

CoLoSL: Concurrent Local Subjective Logic

Azalea Raad, Jules Villard, and Philippa Gardner

Imperial College London
{azalea,j.villard,pg}@doc.ic.ac.uk

Abstract. A key difficulty in verifying shared-memory concurrent programs is reasoning compositionally about each thread in isolation. Existing verification techniques for fine-grained concurrency typically require reasoning about either the entire shared state or disjoint parts of the shared state, impeding compositionality. This paper introduces the program logic CoLoSL, where each thread is verified with respect to its subjective view of the global shared state. This subjective view describes only that part of the state accessed by the thread. Subjective views may arbitrarily overlap with each other, and expand and contract depending on the resource required by the thread. This flexibility gives rise to small specifications and, hence, more compositional reasoning for concurrent programs. We demonstrate our reasoning on a range of examples, including a concurrent computation of a spanning tree of a graph.

1 Introduction

A key difficulty in verifying properties of fine-grained concurrent programs is being able to reason compositionally about each thread in isolation, even though in reality the correctness of the whole system is the collaborative result of intricately intertwined actions of the threads. Such compositional reasoning is essential for: verifying large concurrent systems, since it allows them to be verified component-wise; verifying library code and incomplete programs, since one does not need to know about the context of execution; and replicating a programmer’s intuition about why their implementations are correct, since their informal arguments are typically kept local and do not bring the whole system into the reasoning.

Rely-guarantee [16] and various combinations of rely-guarantee and separation logic reasoning [26,9,10,7,23,3] achieve compositional reasoning for increasingly difficult inter-thread interactions. However, we believe that, despite substantial progress, there are many examples where the specifications and proofs are not as concise as they might be; they are either too coarse or too contrived. We explore a different approach, introducing the program logic CoLoSL, using which we can give small specifications and proofs which correspond to the programmer’s intuition of what shared resource is required by the thread.

Small specifications were emphasised in the work of O’Hearn, Reynolds and Yang on separation logic [19]. The original separation logic [21,15] achieves local, compositional reasoning for sequential heap-manipulating programs by splitting the heap into disjoint heaplets for describing the local resources required by a

program. Compositionality then rests on two powerful mechanisms: a program can be specified using only those resources that it actually accesses; and this specification can be simply reused in any context that contains these resources. By making the specification as small as possible, we can ensure that it can be reused in a large set of possible contexts using the *frame rule*. This plays an important role in achieving scalable compositional reasoning, as each block of code can be proved in isolation and its small specification reused in larger contexts.

Concurrent separation logic extends this compositional reasoning to concurrent programs using the *disjoint concurrency rule*, where individual threads use both local resources private to the thread as well as static shared resources, which can be accessed by all threads through critical sections. Since then, there have been many extensions combining rely-guarantee reasoning with ideas from concurrent separation logic to reason about fine-grained concurrency: RGSep [26] introduced local resource and shared disjoint regions which are stable with respect to a static interference relation stating how the current thread and the environment can affect the region; CAP [7] and its extensions [24,23,3] in addition abstract the regions. This work has achieved substantial success; see the related work section for more details. Yet, we believe these approaches are not always able to provide small specifications for concurrent programs, impeding compositionality.

The problem is due to the rigid nature of the shared disjoint regions and their static interference, which limits how we can work with a concurrent data structure. A disjoint region typically either describes the entire data structure such as a linked list, or contains individual components of the data structure such as the nodes of a linked list. However, threads may have shared access to arbitrary parts of the data structure, which cannot be directly expressed in the reasoning. The interference associated with the region is static in that it is defined when the region is created and is fixed throughout its entire lifetime. However, threads having access to parts of a shared region need only know about the interference on those parts. In addition, it is not always possible to predict all future interactions associated with a shared region at the time of creation. Just as the original separation logic uses the frame rule to obtain small specifications for the *local* state, we seek an analogous framing mechanism on both the shared resource and its interference to obtain small specifications of the *shared* state.

For example, consider a linked list consisting of $n+1$ nodes accessed concurrently by n threads where the i th thread requires access to the i th and $(i+1)$ st nodes. Current approaches cannot provide a small specification for each thread which captures just these two shared nodes and their interference.¹ Now consider a program whose threads manipulate subgraphs of a graph, such as a recursive concurrent spanning tree algorithm. Current approaches cannot give small specifications capturing just the subgraph manipulated by each thread due to intrinsic, unspecified sharing between subgraphs. Finally, consider a concurrent set implementation. In CoLoSL it is possible to describe the interference associated with a new element as it is added to the set. The interference on new elements need

¹ In [13], similar problems were encountered in an attempt to provide small specifications for a doubly-linked list implementation of a concurrent tree.

not be the same as before and may only be known at the time they are added. Existing approaches cannot accommodate such dynamic interference relations.

This paper introduces the program logic CoLoSL, which stands for Concurrent Local Subjective Logic. CoLoSL’s semantic model is based on one global shared state, and each thread is verified with respect to its partial *subjective view* of this state. Each subjective (personalised) view comprises an assertion which describes *parts* of the shared global resource used by a thread, and an interference relation which describes how the thread and the environment may affect these parts. Subjective views may arbitrarily overlap with each other, with both their resources and interference expanding and contracting in accordance with what is required by the thread. Interestingly, this sometimes requires rewriting the interference relation so that the interference on the smaller state captures the same information as the interference on the bigger state. This expansion and contraction of subjective views provides small specifications for individual threads and local reasoning about a thread’s shared state.

We demonstrate CoLoSL reasoning on a range of examples. The first example in §2 is Dijkstra’s token ring mutual exclusion algorithm [5]. Regardless of the size of the ring, we are able to give a small specification to each thread in isolation such that each proof only mentions resources associated with two of its neighbours. This means that the proof can be largely reused when the implementation of the ring is changed to allow dynamic spawning of new participants. In §4, we study two further examples. The first is a concurrent spanning tree algorithm for graphs, where threads are recursively spawned on potentially overlapping subgraphs. We demonstrate that the flexible, overlapping subjective views of CoLoSL are just what we need. The second is a concurrent set module implemented using a hand-over-hand list-locking algorithm. Our CoLoSL reasoning for this set module improves on CAP reasoning [7] in that our small specifications can be dynamically extended to include other behaviour in future whereas the static CAP interference must predict all behaviour from the start.

Most of the technical details have been left out due to space constraints, and are provided in the accompanying technical report [20].

Related work Jones’ rely-guarantee reasoning [16] provided a breakthrough in compositional reasoning about concurrent programs. It described permitted interferences between threads using global rely and guarantee relations on one global shared state. This work is compositional in the sense that the guarantee of one thread is the rely of another. However, O’Hearn [18] demonstrated that, for concurrent reasoning to scale, it is important to work with small specifications based on local state and shared state rather than working with one global state. In particular, he introduced concurrent separation logic based on thread-local state and static invariants for the shared state.

Since then, there has been much recent work on compositional reasoning about fine-grained concurrent algorithms. RGSep [26] combined rely-guarantee reasoning with separation logic, with the state split into thread-local state and disjoint shared regions, and global rely-guarantee relations providing the interference on these regions. Deny-guarantee extends rely-guarantee and RGSep

with permissions that can turn pieces of interference on or off during the proof of a program, as would be typically needed in a programming language with fork and join constructs instead of statically-scoped parallel composition. This has influenced the capabilities of CAP. CAP reasoning [7] increased compositionality by, instead of the global rely-guarantee relations, having static interference relations associated with the disjoint shared regions and capabilities in the local state for dynamically controlling the permitted interference, and concurrent abstract predicates for hiding implementation details. Several program logics have adapted this work to incorporate more abstraction [25], higher-order features [23] and more flexible capabilities [3]. We believe that all this work has limited compositional reasoning in the sense that it is only possible to frame off local state and unused shared regions. It is not possible to frame within regions nor frame inside the interference relations. (It is, of course, possible to weaken the region assertion by the logical implication $P * Q \Rightarrow P * \text{true}$, but this is not the same as being able to shrink regions *and* their interference to just work with P .)

Meanwhile, a different breakthrough in compositionality came from Feng’s local rely-guarantee (LRG) reasoning [10]. The LRG model comprises one global shared state, with assertions describing disjoint but flexible parts of this state and rely-guarantee relations determining how the thread and environment can affect these partial states. With LRG, it is possible to frame off disjoint parts of shared state and their associated disjoint rely-guarantee relations, but, as noted by its author, the strong disjointness restrictions make this approach only applicable for disjoint modules in the code. We took significant inspiration from LRG, combining the flexibility of the LRG approach with our subjective views to reason locally about overlapping shared states.

Finally, Fine-grained Concurrent Separation Logic (FCSL) [17] of Nanevski *et al.* explores a different notion of region called *concurroids*. Like the regions in CAP, concurroids describe disjoint pieces of shared state together with their static interference. It is therefore not possible to frame within concurroids. The term “subjectivity” in FCSL refers to the fact that, unlike CAP, concurroids have three parts: the “joint” part; and the disjoint “self” and “other” parts. Although FCSL did not influence CoLoSL, it does highlight an interesting point regarding resource transfer between regions: in CAP, communication between regions is achieved indirectly via the local state; in FCSL, communication between concurroids is achieved directly through dangling external transitions; and, in CoLoSL, communication between compatible subjective views can be achieved by merging the two views.

2 Informal Development

We sketch a proof of an implementation of Dijkstra’s token ring mutual exclusion algorithm, which pioneered *self-stabilising* distributed algorithms [5]. Our proof highlights the main reasoning principles of CoLoSL and results in *small specifications* for each participant in the ring. Besides their concision, we show how small specifications are robust against non-trivial changes to the program.

$$\begin{array}{c}
\frac{}{\boxed{P * Q}_I \Rightarrow \boxed{P}_I} \text{ (FORGET)} \quad \frac{}{\boxed{P}_{I_1} * \boxed{Q}_{I_2} \Rightarrow \boxed{P \uplus Q}_{I_1 \cup I_2}} \text{ (MERGE)} \\
\frac{}{(P \Rightarrow Q) \Rightarrow \boxed{P}_I \Rightarrow \boxed{Q}_I} \text{ (WEAKEN)} \quad \frac{I \sqsubseteq^P I' \text{ implies } \boxed{P}_I \Rightarrow \boxed{P}_{I'}}{} \text{ (SHIFT)} \\
\frac{}{\boxed{P}_I \Rightarrow \boxed{P}_I * \boxed{P}_I} \text{ (COPY)} \quad \frac{P \odot I \text{ implies } P \Rightarrow \exists A, A'. [A] * \boxed{P * [A']}_I}{} \text{ (EXTEND)} \\
\frac{\{P_1\} \mathbb{C}_1 \{Q_1\} \quad \{P_2\} \mathbb{C}_2 \{Q_2\}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}} \text{ PAR} \quad \frac{P \Rightarrow P' \quad \vdash \{P'\} \mathbb{C} \{Q'\} \quad Q' \Rightarrow Q}{\vdash \{P\} \mathbb{C} \{Q\}} \text{ CONSQ}
\end{array}$$

Figure 1: Main reasoning principles and rules of CoLoSL.

The algorithm assumes a network of $n+1$ machines arranged in a ring, with a designated *master* machine and n slave machines. Each machine maintains a local counter and has access to the value of its left neighbour's counter; the state of the system consists of all $n+1$ counters. Starting in an arbitrary state, the network eventually stabilises to legitimate states [4], with the following global property: in the i th legitimate state, machines 0 to $i-1$ (machine 0 being the master) have some value $v+1$, and all others have value v . In the i th legitimate state, only the i th machine (indicated by \bullet below) can make progress: it increments its counter by 1 and takes the system to the next legitimate state ($(i+1 \bmod n+1)$). For a ring at address x , the i th legitimate state is depicted below (for both $i > 0$ and $i = 0$).

$$\begin{array}{ccccccccc}
x & x+1 & & x+i-1 & x+i & & x+n & & x & x+1 & & x+n \\
\boxed{v+1} & \boxed{v+1} & \cdots & \boxed{v+1} & \boxed{v \bullet} & \cdots & \boxed{v} & & \boxed{v \bullet} & \boxed{v} & \cdots & \boxed{v}
\end{array}$$

In [20], we outline a proof of the token ring's self-stabilisation phase, which shows that the system always converges to a legitimate state. We also provide a proof of the token ring as a mutual exclusion mechanism, where a machine holding \bullet may gain ownership of a shared resource. For simplicity, here we focus on the case where the token ring is in the 0th legitimate state, with all counters holding value 0. Our proof makes use of the CoLoSL principles laid out in Fig. 1, together with the usual concurrency rule of separation logic [18] and the rule of semantic consequence from the views framework [6]. We introduce them informally as needed, and discuss them in more detail in §3.

CoLoSL introduces a new assertion \boxed{P}_I called a *subjective view*, which comprises a *subjective assertion* P describing a *part* of the global shared state and an *interference assertion* I characterising how this partial shared state may be changed by the thread or the environment. Similar to the interference assertions of CAP [7], I declares actions of the form $[a] : Q \rightsquigarrow R$, where a thread in possession of the $[a]$ capability in its local state may carry out its transition and update parts of the shared state satisfying Q to a state satisfying R . Assertions in Hoare triples must be *stable* with respect to the interference from the environment: that is, robust with respect to the interference assertion I .

Consider the program `ring(x)` defined in Fig. 2 which represents a token ring with $n+1$ machines. It is written in pseudo-code resembling C with additional constructs for concurrency: atomic sections $\langle _ \rangle$ which declare that code behaves

```

ring(x)
// { x ⇒ x * [m_x] * [s_{x+1}] * ... * [s_{x+n}] * [n ⇒ n * x ↦ 0 * x+1 ↦ 0 * ... * x+n ↦ 0] }_I
{ master(x) || slave(x+1) || ... || slave(x+n);
} // { x ⇒ x * [m_x] * [s_{x+1}] * ... * [s_{x+n}] * [n ⇒ n * x ↦ 10 * x+1 ↦ 10 * ... * x+n ↦ 10] }_I

master(x)
// { x ⇒ x * [m_x] * [n ⇒ n * x ↦ 0 * x+n ↦ 0] }_{M'_x}
{ while(*x != 10)
// { x ⇒ x * [m_x] * [n ⇒ n * ∃v. x ↦ v * (x+n ↦ v ∨ x+n ↦ v-1)] }_{M'_x}
  { (if (*x == *(x+n)) *x = *x + 1; ) }
} // { x ⇒ x * [m_x] * [n ⇒ n * x ↦ 10 * (x+n ↦ 10 ∨ x+n ↦ 9)] }_{M'_x}

slave(x)
// { x ⇒ x * [s_x] * [x ↦ 0 * (x-1 ↦ 0 ∨ x-1 ↦ 1)] }_{S'_x}
{ while(*x != 10)
// { x ⇒ x * [s_x] * [∃v. x ↦ v * (x-1 ↦ v ∨ x-1 ↦ v+1)] }_{S'_x}
  { (if (*x != *(x-1)) *x = *(x-1); ) }
} // { x ⇒ x * [s_x] * [x ↦ 10 * (x-1 ↦ 10 ∨ x-1 ↦ 11)] }_{S'_x}

s_x ≜ [s_x]: ∃v. x ↦ v * x-1 ↦ v+1      ∼ x ↦ v+1 * x-1 ↦ v+1
s'_x ≜ [s_x]: ∃v. x+1 ↦ v * x ↦ v * x-1 ↦ v+1 ∼ x+1 ↦ v * x ↦ v+1 * x-1 ↦ v+1
m_x ≜ [m_x]: ∃v, n. n ⇒ n * x ↦ v * x+n ↦ v      ∼ n ⇒ n * x ↦ v+1 * x+n ↦ v
m'_x ≜ [m_x]: ∃v, n. n ⇒ n * x+1 ↦ v * x ↦ v * x+n ↦ v ∼
              n ⇒ n * x+1 ↦ v * x ↦ v+1 * x+n ↦ v
l'_x ≜ [s_x]: ∃v, n. n ⇒ n * x ↦ v+1 * x+n ↦ v * x+n-1 ↦ v+1 ∼
              n ⇒ n * x ↦ v+1 * x+n ↦ v+1 * x+n-1 ↦ v+1

I ≜ {m_x, s_{x+1}, ..., s_{x+n}}    M'_x ≜ {m_x, l'_x}    F'_x ≜ {s_x, m'_{x-1}}    S'_x ≜ {s_x, s'_{x-1}}

```

Figure 2: A proof sketch of the token ring in CoLoSL. Assertions in lines starting with `//` describe the local state and the subjective shared state at the relevant program points. The proof of `slave(x)` applies to all slaves expect the first one (called the *foreman*) in the parallel composition, where S'_x is replaced by F'_x .

atomically; and parallel composition `_||_` which spawns threads then waits until they complete. In our example, the $n+1$ threads run in parallel until all counters reach value 10. While the implementation of all the slave threads are identical, we shall see that the proof of the first slave in the ring (at $x+1$) is slightly different from the others. We henceforth refer to the first slave thread as the *foreman*. Let us proceed with the proof of the other slaves.

Proof of the slaves Let us temporarily forget about the proof outline of Fig. 2 and attempt to prove `slave(x)` in isolation, in the spirit of local reasoning. Since `slave(x)` inspects the value of its counter pointed to by `x` and compares it against the counter at `x-1` (its left neighbour in the ring), a tempting precondition for `slave(x)` would be one describing just these two locations, e.g.

$$x \Rightarrow x * [s_x] * [x \mapsto 0 * (x-1 \mapsto 0 \vee x-1 \mapsto 1)]_{S_x} \quad S_x = \{s_x, s_{x-1}\}$$

The above assertion comprises: a) a variable assertion stating that the thread locally owns variable `x` with value x (using the variables-as-resource model [1]);

b) a *capability* $[s_x]$ that allows it to perform the associated s_x action (see below); and c) a subjective view of the shared state: x points to 0, and $x-1$, its left neighbour, points to either 0 or 1 since it might already have incremented its own counter. The \star connective used between assertions is that of separation logic and means that the assertions describe *disjoint* pieces of state. The interference assertion associated with the subjective view is captured by S_x and consists of two actions: s_x and s_{x-1} , where s_x represents an increment of the contents of x under the condition that its value is one less than the value at address $x-1$ (see Fig. 2). Since the current thread owns $[s_x]$ locally, only it can perform s_x . On the other hand, the capability $[s_{x-1}]$ is not locally owned, thus the environment could potentially perform the associated action whenever its precondition (on the left-hand side of \rightsquigarrow) is satisfied. Upon closer inspection, since this subjective view says nothing of the value of the cell at address $x-2$, s_{x-1} could potentially *always* fire. The assertion is thus not *stable*: nothing prevents the counter at $x-1$ from incrementing beyond value 1. A weaker, stable assertion is thus:

$$x \Rightarrow x \star [s_x] \star \boxed{\exists v. x \mapsto 0 \star x-1 \mapsto v}_{S_x}$$

Fortunately, we can do better and obtain a stronger small precondition. Let us first step back and think again at the level of the whole algorithm. As the programmer knows, the situation above cannot happen as $x-2$ can only be at most one ahead of x itself. We can thus replace S_x by S'_x and give a stronger stable precondition that captures just what we want, as in Fig. 2. The proof of $\text{slave}(x)$ is now relatively straightforward. By inspection (or using the rules of §3.3), the invariant of the while loop and the postcondition are also stable. The atomic section temporarily “opens the box” to perform action s_x then “closes back the box”, and preserves the invariant. The final postcondition of $\text{slave}(x)$ follows from the invariant and the boolean expression of the loop.

Proof of the master and foreman The proof sketches of the master and foreman threads given in Fig. 2 are analogous. As the first slave, the foreman has to account for interference from the master instead of another slave. Moreover, the master and its associated action m_x have access to variable n holding the current size of the ring, since they depend on the value of the counter at address $x+n$.

Proof of the ring The precondition of $\text{ring}(x)$ states that it owns all capabilities, which will be distributed amongst the $n+1$ threads; the global variable $n \Rightarrow n$ is shared, as are all $n+1$ counters, initialised to 0. The interference I associated with the subjective view consists of the actions of the $n+1$ threads. Because $\text{ring}(x)$ has a global view of the state of the ring (and moreover all capabilities are held locally), the s_{x+i} actions are enough to guarantee stability.

Let us write \boxed{P}_I for this initial subjective view. This assertion may be freely duplicated using the COPY principle of Fig. 1 and each thread is given a copy together with the appropriate capability using the usual PAR rule of concurrent separation logic. For instance, the thread running $\text{slave}(x+i)$, for $i > 1$, gets $[s_{x+i}] \star \boxed{P}_I$. This assertion does not match the precondition of $\text{slave}(x+i)$ just

yet. Using the principles of Fig. 1, we can weaken the assertion as such:

$$\begin{aligned}
\boxed{P} \quad I &\xRightarrow{\text{(SHIFT)}} \boxed{P} \quad (I \setminus \{s_{x+i-1}\}) \cup \{s'_{x+i-1}\} \\
&\xRightarrow{\text{(FORGET)}} \boxed{x+i-1 \mapsto 0 \star x+i \mapsto 0} \quad (I \setminus \{s_{x+i-1}\}) \cup \{s'_{x+i-1}\} \\
&\xRightarrow{\text{(SHIFT)}} \boxed{x+i-1 \mapsto 0 \star x+i \mapsto 0} \quad \{s'_{x+i-1}, s_{x+i}\} \\
&\xRightarrow{\text{(WEAKEN)}} \boxed{x+i \mapsto 0 \star (x+i-1 \mapsto 0 \vee x+i-1 \mapsto 1)} \quad \{s'_{x+i-1}, s_{x+i}\}
\end{aligned}$$

We start by exchanging the action s_{x+i-1} of I for the stronger action s'_{x+i-1} using the SHIFT principle. In general, SHIFT allows us to replace I with any interference assertion I' that has the same observable effect as far as the subjective assertion P is concerned (written $I \sqsubseteq^P I'$). In this instance, the actions s_{x+i-1} and s'_{x+i-1} have the same effect according to P , as discussed in the proof of slave(x). As such, rewriting s_{x+i-1} as s'_{x+i-1} merely reflects stronger knowledge about how $x+i$ and $x+i-2$ are related through $x+i-1$. In particular, $I \sqsubseteq^P (I \setminus \{s_{x+i-1}\}) \cup \{s'_{x+i-1}\}$ as required.

Next, because subjective views only describe *parts* of the shared state, we can use the FORGET principle to obtain a weaker view of the shared state, in this case a view that ignores all cells in the ring except for those at addresses $x+i-1$ and $x+i$. With all other cells out of the subjective view, their actions no longer have observable effects on the assertion, since they leave $x+i-1$ and $x+i$ unchanged. We can thus apply the SHIFT principle again to *frame off* those actions and obtain $S'_{x+i} = \{s'_{x+i-1}, s_{x+i}\}$.

Finally, we weaken the resulting subjective view so that it is stable with respect to S'_{x+i} , i.e. preserved by those of its actions that the environment may perform (here, s'_{x+i-1}). This yields the precondition of slave(x+i) as in Fig. 2. The preconditions of the master and foreman threads can be derived analogously.

Once all threads have completed their operations, we join up their postconditions using the MERGE principle, which embodies a crucial feature of CoLoSL: different subjective views *overlap*. The overlapping conjunction \bowtie between two assertions means that the two assertions describe potentially *overlapping* pieces of state. In particular, $A \star B \Rightarrow A \bowtie B$ and $A \wedge B \Rightarrow A \bowtie B$. This connective has been used in the past to reason about sharing in data structures [22,11,14]. Since \vee distributes over \bowtie , the subjective view simplifies to $\boxed{x \mapsto 10 \star x+1 \mapsto 10 \star \dots \star x+n \mapsto 10} \quad I'$, where $I' = M'_x \cup F'_{x+1} \cup S'_{x+2} \cup \dots \cup S'_{x+n}$. Finally, since $I' \sqsubseteq^{n \Rightarrow n \star x \mapsto 10 \star x+1 \mapsto 10 \star \dots \star x+n \mapsto 10} I$, we get the postcondition of ring(x) by the SHIFT principle. This concludes our CoLoSL proof of ring(x).

Small specifications and proof reuse Our expansion and contraction of subjective views, in particular with the shifting of interference assertions at key places, enables us to confine the specification and verification of each thread to just the resources they need. Such small specifications make proofs robust against changes to each thread's environment, and provide more opportunities for proof reuse.

For instance, let us now add a thread that dynamically grows the ring by spawning extra slaves to the parallel composition of ring(x) (the details can be

found in [20]). When the ring has $n+1$ machines, we use the EXTEND principle as follows to add a new slave (at $x+n+1$) to the shared state with the associated interference relation and capability.²

$$(x+n+1) \mapsto v \stackrel{\text{(EXTEND)}}{\Rightarrow} \exists s_{x+n+1}. [s_{x+n+1}] * \boxed{x+n+1 \mapsto v}_{\{s_{x+n+1}\}}$$

Here, the *view shift* [6] (or *repartitioning* [7]) $P \Rightarrow Q$ means that an instrumented (logical) state satisfying P may be changed to Q as long as the underlying machine state does not change. In particular, $(P \Rightarrow Q) \Rightarrow (P \Rightarrow Q)$. The point is that our proof changes only minimally to accommodate the new program: the proof of $\text{master}(x)$ accounts for new interference on $n \Rightarrow n$ since the environment can grow the ring, hence mutate n ; the proofs of other threads can be directly reused.

In contrast, in existing CAP-like approaches [7,3], both $n \Rightarrow n$ and the global interference relation are observed by all threads. As such, with the above extension, the global interference relation needs to change (to include the interference on $n \Rightarrow n$), and the proofs of *both* master and slave threads need to be adapted.

3 CoLoSL

We now give a more formal overview of how to use CoLoSL for program verification, eschewing some details of the semantics for lack of space, while still providing enough ingredients to carry program proofs. We describe the underlying model of CoLoSL, give the semantics of CoLoSL assertions, then present proof rules to reduce the various obligations typically encountered in proofs (namely, the side-conditions of SHIFT and EXTEND, and stability checks) to classical separation logic entailments that, in particular, do not mention subjective views.

3.1 CoLoSL worlds

A *world* is a triple (l, g, \mathcal{J}) where l and g are *logical states* and \mathcal{J} is an *action model*. The *local logical state* (or local state) l represents the locally owned resources of a thread, in the standard separation logic sense, while the *global logical state* (or shared state) g represents shared resources. The action model \mathcal{J} records all possible interferences on the shared state.

Logical states have two components: one describes machine states (e.g. stacks and heaps); the other represents *capabilities*. The latter are inspired by the capabilities in CAP [7]: a thread in possession of a given capability is allowed to perform the associated actions (as prescribed by the *action model* component of each world, defined below), while any capability *not* owned by a thread means that the environment can perform the action. This can be seen as a unified

² Unlike CAP [7] and as in iCAP [23], we do not provide an explicit *unsharing* mechanism to claim back shared resources. Instead, this can be simply encoded as an action of the form $\star_{j \in \mathcal{J}}[a_j]: P \rightsquigarrow \star_{j \in \mathcal{J}}[a_j]: a$: a thread holding $\star_{j \in \mathcal{J}}[a_j]$ can move the shared resource P into its local state in exchange for the associated capabilities.

treatment of the rely and guarantee relations in rely-guarantee reasoning [16]: a capability fully owned (resp. fully not owned) during the entire lifetime of a thread represents its guarantee (resp. its rely), while a partially-owned capability means that the corresponding action is both in the rely and the guarantee. Capabilities go beyond the rely-guarantee model [9]; in particular, they may be transferred between a thread and its environment just like any other resource to temporarily block or enable certain actions. See the presentation of CAP [7] and deny-guarantee [9] for further details and motivation.

In general, each component of a logical state is taken from an arbitrary *separation algebra* [2] (i.e. a cancellative, partial commutative monoid) that satisfies the *cross-split* property³[8] (this is needed for \uplus to be associative [14]). As we demonstrate in the examples of §4, our programs often call for a more complex model of machine states and capabilities than that presented here. For instance, we may need our capabilities to be fractionally owned, where ownership of a *fraction* of a capability grants the right to perform the action to both the thread and the environment, while a fully-owned capability by the thread *denies* the right to the environment to perform the associated action. For ease of presentation, the focus of this paper is on the standard stack and heap model for machine states, and finite sets of *tokens* (which are simple names) for capabilities. We assume a set of program values Val , as well as infinite disjoint sets PVar , Loc , and Token of program variables, memory locations, and tokens, respectively.

Definition 1 (Logical states). *A logical state is a tuple $((\sigma, h), \kappa) \in \text{LState}$, also written (σ, h, κ) , of a finite partial stack $\sigma \in \text{Stack}$ associating program variables with values, a heap $h \in \text{Heap}$ associating heap locations with values, and a capability $\kappa \in \mathbb{K}$:*

$$\begin{aligned} \text{Stack} &\triangleq \text{PVar} \rightarrow_{\text{fin}} \text{Val} & \text{Heap} &\triangleq \text{Loc} \rightarrow_{\text{fin}} \text{Val} & \mathbb{M} &\triangleq \text{Stack} \times \text{Heap} \\ \mathbb{K} &\triangleq \mathcal{P}_{\text{fin}}(\text{Token}) & \text{LState} &\triangleq \mathbb{M} \times \mathbb{K} \end{aligned}$$

The local and global logical states of a world are always *compatible*: they can be composed with one another. This captures the intuition that locally-owned resources are disjoint from shared resources. The composition of logical states is defined component-wise as disjoint function union \uplus over stacks and heaps, and disjoint set union \uplus on capabilities.

Definition 2 (Logical state composition). *The composition of logical states $\circ : \text{LState} \times \text{LState} \rightarrow \text{LState}$ is defined as:*

$$((\sigma, h), \kappa) \circ ((\sigma', h'), \kappa') \triangleq ((\sigma \uplus \sigma', h \uplus h'), \kappa \uplus \kappa')$$

We write l to range over arbitrary logical states or just local states, and g to range over logical states representing global shared states. The empty logical state

³ A monoid $(\mathbb{A}, +, \mathbf{1})$ satisfies the cross-split property iff, for all $a, b, c, d \in \mathbb{A}$, if $a + b = c + d$, there exists $x, y, z, w \in \mathbb{A}$ s.t. $a = x + y$, $b = z + w$, $c = x + z$ and $d = y + w$.

$(\emptyset, \emptyset, \emptyset)$ is written $\mathbf{0}_L$. We write $l_1 \leq l_2$ when there exists l such that $l \circ l_1 = l_2$, and write $l_2 - l_1$ to denote the unique such l (by cancellativity) when it exists. When $l_1 \circ l_2$ is defined, we say that l_1 and l_2 are *compatible* and write $l_1 \# l_2$.

An action is a triple (p, q, c) of logical states where p and q are the action *pre-* and *post-states* describing how the shared state is modified by the action; and c is the action *catalyst*. An action catalyst has to be present for the action to take effect, but is left unchanged by the action. It is maximal in the sense that no further, non-empty catalyst c' can be found, which we write $p \perp q: \forall l. l \leq p \wedge l \leq q \implies l = \mathbf{0}_L$.⁴ For instance, as we shall shortly see, s_x in Fig. 2 will be interpreted as the set of actions $S = \{a_v \triangleq ((\emptyset, \{\ell : v\}, \emptyset), (\emptyset, \{\ell : v+1\}, \emptyset), (\emptyset, \{\ell-1 : v+1\}, \emptyset)) \mid v \in \mathbb{N}\}$, where ℓ is the value of x in the current logical environment.

An action model is a partial function from *capabilities* to sets of *actions*. It corresponds to the semantic interpretation of an interference assertion.

Definition 3 (Action models). *The set Action of actions, ranged over by a , and the set AMod of action models, ranged over by \mathcal{J} , are defined as:*

$$\text{Action} \triangleq \text{LState} \times \text{LState} \times \text{LState} \qquad \text{AMod} \triangleq \mathbb{K} \rightarrow \mathcal{P}(\text{Action})$$

Worlds are triples $(l, g, \mathcal{J}) \in \text{LState} \times \text{LState} \times \text{AMod}$ that satisfy several well-formedness conditions: the local and shared states are compatible; the capabilities owned by l and g are in the domain of \mathcal{J} ; and actions in \mathcal{J} are *confined* to g (written $g \odot \mathcal{J}$). An action $a = (p, q, c)$ is confined to g if and only if, whenever it is enabled ($p \circ c$ agrees with g), then its pre-state p is contained in g ($p \leq g$). We motivate the need for the confinement condition in §3.3.

Definition 4 (Well-formedness). *A triple $(l, g, \mathcal{J}) \in \text{LState} \times \text{LState} \times \text{AMod}$ is well-formed, written $\text{wf}(l, g, \mathcal{J})$, iff $l \# g$, $l_\kappa \cup g_\kappa \subseteq \bigcup \text{dom}(\mathcal{J})$ and $g \odot \mathcal{J}$.*

Definition 5 (Worlds). *The set World of worlds consists of all well-formed triples:*

$$\text{World} \triangleq \{w \in \text{LState} \times \text{LState} \times \text{AMod} \mid \text{wf}(w)\}$$

Finally, the composition of two worlds is defined when their local states are compatible, their global shared states and action models are the same, and the resulting tuple is well-formed.

Definition 6 (World composition). *The composition of worlds, $\bullet : \text{World} \times \text{World} \rightarrow \text{World}$, is defined as:*

$$(l, g, \mathcal{J}) \bullet (l', g', \mathcal{J}') \triangleq \begin{cases} (l \circ l', g, \mathcal{J}) & \text{if } g = g', \text{ and } \mathcal{J} = \mathcal{J}', \text{ and } \text{wf}((l \circ l', g, \mathcal{J})) \\ \text{undefined} & \text{otherwise} \end{cases}$$

⁴ Alternatively, the catalyst could be computed a posteriori for each action. However, we often need to isolate the part of the state that is modified by an action, hence our technical choice of recording the catalyst in the model.

3.2 CoLoSL Assertions

Our assertion language extends separation logic with *subjective views* and *capability assertions*.

CoLoSL is parametric in the assertions of machine states and capabilities, and can be instantiated with any assertion language over machine states \mathbb{M} and capabilities \mathbb{K} . In this paper, we use standard heap and stack assertions for machine state assertions, and single token assertions of the form $[a]$ for capability assertions where $a \in \text{Token}$. We write $[A]$ as a shorthand for $\star_{a \in A}[a]$. We assume an infinite set LVar of *logical variables*, disjoint from PVar .

Definition 7 (Assertion syntax). *Given $x \in \text{LVar}$, $x \in \text{PVar}$, and $a \in \text{Token}$, the assertions of CoLoSL, Assn , are described by the grammars below:*

$$\begin{aligned} \text{LAssn} \ni p, q &::= \text{false} \mid E_1 = E_2 \mid \text{emp} \mid x \Rightarrow E \mid E_1 \mapsto E_2 \mid [E] \\ &\quad \mid p \vee q \mid \neg p \mid \exists x. p \mid p \star q \mid p \uplus q \mid p \multimap q \\ \text{Assn} \ni P, Q &::= p \mid \exists x. P \mid P \vee Q \mid P \star Q \mid P \uplus Q \mid \boxed{P}_I \\ I &::= \emptyset \mid \{[A]: \exists \bar{x}. P \rightsquigarrow Q\} \cup I \quad E ::= x \mid a \mid E_1 + E_2 \mid \dots \end{aligned}$$

The syntax and semantics of *local assertions* LAssn are as in standard separation logic with variables-as-resource [1].⁵ Local assertions are interpreted over a world's local state. The empty local state $\mathbf{0}_L$ is denoted by emp . The assertion $x \Rightarrow E$ denotes a singleton stack where x has value E . Similarly, $E_1 \mapsto E_2$ is true of the singleton heap where only address E_1 is allocated and has value E_2 . The capability assertion $[E]$ is true of the singleton capability $\{a\}$ if E evaluates to a . The *separating conjunction* $p \star q$ is true when the local state can be split into two according to \circ such that one state satisfies p and the other satisfies q . The *overlapping conjunction* $p \uplus q$ is true when the local state can be split three-ways according to \circ , such that the \circ -composition of the first two states satisfies p and the \circ -composition of the last two satisfy q [14,12,22]. *Sepraction* (or *existential magic wand*) $p \multimap q$ is true when there exists a local state satisfying p that can be \circ -composed with the current one to yield a state satisfying q . The usual predicates and connectives have their standard classical meaning.

As in RGSep [26], our assertions Assn are defined on top of local assertions. For simplicity, assertions do not include negation nor sepraction. The interpretation of assertions is a simple lift from that of local assertions, with the exception of the subjective view \boxed{P}_I . First, an *interference assertion* I describes actions enabled by a given capability, in the form of a pre- and postcondition. A subjective view \boxed{P}_I is then true of (l, g, J) when $l = \mathbf{0}_L$ and a *subjective state* s can be found in the global shared state g (i.e. $g = s \circ r$ for some r), such that s satisfies P , and J and I agree given the decomposition s and r , written $J \downarrow (s, r, \llbracket I \rrbracket_l)$, in the following sense:

1. every action in I is reflected in J ;

⁵ Note in particular that expressions E do not allow program variables: they can only appear on the left-hand side of $x \Rightarrow E$.

2. every action in \mathcal{J} that has a visible effect on s is reflected in I ;
3. the above holds after any number of actions in \mathcal{J} takes place.

Thus, given a world (l, g, \mathcal{J}) and a subjective view \boxed{P}_I , P describes a subjective state s that is *a part of* g and I describes *all parts* of \mathcal{J} concerning s , while \mathcal{J} describes the overall interference on g . We refer to the above agreement between the action model and the subjective view as the *action model closure property*. We omit its formal definition for lack of space.

The semantics of CoLoSL assertions is given by a forcing relation $w, \iota \vDash P$ between a world w , a logical environment $\iota \in \mathbf{LEnv}$, and an assertion P . We use two auxiliary forcing relations. The first one $l, \iota \vDash_{\mathbf{SL}} p$ interprets local assertions $p \in \mathbf{LAssn}$ in the usual separation logic sense over a logical state l . The second one $s, \iota \vDash_{g, \mathcal{J}} P$ interprets assertions $P \in \mathbf{Assn}$ over a *subjective state* s that is part of the global shared state g , subject to the action model \mathcal{J} . This second relation is needed to deal with the nesting of subjective views.⁶ Since logical connectives are interpreted uniformly in all cases, we write \vDash_{\dagger} for any of the three satisfaction relations, u for elements of either \mathbf{World} or \mathbf{LState} , and \bullet for either \bullet or \circ , as appropriate.

Definition 8 (Assertion semantics). *Given a logical environment $\iota : \mathbf{LVar} \rightarrow \mathbf{Val}$, the semantics of CoLoSL assertions is defined below, using the semantics of interference assertions $\llbracket \cdot \rrbracket_{(\cdot)} : \mathbf{LEnv} \rightarrow \mathbf{AMod}$ also defined below:*

$$\begin{array}{ll}
(l, g, \mathcal{J}), \iota \vDash p & \text{iff } l, \iota \vDash_{\mathbf{SL}} p \\
(l, g, \mathcal{J}), \iota \vDash \boxed{P}_I & \text{iff } l = \mathbf{0}_L \text{ and } \exists s, r. g = s \circ r \text{ and } s, \iota \vDash_{g, \mathcal{J}} P \text{ and } \mathcal{J} \downarrow (s, r, \llbracket I \rrbracket_{\iota}) \\
s, \iota \vDash_{g, \mathcal{J}} p & \text{iff } s, \iota \vDash_{\mathbf{SL}} p \\
s, \iota \vDash_{g, \mathcal{J}} \boxed{P}_I & \text{iff } (s, g, \mathcal{J}), \iota \vDash \boxed{P}_I \\
u, \iota \vDash_{\dagger} \exists x. P & \text{iff } \exists v. u, [\iota \mid x : v] \vDash_{\dagger} P \\
u, \iota \vDash_{\dagger} P \vee Q & \text{iff } u, \iota \vDash_{\dagger} P \text{ or } u, \iota \vDash_{\dagger} Q \\
u, \iota \vDash_{\dagger} P_1 \star P_2 & \text{iff } \exists u_1, u_2. u = u_1 \bullet u_2 \text{ and } u_1, \iota \vDash_{\dagger} P_1 \text{ and } u_2, \iota \vDash_{\dagger} P_2 \\
u, \iota \vDash_{\dagger} P_1 \uplus P_2 & \text{iff } \exists u', u_1, u_2. u = u' \bullet u_1 \bullet u_2 \text{ and} \\
& u' \bullet u_1, \iota \vDash_{\dagger} P_1 \text{ and } u' \bullet u_2, \iota \vDash_{\dagger} P_2 \\
l, \iota \vDash_{\mathbf{SL}} \text{false} & \text{never} \\
l, \iota \vDash_{\mathbf{SL}} \text{emp} & \text{iff } l = \mathbf{0}_L \\
l, \iota \vDash_{\mathbf{SL}} E_1 = E_2 & \text{iff } \llbracket E_1 \rrbracket_{\iota} = \llbracket E_2 \rrbracket_{\iota} \\
l, \iota \vDash_{\mathbf{SL}} x \Rightarrow E & \text{iff } l = (\{x : \llbracket E \rrbracket_{\iota}\}, \emptyset, \emptyset) \\
l, \iota \vDash_{\mathbf{SL}} E_1 \mapsto E_2 & \text{iff } l = (\emptyset, \{\llbracket E_1 \rrbracket_{\iota} : \llbracket E_2 \rrbracket_{\iota}\}, \emptyset) \\
l, \iota \vDash_{\mathbf{SL}} [E] & \text{iff } l = (\emptyset, \emptyset, \{\llbracket E \rrbracket_{\iota}\}) \\
l, \iota \vDash_{\mathbf{SL}} \neg p & \text{iff } l, \iota \not\vDash_{\mathbf{SL}} p \\
l, \iota \vDash_{\mathbf{SL}} p \multimap q & \text{iff } \exists l'. l', \iota \vDash_{\mathbf{SL}} p \text{ and } l \# l' \text{ and } l \circ l', \iota \vDash_{\mathbf{SL}} q
\end{array}$$

$$\llbracket I \rrbracket_{\iota}(A) \triangleq \left\{ (p, q, c) \mid \begin{array}{l} [A]: \exists \bar{x}. P \rightsquigarrow Q \in I \wedge p \perp q \wedge \exists \bar{v}, \mathcal{J}, r, r'. c = r \circ r' \wedge \\ p \circ r, [\iota \mid \bar{x} : \bar{v}] \vDash_{p \circ c, \mathcal{J}} P \wedge q \circ r, [\iota \mid \bar{x} : \bar{v}] \vDash_{q \circ c, \mathcal{J}} Q \end{array} \right\}$$

⁶ This presentation with several forcing relations differs from the usual CAP presentation [7], where assertions are first interpreted over worlds that are not necessarily well-formed, and then cut down to well-formed ones. The advantage of our presentation is that the semantics of assertions is *compositional*, e.g. the semantics of $P \star Q$ follows directly from the semantics of P and Q .

Note that, as in the CAP family [7,23,3], CoLoSL cannot ensure that proved programs do not leak memory. This is because of the following property of the semantics with respect to the shared state (sometimes called “intuitionistic semantics” [21]): if $(l, g, \mathcal{J}), \iota \models P$ then $(l, g \circ g', \mathcal{J}), \iota \models P$.

Five of the principles of Fig. 1 are direct consequences of the semantics.

Lemma 1. *The CoLoSL reasoning principles FORGET, MERGE, SHIFT, WEAKEN, and COPY are valid.*

Proof (sketch). The cases of WEAKEN and COPY are straightforward. For FORGET, MERGE, and SHIFT, we note in [20] that action model closure is preserved by picking a smaller subjective state, taking the union of subjective states and their interference assertions, and shifting the interference assertion, respectively. \square

3.3 Reducing CoLoSL principles to separation logic entailments

We turn to the remaining two principles, EXTEND and SHIFT, and to the stability of assertions. These involve reasoning outside our assertion language, potentially requiring semantic reasoning in the model. Fortunately, it is enough to work with a partial axiomatisation for all three conditions to verify our examples. In this section, we give cut-down versions of these rules for a fragment of the CoLoSL assertion language where the nesting of subjective views is not permitted and interference assertions cannot mention subjective views. This restriction is easily lifted: assertions with nested boxes can always be *flatten* into logically equivalent assertions with no nesting; and interference assertions mentioning other subjective views in their actions may be rewritten into ones that do not. See [20] for the full details.

Confinement The soundness of CoLoSL hinges on the fact that, given a world (l, g, \mathcal{J}) , the action model \mathcal{J} contains all actions that could possibly affect the shared state g . This was captured by a well-formedness condition in the definition of worlds (Def. 5) as $g \textcircled{c} \mathcal{J}$, stipulating that the actions in \mathcal{J} are confined to the shared state g . It is also possible to extend g at any time. Any part l' of the local state can migrate to the shared state under a new set of actions \mathcal{J}' , yielding a new shared state $g \circ l'$ and action model $\mathcal{J} \cup \mathcal{J}'$. This migration is only permitted if $l' \textcircled{c} \mathcal{J}'$. This confinement condition means that the extension does not invalidate the views of the threads.

The technical definition of confinement of an action model on a logical state is given in the technical report [20]. Intuitively, it means that, whenever an action $a = (p, q, c)$ of \mathcal{J}' is enabled, the pre-state p must be a substate of l' . It is possible for some of the catalyst c to lie outside l' , since the fact that it does not change during the course of the action means that it will not have an effect on the views of other threads. For example, recall the interpretation of s_x given by S just before Def. 3. The action $a_0 = ((\emptyset, \{\ell : 0\}, \emptyset), (\emptyset, \{\ell : 1\}, \emptyset), (\emptyset, \{\ell - 1 : 1\}, \emptyset)) \in S$ is confined to the logical state $l_0 = (\emptyset, \{\ell : 0\}, \emptyset)$ because l_0 is the first component of the action, and state $(\emptyset, \{\ell - 1 : 0\}, \emptyset)$ because it is incompatible with the action,

$$\begin{array}{c}
\frac{P \vdash f \quad f \blacktriangleright I}{P \textcircled{C} I} \quad \frac{\forall ([A]: \exists \bar{x}. p \rightsquigarrow q) \in I. f \blacktriangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}}{f \blacktriangleright I} \quad \frac{f \wp p \vdash_{\text{SL}} \text{false}}{f \blacktriangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}} \\
\frac{p \vdash_{\text{SL}} p' * r \quad q \vdash_{\text{SL}} q' * r \quad \text{exact}(r) \quad f \blacktriangleright \{[A]: \exists \bar{x}. p' \rightsquigarrow q'\}}{f \blacktriangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}} \\
\frac{(p \multimap f) * q \vdash_{\text{SL}} f \quad f \Leftrightarrow \bigvee_{i \in J} f_i \quad (\text{precise}(f_i) \wedge f_i \wp p \vdash_{\text{SL}} f_i \text{ for } i \in J)}{f \blacktriangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}}
\end{array}$$

Figure 3: Selected confinement and local fencing rules. We assume that variables in \bar{x} do not appear free in f , and in f_i for any i .

but not state $(\emptyset, \{\ell-1:1\}, \emptyset)$ because the action can potentially affect address ℓ . The definition of confinement also requires that all states resulting from the successive application of actions in \mathcal{J}' themselves confine all actions in \mathcal{J}' . For instance, we also require that $l_1 = (\emptyset, \{\ell:1\}, \emptyset)$ resulting from the application of a_0 on l_0 , $l_2 = (\emptyset, \{\ell:2\}, \emptyset)$ resulting from the application of a_1 on l_1 , and so forth, *all* confine the actions in S . Inspired by the LRG approach [10], we achieve this by first finding a set that is invariant under all actions in \mathcal{J}' (called a *fence*), then checking the confinement condition for each action. We provide the technical details of confinement in [20].

For our examples, it is in fact enough to work with confinement in the logic. We lift the notion of confinement to assertions, and write $P \textcircled{C} I$ when the set of states described by P confines the actions of interface assertion I . In the logic, the shared state can be extended by the resources in P under the interference assertion I via the EXTEND principle, which requires that $P \textcircled{C} I$ be established.

Fig. 3 presents a set of rules which reduces $P \textcircled{C} I$ to a series of entailments in ordinary separation logic. As expressed by the first rule, $P \textcircled{C} I$ holds if there is a weaker local assertion f that acts as a *local fence* for I , written $f \blacktriangleright I$. This relation states that f must be invariant under all actions of I and must confine the actions in f . In [20], we show that it is always possible to weaken an arbitrary assertion into a local assertion. This fencing condition is checked for each action in I (see the second rule). For each action $[A]: \exists \bar{x}. p \rightsquigarrow q$, the three remaining rules of the figure may apply. In the first of these rules, the action cannot possibly fire, because its precondition does not agree with f : no state satisfying f may be extended such that a subpart satisfies p . The second of these rules allows us to trim a *neutral* part r (corresponding to a part of the catalyst in the interpretation of this action) of an action $[A]: \exists \bar{x}. p \rightsquigarrow q$ appearing both in p and q . This only applies when r is *exact*, i.e. satisfied by at most one logical state:⁷ the part of the state denoted by r is then uniquely determined and left unchanged by the action. The last rule finally reduces local fencing to entailment checking, provided the fence f can be expressed as a disjunction of *precise* assertions, i.e. assertions

⁷ $\text{exact}(p) \triangleq \forall \iota, l_1, l_2. l_1, \iota \models_{\text{SL}} p$ and $l_2, \iota \models_{\text{SL}} p$ implies $l_1 = l_2$

$$\begin{array}{c}
\frac{P \vdash f \quad I \sqsubseteq^f I'}{I \sqsubseteq^P I'} \quad \frac{}{I \sqsubseteq^f I} \quad \frac{f \triangleright I \cup I_1 \quad I_1 \sqsubseteq^f I_2}{I \cup I_1 \sqsubseteq^f I \cup I_2} \\
\frac{p \vdash_{\text{SL}} p' * r \quad q \vdash_{\text{SL}} q' * r \quad \text{exact}(r) \quad f \perp p'}{\{[A]: \exists \bar{x}. p \rightsquigarrow q\} \sqsubseteq^f \emptyset} \quad \frac{(p \multimap (f \wp p)) * q \vdash_{\text{SL}} \text{false}}{\{[A]: \exists \bar{x}. p \rightsquigarrow q\} \sqsubseteq^f \emptyset} \\
\frac{f \wp p \vdash_{\text{SL}} \bigvee_{i \in J} f \wp (p * r_i) \quad \text{exact}(r_i) \text{ for } i \in J}{\{[A]: \exists \bar{x}. p \rightsquigarrow q\} \equiv^f \bigcup_{i \in J} \{[A]: \exists \bar{x}. p * r_i \rightsquigarrow q * r_i\}} \\
\frac{\bigcup_{i \in K, j \in J} \{[A]: \exists \bar{x}. p_i \rightsquigarrow q_j\} \equiv^{\text{true}} \{[A]: \exists \bar{x}. \bigvee_{i \in K} p_i \rightsquigarrow \bigvee_{j \in J} q_j\}}{\text{true} \triangleright I} \\
\frac{\forall [A]: \exists \bar{x}. p \rightsquigarrow q \in I. f \triangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}}{f \triangleright I} \quad \frac{p \perp q \quad (p \multimap (f \wp p)) * q \vdash_{\text{SL}} f}{f \triangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}} \\
\frac{f \perp p}{f \triangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}} \quad \frac{p \vdash_{\text{SL}} p' * r \quad q \vdash_{\text{SL}} q' * r \quad \text{exact}(r) \quad f \triangleright \{[A]: \exists \bar{x}. p' \rightsquigarrow q'\}}{f \triangleright \{[A]: \exists \bar{x}. p \rightsquigarrow q\}}
\end{array}$$

Figure 4: Selected action shifting rules. We write $I \equiv^f I'$ for $I \sqsubseteq^f I' \wedge I' \sqsubseteq^f I$, and $p \perp q$ to denote that states satisfying p and q have empty intersections. We assume that variables in \bar{x} do not appear free in f .

satisfied by at most one substate of each logical state.⁸ The first premise states that f is invariant under the action $[A]: \exists \bar{x}. p \rightsquigarrow q$, similar to the way that RGSep encodes stability checks as separation logic entailments. Informally, it reads: for any state in f , remove a part satisfying p , add a state satisfying q , and the result should still be in f . The third premise checks the confinement condition: given a state in local assertion f_i ($l_1 \circ l_2 \vDash f_i$), and a state in p ($l_2 \circ l_3 \vDash p$ with l_1 and l_3 disjoint), the combined state ($l_1 \circ l_2 \circ l_3$) must also be in f_i . Hence, by precision of f_i , we have $l_1 \circ l_2 \circ l_3 = l_1 \circ l_2$, i.e. $l_2 \circ l_3 \leq l_1 \circ l_2$ as required.

Shifting and fencing Fig. 4 presents a partial axiomatisation of the shifting condition $I \sqsubseteq^P I'$ required by the SHIFT principle. As with confinement, we omit the direct semantic definition. Intuitively, the relation $I \sqsubseteq^P I'$ means that interference assertion I can be replaced by I' , because I and I' describe the same interference with respect to the states described by P . The first rule weakens the shifting condition from assertions to local fence assertions. The third rule reduces the shifting judgement $I \cup I_1 \sqsubseteq^f I \cup I_2$ to the simpler $I_1 \sqsubseteq^f I_2$, provided that f is invariant with respect to $I \cup I_1$, written $f \triangleright I \cup I_1$. This invariant fencing condition is necessary: $I_1 \sqsubseteq^f I_2$ only states that I_1 and I_2 have the same effect with respect to f ; $f \triangleright I \cup I_1$ states that f is an *invariant* of the shared state under the combined interferences of I and I_1 .

The next two rules describe situations where it is impossible to apply the action to f : when the precondition of the action is entirely outside f ; or when

⁸ $\text{precise}(p) \triangleq \forall \iota, l, l_1, l_2. l_1 \leq l, l_2 \leq l, l_1, \iota \vDash_{\text{SL}} p$, and $l_2, \iota \vDash_{\text{SL}} p$ implies $l_1 = l_2$

the postcondition is incompatible with that part of f not associated with the precondition. The notation $p \perp q$ asserts that states described by p and q have *empty intersections*: whenever $l_1, \iota \models_{\text{SL}} p$ and $l_2, \iota \models_{\text{SL}} q$, we have $l_1 \perp l_2$ as defined in §3.1. This is expressible in standard separation logic as follows:

$$p \perp q \Leftrightarrow p \vdash_{\text{SL}} \neg(\text{true} \star (\neg \text{emp} \wedge (\text{true} \multimap q)))$$

The next rule is a shifting equivalence that uses the knowledge embodied by f of all the possible states that the subjective shared state may be in to rewrite an action into an equivalent one. More precisely, if whenever the precondition p of the action agrees with f then one of the r_i 's is also true, then adding r_i as a neutral part of the action produces the same behaviour. We can use this rule (with the single $r_0 = x \mapsto v$) to justify the shiftings of §2. The fact that the r_i 's are exact guarantees that no piece of state in r_i is mutated by the “larger” action. In general, it may not be the case that a single exact assertion can be added, but it may be the case that a disjunction of exact facts holds. The last shifting equivalence is straightforward.

The last five rules partially axiomatise the *fencing* relation $f \triangleright I$. Most are similar to those for local fencing $f \blacktriangleright I$. The first two state that *true* fences any interference assertion, and fencing can be checked per action. The last one states that, as for local fencing, neutral parts of actions may be ignored. The one before that states that, contrarily to local fencing, actions are allowed to have effects outside of the fence. If the action precondition does not intersect with the fence, then its effect is entirely outside the fence and the action may be ignored.

Let us now focus on the third of these rules, which states that whenever p and q do not intersect in an action $[A]: \exists \bar{x}. p \rightsquigarrow q$ (e.g. when their common parts have been removed using the last rule), then the application of the action must preserve the fence f . Contrarily to the case of local fencing, the action is allowed to act partly outside of f , hence the state on which the action is applied is $f \uplus p$. However, the whole of the postcondition q of the action is then added, and the resulting state must still be in f . One might instead have expected that only parts of the resulting state need to be represented in f , to mimic the relationship between f and p (and indeed, this is all that is required for stability, as we shall see next). However, we do need the full q . Recall that shifting asserts that two interference assertions have the same effect even after an arbitrary number of steps. Doing otherwise would be unsound because there would be no guarantee that I accounts for all possible actions on that part of q that would be discarded, since it would not be part of f . Hence, we could end up with a new interference assertion I' that breaks the original action model closure property. Finally, the reason why p and q must not intersect stems from similar considerations. If p and q had a non-empty intersection c , such that c is not part of f , then this would force the fence to account for c which would prevent us from forgetting actions associated with it using shifting. For instance, it would make it impossible to forget actions as we do in §2.

Stability Fig. 5 partially axiomatises the stability of assertions of the form $\exists \bar{x}. p \star \star_{i \in J} \boxed{q_i}_{I_i}$. We work with a more restricted form of assertions than our

$$\begin{array}{c}
\frac{}{\text{Stable}(p)} \qquad \frac{\text{Stable}(P)}{\text{Stable}(\exists x. P)} \qquad \frac{\forall j \in J. \text{Stable}(q_j, I_j, p, \biguplus_{i \in J} q_i)}{\text{Stable}(p \star \star_{i \in J} \boxed{q_i}_{I_i})} \\
\frac{\forall [A]: \exists \bar{x}. p_1 \rightsquigarrow p_2 \in I. \text{Stable}(q, \{[A]: \exists \bar{x}. p_1 \rightsquigarrow p_2\}, p, r)}{\text{Stable}(q, I, p, r)} \qquad \frac{[A] \star p \star r \vdash_{\text{SL}} \text{false}}{\text{Stable}(q, \{[A]: \exists \bar{x}. p_1 \rightsquigarrow p_2\}, p, r)} \\
\frac{p \star (r \uplus p_1) \vdash_{\text{SL}} \text{false}}{\text{Stable}(q, \{[A]: \exists \bar{x}. p_1 \rightsquigarrow p_2\}, p, r)} \qquad \frac{(p_1 \text{---} \otimes (q \uplus p_1)) \star p_2 \vdash_{\text{SL}} q \star \text{true}}{\text{Stable}(q, \{[A]: \exists \bar{x}. p_1 \rightsquigarrow p_2\}, p, r)}
\end{array}$$

Figure 5: Selected rules for stability checks. We assume that variables in \bar{x} do not appear free in p , q , and r .

two previous axiomatisations for simplicity; see [20] for the general case. The first two rules state that local assertions are always stable and existentials can be eliminated. The next rule states that checking the stability of $p \star \star_{i \in J} \boxed{q_i}_{I_i}$ boils down to establishing, for each local assertion q_j for $j \in J$, that the four-place predicate $\text{Stable}(q_j, I_j, p, \biguplus_{i \in J} q_i)$ holds: this means that q_i is stable under the interference assertion I_i , a local context p , and a shared context made of the \uplus -combination of all the subjective assertions (including q_j). In turn, checking this fact reduces to checking stability for each action of I_j . The last three rules deal with checking stability for a single action, in a similar way to the fencing rules above.

The first of these rules is unfamiliar. Unlike fencing, stability checking for assertions only has to be checked against actions for which the environment may have the capability. If the capability required by the action cannot exist separately from those held by the assertion (that is, those in $p \star r$) then the environment cannot possibly own the capability to perform the action. Similarly, an action whose precondition is incompatible with the assertion $p \star r$ cannot possibly fire, as stated by the next rule. The last rule checks that q is preserved by the effect of an action. Again, there is a crucial difference with the corresponding check for fencing: in p_2 , the action may bring in some newly-shared state, hence the result $q \star \text{true}$; but the FORGET rule allows us to immediately discard it, if appropriate.

Proof reuse Note that, once the local fencing judgement $f \blacktriangleright I$ has been established, it automatically establishes the weaker fencing judgement $f \triangleright I$, which in turn implies that f is stable under I . Also, fencing is preserved by action shifting. These observations provide the admissible rules:

$$\frac{f \blacktriangleright I}{f \triangleright I} \qquad \frac{p \triangleright I}{\text{Stable}(\boxed{p}_I)} \qquad \frac{f \triangleright I \quad I \sqsubseteq^f I'}{f \triangleright I'}$$

The other directions are not valid in general: fencing lacks the confinement condition required by local fencing; stability of p under I may omit part of the state resulting from an action application to re-establish p , which is not allowed

in fencing; and a fence f for a smaller interference assertion need not be a fence for a larger interference assertion.

3.4 Soundness

We prove the soundness of CoLoSL, parametrised by the underlying models of machine states (\mathbb{M}) and capabilities (\mathbb{K}). We appeal to the general soundness result of the views framework, providing parameters such as the reification function in Def. 9 and proving lemmas required to make the result hold.

The main part of the proof establishes the soundness of the following rule for atomic commands, where \vdash_{SL} represents standard sequential separation logic:

$$\frac{\vdash_{\text{SL}} \{p\} \mathbb{C} \{q\} \quad P \Rrightarrow^{\{p\}\{q\}} Q}{\vdash \{P\} \langle \mathbb{C} \rangle \{Q\}} \text{ATOM}$$

This rule is present in logics arising from CAP [7,23]. For CoLoSL, the *repartitioning* $P \Rrightarrow^{\{p\}\{q\}} Q$ holds if, from any world (l, g, \mathcal{J}) satisfying P , whenever parts of $l \circ g$ satisfies p then substituting that part for any other satisfying q will yield a state $l' \circ g'$ such that there exists \mathcal{J}' such that (l', g', \mathcal{J}') satisfies Q . Moreover, the passage from (l, g, \mathcal{J}) to (l', g', \mathcal{J}') can be achieved via a succession of valid updates from \mathcal{J} and valid extension steps. The details can be found in the accompanying technical report [20].

Semantic validity of Hoare triples depends on the definition of an operational semantics $\mathbb{C}, m \rightarrow^* \mathbb{C}', m'$, where $m, m' \in \mathbb{M}$, and a reification function that relates a CoLoSL world to a concrete machine state.

Definition 9 (Reification). *The reification, $[\cdot]_{\mathbb{W}} : \text{World} \rightarrow \mathbb{M}$ is defined as:*

$$[(\sigma_l, h_l, \kappa_l), (\sigma_g, h_g, \kappa_g), \mathcal{J}]_{\mathbb{W}} \triangleq (\sigma_l \uplus \sigma_g, h_l \uplus h_g)$$

Definition 10 (Valid triple). *A triple is valid, written $\vDash \{P\} \mathbb{C} \{Q\}$, if and only if, for all $\iota \in \text{LEnv}$, $w \in \text{World}$ and $m, m' \in \mathbb{M}$,*

$$(w, \iota \vDash P \text{ and } \mathbb{C}, [w]_{\mathbb{W}} \rightarrow^* \text{skip}, m') \text{ implies } \exists w'. w', \iota \vDash Q \text{ and } m' = [w']_{\mathbb{W}}$$

Theorem 1 (Soundness). *If $\vdash \{P\} \mathbb{C} \{Q\}$ then $\vDash \{P\} \mathbb{C} \{Q\}$.*

4 Examples

4.1 Concurrent Spanning Tree

Programs manipulating arbitrary graphs present a significant challenge for compositional reasoning, because deep, *unspecified* sharing between different components of a graph results in changes to one subgraph affecting other subgraphs pointing into it. This makes it hard to reason about updates to each subgraph in isolation. In a concurrent setting, this difficulty is compounded by the fact that threads working on different parts of the graph may affect each other in ways that are difficult to reason about locally to each subgraph. Using a concurrent spanning

tree algorithm, we demonstrate that CoLoSL reasoning might be just the right approach, as subjective views naturally provide arbitrary overlapping views of the shared state where interferences can be naturally tailored to a given subjective view. With CoLoSL, we have achieved local reasoning about the shared state for this challenging program.

Our example, presented in Fig. 8, operates on a *directed binary graph* (henceforth simply *graph*): that is, a directed graph where each node has at most two successors, called its left and right children. The program concurrently computes an in-place spanning tree of the graph (i.e. a tree that covers all nodes of the graphs from a given root), as follows: each time a new node is encountered, two new threads are created that each prune the edges of its left and right children recursively. A mark bit is associated with each node to keep track of which nodes have already been visited. Each thread returns whether it marked the node it was called on itself or whether somebody else did. In the latter case, the parent thread removes the link from its own root node to the corresponding child. Intuitively, it is allowed to do so because the child has already been reached via some other path in the graph since it was marked by another thread.

We will prove that, given a shared graph as input, the program always returns a tree, i.e. all sharing and cycles have been appropriately removed. Pleasingly, the CoLoSL specification captures just the subgraph manipulated by the thread, instead of the whole graph.

To reason about this program, following [14] we use two representations of graphs. The first is a mathematical representation $\gamma = (V, E)$ where V is a finite set of vertices and $E : V \rightarrow (V \uplus \{\text{null}\}) \times (V \uplus \{\text{null}\})$ is a function associating each vertex with at most two successors, where null denotes the absence of a child. We write $n \in \gamma$ for $n \in V$, $\gamma(n)$ for $E(n)$ which also assumes $n \in \gamma$, $n \sim_\gamma n'$ for $n' \in \gamma(n)$, and \sim_γ^* for the reflexive and transitive closure of \sim_γ .

Mathematical graphs are connected to their in-memory representations by a predicate **graph**(x, γ) shown in Fig. 6, denoting a spatial (in-heap) graph rooted at address x corresponding to the mathematical graph γ . This predicate uses the overlapping conjunction to account for potential sharing between the left and right children and for potential cycles in the graph. Each vertex is represented as three consecutive cells in the heap tracking the mark bit and the addresses of the left and right subgraphs. We write $x \mapsto m, l, r$ for $x \mapsto m * x + 1 \mapsto l * x + 2 \mapsto r$, and $x.l$ and $x.r$ for $x + 1$ and $x + 2$, respectively. When vertex x is in the unmarked state $\mathbf{U}(x, l, r)$, the whole cell $x \mapsto 0, l, r$ resides in the shared state. In the marked state $\mathbf{M}(x)$, only $x \mapsto 1$ is owned. In both cases, the shared state also contains the left and right subgraphs represented by $\mathbf{G}(l, \gamma)$ and $\mathbf{G}(r, \gamma)$.

To understand the difference in ownership between $\mathbf{U}(x, l, r)$ and $\mathbf{M}(x)$, let us look at the interference associated with the graph, which is the union of interferences pertaining each vertex in the graph. For each vertex $n \in \gamma$, the only permitted action is that of marking n . (For simplicity, this action does not require any capability.) When changing the mark field of n from 0 to 1, the current thread also claims ownership of its left and right pointers. Indeed, we observe that other threads need not access the children of n once they see that

$$\begin{aligned} \mathbf{graph}(x, \gamma) &\triangleq \boxed{\mathbf{G}(x, \gamma)}_{I_\gamma} & I_\gamma &\triangleq \bigcup_{n \in \gamma} I(n) & I(n) &\triangleq \{[\emptyset] : \exists l, r. \mathbf{U}(n, l, r) \rightsquigarrow \mathbf{M}(n)\} \\ \mathbf{G}(x, \gamma) &\triangleq (x = \mathbf{null} \wedge \mathbf{emp}) \vee \exists l, r. \gamma(x) = (l, r) \wedge & \mathbf{U}(x, l, r) &\triangleq x \mapsto 0, l, r \\ & (\mathbf{U}(x, l, r) \vee \mathbf{M}(x)) \wp \mathbf{G}(l, \gamma) \wp \mathbf{G}(r, \gamma) & \mathbf{M}(x) &\triangleq x \mapsto 1 \end{aligned}$$

Figure 6: Globally-shared graph predicate.

$$\begin{aligned} \mathbf{g}(x, \gamma) &\triangleq (x = \mathbf{null} \wedge \mathbf{emp}) \vee \exists l, r. \gamma(x) = (l, r) \wedge \\ & \boxed{\mathbf{U}(x, l, r) \vee \mathbf{M}(x)}_{I(x)} * \mathbf{g}(l, \gamma) * \mathbf{g}(r, \gamma) \\ \mathbf{t}(x, \gamma) &\triangleq (x = \mathbf{null} \wedge \mathbf{emp}) \vee (\exists l, r. \gamma(x) = (l, r) \wedge \exists l' \in \{l, \mathbf{null}\}, r' \in \{r, \mathbf{null}\}. \\ & \boxed{\mathbf{M}(x)}_{I(x)} * x.l \mapsto l' * \mathbf{t}(l', \gamma) * x.r \mapsto r' * \mathbf{t}(r', \gamma)) \end{aligned}$$

Figure 7: Locally-shared graph predicate.

it has already been marked. The atomic CAS (compare-and-swap) instruction prevents two threads from concurrently marking the same node and claiming ownership of the same resource.

The $\mathbf{graph}(x, \gamma)$ predicate defined in Fig. 6 is a *global* subjective view of the graph that contains all vertices and the interference associated with γ . However, our spanning tree algorithm operates *locally* as it is called upon recursively for each node. That is, for each $\mathbf{span}(n)$ call (where $n \in \gamma$), the footprint of the call is limited to the subgraph rooted at n . Moreover, in order to reason about the concurrent recursive calls $\mathbf{span}(x \rightarrow l) \parallel \mathbf{span}(x \rightarrow r)$, we need to *split* the state into two using $*$, pass each constituent state to the relevant thread and $*$ -combine the resulting states, as required by the PAR rule. We thus conduct our proof with respect to a *local* description of the graph, $\mathbf{g}(x, \gamma)$ as defined in Fig. 7. The definition of the $\mathbf{g}(x, \gamma)$ predicate is similar to that of $\boxed{\mathbf{G}(x, \gamma)}_{I_\gamma}$ except that the single view $\boxed{\mathbf{G}(x, \gamma)}_{I_\gamma}$ has been broken down into individual views for each vertex n reachable from x . Moreover, the interference assertion of each local view concerning a vertex $n \in \gamma$ has been shifted from I_γ to $I(n)$ so as to reflect only those actions that affect n .

The $\mathbf{t}(x, \gamma)$ predicate, also given in Fig. 7, represents a *tree* rooted at x , as is standard in separation logic [21], and consists, once fully unfolded, of one subjective view for each vertex n reachable from x in γ . The assertion of the subjective view for x reflects the fact that x has been marked, and its left and right pointers have been claimed by the marking thread and moved into its local state. The vertex l' addressed by the left pointer of x corresponds to either the initial value l prior to marking, or to \mathbf{null} when l has more than one predecessor and has been marked by another thread.

Our goal is to prove that the following specification holds using the global graph predicate:

$$\{\mathbf{graph}(x, \gamma)\} \mathbf{b} = \mathbf{span}(x) \{(\mathbf{b} = 0 \wedge \mathbf{emp}) \vee (\mathbf{b} = 1 \wedge \mathbf{t}(x, \gamma))\} \quad (1)$$

```

b = span(x) // {g(x, γ)}
{ // {(x = null ∧ emp) ∨ (∃l, r. γ(x) = (l, r) ∧  $\boxed{U(x, l, r) \vee M(x)}$   $_{I(x)}$  * g(l, γ) * g(r, γ))}
  if (x == null) {
    // {x = null ∧ emp}
    // {t(x, γ)}
    return 1;
  } // {(∃l, r. γ(x) = (l, r) ∧  $\boxed{U(x, l, r) \vee M(x)}$   $_{I(x)}$  * g(l, γ) * g(r, γ))}
  res = ( CAS(x, 0, 1) );
  // {(∃l, r. γ(x) = (l, r) ∧  $\boxed{M(x)}$   $_{I(x)}$  * g(l, γ) * g(r, γ) *
  // ((res = 0 ∧ emp) ∨ (res = 1 ∧ x.l ↦ l * x.r ↦ r))}
  // {(res = 0 ∧ emp) ∨ (res = 1 ∧ ∃l, r. γ(x) = (l, r) ∧
  //  $\boxed{M(x)}$   $_{I(x)}$  * g(l, γ) * g(r, γ) * x.l ↦ l * x.r ↦ r)}
  if (res) {
    // {(res = 1 ∧ ∃l, r. γ(x) = (l, r) ∧  $\boxed{M(x)}$   $_{I(x)}$  * g(l, γ) * g(r, γ) * x.l ↦ l * x.r ↦ r}
    // {g(l, γ)}
    b1 = span(x->l);
    // {(b1 = 0 ∧ emp) ∨ (b1 = 1 ∧ t(l, γ))}
    // {(b1 = 0 ∧ emp) ∨ (b1 = 1 ∧ t(l, γ))} * {(b2 = 0 ∧ emp) ∨ (b2 = 1 ∧ t(r, γ))}
    // {(res = 1 ∧ ∃l, r. γ(x) = (l, r) ∧  $\boxed{M(x)}$   $_{I(x)}$  * x.l ↦ l * x.r ↦ r *
    // ((b1 = 0 ∧ emp) ∨ (b1 = 1 ∧ t(l, γ))) * ((b2 = 0 ∧ emp) ∨ (b2 = 1 ∧ t(r, γ)))}
    if (!b1) { x->l = null; }
    if (!b2) { x->r = null; }
    // {(res = 1 ∧ ∃l, r. γ(x) = (l, r) ∧  $\boxed{M(x)}$   $_{I(x)}$  *
    // {∃l' ∈ {l, null}. x.l ↦ l' * t(l', γ) * ∃r' ∈ {r, null}. x.r ↦ r' * t(r', γ)}
    // {res = 1 ∧ t(x, γ)}
    // {(res = 0 ∧ emp) ∨ (res = 1 ∧ t(x, γ))}
  }
  return res;
} // {(b = 0 ∧ emp) ∨ (b = 1 ∧ t(x, γ))}

```

Figure 8: Code and proof sketch of the concurrent spanning tree program. We omit the obvious variables as resource assertions.

This is achieved below by giving a proof of the analogous specification using the local graph predicate in Fig. 8, and demonstrating that the global graph specification implies the local one using the principles of CoLoSL from Fig. 1. The proof sketched in Fig. 8 is mostly straightforward. One subtlety is the consequence step just before the second **if** statement. There, we first distribute the shared state $\boxed{M(x)}$ $_{I(x)}$ * g(l, γ) * g(r, γ) over the disjunction, then use the fact that it implies emp to discard it in the left disjunct. This proof demonstrates that our CoLoSL reasoning really is compositional, in the sense that we are doing local reasoning on the shared state (the subgraphs).

Let us show how to derive the specification (1) from the one obtained in Fig. 8. We introduce the *iterative star* operator \circledast ; when iterating over the empty set, it denotes emp by convention (needed below, when x is null). We define

$$\begin{aligned}
P(x, \gamma) &\triangleq \circledast_{x \rightsquigarrow^* n} \exists l, r. \gamma(n) = (l, r) \wedge (U(n, l, r) \vee M(n)) \\
Q(x, \gamma) &\triangleq \circledast_{x \rightsquigarrow^* n} \exists l, r. \gamma(n) = (l, r) \wedge \boxed{U(n, l, r) \vee M(n)}_{I(n)}
\end{aligned}$$

From the definitions of $\mathbf{G}(x, \gamma)$ and $\mathbf{g}(x, \gamma)$, one can show that

$$\mathbf{graph}(x, \gamma) \iff \boxed{P(x, \gamma)}_{I_\gamma} \quad \mathbf{g}(x, \gamma) \iff Q(x, \gamma)$$

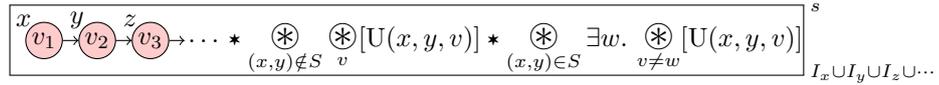
The specification (1) then follows from CONSQ, Fig. 8, and the derivation below:

$$\begin{aligned} \mathbf{graph}(x, \gamma) &\stackrel{\text{(COPY)}}{\Rightarrow} \bigotimes_{x \rightsquigarrow_\gamma^* n} \boxed{P(x, \gamma)}_{I_\gamma} \\ &\stackrel{\text{(FORGET)}}{\Rightarrow} \bigotimes_{x \rightsquigarrow_\gamma^* n} \left(\exists l, r. \gamma(n) = (l, r) \wedge \boxed{U(n, l, r) \vee M(n)}_{I_\gamma} \right) \\ &\stackrel{\text{(SHIFT)}}{\Rightarrow} \bigotimes_{x \rightsquigarrow_\gamma^* n} \left(\exists l, r. \gamma(n) = (l, r) \wedge \boxed{U(n, l, r) \vee M(n)}_{I(n)} \right) \stackrel{\text{def}}{\iff} \mathbf{g}(x, \gamma) \end{aligned}$$

4.2 Set Module

Finally, we give a pictorial description of our reasoning about a concurrent set module implemented as a singly-linked list. We compare our CoLoSL reasoning with the original CAP reasoning of [7], demonstrating that our CoLoSL reasoning provides more concise proofs using our local reasoning about shared state.

Consider the following diagram which illustrates the CAP set predicate of [7]:



The set is represented as a sorted singly-linked list with no duplicate elements. The list starts at address x with value v_1 , points to the next element at address y with value v_2 , and so forth. Hereafter, we write $\mathbf{node}(x, v, y)$ to denote a node at address x , with value v and successor y .

All nodes of the list reside in a single shared region labelled s and the interference on the list is the combined interference associated with each constituent node. Each node at a given address x is associated with a set of update capabilities of the form $[U(x, y, v)]$ for *all* possible addresses y and *all* possible values v . This is to capture all potential successor addresses y and all potential values v that may be stored at address x . In order to modify a node, a thread can acquire the lock associated with the node and subsequently claim the relevant update capability.

Since in CAP the capabilities associated with a region can only be generated upon its creation, the shared region is required to keep track of all possible update capabilities $[U(x, y, v)]$ associated with all addresses x (including those not currently in the domain of the list), all addresses y and all values v . At any one point, given $\mathbf{node}(x, v, y)$, the only update capability that can be claimed by a thread (through locking) is the one that reflects its current status, namely $[U(x, y, v)]$. As a result, an auxiliary mathematical set S is used to track those nodes of the list that are currently locked and thus infer which $[U]$ capabilities have been claimed. The distribution of update capabilities is captured by the two assertions written as the *infinite multiplicative star operator* \bigotimes . The first part of

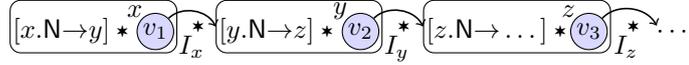
the assertion states that given any node at address x with successor y , if it is not locked, i.e. $(x, y) \notin S$, then all of its update capabilities of the form $[U(x, y, v)]$ lie in the shared region for all values v . Dually, if it is locked, i.e. $(x, y) \in S$, then the update capabilities for all values v but one ($w \neq v$) are in the shared region.

This CAP set predicate is unnecessarily complicated. It is counter-intuitive to have to account for the capabilities associated with addresses not in the domain of the list. Moreover, each thread observes all nodes in the list and thus needs to account for their associated interference.

TaDA [3] took the first steps towards addressing the above shortcomings of the CAP approach. TaDA regions are parametric in the separation algebra of capabilities (called guards). As such, one can choose a more suitable algebra to axiomatise the desired behaviour of capabilities. While TaDA’s approach is much cleaner than that of CAP, it nevertheless requires the foresight of specifying all desired interference associated with the region upon its creation. As such, interference specifications are *static* and cannot be extended with new behaviour even when the existing resources are left untouched. On the other hand, as well as being parametric in its capability separation algebra, the dynamic subjective views of CoLoSL provide local reasoning about the shared resource and its interference.

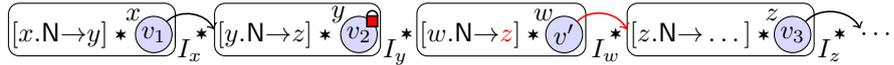
We proceed with the CoLoSL proof of the set implementation. Recall from §3 that CoLoSL is parametric in the separation algebra of capabilities. We thus instantiate it with a heap-like capability separation algebra that is *stateful* and demonstrate that this allows for a more concise proof.

We specify the set predicate as the \star -composition of subjective views associated each node in the singly-linked list as illustrated by:

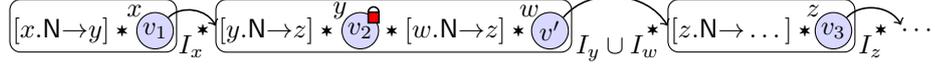


The interference on each subjective view is limited to the node in question. Associated with each node at address x is a “next” capability $[x.N \rightarrow y]$ that tracks its successor y . This is analogous to the $[U(x, y, v)]$ capability of CAP and we shortly demonstrate how it is utilised in our reasoning.

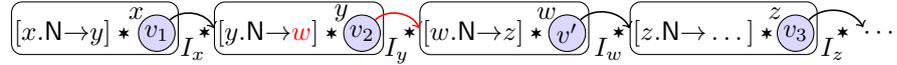
Since CoLoSL allows for dynamic extension of the shared state, we do not need to account for capabilities associated with all addresses. Instead, fresh capabilities are generated dynamically as needed. We demonstrate this by giving a reasoning outline of the $\text{add}(v')$ method that adds value v' to the set by inserting it in the sorted list. Suppose $v_2 < v' < v_3$, and thus a new node w with value v' is to be inserted after node y . The operating thread proceeds by traversing the list by hand-over-hand locking until it reaches node y . It then locks y and claims its next pointer and moves it to its local state, as allowed by I_y . Subsequently, the shared state is extended by the resources associated with the new node and its associated capabilities ($[w.N \rightarrow z]$) are generated on the fly as illustrated by:



Since the locking thread holds the next pointer of y in its local state, it modifies it to point to the new node w . It then unlocks y and returns its next pointer to the shared state. When inserting a new node between y and z , the associated interference assertion I_y allows y to be unlocked only if it has been directed to a new node whose successor is z . As such, the unlocking thread must demonstrate that the new node w does indeed point to z . In order to establish this, we use the MERGE principle to combine the subjective views of y and w as follows:



Finally, y is unlocked; its next pointer is returned to the shared state and its next capability is modified to reflect its new successor. Using the COPY, FORGET and SHIFT principles in order, we obtain the set predicate with w inserted.



We can reason about the remove operation in a similar fashion. The dynamic extension afforded by the EXTEND principle allows us to generate new capabilities only when needed and thus gives way to a concise proof. Moreover, rather than having a distinct capability to modify the element at address x , for each possible successor address y (as with $[U(x, y, v)]$ in CAP), we appeal to a single capability of the form $[x.N \rightarrow y]$ that is modified to $[x.N \rightarrow y']$ whenever x 's successor changes from y to y' . Lastly, using the reasoning principles of MERGE, FORGET, SHIFT and COPY, we can grow and shrink our subjective views as needed. This means that, at any one point, we only view the relevant parts of the shared state. The technical details can be found in [20].

Concluding Remarks We have introduced CoLoSL, a new program logic for reasoning locally about the shared state. We focus on subjective views, which expand and contract to provide a flexible treatment of both the shared resource and its interference. However, CoLoSL is still young, and lacks many features of its various cousins. There are many interesting ideas present in the literature: e.g. abstract states governed by state transition systems [25]; higher-order reasoning [23]; and abstract atomicity [3]. All these ideas require further investigation. Here, our aim was to simply introduce subjective views as a fundamental new way of underpinning such reasoning.

Acknowledgements We are grateful to Aquinas Hobor for providing us with the example of §2, and to Matthew Parkinson for suggesting to record catalysts in the model. We would also like to thank Pedro da Rocha Pinto for his continuous feedback on earlier versions of this paper. This research was funded by EPSRC grants K008528/1 and H008373/2.

References

1. R. Bornat, C. Calcagno, and H. Yang. Variables as resource in separation logic. *ENTCS*, 155, 2006.
2. C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *LICS*, 2007.
3. P. da Rocha Pinto, T. Dinsdale-Young, and P. Gardner. TaDA: A logic for time and data abstraction. In *ECOOP*, 2014.
4. E. Dijkstra. A belated proof of self-stabilization. *Distributed Computing*, 1(1), 1986.
5. E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11), 1974.
6. T. Dinsdale-Young, L. Birkedal, P. Gardner, M. J. Parkinson, and H. Yang. Views: compositional reasoning for concurrent programs. In *POPL*, 2013.
7. T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP*. Springer, 2010.
8. R. Dockins, A. Hobor, and A. W. Appel. A fresh look at separation algebras and share accounting. In *APLAS*. Springer, 2009.
9. M. Dodds, X. Feng, M. J. Parkinson, and V. Vafeiadis. Deny-guarantee reasoning. In *ESOP*. Springer, 2009.
10. X. Feng. Local rely-guarantee reasoning. In *POPL*. ACM, 2009.
11. P. Gardner, S. Maffeis, and G. D. Smith. Towards a program logic for JavaScript. In *POPL*. ACM, 2012.
12. P. Gardner, S. Maffeis, and G. D. Smith. Towards a program logic for JavaScript. In *POPL*. ACM, 2012.
13. P. Gardner, A. Raad, M. Wheelhouse, and A. Wright. Abstract local reasoning for concurrent libraries: mind the gap. In *MFPS*, 2014.
14. A. Hobor and J. Villard. The ramifications of sharing in data structures. In *POPL*, 2013.
15. S. S. Ishtiaq and P. W. O’Hearn. BI as an assertion language for mutable data structures. In *POPL*, 2001.
16. C. B. Jones. Specification and design of (parallel) programs. In *IFIP Cong.*, 1983.
17. A. Nanevski, R. Ley-Wild, I. Sergey, and G. A. Delbianco. Communicating state transition systems for fine-grained concurrent resources. In *ESOP*. Springer, 2014.
18. P. W. O’Hearn. Resources, concurrency, and local reasoning. *TCS*, 375(1-3), 2007.
19. P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL*. Springer, 2001.
20. A. Raad, J. Villard, and P. Gardner. CoLoSL: Concurrent Local Subjective Logic. <http://www.doc.ic.ac.uk/~azalea/ESOP2015/CoLoSL-TR.pdf>, 2014.
21. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*. IEEE Computer Society, 2002.
22. J. C. Reynolds. A short course on separation logic. <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/fox-19/member/jcr/wwwaac2003/notes7.ps>, 2003.
23. K. Svendsen and L. Birkedal. Impredicative concurrent abstract predicates. In *ESOP*. Springer, 2014.
24. K. Svendsen, L. Birkedal, and M. Parkinson. Modular reasoning about separation of concurrent data structures. In *ESOP*, 2013.
25. A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. *SIGPLAN Not.*, 48(9), Sept. 2013.
26. V. Vafeiadis and M. J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*. Springer, 2007.