bowtie T_{\{\text{tlock}, \text{own}, \text{next}, I\}}$.

Definition 12 (try action). $\forall n : \text{nat}$. `try`(n) = $(P \bowtie T_{\{\text{tlock}, \text{own}, \text{next}, I\}}, \text{bool}, \sigma, \mu)$, where

$$\begin{aligned} w \in \sigma & \iff \\ & w = \text{priv} \rightarrow [\text{empty} \mid \text{empty} \mid h_0] \cup \\ & \quad \text{tlock} \rightarrow [(\text{ts} \cup \{n\}, \text{as}) \mid \text{own} \rightarrow \langle n_1, b \rangle \cup \text{next} \rightarrow n_2 \cup h \mid (\text{to}, \text{ao})] \\ (w, w', \text{res}) \in \mu & \iff \\ & w = \text{priv} \rightarrow [\text{empty} \mid \text{empty} \mid h_0] \cup \\ & \quad \text{tlock} \rightarrow [(\text{ts} \cup \{n\}, \text{as}) \mid \text{own} \rightarrow \langle n_1, b \rangle \cup \text{next} \rightarrow n_2 \cup h \mid (\text{to}, \text{ao})] \wedge \\ & \text{if } n = n_1 \\ & \text{then } w' = \text{priv} \rightarrow [h \mid \text{empty} \mid h_0] \cup \\ & \quad \text{tlock} \rightarrow [(\text{ts} \cup \{n\}, \text{as}) \mid \text{own} \rightarrow \langle n_1, \text{true} \rangle \cup \text{next} \rightarrow n_2 \mid (\text{to}, \text{ao})] \wedge \\ & \quad \text{res} = \text{true} \wedge I(\text{as} \bullet \text{ao}) h \\ & \text{else } w' = w \wedge \text{res} = \text{false} \end{aligned}$$

One can see that the action `try`(n) corresponds to reading from the heap cell `own` \rightarrow – (hence, it’s operational), but it has a interesting effect on the auxiliary state. In particular, if the comparison succeeds (then-branch), the ownership over heap h is transferred to the *self*-part of `priv`, and logical flag is flipped, which corresponds to the internal transition of the entanglement $P \bowtie T_{\{\text{tlock}, \text{own}, \text{next}, I\}}$. Otherwise, the thread has approached the counter with a ticket that is not called, and thus is denied service for now. The state therefore doesn’t change, which corresponds to the reflexivity of the internal transition.

The action can only be performed if the ticket n , which is being checked for the service, is present in the *self* part, hence $t_S \cup \{n\}$ (that is the thread indeed has right to *try* claim the service).

Definition 13 (droptkt action). $\text{droptkt} = (T_{\{\text{tlock}, \text{own}, \text{nxt}, I\}}, \text{unit}, \sigma, \mu)$, where

$$\begin{aligned}
 w \in \sigma & \iff w = \text{priv} \rightarrow [h \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{tlock} \rightarrow [(t_S \cup \{n_1\}, a_S) \mid \text{own} \rightarrow \langle n_1, b \rangle \cup \text{nxt} \rightarrow n_2 \mid (t_O, a_O)] \wedge \\
 & \quad I(f(a_S) \bullet a_O) h \wedge f \text{ is local} \\
 (w, w', \text{res}) \in \mu & \iff w = \text{priv} \rightarrow [h \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{tlock} \rightarrow [(t_S \cup \{n_1\}, a_S) \mid \text{own} \rightarrow \langle n_1, b \rangle \cup \text{nxt} \rightarrow n_2 \mid (t_O, a_O)] \wedge \\
 & \quad I(f(a_S) \bullet a_O) h \wedge f \text{ is local} \wedge \\
 & \quad w' = \text{priv} \rightarrow [\text{empty} \mid \text{empty} \mid h_O] \cup \\
 & \quad \text{tlock} \rightarrow [(t_S, a_S) \mid \text{own} \rightarrow \langle n_1 + 1, \text{false} \rangle \cup \text{nxt} \rightarrow n_2 \cup h \mid (t_O, a_O)]
 \end{aligned}$$

In the definition of droptkt , the *self*-contribution with respect to the abstract state a_S is captured via some *auxiliary function* f , which is required to be *local* [10]. Another important change is removing a ticket n_1 from the *self*-set t_S , once the service is finished (*i.e.*, the lock is unlocked).

E.5 Lock predicates

In order to verify the `lock` and `unlock` procedures, we first introduce a number of auxiliary predicates to simplify the presentation of the reasoning.

Definition 14 ($T_{\{\text{tlock}, \text{own}, \text{nxt}, I\}}$ predicates).

$$\begin{aligned}
 \text{isLock}_{\text{tlock}}(a_S) & \hat{=} \text{tlock} \xrightarrow{s} (t_S, a_S) \\
 \text{canTry}_{\text{tlock}}(n, a_S) & \hat{=} \text{tlock} \xrightarrow{s} (t_S \cup \{n\}, a_S) \wedge \text{tlock} \xrightarrow{j} (\text{own} \rightarrow \langle n_1, b \rangle \cup h) \wedge \\
 & \quad (n = n_1) \implies \neg b \\
 \text{Locked}_{\text{tlock}}(a_S, a_O) & \hat{=} \text{tlock} \xrightarrow{s} (t_S \cup \{n_1\}, a_S) \wedge \text{tlock} \xrightarrow{j} (\text{own} \rightarrow \langle n_1, \text{true} \rangle \cup h)
 \end{aligned}$$

In Definition 14 and further we implicitly treat all non-bound logical variables as existentially quantified. The predicate $\text{isLock}_{\text{tlock}}$ identifies a state w as one satisfying to the structure of the $T_{\{\text{tlock}, \text{own}, \text{nxt}, I\}}$ concurroid. The predicate $\text{canTry}_{\text{tlock}}$ asserts that the current thread can make an attempt to acquire the lock with the ticket n , but it might fail. The auxiliary condition $(n = n_1) \implies \neg b$ is essential for verifying the body of `lock` in order to exclude the bogus situation when the same thread tries to re-acquire the lock which it already owns.

Curiously, the $\text{isLock}_{\text{tlock}}$ and $\text{Locked}_{\text{tlock}}$ predicates can be considered as implementation of the *abstract* interface for locks, as presented by the work on CAP [4]: we deliberately picked the same name for ours. In particular, we can prove the similar properties for ours predicates with respect to the locking intuition:

Lemma 3. $\text{Locked}_{\text{tlock}}(a_S, a_O) \implies \text{isLock}_{\text{tlock}}(a_S)$.

Lemma 4 (Mutual exclusion for Ticketed Lock).

$$w \models \text{Locked}_{\text{tlock}}(a_S, a_O) \wedge w^\top \models \text{Locked}_{\text{tlock}}(a_O, a_S) \implies \text{False}.$$

Lemma 4 is a “subjective” formulation of the mutual exclusion property, stating that the lock can be simultaneously acquired by at most one thread. We address the interested reader to our formal development in Coq [11] for the proofs of these and some other lemmas about locks as well as for implementation of the same abstract interface for the CAS-based spin lock.

E.6 Verifying locking procedures

We culminate the development of this appendix with the verification of the `lock` procedure for the ticketed lock (Section E.1).⁸

$$\frac{\left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{isLock}_{\text{tlock}}(a_S) \right\}}{\text{lock}} \left\{ \exists h, a_O, \text{priv} \xrightarrow{s} h * \text{Locked}_{\text{tlock}}(a_S, a_O) \wedge I(a_S \bullet a_O) h \right\} \quad (6)$$

Since in our logic imperative `while`-loops are implemented as a fixed point iteration of a tail-recursive function, for such functions a specification should be defined by the programmer (same is true for loop invariants in the standard Hoare logic). We begin by assuming for each auxiliary element of PCM a_S and a ticket number n a family of carried functions $\text{loop}(n : \text{nat}) : \text{unit} \rightarrow \text{unit}$ with the following FCSL specification:

$$\frac{\left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{canTry}_{\text{tlock}}(n, a_S) \right\}}{\text{loop}(n)(-)} \left\{ \exists h, a_O, \text{priv} \xrightarrow{s} h * \text{Locked}_{\text{tlock}}(a_S, a_O) \wedge I(a_S \bullet a_O) h \right\} \quad (7)$$

We proceed by unfolding the definition of the locking procedure and presenting the proof outline.

⁸ Verification of `unlock` is significantly simpler, and we address the reader to our formal development for the details [11].

$$\begin{aligned}
 & \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{isLock}_{\text{tlock}}(\mathbf{a}_S) \right\} \\
 & n \leftarrow \text{inject}'(\text{take}); \\
 & \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{tlock} \xrightarrow{s} (\mathbf{t}_S \cup \{n_2\}, \mathbf{a}_S) \wedge \text{tlock} \xrightarrow{j} (\text{own} \rightarrow - \cup \text{next} \rightarrow n'_2 \cup h') \wedge \right. \\
 & \left. \left\{ n = n_2 \wedge n_2 \leq n'_2 \right\} \right\} \\
 & \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{canTry}_{\text{tlock}}(n, \mathbf{a}_S) \right\} \\
 & \text{fix.} \text{loop}(n). _ . (\\
 & \text{if } (\text{try}(n)) \text{ then} \\
 & \quad \left\{ \text{priv} \xrightarrow{s} h * \text{tlock} \rightarrow [(\mathbf{t}_S, \mathbf{a}_S) \mid \text{own} \rightarrow \langle n, \text{true} \rangle \cup \text{next} \rightarrow n''_2 \mid (\mathbf{t}'_O, \mathbf{a}_O)] \wedge I(\mathbf{a}_S \bullet \mathbf{a}_O) h \right\} \\
 & \quad \text{return } (); \\
 & \quad \left\{ \text{priv} \xrightarrow{s} h * \text{Locked}_{\text{tlock}}(\mathbf{a}_S, \mathbf{a}_O) \wedge I(\mathbf{a}_S \bullet \mathbf{a}_O) h \right\} \\
 & \text{else} \\
 & \quad \left\{ \text{priv} \xrightarrow{s} \text{empty} * \text{canTry}_{\text{tlock}}(n, \mathbf{a}_S) \right\} \\
 & \quad \text{loop}(n)(_); \\
 & \quad \left\{ \exists h, \mathbf{a}_O, \text{priv} \xrightarrow{s} h * \text{Locked}_{\text{tlock}}(\mathbf{a}_S, \mathbf{a}_O) \wedge I(\mathbf{a}_S \bullet \mathbf{a}_O) h \right\} \\
 & _)(_) \\
 & \left\{ \exists h, \mathbf{a}_O, \text{priv} \xrightarrow{s} h * \text{Locked}_{\text{tlock}}(\mathbf{a}_S, \mathbf{a}_O) \wedge I(\mathbf{a}_S \bullet \mathbf{a}_O) h \right\}
 \end{aligned}$$

In fact, the proof has been implicitly split into two parts: (a) proving the correctness of the *loop* specification (7) in the body of the locking procedure by the rule **FIX** (Figure 2) and (b) proving the specification of the locking itself, assuming the *loop* spec. In the outline we have blended them both together, as the proof of (7) is straightforward: just by reading the specification of the *try* action, we know that in *then*-branch the heap *h* has been transferred to *self*-part of *priv*. Moreover, in the proof outline we omitted some reasoning about stability, which, indeed, has been taken into account. In particular, in the assertion right after the line “if (try(*n*)) then”, we mention n''_2 rather than n'_2 , since this counter could have been changed by the *other* threads. These subtle details become clear in the mechanized proofs and are omitted above for the sake of brevity.

Finally, we have employed a slightly different version of injection, which we emphasized by using the *inject'* command. It differs from one we described in Section 3 by the form of entanglement it handles. The original rule **INJECT** enabled concurroid framing with respect to the right operand of \bowtie . In contrast, the following rule enables concurroid framing with respect to the left operand of \bowtie , or right operand of the “reversed” operator \bowtie :

$$\frac{\{p\} c \{q\} @ U \quad r \text{ stable under } V}{\{p * r\} \text{inject}' c \{q * r\} @ U \bowtie V} \text{INJECT}'$$

The rule was, indeed, proved sound as well [11]. In general, we have actually proved the injection rule sound for any operator satisfying a number of general properties, and then shown that \bowtie and \bowtie (as well as some other combinators), satisfy these general properties.

F Semantics and Soundness

The semantic model for FCSL programs largely relies on the denotational semantic of *action trees* [10], which implement finite, partial approximations of the behavior of FCSL commands. Action trees are a generalization of the Brookes’ action traces in the following sense. Where action trace semantics approximate a program by a set of traces, we approximate with a set of trees. A tree differs from a trace in the following: a trace is a sequence of actions and their results, whereas a tree contains an action followed by a *continuation* which itself is a tree parametrized wrt. the output of the action. In this appendix we introduce a number of notions in order to formalize the safety of FCSL.

The high-level view of the development is as follows. We first define semantic behavior (*i.e.*, operational semantics) for action trees wrt. a program state (Section F.1). We consider program state that, in addition to concrete heaps includes the logical information such as auxiliary state, or division of state between concurroid labels. Thus, the operational semantics of trees is *instrumented*. In Section F.4, we prove the standard erasure result of shared-memory concurrency logics, whereby logical information in program states doesn’t influence the result of the computation.

Taking the low-level operational semantics of trees from Section F.1 as a base for our soundness result, we relate it to the high-level transitions of a concurroid by an *always* predicate (Section F.5) that ensures a tree is resilient to any amount of a concurroid’s rely-interference, and that all operational steps by a tree are *safe* (*i.e.*, providing an analogue to *preservation* property for the *action safety* predicate) and correspond to concurroid interference.

The denotational semantics (Section F.6) interprets Hoare tuple judgments by the *monadic Hoare type* $\llbracket p \rrbracket A \llbracket q \rrbracket @ U$, which is a *complete lattice* of trees that are *always-safe* to run from any initial configuration that satisfies precondition p and if they terminate produce a final configuration that satisfies postcondition q (under possible interference from programs respecting the transitions of the concurroid U). The complete lattice structure makes the semantic domain suitable for modeling recursion. A command is denoted by the set of its tree approximations, and a procedure is denoted by a function into a set of trees. The *soundness* of FCSL (Section F.7) follows from showing that the denotations of the commands listed in Section 5 satisfy the appropriate *always* predicate (*i.e.*, adhere to the *progress-and-preservation* property), and that *always* satisfies certain closure conditions.

We choose the Calculus of Inductive Constructions (CiC) [1, 16] as our meta logic. This has several important benefits. First, we can define a *shallow embedding* of FCSL into CiC that allows us to program and prove directly *with the semantic objects*, thus immediately lifting FCSL to a full-blown programming language and verification system with higher-order functions, abstract types, abstract predicates, and a module system. We also gain a powerful dependently-typed λ -calculus, which we use to formalize all semantic definitions and metatheory, including the definition of action trees by *iterated inductive definitions* [16], specification-level functions, and programming-level higher-order procedures. Finally, we were able to mechanize the entire semantics and metatheory [11] in the Coq proof assistant implementation of CiC.

Following [10], we use set-theoretic and type-theoretic notation as appropriate. The reader unconcerned with the fine points of the type theory, may read a typing judgment $x : A$ as a set membership predicate $x \in A$.

F.1 Action trees and their operational semantics

The type family $\text{tree } U A$ of A -returning action trees is defined by an iterated inductive definition, as follows.

Definition 15 (Action trees).

Unfinished	: $\text{tree } U A$
$\text{Ret } (v : A)$: $\text{tree } U A$
$\text{Act } (a : (U, A, \sigma, \mu))$: $\text{tree } U A$
$\text{Seq } (t : \text{tree } U B) (k : B \rightarrow \text{tree } U A)$: $\text{tree } U A$
$\text{Par } (t_1 : \text{tree } U B_1) (t_2 : \text{tree } U B_2) (k : B_1 \times B_2 \rightarrow \text{tree } U A)$: $\text{tree } U A$
$\text{Inject } (t : \text{tree } V A)$: $\text{tree } U A$
$\text{Hide}_{\widehat{\Phi}, g} (t : \text{tree } V A)$: $\text{tree } U A$

Most of the constructors in Definition 15 are self-explanatory. Since trees have finite depth, they can only *approximate* divergent computations, thus the Unfinished tree indicates an incomplete approximation. $\text{Ret } v$ is a terminal computation that returns value $v : A$. The constructor Act takes as a parameter an action (U, A, σ, μ) , defined in concurroid U (Section 5). $\text{Seq } t k$ sequentially composes a B -returning tree t with a continuation k that takes t 's return value and generates the rest of the approximation. $\text{Par } t_1 t_2 k$ is the parallel composition of trees t_1 and t_2 , and a continuation k that takes the pair of their results when they join. CiC's iterated inductive definition permits the recursive occurrences of tree to be *nonuniform* (e.g., $\text{tree } B_i$ in Par) and *nested* (e.g., the *positive* occurrence of $\text{tree } A$ in the continuation). Since the CiC function space \rightarrow includes case-analysis, the continuation may branch upon the argument, which captures the pure computation of conditionals. This closely corresponds to the operational intuition and leads to a straightforward denotational semantics.

The Inject constructor embeds a tree $t : \text{tree } V A$ (of a *different*, i.e., smaller, concurroid V) for the underlying computation and generates a tree in the concurroid U . V and U are not arbitrary, but it has to be possible to *inject* V into U , as we shall define shortly. In particular, we will be able to prove that we can inject V into $U = V \times W$ for any W .

The Hide constructor embeds a *generalized refinement function* $\widehat{\Phi}$, an abstract state g and a tree $t : \text{tree } V A$ for the underlying computation and generates a tree in the concurroid U . Again, U and V are not arbitrary but it has to be possible to *refine* V into U . In particular, we will be able to prove that we can refine $U = P \times V_1 \times V_2$ into $V = P \times V_1$.

In general, in Coq, the proof that V injects into U , or that V refines U , has to be an annotation on the Inject and Hide constructors, but we omit the annotations for brevity.

F.2 Injection and Refinement

We define the predicate $\text{injects } V \ U$, which intuitively means that the “larger” concurroid U can be considered as an entanglement of the “smaller” concurroid V with some additional concurroid W .

Definition 16 (Injection predicate).

$\text{injects } V \ U \iff$ *there exists W , such that the following three statements hold:*

1. $w \in \mathcal{W}_U \iff w = w_1 \cup w_2 \wedge w_1 \in \mathcal{W}_V \wedge w_2 \in \mathcal{W}_W$;
2. $\forall w_1 \ w_2 \ w. w_1 \cup w_2 \in \mathcal{W}_U \wedge w_2 \cup w \in \mathcal{W}_U \wedge (w_1, w_2) \in \tau_V \implies (w_1 \cup w, w_2 \cup w) \in \tau_U$;
3. $\forall w_1 \ w'_1 \ w_2 \ w'_2.$

$$s_1 \in \mathcal{W}_V \wedge s_2 \in \mathcal{W}_V \wedge (s_1 \cup s'_1, s_2 \cup s'_2) \in \mathcal{G}_U \implies \\ (s_1, s_2) \in \mathcal{G}_V \wedge (s'_1, s'_2) \in \mathcal{G}_W, \quad \text{where}$$

\mathcal{G}_X is a guarantee relation of a concurroid X (Definition 4).

The first requirement ensures that the states of U can be constructed as Cartesian products of states of V and W . The second requirement ensures that internal transitions of the smaller concurroid V do not break the coherence of it enclosing concurroid U . The third requirement ensures that the transitions within U can be always uniquely split to the transitions done by V and by the “addition” W . Notably, the first requirement ensures that if $\text{injects } V \ U$ holds, there always exists a “smaller” state w_V , such that $w_V \in \mathcal{W}_V$ for any $w_U \in \mathcal{W}_U$.

We will also introduce another auxiliary definition in order to have a handle to the existential concurroid W in Definition 16.

Definition 17. *We will write $U = V \blacktriangleright W$, when $\text{injects } V \ U$ and W is an additional concurroid from Definition 16.*

The notation \blacktriangleright hints that already familiar entanglement operators \bowtie and \bowtie can be seen as particular cases of \blacktriangleright .

Lemma 5. $\text{injects } V \ (V \bowtie W)$ for every V and W with disjoint label sets.

We next define the refinement relation on concurroids. The concurroid U is refined into V , if we can, intuitively, elaborate the states of U into those of V . In other words, if V 's states can be seen to contain the states of U , plus some other additional state. Abstractly, we capture the dependence by positing an elaboration predicate $\widehat{\Phi}$ between the state spaces of U and V , with a number of properties shown below. Additionally, $\widehat{\Phi}$ takes a value $g : G$ of user-specified type, which is an abstraction of the mentioned additional state. Thus, we read $\widehat{\Phi}(g)(w, w')$ to say that g elaborates w into w' (via $\widehat{\Phi}$).

Definition 18 (Elaboration predicate).

Given a type G of abstract states and a function $\widehat{\Phi} : G \rightarrow (\mathcal{W}_U, \mathcal{W}_V) \rightarrow \text{prop}$, refines $U \ V \ \widehat{\Phi} \iff$ the following statements hold:

1. $\widehat{\Phi}(g)(w, w_1) \wedge \widehat{\Phi}(g)(w, w_2) \implies w_1 = w_2$;

2. $\widehat{\Phi}(g_1)(w_1, w) \wedge \widehat{\Phi}(g_1)(w_2, w) \implies g_1 = g_2 \wedge w_1 = w_2$;
3. $\widehat{\Phi}(g)(w, w') \implies w \in \mathcal{W}_U \wedge w' \in \mathcal{W}_V$;
4. $\widehat{\Phi}(g)(w_1, w'_1) \wedge (w_1, w_2) \in \mathcal{R}_U \implies \exists w'_2, \widehat{\Phi}(g)(w_2, w'_2) \wedge (w'_1, w'_2) \in \mathcal{R}_V$;
5. $\widehat{\Phi}(g_1)(w_1, w'_1) \wedge (w'_1, w'_2) \in \tau_V \implies \exists g_2, w_2. \widehat{\Phi}(g_2)(w_2, w'_2) \wedge (w_1, w_2) \in \tau_U$;
6. $\widehat{\Phi}(g_1)(w_1, w'_1) \wedge \widehat{\Phi}(g_2)(w_2, w'_2) \implies (w_1 \cdot o = w_1 \cdot o \iff w'_1 \cdot o = w'_2 \cdot o)$;
7. $\forall p. \exists q.$

$$\forall g, w, w'. \widehat{\Phi}(g)(w \triangleright p, w') \implies \left(\exists w'_1, w' = w'_1 \triangleright q \wedge \widehat{\Phi}(g)(w \triangleleft p, w'_1 \triangleleft q) \right) \wedge$$

$$\forall g, w, w'. \widehat{\Phi}(g)(w, w' \triangleleft q) \implies \left(\exists w_1, w = w_1 \triangleleft p \wedge \widehat{\Phi}(g)(w_1 \triangleright p, w' \triangleright q) \right);$$
8. $\widehat{\Phi}(g)(w, w') \implies \lfloor w \rfloor = \lfloor w' \rfloor$;

If the properties in Definition 18 are satisfied, we call $\widehat{\Phi}$ and elaboration predicate for U and V . In Definition 18, (1) states that $\widehat{\Phi}$ uniquely determines refined states (however, some states in U need not refine into anything in V , *i.e.*, $\widehat{\Phi}$ is a partial function); (2) states that the $\widehat{\Phi}$ is injective: abstracting back from a refinement is unique. Thus, (1) and (2) state that $\widehat{\Phi}$ is a partial bijection. (3) ensures that the elaboration maps well-formed states to well-formed states; (4) specifies that environment steps in the “coarse” world do not change introduced abstractions in the “fine” world, as environment steps in the coarse world can’t see the refinement; (5) states that “abstracting back” preserves internal transitions. The properties (4) and (5) can be seen as stating commutativity of the apparent categorical diagrams. They may also be seen as stating simulation properties. In (4), a rely transitions of U can always be matched by a rely transition of V . In (5), an internal transition of V may always be matched by an internal transition of U . (6) postulates that refinement realigns *self* and *joint* parts, but doesn’t move things into/out of *other*; (7) is a version of framing, namely, every every extension to *other* of a coarse state is uniquely refined into an addition to fine state; finally, (8) ensures that refinement only deals with abstract parts of the state and does not change the heap.

Lemma 6 (Hiding and refinement).

Let U and V be concurroids, and let Φ be an abstraction function from a HIDE; that is, Φ satisfies the properties about concurroid V from Section C. Let $\widehat{\Phi}$ be constructed as follows.

$$\begin{aligned} \widehat{\Phi}(g)(w, w') \iff \exists h, w_1, w_2, w &= (\text{priv} \rightarrow (h \cup \lfloor w_1 \rfloor)) \cup w_2 \wedge \\ w' &= (\text{priv} \rightarrow h) \cup w_1 \cup w_2 \wedge \\ w_1 &\in \Phi(g) \wedge \\ w &\in \mathcal{W}_{P \times U} \wedge w' \in \mathcal{W}_{P \times U \times V} \end{aligned}$$

Then $\widehat{\Phi}$ is a valid elaboration predicate, that is refines $(P \times U) (P \times U \times V) \widehat{\Phi}$.

F.3 Operational semantics of action trees

The judgment for small-step operational semantics of action trees has the form $w; t \xrightarrow{\pi}_t w', t'$ (Figure 3). It operates on program states w and paths π to step the tree t from initial state

Fig. 3 Tree stepping $w; t \xrightarrow{\pi}_t w'; t'$ with respect to the path π

$$\begin{array}{c}
\frac{(w, w', v) \in \mu}{w; \text{Act } (U, A, \sigma, \mu) \xrightarrow{\text{ChoiceAct}}_t w'; \text{Ret } v} \quad \frac{}{w; \text{Par } (\text{Ret } v_1) (\text{Ret } v_2) k \xrightarrow{\text{ParRet}}_t w; k (v_1, v_2)} \\
\frac{w; t_1 \xrightarrow{\pi}_t w'; t'_1}{w; \text{Par } t_1 t_2 k \xrightarrow{\text{ParL } \pi}_t w'; \text{Par } t'_1 t_2 k} \quad \frac{w; t_2 \xrightarrow{\pi}_t w'; t'_2}{w; \text{Par } t_1 t_2 k \xrightarrow{\text{ParR } \pi}_t w'; \text{Par } t_1 t'_2 k} \\
\frac{}{w; \text{Seq } (\text{Ret } v) k \xrightarrow{\text{SeqRet}}_t w; (k v)} \quad \frac{w; t_1 \xrightarrow{\pi}_t w'; t_2}{w; \text{Seq } t_1 k \xrightarrow{\text{SeqStep } \pi}_t w'; \text{Seq } t_2 k} \\
\frac{}{w; \text{Inject } (\text{Ret } v) \xrightarrow{\text{InjRet}}_t w; \text{Ret } v} \\
\frac{\text{injects } V U \quad w_U = w_V \cup w_0 \quad w_V; t_1 \xrightarrow{\pi}_t w'_V; t_2 \quad w'_U = w'_V \cup w_0}{w_U; \text{Inject } t_1 \xrightarrow{\text{InjStep } \pi}_t w'_U; \text{Inject } t_2} \\
\frac{}{w; \text{Hide } \widehat{\Phi} g (\text{Ret } v) \xrightarrow{\text{HideRet}}_t w; \text{Ret } v} \\
\frac{\text{hides } U V \widehat{\Phi} \quad \widehat{\Phi}(g)(w_U, w_V) \quad w_V; t_1 \xrightarrow{\pi}_t w'_V; t_2 \quad \widehat{\Phi}(g')(w'_U, w'_V)}{w_U; \text{Hide } \widehat{\Phi} g t_1 \xrightarrow{\text{HideStep } \pi}_t w'_U; \text{Hide } \widehat{\Phi} g' t_2}
\end{array}$$

w to a reduced tree t' in ending state w' . Intuitively, the path π identifies the position in the tree to be reduced, which may be an action, or a beta redex.

$$\begin{array}{l}
\pi ::= \text{ChoiceAct} \quad | \quad \text{SeqRet} \quad | \quad \text{SeqStep } \pi \quad | \\
\text{ParRet} \quad \quad | \quad \text{ParL } \pi \quad | \quad \text{ParR } \pi \quad \quad | \\
\text{HideStep } \pi \quad | \quad \text{HideRet} \quad | \quad \text{InjStep } \pi \quad | \quad \text{InjRet}.
\end{array} \tag{8}$$

Stepping is undefined for the Unfinished and Ret trees. For the Seq and Par trees, the path π selects a β -redex and performs the appropriate reduction. For the Inject tree in the non-Ret case the reduction of the tree t_1 is performed for a smaller state w_V instead of the large state w_U , and the resulted state w'_V is “plugged” back using the addition w_0 in order to obtain the resulting large state w'_U . The situation for Hide is symmetric: the coarse-grained state w_U and the fine-grained state w_V are related by the elaboration $\widehat{\Phi}$. Note that the modifications of the changes to the abstract values g are recorded into g' .

The decorated semantics in Figure 3 may not make a step for two different reasons. The first, benign, reason is that the chosen path π doesn't actually determine an action or a redex in the tree t . For example, we may have $t = \text{Unfinished}$ and $\pi = \text{ParR}$. But we can choose the right side of a parallel composition only in a tree whose head

constructor is `Par`, which is not the case with `Unfinished`. We consider such paths that don't determine an action or a redex in a tree ill-formed.

The second reason arises when π is actually well-formed. In that case, the constructors of the path uniquely determine a number of rules of the operational semantics that should be applied to step the tree. However, what is some of the rules can't be applied because their side-conditions are not satisfied by the current state.

We next state a form of the progress property, showing that the latter case can't occur in trees that approximate well-proved programs. We first formally define paths which are well-formed wrt. a tree.

Definition 19 (Good paths). *Let $t : \text{tree } U \ A$ and π be a path. Then the predicate $\text{good } t \ \pi$ is defined as follows:*

$\text{good } (\text{Act } (U, A, \sigma, \mu))$	<code>ChoiceAct</code>	$\hat{=}$	<code>true</code>
$\text{good } (\text{Seq } (\text{Ret } v) _)$	<code>SeqRet</code>	$\hat{=}$	<code>true</code>
$\text{good } (\text{Seq } t _)$	<code>SeqRet</code>	$\hat{=}$	$\text{good } t \ \pi$
$\text{good } (\text{Par } (\text{Ret } _) (\text{Ret } _) _)$	<code>ParRet</code>	$\hat{=}$	<code>true</code>
$\text{good } (\text{Par } t_1 \ t_2 \ _)$	<code>ParL</code>	$\hat{=}$	$\text{good } t_1 \ \pi$
$\text{good } (\text{Par } t_1 \ t_2 \ _)$	<code>ParR</code>	$\hat{=}$	$\text{good } t_2 \ \pi$
$\text{good } (\text{Inject } (\text{Ret } _))$	<code>InjRet</code>	$\hat{=}$	<code>true</code>
$\text{good } (\text{Inject } t)$	<code>InjStep</code>	$\hat{=}$	$\text{good } t \ \pi$
$\text{good } (\text{Hide } _ _ (\text{Ret } _))$	<code>HideRet</code>	$\hat{=}$	<code>true</code>
$\text{good } (\text{Hide } _ _ t)$	<code>HideStep</code>	$\hat{=}$	$\text{good } t \ \pi$
$\text{good } t$	π	$\hat{=}$	<code>false</code> <i>otherwise</i>

We next require a notion of safety for a tree and a path.

Definition 20 (Safe path for an action tree). *Let $t : \text{tree } U \ A$, π be a path π , and state $w \in \mathcal{W}_U$. $\text{safe } t \ \pi \ w$ is defined as follows:*

$\text{safe } (\text{Act } (U, A, \sigma, \mu))$	<code>ChoiceAct</code>	$w \hat{=}$	$(w \in \sigma)$
$\text{safe } (\text{Seq } (\text{Ret } v) _)$	<code>SeqRet</code>	$w \hat{=}$	<code>true</code>
$\text{safe } (\text{Seq } t _)$	<code>SeqStep</code>	$\pi \ w \hat{=}$	$\text{safe } t \ \pi \ w$
$\text{safe } (\text{Par } (\text{Ret } _) (\text{Ret } _) _)$	<code>ParRet</code>	$w \hat{=}$	<code>true</code>
$\text{safe } (\text{Par } t_1 \ t_2 \ _)$	<code>ParL</code>	$\pi \ w \hat{=}$	$\text{safe } t_1 \ \pi \ w$
$\text{safe } (\text{Par } t_1 \ t_2 \ _)$	<code>ParR</code>	$\pi \ w \hat{=}$	$\text{safe } t_2 \ \pi \ w$
$\text{safe } (\text{Inject } (\text{Ret } _))$	<code>InjRet</code>	$w \hat{=}$	<code>true</code>
$\text{safe } (\text{Inject } (t : \text{tree } V \ A))$	<code>InjRet</code>	$\pi \ w \hat{=}$	$\text{injects } V \ U \wedge w = w_V \cup w_0 \wedge \text{safe } t \ \pi \ w_V$
$\text{safe } (\text{Hide } _ _ (\text{Ret } _))$	<code>HideRet</code>	$w \hat{=}$	<code>true</code>
$\text{safe } (\text{Hide } \Phi \ g \ (t : \text{tree } V \ A))$	<code>HideStep</code>	$\pi \ w \hat{=}$	$\text{hides } U \ V \ \Phi \wedge \Phi(g)(w, w_V) \wedge \text{safe } t \ \pi \ w_V$
$\text{safe } t$	π	$w \hat{=}$	<code>true</code> <i>otherwise</i>

The definition traverses the path and the tree, and in those cases when the tree and the path match, collects the side-conditions on the *initial states* of the operational semantics rules.

Lemma 7 (Progress on action trees). *For any concurroid U , state $w \in \mathcal{W}_U$, action tree $t : \text{tree } U \ A$ and a path π , if safe $t \ \pi \ w$, then either of the following holds*

1. $\neg \text{good } t \ \pi$;
2. $\exists w' \ t', w; t \xrightarrow{\pi}_t w'; t'$.

The lemma shows that if we can match the tree t with a path π and thus manage to pick an action or a redex (*i.e.*, $\text{good } t \ \pi$), then the safety predicate implies the side conditions of the inference rule determined by π .

The import of the lemma is in allowing that safety and stepping relations may be defined independently of each other. The difficulty of the proof consists in showing that the various conditions from the reduction rules, that have *not* been incorporated into the safety predicate, are nevertheless satisfied. For example, the reduction rules for `Inject` and `Hide` contain such side conditions on the *ending* states of their premisses. We have proved that even though these conditions are not incorporated by the **safe** predicate, they are implied by it. The proof relies on the general properties of injection and refinement listed in Section F.2.

Defining safety independently from stepping will enable us to develop a semantics in the style of Milner and Wright and Felleisen, whereby we only give meaning to programs that are well-proved in FCSL, analogously to how Milner and Wright-Felleisen semantics give meaning to only well-typed programs. As the type of the program, we will use its Hoare-tuple specification, and the denotation of a program wrt. a specification with precondition p , will only contain approximating action trees whose safety predicate is implied by p . The fact that the safety predicate actually means safety, *i.e.*, it allows stepping, is a consequence of the Progress lemma above.

F.4 Erased FCSL action trees

We next define an erased version of action trees that remove the various logical annotations. One example of such erased logical information is the proofs of injection and refinement that we implicitly elided from the `INJECT` and `HIDE` constructors. Another example is the logical information from the actions embedded into action trees, which, as we have shown, operate on *decorated* states.

We then prove that such logical information doesn't influence the stepping of the tree. In other words, states that erase to equal concrete heaps, produce output states that also erase to equal concrete heaps.

Definition 21 (Tree erasure).

Given a tree $t : \text{tree } U \ A$, its erasure $[t] : \text{etree } A$ is defined as follows:

$[Unfinished]$	\equiv	<code>Unfinished</code>
$[Ret \ v]$	\equiv	<code>Ret \ v</code>
$[Act \ (U, A, \sigma, \mu)]$	\equiv	<code>Act \ (A, [\sigma], [\mu])</code>
$[Seq \ t \ k]$	\equiv	<code>Seq \ [t] \ (\lambda v. [k \ v])</code>
$[Par \ t_1 \ t_2 \ k]$	\equiv	<code>Par \ [t_1] \ [t_2] \ (\lambda v. [k \ v])</code>
$[Inject \ t]$	\equiv	<code>Inject \ [t]</code>
$[Hide \ _ _ \ t]$	\equiv	<code>Hide \ [t]</code>

Theorem 2 (Erasure irrelevance). *For a concurroid U , states $w_1, w_2 \in \mathcal{W}_U$ and action trees $t_1, t_2 : \text{tree } U \ A$, if $\llbracket w_1 \rrbracket = \llbracket w_2 \rrbracket$, $\llbracket t_1 \rrbracket = \llbracket t_2 \rrbracket$, $w_1; t_1 \xrightarrow{\pi} w'_1; t'_1$ and $w_2; t_2 \xrightarrow{\pi} w'_2; t'_2$, then $\llbracket w'_1 \rrbracket = \llbracket w'_2 \rrbracket$ and $\llbracket t'_1 \rrbracket = \llbracket t'_2 \rrbracket$.*

Ideally, we will need to prove a stronger result, stating that if the tree and state evaluate under semantics from Figure 3, so do their erased counterparts under some sort of erased semantics. However, we have not carried out this exercise due to time constraints and decided to leave it for the future work.

F.5 Auxiliary modal predicates

In this section we next define a number of *modal* predicates over *all* possible steps of execution to relate the operational semantics over heaps and concurroid transitions over admissible states.

Definition 22 (Modal predicates). *The predicates $\text{always}_U^\zeta w \ t \ P$, $\text{always}_U w \ t \ P$ and $\text{after}_U w \ t \ Q$ are defined relative to a schedule ζ , concurroid U , state w , A -returning rich tree $t : \text{tree } U \ A$ and arbitrary predicates $P : \text{state} \rightarrow \text{tree } U \ A \rightarrow \text{prop}$ and $Q : A \rightarrow \text{state} \rightarrow \text{prop}$:*

$$\begin{aligned} \text{always}_U^\zeta w \ t \ P &\hat{=} \text{if } \zeta = \pi :: \zeta' \\ &\quad \text{then } \forall w_2. (w, w_2) \in \mathcal{R}_U \implies \\ &\quad \quad \text{safe } t \ \pi \ w_2 \ \wedge \\ &\quad \quad \left(\forall w_3 \ t_2. w_2; t \xrightarrow{\pi} w_3; t_2 \implies \text{always}_U^{\zeta'} w_3 \ t_2 \ P \right) \\ &\quad \text{else } \forall w_2. (w, w_2) \in \mathcal{R}_U \implies P \ w_2 \ t \\ \\ \text{always}_U w \ t \ P &\hat{=} \forall \zeta. \text{always}_U^\zeta w \ t \ P \\ \\ \text{after}_U w \ t \ Q &\hat{=} \text{always}_U w \ t \ (\lambda w' \ t. \forall v'. (t' = \text{Ret } v') \implies Q \ v' \ w') \end{aligned}$$

The predicate $\text{always}_U w \ t \ P$ expresses the fact that starting from the state w , the tree t remains memory-safe and the user-chosen predicate P holds of all intermediate configurations and trees, for any schedule ζ and under any environment of the concurroid U . The helper predicate $\text{always}_U^\zeta w \ t \ P$ is defined by induction on ζ : the concurroid U is allowed to make arbitrary rely-transitions from w to w_2 , in the resulting configuration the predicate $P \ w_2 \ t$ holds; moreover, if the schedule is $\pi :: \zeta'$, then the resulting configuration must have a state that is *safe* for t , π and w_2 (so, according to Theorem 7, it can step), and, if t steps to w_3 and t_2 , then the predicate recurses on ζ' , w_3 and t_2 . One can notice that if the predicate $\text{always}_U^\zeta w \ t \ P$, it automatically implies “preservation” of the safe predicate at each step of reduction of the tree t .

The predicate $\text{after}_U w \ t \ Q$ encodes that t is *safe*; however, $Q \ v' \ w'$ only holds if t steps completely to $\text{Ret } v'$ in configuration w' .

Fig. 4 Denotational semantics of FCSL judgments.

$$\begin{aligned}
\llbracket \cdot \rrbracket &\hat{=} \cdot \\
\llbracket T, x : A \rrbracket &\hat{=} \llbracket T \rrbracket, x : A \\
\llbracket T, \forall x : B. \{p\} f(x) : A \{q\} @ U \rrbracket &\hat{=} \llbracket T \rrbracket, f : \forall x : B. \{p\} A \{q\} @ U \\
\llbracket T \vdash \{p\} c : A \{q\} @ U \rrbracket &\hat{=} \llbracket T \rrbracket \vdash_{\text{CiC}} \llbracket c \rrbracket : \{p\} A \{q\} @ U \\
\llbracket T \vdash \forall x : B. \{p\} F(x) : A \{q\} @ U \rrbracket &\hat{=} \llbracket T \rrbracket \vdash_{\text{CiC}} \llbracket F \rrbracket : \forall x : B. \{p\} A \{q\} @ U \\
\llbracket T \vdash e : A \rrbracket &\hat{=} \llbracket T \rrbracket \vdash_{\text{CiC}} e : A \\
\llbracket T \vdash (p_1, q_1) \sqsubseteq (p_2, q_2) \rrbracket &\hat{=} \\
&\llbracket T \rrbracket \vdash_{\text{CiC}} \forall w w'. (w \models \exists \bar{v}_2. p_2 \implies w \models \exists \bar{v}_1. p_1) \wedge \\
&\quad ((\forall \bar{v}_1 \text{ res. } w \models p_1 \implies w' \models q_1) \implies (\forall \bar{v}_2 \text{ res. } w \models p_2 \implies w' \models q_2)) \\
&\text{where } \bar{v}_i \hat{=} \text{FLV}(p_i, q_i)
\end{aligned}$$

F.6 Denotational semantics of FCSL

In this section, we define the denotational semantics of FCSL as a shallow embedding into CiC. For the sake of simplicity, one can think of the embedding as of a macro-expansion, which compositionally replaces FCSL programs by expressions and values of CiC.

The semantics of commands and judgments are defined simultaneously. The mutual recursion is necessary because the denotation of judgments depends on the denotation of commands and procedures, while the denotation of a fixed point procedure depends on the denotation of its procedure triple to determine the lattice in which to take the fixed point.

We denote FCSL programs as *sets* T of trees of increasing precision including the Unfinished tree, which is the coarsest possible approximation of any program:

$$\text{prog } U A \hat{=} \{T \subseteq \mathcal{P}(\text{tree } U A) \mid \text{Unfinished} \in T\}.$$

To model recursion, we construct a complete lattice of Hoare types to get fixed points. We use the *after* predicate (Definition 22) to ensure the tree approximations are memory safe, respect mutual exclusion, and satisfy their FCSL specifications.

Definition 23 (FCSL Hoare types). For a concurroid U and type A , fix a precondition $p : \text{state} \rightarrow \text{prop}$ and a postcondition $q : A \rightarrow \text{state} \rightarrow \text{prop}$, with free logical variables $\text{FLV}(p, q)$. The FCSL Hoare type $\{p\} A \{q\} @ U$ is defined as follows:

$$\{p\} A \{q\} @ U \hat{=} \left\{ T \in \text{prog } U A \mid \begin{array}{l} \forall \text{FLV}(p, q) w (t \in T). \\ w \models p \wedge w \in \mathcal{W}_U \implies \text{after}_U w t q \end{array} \right\}$$

Intuitively, the denotation of a FCSL judgment $\{p\} c : A \{q\} @ U$ is the set of trees T denoting the command c , together with a proof that for any initial configuration w that satisfies the precondition p , then *after* executing any tree $t \in T$ from c produces some result value and final configuration that satisfy postcondition q . The definition quantifies

Fig. 5 Denotational semantics of FCSL commands.

$$\begin{aligned}
\llbracket \text{return } (v : A) \rrbracket &\triangleq \{\text{Unfinished}, \text{Ret } v\} \\
\llbracket x \leftarrow c_1; c_2 \rrbracket &\triangleq \{\text{Unfinished}\} \cup \{\text{Seq } t_1 k \mid t_1 \in \llbracket c_1 \rrbracket, \forall x. k x \in \llbracket c_2 \rrbracket\} \\
\llbracket c_1 \parallel c_2 \rrbracket &\triangleq \{\text{Unfinished}\} \cup \{\text{Par } t_1 t_2 (\lambda v. \text{Ret } v) \mid t_1 \in \llbracket c_1 \rrbracket, t_2 \in \llbracket c_2 \rrbracket\} \\
\llbracket \text{if } e \text{ then } c_1 \text{ else } c_2 \rrbracket &\triangleq \text{if } e \text{ then } \llbracket c_1 \rrbracket \text{ else } \llbracket c_2 \rrbracket \\
\llbracket \text{act } a \rrbracket &\triangleq \{\text{Unfinished}, \text{Act } a\} \\
\llbracket \text{inject } c \rrbracket &\triangleq \{\text{Unfinished}\} \cup \{\text{Inject } t \mid t \in \llbracket c \rrbracket\} \\
\llbracket \text{hide}_{\phi, g} c \rrbracket &\triangleq \{\text{Unfinished}\} \cup \{\text{Hide } \widehat{\Phi} g t \mid t \in \llbracket c \rrbracket\}, \text{ where} \\
&\quad \widehat{\Phi} \text{ is constructed from } \Phi \text{ as in Lemma 6} \\
\llbracket F(e) \rrbracket &\triangleq \llbracket F \rrbracket(e) \\
\llbracket \text{fix } f^{B.A.p.q}. x : A. c \rrbracket &\triangleq \text{lfp}_{\forall x.B. \llbracket p \rrbracket A \llbracket q \rrbracket @ U} (\lambda f. \lambda x. \llbracket c \rrbracket)
\end{aligned}$$

over the free logical variables of p and q in order to give these variables local scope, as stipulated in Section 5. FCSL assertions (Section 3) are arbitrary CiC predicates of type $\text{state} \rightarrow \text{prop}$; and we model the relation $w \models p$ of FCSL as the application $p w$ in CiC.

Lemma 8. *The type $\llbracket p \rrbracket A \llbracket q \rrbracket @ U$ is a complete lattice, with set union as the join operator, and $\{\text{Unfinished}\}$ as the unit element.*

The type $\forall x : B. \llbracket p \rrbracket A \llbracket q \rrbracket @ U$ of functions mapping $x : B$ into $\llbracket p \rrbracket A \llbracket q \rrbracket @ U$, where A, p, q may depend on x , is also a complete lattice, with the join operator on functions defined point-wise, and the constant $\{\text{Unfinished}\}$ function as the unit element.

The **denotation of judgments** $\llbracket \Gamma \vdash J \rrbracket$ (Figure 4) turns FCSL judgments into CiC typing judgments (\vdash_{CiC}). A command specification $\{p\} - : A \{q\} @ U$ is denoted by the CiC type $\llbracket p \rrbracket A \llbracket q \rrbracket @ U$, and a procedure specification $\forall x : B. \{p\} - : A \{q\} @ U$ is denoted by the CiC dependent function type $\forall x : B. \llbracket p \rrbracket A \llbracket q \rrbracket @ U$. The Hoare ordering $(p_1, q_1) \sqsubseteq (p_2, q_2)$ judgment is explained in Section D, so here we just literally repeat the definition given there.

The **denotation of commands and procedures** (Figure 5) is subsidiary to that of judgments because the fixed-point construction is indexed by the argument and return types, and the pre- and postconditions. An A -returning command c is denoted by a set of approximating trees in $\text{prog } U A$, and an A -returning procedure F with argument B is denoted by a set of trees in $B \rightarrow \text{prog } U A$.

In the semantic translation for the sequential composition $\llbracket x \leftarrow c_1; c_2 \rrbracket$, we employ the fact that the translation $\llbracket c_2 \rrbracket$ can produce *open* values (*i.e.*, trees t_2), where a variable x of CiC is not bound. We subsequently close them by constructing a continuation k as a CiC function, in an indirect way, that is, quantifying over all possible inputs, hence $\forall x, k x \in \llbracket c_2 \rrbracket$. For the denotational semantics of the hiding statement $\llbracket \text{hide}_{\phi, g} c \rrbracket$, we appeal to Lemma 6 (Section F.2), which give a constructive way to construct an elaboration predicate $\widehat{\Phi}$ from a given abstraction function Φ . Since all FCSL program constructors preserve *monotonicity*, the $\text{fix } f.x.c$ procedure can take the least fixed point lfp of the function $\lambda f. \lambda x. \llbracket c \rrbracket$ by the Knaster-Tarski theorem.

F.7 Soundness Result

We culminate this section with the proof of soundness of the interpretation. The main soundness theorem (Theorem 3), which we present at the end of this section, is a key result of our formal development. It states that FCSL is sound as a logic with respect to its translation into CiC (*i.e.* shallow embedding), that is, for any judgment $J, \Gamma \vdash J$ implies $\llbracket \Gamma \vdash J \rrbracket$. In particular, it means that if a judgment of the form $\Gamma \vdash \{p\}c : A \{q\}@U$ is derivable in FCSL, then its translation $\llbracket \Gamma \rrbracket \vdash_{\text{CiC}} \llbracket c \rrbracket : \{p\}A\{q\}@U$ is derivable in CiC. Next, let us apply Definition 23 of a Hoare type to $\{p\}A\{q\}@U$ and recall that $\llbracket c \rrbracket$ is just a set of action trees, corresponding to the command c (Figure 5). Therefore, by the derivation under \vdash_{CiC} , all action trees from $\llbracket c \rrbracket$ will belong to $\{p\}A\{q\}@U$, which automatically implies that for any state w and each tree from $\llbracket c \rrbracket$, if w satisfies p and $w \in \mathcal{W}_U$, then $\text{after}_U w t q$ also holds. The later means (Definition 22) that t is safe and $q \vee w'$ eventually holds if t steps completely to $\text{Ret } v$ in some final configuration w' . Since after employs the **always** predicate, it also implies the progress property (Lemma 7) for all action trees from $\llbracket c \rrbracket$.

Auxiliary modal lemmas The proof of the soundness result will rely on a number of auxiliary modal lemmas, which are listed below. For lack of space, we omit their proofs. They are relatively straightforward, usually by an induction on the schedule ζ used by **always** predicate. We have carried all of proofs for these lemmas in Coq [11], they usually proceed by an induction on the schedule ζ . In the statements of the lemmas below, we always assume U to be some fixed concurroid, t ranges over action trees and w ranges over states.

The Universal lemma states that the modal predicate **always** commutes with universal quantification, which yields to the soundness of an infinitary **CONJUNCTION** rule (Figure 2). The assumption $\text{always}_U w t (\lambda w' t'. \text{True})$ makes the lemma hold when the quantification over x is vacuous.

Lemma 9 (Universal). *If $\text{always}_U w t (\lambda w' t'. \text{True})$, then **always** commutes with universal quantification:*

$$\text{always}_U w t (\lambda w' t'. \forall x. P x w' t') \iff \forall x. \text{always}_U w t (\lambda w' t'. P x w' t')$$

The implication lemma corresponds to weakening the postcondition, which is necessary for the proof of **CONSEQ** rule.

Lemma 10 (Implication for after). *If $\text{after}_U w t Q_1$ and $Q_1 \vee w \implies Q_2 \vee w$, for all v and $w \in \mathcal{W}_U$, then $\text{after}_U w t Q_2$.*

Closure under sequential composition justifies the **Seq** rule: q holds at the end of the composed tree if final configuration of the prefix t_1 can be used as an initial configuration for the suffix to show q holds after.

Lemma 11 (Closure under sequential composition). *For some action tree $t_1 : \text{tree } U B$ and $K : B \rightarrow \mathcal{P}(\text{tree } U A)$, if $t_2 \in \{\text{Seq } t_1 k \mid \forall x. k x \in K x\}$ and*

$$\text{after}_U w_1 t_1 (\lambda v w. \forall t. (t \in K v) \implies \text{after}_U w t q),$$

then $\text{after}_U w_1 t_2 q$.

Closure under parallel composition justifies the PAR rule. Intuitively, it holds because when (an approximation t_2 of) c_2 takes a step over its private and shared state, it amounts to \mathcal{R}_U environment interference on (an approximation t_1 of) c_1 , and vice versa. Note that the pattern of rearranging subjective *self/other* components recurs at the level of triples $w = [s \mid j \mid o]$: the parallel composition uses $[s_1 \circ s_2 \mid j \mid o]$ and the left and right child threads use $[s_1 \mid j \mid o \circ s_2]$ and $[s_2 \mid j \mid o \circ s_1]$, respectively.

Lemma 12 (Closure under parallel composition). *If $\text{after}_U ([s_1 \circ s_2 \mid j \mid o]) t_1 Q_1$ and $\text{after}_U ([s_2 \circ s_1 \mid j \mid o]) t_2 Q_2$, then*

$$\begin{aligned} \text{after}_U ([s_1 \circ s_2 \mid j \mid o]) (\text{Par } t_1 t_2 (\lambda x. \text{Ret } x)) \\ (\lambda v' w'. \exists s'_1 s'_2 j' o'. w' = [s'_1 \circ s'_2 \mid j' \mid o']) \quad \wedge \\ P_1 \pi_1 v' [s'_1 \mid j' \mid o' \circ s'_2] \quad \wedge \\ P_2 \pi_2 v' [s'_2 \mid j' \mid o' \circ s'_1]). \end{aligned}$$

Closure under injection is targeted to justify the INJECT rule. It states that if *after* for some predicate Q holds on a state w_U from a “small” concurroid V , it will hold in a large concurroid $U = V \blacktriangleright W$ with an additional state, reachable by the \mathcal{R}_W relation of an addition W (Definition 17).

Lemma 13 (Closure under injection). *For the states w_V and w_0 , such that $U = V \blacktriangleright W$, $w_V \in \mathcal{W}_U$, $w_0 \in \mathcal{W}_W$ and $w_V \cup w_0 \in \mathcal{W}_V$. If $\text{after}_V w_V t Q$, then*

$$\begin{aligned} \text{after}_U (w_V \cup w_0) (\text{Inject } t) (\lambda v' w'. \exists w'_V w'_0. w' = w'_V \cup w'_0 \wedge w'_V \in \mathcal{W}_U \wedge \\ (w_0, w'_0) \in \mathcal{R}_W \wedge Q v' w'_0). \end{aligned}$$

Closure under hiding justifies the rule HIDE. It is important to notice that in the final state w'_V of the refined execution in V is related to the final state w'_U of the coarse execution in U by the same refinement function $\widehat{\Phi}$ (Definition 18), but with different auxiliaries g' , hence the conjunct $\widehat{\Phi}(g')(w'_U, w'_V)$.

Lemma 14 (Closure under hiding). *If $\widehat{\Phi}(g)(w_U, w_V)$, such that refines $U V \widehat{\Phi}$, and $\text{after}_V w_V t Q$, then*

$$\text{after}_U w_U (\text{Hide } \Phi g t) (\lambda v w'_U. \exists w'_V g'. \widehat{\Phi}(g')(w'_U, w'_V) \wedge Q v w'_V)$$

Finally, we can formulate and prove the main theorem. The theorem essentially states that if a judgment $\Gamma \vdash J$ can be derived in FCSL, its translation (Figure 4) can be also derived in CiC.

Theorem 3 (Soundness). *If $\Gamma \vdash J$, then $\llbracket \Gamma \vdash J \rrbracket$.*

Proof. The proof goes by induction on the derivation of J . The proof for pure expression is straightforward, since $\llbracket \Gamma \vdash e : A \rrbracket$ is just $\llbracket \Gamma \rrbracket \vdash_{\text{CiC}} e : A$. So let us focus on the judgments of the form $\Gamma \vdash \{p\} c : A \{q\} @ U$ that check the specifications of FCSL commands (Figure 2).

Each basic command (e.g., *return*, *act*) is sound because its pre- and postconditions are stable under environment interference, the precondition implies the command

is safe (by Definition 23 and through the after predicate), and the resulting configuration satisfies the postcondition. The SEQ, INJECT and HIDE rules are sound by Lemmas 11, 13 and 14. The PAR rule is sound by Lemma 12, as the *self/other* exchange in the lemma accounts exactly for the interpretation of the subjective separate conjunction \otimes (Figure 1), and the rule FRAME is just a particular case of PAR with c_2 being an idle thread (e.g., `return ();`), hence the stability requirement. The fix rule is sound by the Knaster-Tarski theorem. The CONJ rule is sound by Lemma 9, and the CONSEQ rule by Lemma 10. The EXISTential rule and IF rules are derivable. Since SCSL procedures are interpreted as (monadic) CiC functions, the procedure APPLICATION and HYPOTHESIS rules are sound by the function application and hypothesis rules of CiC. \square