

Communicating State Transition Systems for Fine-Grained Concurrent Resources

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 extended version [11].
- 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

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

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* concurroid 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 concurroid $CSL_{lock, lk, I}$ (i.e., private state and one lock). The lock lk protects a shared integer pointer x , that is, the resource invariant is $I (a : nat) (h : heap) \hat{=} h = x \rightarrow a$. For the nat argument, we chose the PCM structure under addition; thus, an assertion $lock \mapsto^s (-, a_S)$ expresses that the *self* thread has added a_S to x , and dually for $lock \mapsto^o (-, a_O)$. 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 $incr(n)$ acquires the lock to ensure exclusive access to x , increments x by n , and releases the lock. In FCSL, it has the following specification:

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

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 $incr(i) \parallel incr(j)$ 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.

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

If c is verified wrt. concurroid U , it can be *injected* (i.e. coerced) into a larger concurroid $U \times V$. In programs, we use the explicit coercion `inject` to describe the change of “type” from U to $U \times V$. 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* [11] 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 $CSL_{\text{lock},lk,l} = P \times L_{\text{lock},lk,l}$, which we have already used before.

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

`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. $CSL_{\text{lock},lk,l}$.

`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.

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

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 $\Psi(\text{Own}, a_S) \text{ empty} = \text{priv} \mapsto^s lk \rightarrow \text{false} \cup x \rightarrow 0$. Also, $\Phi(m_S, a_S) \multimap \text{lock} \mapsto^s (m'_S, a'_S)$ is equivalent to $(m_S, a_S) = (m'_S, a'_S)$ in the label-free state.

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

The soundness of HIDE depends on a number of semantic properties of Φ [11]. 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 a'_S , where a'_S is some value that is coherent with the acquired heap h , *i.e.*, which makes the resource invariant $I(a_S \bullet a_O) h$ 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 .

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

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) [11]. 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) \multimap q)\} \text{hide}_{\Phi, g} c \{\exists g'. \Psi g' h' * (\Phi(g') \multimap q)\} @ P \times U} \text{HIDE} \\
\\
\text{where } \Psi g h = \exists k: \text{heap}. \text{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 [11]. 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/fcsl/>.
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*.