

An Asynchronous Soundness Theorem for Concurrent Separation Logic

Paul-André Melliès

IRIF, CNRS & Université Paris Diderot

Léo Stefanescu

IRIF, Université Paris Diderot & CNRS

Abstract

Concurrent separation logic (CSL) is a specification logic for concurrent imperative programs with shared memory and locks. In this paper, we develop a concurrent and interactive account of the logic inspired by asynchronous game semantics. To every program C , we associate a pair of asynchronous transition systems $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ which describe the operational behavior of the code when confronted to its Environment or Frame — both at the level of machine states (S) and of machine instructions and locks (L). We then establish that every derivation tree π of a judgment $\Gamma \vdash \{P\}C\{Q\}$ defines a winning and asynchronous strategy $\llbracket \pi \rrbracket_{Sep}$ with respect to both asynchronous semantics $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$. From this, we deduce an asynchronous soundness theorem for CSL, which states that the canonical map $\mathcal{L} : \llbracket C \rrbracket_S \rightarrow \llbracket C \rrbracket_L$ from the stateful semantics $\llbracket C \rrbracket_S$ to the stateless semantics $\llbracket C \rrbracket_L$ satisfies a basic fibrational property. We advocate that this provides a clean and conceptual explanation for the usual soundness theorem of CSL, including the absence of data races.


1 Introduction

A simple way to understand an imperative (possibly nondeterministic) program C is to interpret it as a binary relation $[C] \subseteq S \times S$ between machine states $s, s' \in S$. In that approach, the statement $s[C]s'$ indicates that one execution trace (at least) of the program C has initial state $s \in S$ and final state $s' \in S$. One practical advantage of this description is that the binary relation $[C]$ abstracts away from the execution traces of the program C , and only retains their initial and final states. However crude, this abstraction is generally sufficient to analyze the properties of sequential imperative programs, and to establish the soundness of Hoare logic. Unfortunately, the abstraction becomes too coarse when one decides to shift to concurrent imperative programs with shared memory and locks, and to establish the soundness of a specification logic like Concurrent Separation Logic (CSL). To that purpose, it has long been recognized that one needs a proper account of the execution traces of the program C , see Brookes [5]. In this paper, we go one step further, and advocate that the soundness theorem of CSL, and more specifically the absence of data races, is intrinsically related to the asynchronous structure of the execution paths of C . Inspired by asynchronous game semantics, we interpret every concurrent imperative program C as a pair of asynchronous graphs $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ related by an asynchronous graph homomorphism

$$\mathcal{L}_C \quad : \quad \llbracket C \rrbracket_S \longrightarrow \llbracket C \rrbracket_L \quad (1)$$

We thus start by recalling the notion of *asynchronous graph* [14, 15] before discussing the relationship between time and space separation.

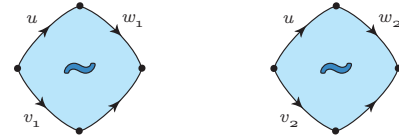
Asynchronous graphs A graph $G = (V, E, \partial^-, \partial^+)$ consists of a set V of vertices or nodes, a set of E of edges or transitions, and a source and a target function $\partial^-, \partial^+ : E \rightarrow V$. An *asynchronous graph* (G, \diamond) is a graph G equipped with a binary relation \diamond between paths $f, g : P \rightarrow Q$ of length 2, with the same source and target vertices. A pair (f, g) such that $f \diamond g$ is called a permutation tile and is depicted as a 2-dimensional tile between the paths $f = u \cdot v'$ and $g = v \cdot u'$ as follows:



The intuition conveyed by such a permutation tile $u \cdot v' \diamond v \cdot u'$ is that the two transitions u and v are independent. For that reason, the two paths $u \cdot v'$ and $v \cdot u'$ may be seen as equivalent up to scheduling. The binary relation \diamond is required to satisfy the two axiomatic properties below.

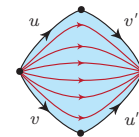
Axiom 1. The permutation relation \diamond is symmetric, in the sense that $u \cdot v' \diamond v \cdot u'$ implies $v \cdot u' \diamond u \cdot v'$ for all transitions u, v, u', v' .

Axiom 2. In the situation below where $u \cdot w_1 \diamond v_1 \cdot u_1$ and $u \cdot w_2 \diamond v_2 \cdot u_2$, one has that $v_1 = v_2$ if and only if $w_1 = w_2$.



Two paths $f, g : M \rightarrow N$ of an asynchronous graph are equivalent modulo one permutation tile $h_1 \diamond h_2$ when f and g factor as $f = d \cdot h_1 \cdot e$ and $g = d \cdot h_2 \cdot e$ for two paths $d : M \rightarrow P$ and $e : Q \rightarrow N$. We write $f \sim g$ when the path $f : M \rightarrow N$ is equivalent to the path $g : M \rightarrow N$ modulo a number of such permutation tiles. Note that the relation \sim is symmetric, reflexive and transitive, and thus defines an equivalence relation, closed under composition.

Separation in space and time The 2-dimensional permutation tiles $f \diamond g$ provide a topological means to reflect the *temporal* nature of independence in concurrency theory. Every permutation tile (2) indicates that the two transitions u and v are independent in time: they may be equivalently executed in the sequential order $u \cdot v'$ or in the sequential order $v \cdot u'$. Although all the asynchronous graphs considered in this paper are discrete, it is enlightening to take the topological intuition of “homotopy” seriously, and to imagine that the path $u \cdot v'$ could be transformed “continuously” into the path $v \cdot u'$ by a sequence of local deformations of the form



as it would be possible if one embedded our asynchronous graphs (G, \diamond) in the topological framework of directed homotopy, see [6].

Interestingly, in practical situations, the *temporal* independence of two transitions u and v is not primitive: it is a consequence of their *spatial* separation. In that respect, the idea of temporal independence may be seen as a layer of abstraction above the more concrete and machine-dependent idea of spatial separation. We illustrate this basic but important point by constructing an asynchronous graph (G, \diamond_G) based on a very simple machine model, consisting of

- a countable set \mathbf{Var} of variables, written x, y, \dots ,
- a countable set \mathbf{Val} of values, written v, w, \dots ,
- a countable set $\mathbf{Loc} \subseteq \mathbf{Val}$ of memory locations, written ℓ .

A *memory state* $\mu = (s, h)$ of the machine is defined as a pair consisting of two partial functions

$$s : \mathbf{Var} \rightarrow_{\text{fin}} \mathbf{Val} \quad h : \mathbf{Loc} \rightarrow_{\text{fin}} \mathbf{Val} \quad (3)$$

with finite domains, called the *stack* s and the *heap* h of the memory state μ . The instructions m of our machine are of three kinds:

$$x := v \quad x := [\ell] \quad [\ell] := x \quad (4)$$

where (1) the instruction $x := v$ assigns a value v to the variable x , (2) the instruction $x := [\ell]$ loads the value $h(\ell)$ at location ℓ and assigns it to the variable x , and (3) the instruction $[\ell] := x$ stores at location ℓ the current value $s(x)$ of the variable x . The asynchronous graph (G, \diamond_G) is defined in the following way. Its nodes are the memory states (3) of the machine, and its transitions are of the form

$$\begin{aligned} (s, h) &\xrightarrow{x:=v} (s', h) && \text{when } s' = s \otimes \{x \mapsto v\}, \\ (s, h) &\xrightarrow{x:=[\ell]} (s', h) && \text{when } h(\ell) \text{ is defined and} \\ &&& s' = s \otimes \{x \mapsto h(\ell)\}, \\ (s, h) &\xrightarrow{[\ell]:=x} (s, h') && \text{when } s(x) \text{ is defined and} \\ &&& h' = h \otimes \{\ell \mapsto s(x)\}. \end{aligned}$$

Here, we use the following convenient notation: given a partial function $f : X \rightarrow_{\text{fin}} Y$ with finite domain between two sets X and Y , and an element $y \in Y$, we write $f \otimes \{x \mapsto y\} : X \rightarrow_{\text{fin}} Y$ for the partial function with finite domain defined as

$$f \otimes \{x \mapsto y\} : x' \mapsto \begin{cases} f(x) & \text{when } x' \neq x, \\ y & \text{when } x' = x. \end{cases}$$

In order to define the permutation tiles of the asynchronous graph (G, \diamond_G) , one observes that every transition

$$u : (s, h) \xrightarrow{m} (s', h')$$

performed by an instruction m reads and writes on a specific area

$$\text{rd}(u) \subseteq \mathbf{Var} + \mathbf{Loc} \quad \text{wr}(u) \subseteq \mathbf{Var} + \mathbf{Loc}$$

of the memory of the machine, which we shall call its *footprint*. This footprint may be computed from the instruction m performing the transition $u = (\mu, m, \mu')$ in the following way:

$\begin{aligned} \text{rd}(x := v) &= \emptyset \\ \text{wr}(x := v) &= \{x\} \\ \text{rd}(x := [\ell]) &= \{\ell\} & \text{rd}([\ell] := x) &= \{x\} \\ \text{wr}(x := [\ell]) &= \{x\} & \text{wr}([\ell] := x) &= \{\ell\} \end{aligned}$
--

Now, suppose given two transitions $u : \mu \rightarrow \mu_1$ and $v : \mu \rightarrow \mu_2$ starting from the same memory state μ in the graph G . The two transitions u and v are declared *independent* when

$$(\text{rd}(u) \cup \text{wr}(u)) \cap \text{wr}(v) = \emptyset \quad \text{and} \quad \text{wr}(u) \cap (\text{rd}(v) \cup \text{wr}(v)) = \emptyset.$$

Note that the independence of the transitions u and v is a *consequence* of their spatial separation. It is not difficult to see that for every pair of such independent transitions

$$u : \mu_1 \xrightarrow{m_1} \mu_2 \quad v : \mu_2 \xrightarrow{m_2} \mu_3$$

there exists a unique memory state μ'_2 such that

$$u' : \mu'_2 \xrightarrow{m_1} \mu_3 \quad v' : \mu_1 \xrightarrow{m_2} \mu'_2$$

are transitions of the graph G . In that case, we say that u' is the residual of u after v and, symmetrically, that v' is the residual of v after u . This basic confluence property leads us to the following definition. A permutation tile of the form (2)

$$u \cdot v' \diamond_G v \cdot u'$$

in the asynchronous graph (G, \diamond_G) is defined as a pair of independent transitions u and v where the transition u' is defined as the residual of u after v , and the transition v' is defined as the residual of v after u . It is not difficult to see that the graph $G = (V, E)$ of memory states and transitions between them, together with the notion of permutation tile $u \cdot v' \diamond_G v \cdot u'$ just defined, satisfy the axioms required of an asynchronous graph (G, \diamond_G) .

Stateful vs. stateless semantics Along the *stateful* description of the machine provided by the asynchronous graph (G, \diamond_G) , comes a *stateless* description of the same machine, conveyed this time by an asynchronous graph (H, \diamond_H) where only the instructions are considered, not their action on the machine states. Accordingly, the graph H has a single node $*$ and a transition

$$a : * \xrightarrow{m} *$$

for each instruction m of the machine displayed in (4) parametrized by $x \in \mathbf{Var}$, $v \in \mathbf{Val}$ and $\ell \in \mathbf{Loc}$. The graph H is moreover equipped with a permutation tile

$$a \cdot b' \diamond_H b \cdot a'$$

for every pair $a = a'$ and $b = b'$ of instructions of the machine. The two asynchronous transition graphs (G, \diamond_G) and (H, \diamond_H) are related by an asynchronous graph homomorphism

$$\mathcal{L} : (G, \diamond_G) \longrightarrow (H, \diamond_H) \quad (5)$$

which maps every memory state μ to the node $*$, and every instruction to itself. We recall the definition of such a homomorphism:

Definition 1.1 (homomorphism). An asynchronous graph homomorphism

$$\mathcal{F} : (G, \diamond_G) \longrightarrow (H, \diamond_H) \quad (6)$$

is a graph homomorphism $\mathcal{F} : G \rightarrow H$ between the underlying graphs, such that

$$u \cdot v' \diamond_G v \cdot u' \Rightarrow \mathcal{F}(u) \cdot \mathcal{F}(v') \diamond_H \mathcal{F}(v) \cdot \mathcal{F}(u')$$

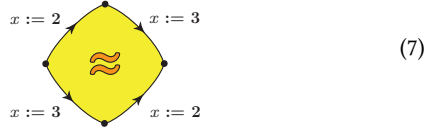
for all transitions u, u', v, v' of the asynchronous graph G .

Note that, in that situation, one has

$$f \sim g \Rightarrow \mathcal{L}(f) \approx \mathcal{L}(g)$$

for all paths $f, g : M \rightarrow N$ in G , where \approx denotes the permutation equivalence in the asynchronous graph (H, \diamond_H) .

Data races as topological obstructions The reason for the liberal definition of \diamond_H is that nothing should forbid two instructions m_1 and m_2 to commute at the *stateless* level of abstraction. By way of illustration, there exists a permutation tile in H (depicted below in light yellow) which permutes the two instructions $x := 2$ and $x := 3$ in the following way:



This permutation tile (7) should be understood as a basic example of *data race* in the machine, where the two instructions $x := 2$ and $x := 3$ compete for the same variable x . As a matter of fact, one key observation and guiding idea of the paper is that such a data race may be detected by the fact that it defines a permutation tile in the stateless semantics (H, \diamond_H) which does *not* lift along \mathcal{L} to a permutation tile in the stateful semantics (G, \diamond_G) . This line of thought leads us to the following definitions of 1-fibration and 2-fibration.

Definition 1.2 (1-fibration). An asynchronous graph homomorphism $\mathcal{F} : (G, \diamond_G) \rightarrow (H, \diamond_H)$ is called a 1-fibration when for every node x of G and transitions $v : \mathcal{F}(x) \rightarrow z$, there exists a transition $u : x \rightarrow y$ such that $\mathcal{F}(u) = v$.

Definition 1.3 (2-fibration). An asynchronous graph homomorphism $\mathcal{F} : (G, \diamond_G) \rightarrow (H, \diamond_H)$ is called a 2-fibration when for every pair of transitions u and v' defining a path $u \cdot v'$ of length 2 in G and for every permutation tile

$$\mathcal{F}(u) \cdot \mathcal{F}(v') \diamond_H b \cdot a'$$

in H , there exists a pair of transitions v and u' in G such that

$$u \cdot v' \diamond_G v \cdot u' \quad \text{and} \quad \mathcal{F}(v) = b \quad \text{and} \quad \mathcal{F}(u') = a'.$$

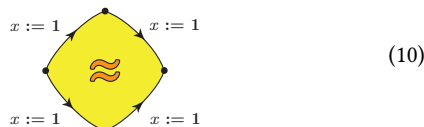
Coming back to our construction, our point is that the asynchronous graph homomorphism \mathcal{L} defined in (5) is *not* a 2-fibration because of the presence of data races such as (7) in the stateless semantics. Typically, any sequence of transitions in (G, \diamond_G)

$$\mu_1 \xrightarrow{x:=2} \mu_2 \xrightarrow{x:=3} \mu_3 \quad (8)$$

mapped by \mathcal{L} to the upward border $* \xrightarrow{x:=2} * \xrightarrow{x:=3} *$ of the permutation tile (7) in the asynchronous graph (H, \diamond_H) satisfies $\mu_2(x) = 2$ and $\mu_3(x) = 3$. For that reason, there exists no way to lift the permutation tile (7) along \mathcal{L} and to permute the sequence of instructions (8) accordingly in (G, \diamond_G) as follows:

$$\mu_1 \xrightarrow{x:=3} \mu'_2 \xrightarrow{x:=2} \mu_3 \quad (9)$$

because this would mean that $\mu_3(x) = 2$, and this would contradict the fact that $\mu_3(x) = 3$. More generally, every data race in the machine may be detected as a topological obstruction to the fact that the stateful-to-stateless homomorphism \mathcal{L} is a 2-fibration. Note that, in the same way but for different reasons, the data race between the two instructions $x := 1$ described by the permutation tile in H below



does *not* lift along \mathcal{L} to a permutation tile in G . Indeed, the instruction $x := 1 : \mu \rightarrow \mu'$ starting from any memory state μ has the nontrivial footprint $\text{wr}(x := 1) = \{x\}$, and is thus not independent of itself in the asynchronous graph (G, \diamond_G) .

An asynchronous semantics of code The machine just considered is a very elementary toy model, which can be easily extended with locks and with memory allocation and deallocation. Also, more than in the machine itself, we are interested in the asynchronous description of the code C we want to analyse. We thus need to explain how we shift from the machine to the code. Interestingly, the story remains essentially the same. To every program C , we associate a stateful interpretation $\llbracket C \rrbracket_S$ and a stateless interpretation $\llbracket C \rrbracket_L$ which reflect the interactive behavior of the program C when confronted to its Environment, called Frame in that context. The two interpretations $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ are formulated as *asynchronous transition systems* (ATS) related by a homomorphism

$$\mathcal{L}_C : \llbracket C \rrbracket_S \longrightarrow \llbracket C \rrbracket_L \quad (11)$$

mentioned in (1) which plays the same role for the code C as the homomorphism (5) for the machine model. The two ATSs $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ are defined uniformly by structural induction on the program C . Their construction – and more specifically the interpretation of the parallel product $C_1 \parallel C_2$ – requires to develop a number of new techniques, in particular an asynchronous parallel product of two ATSs based on the same machine model.

The asynchronous soundness theorem As in the case of the machine model, the data races produced by the program C will be detected as *obstructions* to the fact that \mathcal{L} is a 2-fibration. Typically, the program C defined as $x := 2 ; x := 3$ is data-race-free because the permutation tile (7) does not appear in the stateless semantics $\llbracket C \rrbracket_L$, while the program C' defined as $x := 2 \parallel x := 3$ produces a data race reflected by the fact that the permutation tile (7) appears in the stateless interpretation $\llbracket C' \rrbracket_L$ and cannot be lifted along \mathcal{L} to the stateful interpretation $\llbracket C' \rrbracket_S$.

In the present paper, we carry on our game-theoretic investigation of Concurrent Separation Logic (CSL) initiated in [16] and establish that *well-specified* programs are data-race-free. We achieve this by interpreting every derivation tree

$$\frac{\vdots \pi}{\Gamma \vdash \{P\}C\{Q\}} \quad (12)$$

of CSL as an asynchronous strategy $\llbracket \pi \rrbracket_{\text{Sep}}$ playing on the asynchronous game of separated states. Our asynchronous version of the Soundness Theorem is then formulated in the following fibrational way. Suppose that a code C comes equipped with a proof of the Hoare triple $\Gamma \vdash \{P\}C\{Q\}$ in CSL, and consider the asynchronous subgraph $\llbracket \{P\}C \rrbracket_S^\tau$ obtained by restricting $\llbracket C \rrbracket_S$ to the nodes reachable from an initial node satisfying the precondition P . In that situation, we establish (see Thm. 8.3 in §8 for details) that

Asynchronous Soundness Theorem. *The stateful-to-stateless homomorphism $\mathcal{L}_C : \llbracket C \rrbracket_S \rightarrow \llbracket C \rrbracket_L$ is a 2-fibration when restricted to the asynchronous subgraph $\llbracket \{P\}C \rrbracket_S^\tau$.*

The 2-fibrational property is conceptually new and provides the first structural explanation for the absence of data races in concurrent programs specified by CSL.

Related works Stephen Brookes established the first proof of soundness of CSL in [5], using a stateless trace semantics similar to $\llbracket C \rrbracket_L$ for the concurrent imperative programs. More recently, Viktor Vafeiadis [20] gave a new proof of soundness, based this time on a stateful operational semantics, similar to $\llbracket C \rrbracket_S$. Our approach can be seen as unifying the two schools of semantics, by revealing the asynchronous graph morphism (11) between them. Also, one main benefit of our asynchronous approach is that we can directly describe and analyze the concurrent execution of two instructions.

In the same way as we do here, Jonathan Hayman and Glynn Winskel [9] establish the soundness of CSL in a “truly concurrent” setting. They interpret programs as Petri nets, where the interference of the environment is modeled by adding events to the Petri net. In contrast to our work, precision of the invariants is necessary for their semantics to work whereas [7] has shown that precision is only needed in order to interpret properly the conjunction rule of CSL. Another difference is that they consider a language somewhat different from Brookes’ original language [5], without local variables, but with dynamic binding of resources.

We give in [16] a game-theoretic interpretation of CSL, where every Hoare triple is interpreted as a *game* between Adam and Eve, and every derivation tree π as a *winning strategy* for Eve in that game. Every program is interpreted there as a set of purely sequential traces. For that reason, it is not possible to establish in this framework the absence of data races, at least in a nice and conceptual way. One main achievement of the paper is thus to define a properly asynchronous game semantics of CSL, and to derive for the first time the absence of data races from purely semantic considerations on the model.

Synopsis of the paper After the machine states and instructions are described in §2, we construct in §3 the two asynchronous graphs \mathfrak{A}_S and \mathfrak{A}_L defining our stateful and stateless machine models. We then explain in §4 how to interpret every code C as a pair $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ of asynchronous transition systems (ATS) with respective machine models \mathfrak{A}_S and \mathfrak{A}_L . Once the notions of logical state and of separated state are recalled in §5 and in §6, we explain in §7 how to interpret every proof π of CSL as an asynchronous strategy $\llbracket \pi \rrbracket_{Sep}$ playing on the machine model \mathfrak{A}_{Sep} of separated states. From this, we establish our asynchronous soundness theorem in §8, and conclude in §9.

2 Machine states and machine instructions

We introduce below the notions of *machine state* and of *machine instruction* which will be used throughout the paper. We suppose given countable sets \mathbf{Var} of *variable names*, \mathbf{Val} of *values*, $\mathbf{Loc} \subseteq \mathbf{Val}$ of *memory locations*, and $\mathbf{LockName}$ of *resources*. In practice, we consider the case where $\mathbf{Loc} = \mathbb{N}$ and $\mathbf{Val} = \mathbb{Z}$.

Definition 2.1 (*Memory states*). A *memory state* μ is a pair (s, h) of partial functions with finite domains $s : \mathbf{Var} \rightarrow_{fin} \mathbf{Val}$ and $h : \mathbf{Loc} \rightarrow_{fin} \mathbf{Val}$ called the *stack* s and the *heap* h of the memory state μ . The set of memory states is denoted by \mathbf{State} . The domains of the partial function s and of h are denoted by $\text{vdom}(\mu)$ and $\text{hdom}(\mu)$ respectively, and we write $\text{dom}(\mu)$ for their disjoint union.

Definition 2.2 (*Machine states*). A *machine state* is either a pair $\mathfrak{s} = (\mu, L)$ consisting of a memory state μ and a subset of resources $L \subseteq \mathbf{LockName}$, called the *lock state*, which describes the subset of locked resources in \mathfrak{s} ; or an error state $\frac{1}{2}$. The set of machine states

is denoted by \mathbf{MState} . Formally:

$$\mathbf{MState} = \mathbf{State} \times \mathcal{P}(\mathbf{LockName}) + \{\frac{1}{2}\}$$

A machine step is defined as a labeled transition between machine states. There are two kinds of transitions:

$$(\mu, L) \xrightarrow{m} (\mu', L') \quad (\mu, L) \xrightarrow{m} \frac{1}{2} \quad (13)$$

depending on whether the instruction $m \in \mathbf{Instr}$ has been executed successfully (on the left) or has produced a runtime error (on the right). In particular, $\frac{1}{2}$ has no successor. The machine instructions $m \in \mathbf{Instr}$ which label the machine steps are of the following form:

$$\begin{aligned} m ::= & x := E \mid x := [E] \mid [E] := E' \mid \text{nop} \\ & \mid x := \text{alloc}(E, \ell) \mid \text{dispose}(E) \mid P(r) \mid V(r) \end{aligned}$$

where $x \in \mathbf{Var}$ is a variable, $r \in \mathbf{LockName}$ is a resource name, ℓ is a location, and E, E' are arithmetic expressions, possibly with “free” variables in \mathbf{Var} . For example, the instruction $x := E$ executed in a machine state $\mathfrak{s} = (\mu, L)$ assigns to the variable x the value $E(\mu) \in \mathbf{Val}$ when the value of the expression E can be evaluated in the memory state μ , and produces the runtime error $\frac{1}{2}$ otherwise. The instruction $P(r)$ acquires the resource variable r when it is available, while the instruction $V(r)$ releases it when r is locked, as described below:

$$\begin{array}{c} \frac{E(\mu) = v}{(\mu, L) \xrightarrow{x:=E} (\mu[x \mapsto v], L)} \quad \frac{E(\mu) \text{ not defined}}{(\mu, L) \xrightarrow{x:=E} \frac{1}{2}} \\ \frac{r \notin L}{(\mu, L) \xrightarrow{P(r)} (\mu, L \uplus \{r\})} \quad \frac{r \notin L}{(\mu, L \uplus \{r\}) \xrightarrow{V(r)} (\mu, L)} \end{array}$$

The inclusion $\mathbf{Loc} \subseteq \mathbf{Val}$ means that an expression E may also denote a location. In that case, $[E]$ refers to the value stored at location E in the heap. The instruction $x := \text{alloc}(E, \ell)$ allocates some memory space on the heap at address $\ell \in \mathbf{Loc}$, initializes it with the value of the expression E , and assigns the address ℓ to the variable $x \in \mathbf{Var}$ if *location* was free, otherwise there is no transition. $\text{dispose}(E)$ deallocates the location denoted by E when it is allocated, and returns $\frac{1}{2}$ otherwise. Finally, the instruction nop (for no-operation) does not alter the state.

3 Asynchronous Machine Models

As explained in the introduction, *machine models* are described using asynchronous graphs. Since we consider *stateful* as well as *stateless* descriptions of the machine and of the code, we will consider two kinds of machine models, organized into a pair of asynchronous graphs: the *stateful model* \mathfrak{A}_S based on machine states, and the *stateless model* \mathfrak{A}_L based on locks. Their tiles will be defined using the notion of footprint, which summarizes which area of the state (memory, locks) an instruction relies on, and how it uses it. In both cases, we write $\text{footprint}_s(m)$ for the footprint of an instruction m in state s , omitting the subscript when it is clear from the context. Our machine models \mathfrak{A}_S and \mathfrak{A}_L are parameterized over the finite set $\mathbf{Locks} \subseteq \mathbf{LockName}$ of locks, or resources, which are considered well-defined. We sometimes write $\mathfrak{A}_S(\mathbf{Locks})$ or $\mathfrak{A}_L(\mathbf{Locks})$ to make it explicit.

The stateful model A machine state footprint

$$\rho \in \mathcal{P}(\mathbf{Var} + \mathbf{Loc}) \times \mathcal{P}(\mathbf{Var} + \mathbf{Loc}) \times \mathcal{P}(\mathbf{Locks}) \times \mathcal{P}(\mathbf{Loc})$$

is, made of: (i) $\text{rd}(\rho)$, the part of the memory that is *read*, (ii) $\text{wr}(\rho)$, the part of the memory that is *written*, (iii) $\text{lock}(\rho)$, the locks that

are *touched*, and (iv) $\text{mem}(\rho)$ the addresses that are *allocated* or *deallocated*. Two footprints ρ and ρ' are declared *independent* when:

$$\begin{aligned} (\text{rd}(\rho) \cup \text{wr}(\rho)) \cap \text{wr}(\rho') &= \emptyset & \text{lock}(\rho) \cap \text{lock}(\rho') &= \emptyset \\ (\text{rd}(\rho') \cup \text{wr}(\rho')) \cap \text{wr}(\rho) &= \emptyset & \text{mem}(\rho) \cap \text{mem}(\rho') &= \emptyset \end{aligned}$$

The stateful model \mathfrak{A}_S is the following asynchronous graph: its nodes are the machine states in **MState**, its transitions are of the form

$$(\mu, L) \xrightarrow{m} (\mu', L') \quad \text{or} \quad (\mu, L) \xrightarrow{m} \zeta$$

corresponding to the machine steps, defined in §2. The asynchronous tiles of \mathfrak{A}_S are the squares of the form

$$s \xrightarrow{m} s_1 \xrightarrow{m'} s' \quad \sim \quad s \xrightarrow{m'} s_2 \xrightarrow{m} s'$$

where their footprints are independent in the sense above.

The stateless model A lock footprint

$$\rho \in \wp(\text{Locks}) \times \wp(\text{Loc})$$

is made of a set of locks $\text{lock}(\rho)$ and a set of locations $\text{mem}(\rho)$. Two such footprints are *independent* when their sets are componentwise disjoint. The stateless model \mathfrak{A}_L is defined in the following way: its nodes are the subsets of *Locks*, and its transitions are all the edges of the form (note the non-determinism)

$$\begin{array}{ccc} L \xrightarrow{P(r)} L \uplus \{r\} & L \xrightarrow{\text{alloc}(\ell)} L & L \xrightarrow{\tau} L \\ L \uplus \{r\} \xrightarrow{V(r)} L & L \xrightarrow{\text{dispose}(\ell)} L & L \xrightarrow{m} \zeta \end{array}$$

where m is a *lock instruction* of the form:

$$P(r) \mid V(r) \mid \text{alloc}(\ell) \mid \text{dispose}(\ell) \mid \tau$$

for $\ell \in \text{Loc}$ and $r \in \text{Locks}$. The purpose of these transitions is to extract from each instruction of the machine its synchronization behavior. An important special case, the transition τ represents the absence of any synchronization mechanism in an instruction like $x := E$, $x := [E]$ or $[E] := E'$. The asynchronous tiles of \mathfrak{A}_L are the squares of the form

$$L \xrightarrow{x} L_1 \xrightarrow{y} L' \quad \sim \quad L \xrightarrow{y} L_2 \xrightarrow{x} L'$$

when the lock footprints of x and y are independent. It is worth noting that L' may be equal to ζ in such an asynchronous tile. Note that the asynchronous graph \mathfrak{A}_L is more liberal than \mathfrak{A}_S about which footprints commute, because it only takes into account the locks as well as the allocated and deallocated locations. As explained in the introduction, this mismatch enables us to detect *data races* in the machine as well as in the code.

Remark The last component $\text{mem}(\rho)$ in the machine state footprint as well as in the lock footprint enables us to forbid a deallocation followed by an allocation to happen at the same address without some kind of synchronization, both at the stateful and stateless level. This is consistent with practice, since the malloc implementation would typically synchronize its accesses to the free-list(s) of the different threads.

4 Asynchronous Semantics of Code

In this section, we associate to every program C a pair of asynchronous transition systems $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ over the machine models \mathfrak{A}_S and \mathfrak{A}_L introduced in the previous section. The first interpretation $\llbracket C \rrbracket_S$ is *stateful* and describes how each instruction of the program C acts on the memory states and on the locks. The second

interpretation $\llbracket C \rrbracket_L$ is *stateless* and only remembers the action of the instructions on the locks.

4.1 Asynchronous transition systems (ATSs)

Asynchronous transition systems (ATSs) are specific asynchronous graphs where every transition is either executed by Code or by Frame. We thus start by introducing the following notion:

Definition 4.1 (Asynchronous graph with polarities). An asynchronous graph with polarities is an asynchronous graph (G, \diamond_G) where every transition is assigned a *polarity* Code or Frame. One requires that in every permutation tile $u \cdot v' \diamond_G v \cdot u'$, the two transitions u and u' (symmetrically v and v') have the same polarity.

A path in an asynchronous graph G with polarities is called *Code-proper* when it contains (at least) one Code transition. A node x is called *initial* in G when there are no Code-proper incoming paths into x , and *final* when there are no Code-proper outgoing paths from x . The sets of initial and final nodes in G are denoted $\partial_0 G$ and $\partial_1 G$, respectively. The graph G is called *Code-acyclic* when there are no Code-proper cycles, that is, every cycle of the graph G contains only Frame transitions. A set S of nodes of a graph is *forward-closed* when $x \in S$ and $x \rightarrow y$ implies that $y \in S$.

Definition 4.2 (ATS). An asynchronous transition system (ATS) is a Code-acyclic asynchronous graph with polarities (G, \diamond_G) equipped with a forward-closed subset $|G| \subseteq \partial_1(G)$ of final nodes. A final node in $|G|$ is called a *returning node* of the ATS.

Definition 4.3. An ATS with machine model (\mathfrak{A}, \diamond) is defined as an ATS $(G, |G|)$ equipped with an asynchronous graph homomorphism

$$\lambda_G : (G, \diamond_G) \longrightarrow (\mathfrak{A}, \diamond)$$

One requires moreover that

1. the map λ_G defines a bijection between the set $\partial_0 G$ of initial nodes and the set of nodes of \mathfrak{A} , and an injection from the set $|G|$ of returning nodes into the set of nodes of \mathfrak{A} .
2. the map λ_G is a Frame 1-fibration, in the sense that for every transition $v : \lambda_G(x) \rightarrow z$ in the machine model \mathfrak{A} , there exists a unique Frame transition $u : x \rightarrow y$ in G such that $\lambda_G(u) = v : \lambda_G(x) \rightarrow \lambda_G(y)$,
3. the map λ_G is a Code-Frame and Frame-Frame 2-fibration, in the sense that for every sequence of transitions $x \xrightarrow{u} y \xrightarrow{v'} z$ in G where $v' : y \rightarrow z$ is a Frame transition, and for every permutation tile in \mathfrak{A} of the form:

$$\begin{array}{ccc} & \lambda_G(u) & \\ & \curvearrowright & \lambda_G(v') \\ & \curvearrowleft & \\ b & & a \end{array} \quad (14)$$

there exists a sequence of transitions $x \xrightarrow{v} y' \xrightarrow{u'} z$ and a permutation tile $u \cdot v' \diamond_G v \cdot u'$ in G transported by λ_G to the permutation tile (14) in the sense that

$$\lambda_G(x) \xrightarrow{\lambda_G(v)} \cdot \xrightarrow{\lambda_G(u')} \lambda_G(z) = \lambda_G(x) \xrightarrow{b} \cdot \xrightarrow{a} \lambda_G(z)$$

Notation We often find convenient to label the transitions $u : x \rightarrow y$ in G with the instruction or lock instruction m which labels the transition $\lambda_G(u)$ in the underlying asynchronous graph \mathfrak{A}_S or \mathfrak{A}_L . We also write $m : C$ or $m : F$ to mean that the transition $u : x \rightarrow y$ has the polarity Code or Frame in G , respectively.

4.2 Basic constructions on ATSs

The asynchronous interpretations $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ of the program C are performed by structural induction, using a number of primitive operations on ATSs defined below. Note that whenever a construction makes some nodes unreachable from the initial nodes, they are implicitly removed.

Sum The sum of two ATSs G_1 and G_2 with same machine model \mathfrak{A} , written $G_1 \oplus G_2$, is the disjoint union of the two asynchronous graphs G_1 and G_2 , where we identify their respective initial and returning states together, when they have the same image under λ_{G_1} and λ_{G_2} . This means that for the case of the returning states, there are three cases. If they both have returning states, we identify $\partial_1(G_1)$ with $\partial_1(G_2)$; if only one of G_1 and G_2 has returning states, we keep this one as our returning states; otherwise the juxtaposition has no returning states.

Sequential composition The sequential composition $G; G'$ of two ATSs G and G' is the disjoint union of G and G' where we identify the returning nodes of G and the initial nodes of G' with the same underlying image under λ_G and $\lambda_{G'}$. Because we remove the inaccessible nodes, when G has no returning nodes, $G; G' = G$.

Parallel product The parallel product $G_1 \parallel G_2$ of two ATSs G_1 and G_2 over the same machine model \mathfrak{A} is defined as follows. The nodes of $G_1 \parallel G_2$ are the pairs of nodes $x_1 | x_2 \in G_1 \times G_2$ such that $\lambda_{G_1}(x_1) = \lambda_{G_2}(x_2)$ and $\lambda_{G_1 \parallel G_2}(x_1 | x_2)$ is defined to be that common value. The transitions of $G_1 \parallel G_2$ are of three kinds:

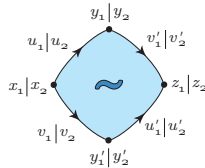
1. the Code transitions $x_1 | x_2 \xrightarrow{m:C} x'_1 | x'_2$ where

$$x_1 \xrightarrow{m:C} x'_1 \text{ in } G_1 \quad \text{and} \quad x_2 \xrightarrow{m:F} x'_2 \text{ in } G_2.$$
2. the Code transitions $x_1 | x_2 \xrightarrow{m:C} x'_1 | x'_2$ where

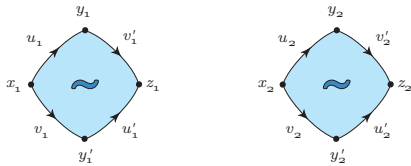
$$x_1 \xrightarrow{m:F} x'_1 \text{ in } G_1 \quad \text{and} \quad x_2 \xrightarrow{m:C} x'_2 \text{ in } G_2.$$
3. the Frame transitions $x_1 | x_2 \xrightarrow{m:F} x'_1 | x'_2$ where

$$x_1 \xrightarrow{m:F} x'_1 \text{ in } G_1 \quad \text{and} \quad x_2 \xrightarrow{m:F} x'_2 \text{ in } G_2.$$

Note that every transition $u : x_1 | x_2 \rightarrow y_1 | y_2$ in the graph $G_1 \parallel G_2$ is a pair $u = (u_1, u_2)$ also written $u = u_1 | u_2$ of a transition $u_1 : x_1 \rightarrow y_1$ in G_1 and $u_2 : x_2 \rightarrow y_2$ in G_2 . A permutation tile in $G_1 | G_2$



is then defined as a square whose projections



define permutation tiles in (G_1, \diamond_1) and (G_2, \diamond_2) , respectively. Finally, the returning nodes $x_1 | x_2 \in |G_1 \parallel G_2|$ are defined as the pairs $x_1 | x_2$ of returning nodes $x_1 \in |G_1|$ and $x_2 \in |G_2|$.

The parallel product of G_1 and G_2 is *asynchronous* in the sense that every Code transition in $G_1 \parallel G_2$ is a Code transition performed by G_1 and seen as a Frame transition by G_2 , or symmetrically, a Code transition performed by G_2 and seen as a Frame transition by G_1 . In particular, by definition, the two components G_1 and G_2 never execute (or “fire”) a Code transition simultaneously in $G_1 \parallel G_2$. At the level of permutation tiles, a Code transition $u_1 | u_2 : x_1 | x_2 \rightarrow y_1 | y_2$ performed in $G_1 \parallel G_2$ by the component G_1 and a Code transition $v'_1 | v'_2 : y_1 | y_2 \rightarrow z_1 | z_2$ performed in $G_1 \parallel G_2$ by the component G_2 define a permutation tile precisely when the transitions $\lambda_{G_1 \parallel G_2}(u_1 | u_2) = \lambda_{G_1}(u_1) = \lambda_{G_2}(u_2)$ and $\lambda_{G_1 \parallel G_2}(v'_1 | v'_2) = \lambda_{G_1}(v'_1) = \lambda_{G_2}(v'_2)$ define a permutation tile in the underlying machine model \mathfrak{A} . As a matter of fact, one purpose of the machine model (\mathfrak{A}, \diamond) is precisely to provide that piece of information necessary to construct the parallel product of G_1 and G_2 .

Resource hiding In order to interpret the resource introduction construct resource r do C , we introduce a *hiding* operator $\text{hide}[r]$ on ATSs which hides the new resource r , similarly to the operator ν in the π -calculus. Formally, if G is an ATS over $\mathfrak{A}(\text{Locks} \uplus \{r\})$, then $\text{hide}[r](G)$ is the ATS over $\mathfrak{A}(\text{Locks})$ where: (1) the resource r has been removed from the sets of locked resources of all states, (2) the Code transitions $P(r)$ and $V(r)$ are replaced with nops, (3) the Frame transitions $P(r)$ and $V(r)$ are removed from the graph G , and (4) the remaining permutation tiles are preserved. Moreover, we only keep as initial and returning states the initial and returning states x of G such that the resource r is not held in $\lambda_G(x)$.

Critical sections Dually, inside critical sections, we need to “lift” ATSs over some set *Locks* of locks to ATSs over $\text{Locks} \uplus \{r\}$. This can be done naturally in this case because we know that, during the critical section, the resource r is held by the Code. Formally, $\text{when}[r](G)$ has the same underlying asynchronous graph as G , where $\lambda' := \lambda_{\text{when}[r](G)}$ is defined by:

$$\begin{aligned} \lambda'(x) &:= L \uplus \{r\} && \text{if } \lambda_G(x) = L \\ \lambda'(x) &:= (\mu, L \uplus \{r\}) && \text{if } \lambda_G(x) = (\mu, L). \end{aligned}$$

This does not define an ATS yet, for condition 3 is not satisfied: there are not enough Environment transitions. This is why we must freely add Frame transitions and new nodes to make it an ATS. The returning nodes are defined to be the same as G .

Other constructions on ATSs Given an ATS G and a Boolean formula B , we define $\text{whentru}[B](G)$ as the graph G where, among the Code transitions out of initial nodes, we only keep those where B holds on $\lambda_G(x)$. Then, we remove the nodes made unreachable by this edge removal. Similarly, we define $\text{whenfalse}[B]$ for when B does not hold. Finally, $\text{whenabort}[B]$ is the graph with transitions from the initial states where B errors out, because it tries to read undefined variables, to \perp . Note that in the case of the $\llbracket C \rrbracket_L$ semantics, since the nodes do not contain information on the state, the first two constructions above are the identity. This means that we sometimes consider impossible branches.

4.3 Asynchronous semantics of the code

We explain how to give a semantics to any code C as an ATS $\llbracket C \rrbracket$, by induction on its structure. This lets us build $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ in the same way. First, we give the syntax of our imperative concurrent

language, which we borrow from [5, 20].

$$\begin{aligned}
B &::= \mathbf{true} \mid \mathbf{false} \mid B \wedge B' \mid B \vee B' \mid E = E' \\
E &::= 0 \mid 1 \mid \dots \mid x \mid E + E' \mid E * E' \\
C &::= x := E \mid x := [E] \mid [E] := E' \mid C; C' \mid C_1 \parallel C_2 \mid \text{skip} \\
&\mid \text{while } B \text{ do } C \mid \text{resource } r \text{ do } C \mid \text{with } r \text{ when } B \text{ do } C \\
&\mid \text{if } B \text{ then } C_1 \text{ else } C_2 \mid x := \text{malloc}(E) \mid \text{dispose}(E)
\end{aligned}$$

Semantics of instructions To every instruction $m \in \mathbf{Instr}$, we associate the ATS $\llbracket m \rrbracket$ with machine model \mathfrak{A} defined as two copies $\mathfrak{A}_0 + \mathfrak{A}_1$ (called *source* and *target*) of the asynchronous graph \mathfrak{A} . Every transition in $\mathfrak{A}_0 + \mathfrak{A}_1$ is assigned the Frame polarity. To this, one adds a Code transition $x_0 \rightarrow y_1$ for every transition of the form (13) labeled by m in the small step semantics. Here, x_0 and y_1 are the nodes x and y of \mathfrak{A} taken in the source and target components \mathfrak{A}_0 and \mathfrak{A}_1 of $\llbracket m \rrbracket$, respectively. The transition $x_0 \rightarrow y_1$ is mapped by $\lambda_{\llbracket m \rrbracket}$ to the transition associated to the small step transition (13) in $\mathfrak{A} = \mathfrak{A}_S$ or $\mathfrak{A} = \mathfrak{A}_S$. Finally, one adds a Code-Frame permutation tile in $\llbracket m \rrbracket$ for each Code-Frame permutation tile in \mathfrak{A} , in such a way that $\lambda_{\llbracket m \rrbracket} : \llbracket m \rrbracket \rightarrow \mathfrak{A}$ defines a Code-Frame 2-fibration.

Leaf codes For leaf codes that correspond to instructions (all, except for `malloc`), their semantics is the same as that of the instruction. For `malloc(E)`, we take the non-deterministic union of all the `alloc(E, ℓ)`:

$$\llbracket \text{malloc}(E) \rrbracket := \bigoplus_{\ell \in \text{Loc}} \llbracket \text{alloc}(E, \ell) \rrbracket$$

Conditionals Conditional branching is interpreted as

$$\begin{aligned}
\llbracket \text{if } B \text{ then } C \text{ else } C_1 \rrbracket &= \text{whentrue}[B](\llbracket \text{nop} \rrbracket); \llbracket C_1 \rrbracket \\
&\oplus \text{whenfalse}[B](\llbracket \text{nop} \rrbracket); \llbracket C_2 \rrbracket \\
&\oplus \text{whenabort}[B]
\end{aligned}$$

The nops are needed because the environment can interfere between the evaluation of B and the beginning of the execution of C_i .

Sequential and parallel compositions We use the sequential and parallel product of ATSs with machine models defined in §4.2, in the following way:

$$\llbracket C_1 \parallel C_2 \rrbracket = \llbracket C_1 \rrbracket \parallel \llbracket C_2 \rrbracket \quad \llbracket C_1; C_2 \rrbracket = \llbracket C_1 \rrbracket; \llbracket C_2 \rrbracket.$$

Resource introduction The interpretation of resource r do C is defined as

$$\llbracket \text{resource } r \text{ do } C \rrbracket = \text{hide}[r](\llbracket C \rrbracket)$$

Critical sections The semantics $\llbracket \text{with } r \text{ when } B \text{ do } C \rrbracket$ is defined using the sequential composition above and `whentrue`:

$$\text{whentrue}[B](\llbracket P(r) \rrbracket; \text{when}[r](\llbracket C \rrbracket); \llbracket V(r) \rrbracket) \oplus \text{whenabort}[B].$$

Loops For loops, the interpretation of $C' = \text{while } B \text{ do } C$ is defined as the (possibly infinite) least fixpoint of the function F :

$$\begin{aligned}
F(G) &= \text{whentrue}[B](\llbracket \text{nop} \rrbracket); \llbracket C \rrbracket; G \oplus \text{whenfalse}[B](\llbracket \text{nop} \rrbracket) \\
&\oplus \text{whenabort}[B].
\end{aligned}$$

Remark The map λ_G is a 2-fibration for Code-Frame and Frame-Frame permutations, but not for Code-Code permutations in general. Consider for instance the interpretation of the program

$$C = \text{resource } r \text{ do } \{ (P(r); V(r)) \parallel (P(r); V(r)) \}$$

where $r \in \mathbf{LockName}$ is a resource name. Since the resource introduction performed by resource r do C is interpreted by hiding the resource r , the two instructions $P(r)$ and $V(r)$ are both transformed in nops instructions. However the two nops *do not* form a tile! Another example is, of course, the sequential composition $C_1; C_2$ of two codes C_1 and C_2 .

4.4 Comparing the stateful and the stateless semantics

We construct a category of ATSs with machine models, in the following way. A *morphism* between ATSs with machine models

$$\lambda_{G_1} : G_1 \rightarrow \mathfrak{A}_1 \quad \lambda_{G_2} : G_2 \rightarrow \mathfrak{A}_2$$

is a pair of asynchronous graph morphisms $\mathcal{F} : \mathfrak{A}_1 \rightarrow \mathfrak{A}_2$ and $\mathcal{G} : G_1 \rightarrow G_2$ such that the diagram below commutes:

$$\begin{array}{ccc}
G_1 & \xrightarrow{\mathcal{G}} & G_2 \\
\lambda_{G_1} \downarrow & & \downarrow \lambda_{G_2} \\
\mathfrak{A}_1 & \xrightarrow{\mathcal{F}} & \mathfrak{A}_2
\end{array}$$

One requires moreover that \mathcal{G} send initial (resp. returning) nodes of G_1 to initial (resp. returning) nodes of G_2 . This defines a category noted **ATS**. Let $\mathcal{F} : \mathfrak{A}_S \rightarrow \mathfrak{A}_L$ denote the asynchronous graph morphism which transports every machine state $s = (\mu, L)$ to the underlying subset $L \subseteq \text{Locks}$ of locks held in s . Every instruction $m \in \mathbf{Instr}$ comes equipped with an ATS morphism

$$\mathcal{L}_m = (\mathcal{F}, \mathcal{G}_m) : \llbracket m \rrbracket_S \longrightarrow \llbracket m \rrbracket_L$$

where the asynchronous graph morphism \mathcal{G}_m is defined as

$$(\mu, L) \xrightarrow{m} (\mu', L') \longmapsto L \xrightarrow{m} L'$$

Because the stateful and stateless interpretations $\llbracket - \rrbracket_S$ and $\llbracket - \rrbracket_L$ are defined using the same functorial operations over \mathcal{F} , we can associate to every code C a morphism of **ATS**

$$\mathcal{L}_C = (\mathcal{F}, \mathcal{G}_C) : \llbracket C \rrbracket_S \longrightarrow \llbracket C \rrbracket_L$$

starting from the family of morphisms \mathcal{L}_m associated to instructions. Note that this morphism $\mathcal{L}_C = (\mathcal{F}, \mathcal{G}_C)$ living in the category **ATS** plays a fundamental role in the present work, since our asynchronous refinement of the original Soundness Theorem for CSL relies on it, see §8 for details.

5 Logical States

As discussed in [16], reasoning about concurrent programs in separation logic requires to introduce an appropriate notion of *logical state*, including information about permissions. The version of concurrent separation logic we consider is almost the same as its original formulation by O'Hearn and Brookes [5, 17]. One difference is that we benefit from the work of Bornat, Calcagno, O'Hearn, Parkinson and Yang in [2, 3, 18] and use permissions p and the predicate $\text{Own}_p(x)$ in order to handle the heap as well as variables in the stack. We suppose given an arbitrary partial cancellative commutative monoid **Perm** which we call the permission monoid, following [2]. The element \top will be used as the permission required for a program to write somewhere in memory. We thus require that \top does not admit any multiples, ie. $\forall x \in \mathbf{Perm}, \top \cdot x$ is not defined. The intuition (which we will need to turn into a theorem) is that we prevent in this way concurrent mutation and observation of the same location, that is, data races. The set **LState**

of *logical states* is defined in much the same way as the set **State** of memory states, with the addition of permissions:

$$\mathbf{LState} = (\mathbf{Var} \rightarrow_{\text{fin}} (\mathbf{Val} \times \mathbf{Perm})) \times (\mathbf{Loc} \rightarrow_{\text{fin}} (\mathbf{Val} \times \mathbf{Perm}))$$

The main benefit of permissions is that they enable us to define a *separation product* $\sigma * \sigma'$ between two logical states σ and σ' , which generalizes the disjoint union. When it is defined, the logical state $\sigma * \sigma'$ is defined as a partial function with domain

$$\text{dom}(\sigma * \sigma') = \text{dom}(\sigma) \cup \text{dom}(\sigma')$$

in the following way: for $a \in \mathbf{Var} \amalg \mathbf{Loc}$,

$$\sigma * \sigma'(a) = \begin{cases} \sigma(a) & \text{if } a \in \text{dom}(\sigma) \setminus \text{dom}(\sigma') \\ \sigma'(a) & \text{if } a \in \text{dom}(\sigma') \setminus \text{dom}(\sigma) \\ (v, p \cdot p') & \text{if } \sigma(a) = (v, p) \text{ and } \sigma'(a) = (v, p') \end{cases}$$

The separation product $\sigma * \sigma'$ of the two logical states σ and σ' is not defined otherwise. In particular, the memory states underlying σ and σ' agree on the values of the shared variables and heap locations when the separation product is well defined. The syntax and the semantics of the *formulas* of Concurrent Separation Logic is the same as in Separation Logic. The grammar of formulas is:

$$P, Q, R, J ::= \mathbf{emp} \mid \mathbf{true} \mid \mathbf{false} \mid P \vee Q \mid P \wedge Q \mid \neg P \mid \forall v. P \mid \exists v. P \\ \mid P * Q \mid v \xrightarrow{L} w \mid \text{Own}_p(x) \mid E_1 = E_2$$

where $x \in \mathbf{Var}$, $p \in \mathbf{Perm}$, $v, w \in \mathbf{Val}$. Given a logical state $\sigma = (s, h)$ consisting of a logical stack s and of a logical heap h , the semantics of the formulas, expressed as the predicate $\sigma \models P$, is standard:

$$\begin{aligned} \sigma \models v \xrightarrow{L} w &\iff v \in \mathbf{Loc} \wedge s = \emptyset \wedge h = [v \mapsto (w, p)] \\ \sigma \models \text{Own}_p(x) &\iff \exists v \in \mathbf{Val}, s = [x \mapsto (v, p)] \wedge h = \emptyset \\ \sigma \models E_1 = E_2 &\iff \llbracket E_1 \rrbracket = \llbracket E_2 \rrbracket \wedge \text{fv}(E_1 = E_2) \subseteq \text{vdom}(\sigma) \\ \sigma \models P \wedge Q &\iff \sigma \models P \text{ and } \sigma \models Q \\ \sigma \models P * Q &\iff \exists \sigma_1 \sigma_2, \sigma = \sigma_1 * \sigma_2 \text{ and } \sigma_1 \models P \text{ and } \sigma_2 \models Q. \end{aligned}$$

The *proof system* underlying concurrent separation logic is a sequent calculus, whose sequents are *Hoare triples* of the form

$$\Gamma \vdash \{P\} C \{Q\}$$

where $C \in \mathbf{Code}$, P, Q are predicates, and Γ is a context, defined as a partial function with finite domain from the set **LockName** of resource variables to predicates. Intuitively, the context $\Gamma = r_1 : J_1, \dots, r_k : J_k$ describes the *invariant* J_i satisfied by the resource variable r_i . The purpose of these resources is to describe the fragments of memory shared between the various threads during the execution.

The more interesting inference rules of CSL are given in Figure 1, and the reader will find the full system in the Appendix. The inference rule **RES** associated to resource r do C moves a piece of logical state which is owned by the Code into the shared context Γ , which means that it can be accessed concurrently inside the code C . However, the access to that piece of state is mediated by the *with* construct, which grants temporary access under the condition that one must give it back (rule **WITH**). Note that the rule **WITH** has the side condition $P \Rightarrow \text{def}(B)$. This means that if P is true in some logical state, then it implies, for each free variable x of B , that there exists some permission p such that $\text{Own}_p(x)$ holds.

6 The machine model of separated states

We recall the notion of *separated state* formulated in [16] whose purpose is to separate the logical memory state into one region controlled by the Code, one region controlled by the Frame, and one independent region for each unlocked resource. In order to define the notion, we suppose given a finite set $\mathbf{Locks} \subseteq \mathbf{LockName}$ of resource variables, or locks.

Definition 6.1. A separated state is a triple

$$(\sigma_C, \sigma, \sigma_F) \in \mathbf{LState} \times (\mathbf{Locks} \rightarrow \mathbf{LState} + \{C, F\}) \times \mathbf{LState}$$

such that the logical state below is defined:

$$\sigma_C * \left\{ \bigotimes_{r \in \text{dom}(\sigma)} \sigma(r) \right\} * \sigma_F \in \mathbf{LState} \quad (15)$$

where $\text{dom}(\sigma) = \{r \in \mathbf{Locks} \mid \sigma(r) \in \mathbf{LState}\}$,

$$\text{dom}_C(\sigma) = \{r \in \mathbf{Locks} \mid \sigma(r) = C\},$$

$$\text{dom}_F(\sigma) = \{r \in \mathbf{Locks} \mid \sigma(r) = F\}.$$

We say that a separated state $(\sigma_C, \sigma, \sigma_F)$ combines into a machine state $\mathfrak{s} = (\mu, L)$ precisely when $L = \text{dom}_C(\sigma) \uplus \text{dom}_F(\sigma)$ and when the function $U : \mathbf{LState} \rightarrow \mathbf{State}$ which forgets the permissions transports the logical state (15) into the memory state $\mu \in \mathbf{State}$. Note that, by definition, every separated state $(\sigma_C, \sigma, \sigma_F)$ combines into a unique machine state, which we write for concision

$$(\mu, L) = \bigotimes(\sigma_C, \sigma, \sigma_F). \quad (16)$$

Interestingly, the notion of separated state comes with the same notion of footprint as the machine states, defined as elements of

$$\rho \in \wp(\mathbf{Var} + \mathbf{Loc}) \times \wp(\mathbf{Var} + \mathbf{Loc}) \times \wp(\mathbf{Locks}) \times \wp(\mathbf{Loc}).$$

which describes the footprint of a transition by Eve or Adam.

Definition 6.2. The machine model of separated states $\mathfrak{s}_{\text{Sep}}$ is the asynchronous graph whose nodes are the separated states and whose edges are either Adam or Eve transitions:

- Eve transitions are of the form

$$(\sigma_C, \sigma, \sigma_F) \xrightarrow{m:C} (\sigma'_C, \sigma', \sigma'_F)$$

where $m \in \mathbf{Instr}$ is an instruction such that

$$\bigotimes(\sigma_C, \sigma, \sigma_F) \rightsquigarrow^m \bigotimes(\sigma'_C, \sigma', \sigma'_F)$$

and such that the following conditions are satisfied:

$$\begin{aligned} \forall \ell \notin \text{wr}(m), \sigma_C(\ell) = \sigma'_C(\ell) \quad \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_C) \\ \text{lock}(m) \subseteq \text{dom}(\sigma) \cup \text{dom}_C(\sigma) \quad \forall r \notin \text{lock}(m), \sigma(r) = \sigma'(r). \end{aligned}$$

- Adam moves of the form

$$(\sigma_C, \sigma, \sigma_F) \xrightarrow{m:F} (\sigma_C, \sigma', \sigma'_F)$$

where $m \in \mathbf{Instr}$ is an instruction, such that

$$\bigotimes(\sigma_C, \sigma, \sigma_F) \rightsquigarrow^m \bigotimes(\sigma_C, \sigma', \sigma'_F)$$

and such that the following conditions are satisfied:

$$\begin{aligned} \forall \ell \notin \text{wr}(m), \sigma_F(\ell) = \sigma'_F(\ell) \quad \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_F) \\ \text{lock}(m) \subseteq \text{dom}(\sigma) \cup \text{dom}_F(\sigma) \quad \forall r \notin \text{lock}(m), \sigma(r) = \sigma'(r). \end{aligned}$$

The permutation tiles of the asynchronous graph $\mathfrak{s}_{\text{Sep}}$ are defined in the expected way, see the Appendix for details. The resulting definition of the machine model $\mathfrak{s}_{\text{Sep}}$ of separated states ensures that the operation (16) defines a morphism $\bigotimes : \mathfrak{s}_{\text{Sep}} \rightarrow \mathfrak{s}_S$ of asynchronous graphs from $\mathfrak{s}_{\text{Sep}}$ to the stateful model \mathfrak{s}_S .

$$\begin{array}{c}
\frac{}{\Gamma \vdash \{E \mapsto -\}[E] := E' \{E \mapsto E'\}} \text{STORE} \quad \frac{\Gamma \vdash \{P\}C\{Q\}}{\Gamma \vdash \{P * R\}C\{Q * R\}} \text{FRAME} \quad \frac{\Gamma, r : J \vdash \{P\}C\{Q\}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C\{Q * J\}} \text{RES} \\
\frac{\Gamma \vdash \{P_1\}C\{Q_1\} \quad \Gamma \vdash \{P_2\}C\{Q_2\}}{\Gamma \vdash \{P_1 \vee P_2\}C\{Q_1 \vee Q_2\}} \text{DISJ} \quad \frac{\Gamma \vdash \{P_1\}C_1\{Q_1\} \quad \Gamma \vdash \{P_2\}C_2\{Q_2\}}{\Gamma \vdash \{P_1 * P_2\}C_1 \parallel C_2\{Q_1 * Q_2\}} \text{PAR} \quad \frac{P \Rightarrow \text{def}(B) \quad \Gamma \vdash \{(P * J) \wedge B\}C\{Q * J\}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ when } B \text{ do } C\{Q\}} \text{WITH}
\end{array}$$

Figure 1. Inference rules of Concurrent Separation Logic

7 An asynchronous semantics of proofs

In this section, we interpret derivation trees (or proofs) of CSL in our asynchronous semantics. In the same way as we did for the Code in §4, we interpret every proof π of a Hoare triple $\Gamma \vdash \{P\}C\{Q\}$ as an asynchronous transition system (ATS, Definition 4.2). The underlying asynchronous machine model is \mathfrak{A}_{Sep} , the graph of separated states. As in the previous case, our ATSs have the 1-fibration property, and moreover the initial states are all the states that satisfy P , and all the final states satisfy Q . The interpretation $\llbracket \pi \rrbracket_{Sep}$ also satisfies that the second component σ of all its nodes of $\llbracket \pi \rrbracket_{Sep}$ satisfies the invariants of Γ pointwise. In order to define the interpretation of a proof π by induction on its structure, we start by defining a small number of new constructions on ATSs.

The parallel product with separated states In order to define the parallel product $G_1 \parallel G_2$ of two ATSs on the model \mathfrak{A}_{Sep} of separated states, we need to adapt the compatibility condition given by the equality $\lambda_G(x_1) = \lambda_G(x_2)$ in the case of the stateful and stateless models \mathfrak{A}_S and \mathfrak{A}_L . In the case of \mathfrak{A}_{Sep} , two nodes of G_1 and G_2 should be declared compatible when they describe the two “subjective” (and dual) views of the same situation, provided in this case by a separated state of $G_1 \parallel G_2$. This leads us to the notion of *three-party* separated state, defined as a tuple $(\sigma_1, \sigma_2, \sigma^*, \sigma_F)$ where $\sigma_1, \sigma_2, \sigma_F \in \mathbf{LState}$ are logical states, where $\sigma^* : \mathbf{Locks} \rightarrow \mathbf{LState} + \{C_1, C_2, F\}$, and where the product $\otimes(\sigma_1, \sigma_2, \sigma^*, \sigma_F)$ immediately adapted from (15) is well-defined.

We define three functions on these new separated states: the “objective” projection, which corresponds to the view of the whole program $C_1 \parallel C_2$, is defined by:

$$(\sigma_1, \sigma_2, \sigma^*, \sigma_C) \mapsto (\sigma_1 * \sigma_2, \sigma^*[C_i \mapsto C], \sigma_C)$$

the left and right “subjective” projections

$$\lambda_1 : (\sigma_1, \sigma_2, \sigma^*, \sigma_C) \mapsto (\sigma_1, \sigma^*[C_1 \mapsto C, C_2 \mapsto F], \sigma_C * \sigma_2)$$

$$\lambda_2 : (\sigma_1, \sigma_2, \sigma^*, \sigma_C) \mapsto (\sigma_2, \sigma^*[C_1 \mapsto F, C_2 \mapsto C], \sigma_C * \sigma_1)$$

which give the state of the program from the point of view of one of the programs in parallel. This leads us to the following definition:

Definition 7.1. Two separated states $x_1, x_2 \in \mathbf{SState}$ are *compatible* when there exists a three-party separated state y such that $\lambda_1(y) = x_1$ and $\lambda_2(y) = x_2$. Note that the three-party separated state y is unique in that case.

Framing To handle framing (FRAME rule), we need to be able to move a piece of (logical) heap from the Frame side to the Code side. First, we define the framing of a logical state σ_R . Given an ATS G over \mathfrak{A}_{Sep} , we define $\text{frame}[\sigma_R](G)$ pointwise with:

$$\lambda'(x) = \begin{cases} (\sigma_C * \sigma_R, \sigma, \sigma_F) & \text{if } \lambda(x) = (\sigma_C, \sigma, \sigma_F * \sigma_R) \\ \text{undefined} & \text{otherwise} \end{cases}$$

Given such a graph G and a predicate R , we define $\text{frame}[R](G)$ as the following union of ATSs (defined in §A.3):

$$\text{frame}[R](G) = \bigcup_{\sigma_R \models R} \text{frame}[\sigma_R](G)$$

Resource introduction To give a semantics to the resource introduction rule, we need to extend the hiding operator to separated states. More precisely, the action of $\text{hide}[r]$, in addition to replacing $P(r)$ and $V(r)$ with nop , is:

$$\begin{aligned}
(\sigma_C, \sigma \uplus [r \mapsto \sigma], \sigma_F) &\mapsto (\sigma_C * \sigma, \sigma, \sigma_F) \\
(\sigma_C, \sigma \uplus [r \mapsto C \text{ or } F], \sigma_F) &\mapsto (\sigma_C, \sigma, \sigma_F)
\end{aligned}$$

Critical sections Similarly, we extend the definition of $\text{when}[r](G)$ to separated states: to the same underlying graph we associate the asynchronous morphism λ' defined as

$$\lambda'(x) = (\sigma_C, \sigma \uplus [r \mapsto C], \sigma_F) \quad \text{if } \lambda_G(x) = (\sigma_C, \sigma, \sigma_F)$$

We also need to take and release locks in the semantics of the proofs: the ATS $\text{acquire}[r]$ is defined by its Eve moves:

$$(\sigma_C, \sigma \uplus [r \mapsto \sigma], \sigma_F) \xrightarrow{P(r):C} (\sigma_C * \sigma, \sigma \uplus [r \mapsto C], \sigma_F)$$

and $\text{release}[r]$ by (for all $\sigma \in \mathbf{LState}$ satisfying r 's invariant in Γ):

$$(\sigma_C * \sigma, \sigma \uplus [r \mapsto C], \sigma_F) \xrightarrow{V(r):C} (\sigma_C, \sigma \uplus [r \mapsto \sigma], \sigma_F).$$

Machine instructions The rules that correspond to machine instructions $m \in \mathbf{Instr}$ (such as **LOAD**) are interpreted in the obvious way, always preserving the permission associated to affected locations.

Semantics of proofs

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P\}C_1\{Q\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{Q\}C_2\{R\} \end{array}}{\Gamma \vdash \{P\}C_1; C_2\{R\}} \right] \right]_{Sep} = \llbracket \pi_1 \rrbracket_{Sep}; \llbracket \pi_2 \rrbracket_{Sep}$$

For the parallel product rule **PAR**, we use the parallel product of ATSs using the above notion of *compatibility*:

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\}C_1\{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\}C_2\{Q_2\} \end{array}}{\Gamma \vdash \{P_1 * P_2\}C_1 \parallel C_2\{Q_1 * Q_2\}} \right] \right]_{Sep} = \llbracket \pi_1 \rrbracket_{Sep} \parallel \llbracket \pi_2 \rrbracket_{Sep}$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma, r : J \vdash \{(P * J) \wedge B\}C\{Q * J\} \end{array}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ when } B \text{ do } C\{Q\}} \right] \right]_{Sep} = \text{whentrue}[B](\text{acquire}[r]); \text{when}[r](\llbracket \pi \rrbracket_{Sep}); \text{release}[r]$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\}C\{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\}C\{Q_2\} \end{array}}{\Gamma \vdash \{P_1 \vee P_2\}C\{Q_1 \vee Q_2\}} \right] \right]_{Sep} = \llbracket \pi_1 \rrbracket_{Sep} \cup \llbracket \pi_2 \rrbracket_{Sep}$$

The semantics of FRAME and RES are defined in the obvious way using $\text{frame}[R]$ and $\text{hide}[r]$; and the semantics of IF and WHILE are interpreted the same way as for the code. One can give a semantics to the proof rule for conjunction, which requires precision of the invariants (see §A.1).

8 An asynchronous soundness theorem

At this stage, we are ready to state our soundness theorem for Concurrent Separation Logic. We start by observing that every proof π in CSL of a Hoare triple of the form $\Gamma \vdash \{P\}C\{Q\}$ comes equipped with a morphism of asynchronous graphs

$$\mathcal{S}_\pi : \llbracket \pi \rrbracket_{\text{Sep}} \longrightarrow \llbracket C \rrbracket_S$$

which makes the diagram below commute

$$\begin{array}{ccc} \llbracket \pi \rrbracket_{\text{Sep}} & \xrightarrow{\mathcal{S}_\pi} & \llbracket C \rrbracket_S \\ \lambda_\pi \downarrow & & \downarrow \lambda_C \\ \mathfrak{A}_{\text{Sep}} & \xrightarrow{\otimes} & \mathfrak{A}_S \end{array}$$

The morphism \mathcal{S}_π thus defines a morphism of ATS which relates the interpretation of the proof π with the stateful interpretation of C . For every such proof π of a Hoare triple $\Gamma \vdash \{P\}C\{Q\}$, we write

$$\mathcal{L}_\pi : \llbracket \pi \rrbracket_{\text{Sep}} \longrightarrow \llbracket C \rrbracket_L$$

for the composite $\mathcal{L}_\pi = \mathcal{L}_C \circ \mathcal{S}_\pi$ below:

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}_\pi} \llbracket C \rrbracket_S \xrightarrow{\mathcal{L}_C} \llbracket C \rrbracket_L$$

Our soundness theorem follows from two properties of the asynchronous strategy $\llbracket \pi \rrbracket_{\text{Sep}}$ associated to a CSL proof tree π . The first property (*1-soundness*) implies that a well-specified program does not crash during a *valid execution*, that is, an execution which starts from a state satisfying the precondition P and where the Frame performs only legal transitions. The second property (*2-soundness*) implies that such a program does not encounter any data race.

Theorem 8.1 (1-soundness). \mathcal{S}_π is a Code 1-fibration.

A *Code 1-fibration* is a 1-fibration (Def. 1.2) where we only ask that Code transitions can be lifted, similarly to axiom 3 of Def. 4.3. This lifting property reflects the fact that the strategy $\llbracket \pi \rrbracket_{\text{Sep}}$ interpreting the proof π is *winning*, in the sense that every transition performed by the Code on machine states can be lifted (and thus logically justified) by the strategy into a transition between separated states, see [16] for a discussion. This implies in particular that *well specified programs do not go wrong*, because the error state \perp cannot be lifted to a separated state. The next statement is of a different nature: it says that the strategy $\llbracket \pi \rrbracket_{\text{Sep}}$ adapts at the *separated* level to the possible reorderings of scheduling performed at the *stateless* level:

Theorem 8.2 (2-soundness). \mathcal{L}_π is a 2-fibration.

This property implies (in particular) that valid executions of C never produce data races. More deeply, it says that two executions which are equivalent modulo \approx at the stateless level, in the sense that they behave in the same way with respect to the locks (each thread acquires and releases each lock in the same order), are also equivalent modulo \sim at the stateful level. To make this statement formal, consider a well-specified program $\theta \vdash \{P\}C\{Q\}$ and define $\llbracket \{P\}C \rrbracket_S^\tau$ as the subgraph of $\llbracket C \rrbracket_S$ obtained by removing every Frame transition, and keeping only the states which can be reached from an

initial node satisfying the precondition P . We are interested in the morphism

$$\mathcal{L}_C^P : \llbracket \{P\}C \rrbracket_S^\tau \longrightarrow \llbracket C \rrbracket_L$$

obtained by restricting \mathcal{L}_C to the asynchronous subgraph $\llbracket \{P\}C \rrbracket_S^\tau$ of $\llbracket C \rrbracket_S$. This enables us to establish the soundness theorem announced in the introduction:

Theorem 8.3 (Soundness). \mathcal{L}_C^P is a 2-fibration.

9 Conclusion and future works

For the first time, we devise and establish a properly asynchronous version of the Soundness Theorem for Concurrent Separation Logic (CSL). In our formulation, the absence of data races follows from a more fundamental lifting property of scheduling along the stateful-to-stateless translation $\llbracket C \rrbracket_S \rightarrow \llbracket C \rrbracket_L$. The proof of the theorem itself is original in design, and relies on the construction of an asynchronous game semantics of CSL, building on the foundations set in [16]. In future work, we wish to adapt this asynchronous semantics of CSL to weak memory models [4, 11–13, 19] and to distributed algorithms [8]. In another direction of investigation, we want to extend our version of the soundness theorem to a higher-order and axiomatic setting like Iris [10]. Also, now that the asynchronous soundness theorem has been established by semantic means, a nice and instructive challenge will be to prove it again using purely syntactic techniques, in the line adopted for Mezzo [1].

Acknowledgments The authors are grateful to Richard Bornat, Stephen Brookes, Tony Hoare, François Pottier and Viktor Vafeiadis for discussions at an early stage of this work.

References

- [1] T. Balabonski, F. Pottier, and J. Protzenko. Type soundness and race freedom for mezzo. In *FLOPS*, 2014.
- [2] R. Bornat, C. Calcagno, P. O'Hearn, and M. Parkinson. Permission accounting in separation logic. In *POPL*, 2005.
- [3] R. Bornat, C. Calcagno, and H. Yang. Variables as resource in separation logic. *ENTCS*, 155, 2006.
- [4] G. Boudol and G. Petri. Relaxed memory models: An operational approach. *POPL*, 2009.
- [5] S. Brookes. A semantics for concurrent separation logic. In *CONCUR*, 2004.
- [6] L. Fajstrup, E. Goubault, E. Haucourt, S. Mimram, and M. Raussen. *Directed Algebraic Topology and Concurrency*. Springer, 2016.
- [7] A. Gotsman, J. Berdine, and B. Cook. Precision and the conjunction rule in concurrent separation logic. 2011. MFPS.
- [8] A. Gotsman, H. Yang, C. Ferreira, M. Najafzadeh, and M. Shapiro. 'cause i'm strong enough: reasoning about consistency choices in distributed systems. In *POPL*, 2016.
- [9] J. Hayman and G. Winskel. Independence and concurrent separation logic. *LMCS*, 2008.
- [10] R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*, 2015.
- [11] J. Kaiser, H. Dang, D. Dreyer, O. Lahav, and V. Vafeiadis. Strong logic for weak memory: Reasoning about release-acquire consistency in iris. In *ECOOP*, 2017.
- [12] J. Kang, C. Hur, O. Lahav, V. Vafeiadis, and D. Dreyer. A promising semantics for relaxed-memory concurrency. In *POPL*, 2017.
- [13] R. Kavanagh and S. Brookes. A denotational semantics for SPARC TSO. MFPS, 2017.
- [14] P. Mellès and S. Mimram. Asynchronous games: Innocence without alternation. In *CONCUR*, 2007.
- [15] P.-A. Mellès. *Une étude micrologique de la négation*. HDR, U. Paris Diderot, 2017.
- [16] P.-A. Mellès and L. Stefanescu. A game semantics for concurrent separation logic. In *MFPS*, 2017.
- [17] P. W. O'Hearn. Resources, concurrency, and local reasoning. *TCS*, 375, 2007.
- [18] M. J. Parkinson, R. Bornat, and C. Calcagno. Variables as resource in hoare logics. In *LICS*, 2006.
- [19] K. Svendsen, J. Pichon-Pharabod, M. Doko, O. Lahav, and V. Vafeiadis. A separation logic for a promising semantics. 2018. ESOP.
- [20] V. Vafeiadis. Concurrent separation logic and operational semantics. *ENTCS*, 276, 2011.

A Definitions omitted in the core of the paper

A.1 The proof system, and the conjunction

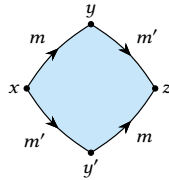
The proof system The inference rules are listed in Figure 2. Notice that the context $\Gamma = r_1 : J_1, \dots, r_k : J_k$ is required to be *precise* in the rule CONJ. This means that each of the predicates J_i is precise in the following sense:

Definition A.1 (Precise predicate). A predicate P is precise when, for every logical state $\sigma \in \mathbf{LState}$, there exists at most one logical state $\sigma' \in \mathbf{LState}$ such that $\sigma' \models P$ and

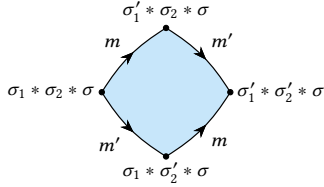
$$\exists \sigma'' \in \mathbf{LState}, \quad \sigma = \sigma' * \sigma''.$$

A.2 The tiles of the machine model of separated states

Like the two other machine models \mathfrak{A}_S and \mathfrak{A}_L , the tiles of \mathfrak{A}_{Sep} are the squares



where the footprints of m and m' at state x are independent. More concretely, an Eve-Eve tile is of the following form (where we only write the first component of each separated state):



For example, the first state σ_C is split into $\sigma_1 * \sigma_2 * \sigma$, where the domain of σ_1 is $\text{wr}(m)$, that of σ_2 is $\text{wr}(m')$ and σ is the rest of σ_C .

A.3 Union and intersection of ATSS

It is sometimes necessary to combine in a single asynchronous strategy two asynchronous strategies G_1 and G_2 whose purpose is to justify the transitions performed by the very same Code C . This is precisely the purpose of the disjoint union $G_1 \cup G_2$.

Union The disjoint union $G_1 \cup G_2$ of two ATSSs G_1 and G_2 over the model \mathfrak{A}_{Sep} of separated states is defined in essentially the same way as the disjoint sum $G_1 \oplus G_2$. The only difference with the disjoint sum $G_1 \oplus G_2$ is that the transitions (or moves) in $G_1 \cup G_2$ are meant to justify the transitions of the *same* program C , whereas the transitions (or moves) of the disjoint sum $G_1 \oplus G_2$ are meant to justify the transitions of two different programs C_1 and C_2 . So, given two derivation trees in CSL

$$\frac{\vdots \pi_1}{\Gamma \vdash \{P_1\}C\{Q_1\}} \quad \frac{\vdots \pi_2}{\Gamma \vdash \{P_2\}C\{Q_2\}}$$

the asynchronous homomorphism associated to $\llbracket \pi_1 \rrbracket_{Sep} \cup \llbracket \pi_2 \rrbracket_{Sep}$ is of the form

$$\mathcal{S}_{\pi_1 \cup \pi_2} : \llbracket \pi_1 \rrbracket_{Sep} \cup \llbracket \pi_2 \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_S$$

whereas the homomorphism associated to the sum is of the form:

$$\mathcal{S}_{\pi_1 \oplus \pi_2} : \llbracket \pi_1 \rrbracket_{Sep} \oplus \llbracket \pi_2 \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_S \oplus \llbracket C \rrbracket_S$$

The disjoint union $\mathcal{S}_{\pi_1 \cup \pi_2}$ may be thus obtained by postcomposing $\mathcal{S}_{\pi_1 \oplus \pi_2}$ with the codiagonal of $\llbracket C \rrbracket_S$:

$$\begin{aligned} & \llbracket \pi_1 \rrbracket_{Sep} \oplus \llbracket \pi_2 \rrbracket_{Sep} \xrightarrow{\mathcal{S}_{\pi_1 \cup \pi_2}} \llbracket C \rrbracket_S \\ = & \llbracket \pi_1 \rrbracket_{Sep} \oplus \llbracket \pi_2 \rrbracket_{Sep} \xrightarrow{\mathcal{S}_{\pi_1 \oplus \pi_2}} \llbracket C \rrbracket_S \oplus \llbracket C \rrbracket_S \xrightarrow{\text{codiagonal}} \llbracket C \rrbracket_S \end{aligned}$$

Intersection In order to interpret the introduction rule of the conjunction, we need to define the intersection of two ATSSs G_1 and G_2 over the same machine model \mathfrak{A} . The definition of $G_1 \cap G_2$ is very similar to the definition of the parallel product. The only difference lies in the fact that the two ATSSs G_1 and G_2 must behave synchronously. The nodes of $G_1 \cap G_2$ are defined in the same way as the nodes of $G_1 \parallel G_2$, that is, as the pairs $x_1 \& x_2$ consisting of a node x_1 of G_1 and of a node x_2 of G_2 , such that

$$\lambda_{G_1}(x_1) = \lambda_{G_2}(x_2)$$

There are two types of transitions. The Code transitions which are pairs of Code transitions:

$$x_1 \& x_2 \xrightarrow{m:C} y_1 \& y_2$$

when $x_1 \xrightarrow{m:C} y_1 \in G_1$ and $x_2 \xrightarrow{m:C} y_2 \in G_2$ and similarly for Frame transitions. A square is a tile in $G_1 \cap G_2$ precisely when it is the superposition of two tiles of G_1 and of G_2 .

A.4 Semantics of the conjunction

The interpretation of the conjunction rule relies on the intersection of ATSSs just defined:

$$\left[\frac{\vdots \pi_1 \quad \vdots \pi_2}{\Gamma \vdash \{P_1\}C\{Q_1\} \quad \Gamma \vdash \{P_2\}C\{Q_2\}} \right] = \llbracket \pi_1 \rrbracket_{Sep} \cap \llbracket \pi_2 \rrbracket_{Sep}$$

Note that we do not rely on the preciseness of the context Γ to define the interpretation of the CONJ proof rule, though precision is needed in order to establish the soundness of the rule.

B Proof of the 1-soundness theorem (Thm. 8.1)

In this section we prove the 1-soundness theorem of CSL. The proof is done by induction on the structure of the proof tree. Each case of the induction is given its own lemma. We focus on the non-leaf rules of CSL.

B.1 Parallel composition

We begin with the rule for parallel composition. The corresponding case in the induction is the following.

Lemma B.1. *Suppose that π is the derivation tree*

$$\frac{\vdots \pi_1 \quad \vdots \pi_2}{\Gamma \vdash \{P_1\}C_1\{Q_1\} \quad \Gamma \vdash \{P_2\}C_2\{Q_2\}} \text{PAR}$$

and that the interpretation

$$\mathcal{S}_i : \llbracket \pi_i \rrbracket_{Sep} \longrightarrow \llbracket C_i \rrbracket_S$$

is a 1-fibration on Code transitions, for $i = 1, 2$. In that case, the asynchronous graph morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{Sep} \longrightarrow \llbracket C_1 \parallel C_2 \rrbracket_S$$

is also a Code 1-fibration.

$$\begin{array}{c}
\frac{}{\Gamma \vdash \{(\text{Own}_\top(x) * P) \wedge E = v\}x := E\{(\text{Own}_\top(x) * P) \wedge x = v\}} \text{AFF} \qquad \frac{}{\Gamma \vdash \{E \mapsto -\}[E] := E'\{E \mapsto E'\}} \text{STORE} \\
\frac{x \notin \text{fv}(E)}{\Gamma \vdash \{E \mapsto^P v * \text{Own}_\top(x)\}x := [E]\{E \mapsto^P v * \text{Own}_\top(x) * x = v\}} \text{LOAD} \qquad \frac{\Gamma \vdash \{P\}C_1\{Q\} \quad \Gamma \vdash \{Q\}C_2\{R\}}{\Gamma \vdash \{P\}C_1; C_2\{R\}} \text{SEQ} \\
\frac{P \Rightarrow \text{def}(B) \quad \Gamma \vdash \{P \wedge B\}C_1\{Q\} \quad \Gamma \vdash \{P \wedge \neg B\}C_2\{Q\}}{\Gamma \vdash \{P\} \text{if } B \text{ then } C_1 \text{ else } C_2\{Q\}} \text{IF} \qquad \frac{\Gamma \text{ is precise} \quad \Gamma \vdash \{P_1\}C\{Q_1\} \quad \Gamma \vdash \{P_2\}C\{Q_2\}}{\Gamma \vdash \{P_1 \wedge P_2\}C\{Q_1 \wedge Q_2\}} \text{CONJ} \\
\frac{\Gamma \vdash \{P_1\}C\{Q_1\} \quad \Gamma \vdash \{P_2\}C\{Q_2\}}{\Gamma \vdash \{P_1 \vee P_2\}C\{Q_1 \vee Q_2\}} \text{DISJ} \qquad \frac{\Gamma, r : J \vdash \{P\}C\{Q\}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C\{Q * J\}} \text{RES} \qquad \frac{P \Rightarrow \text{def}(B) \quad \Gamma \vdash \{(P * J) \wedge B\}C\{Q * J\}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ when } B \text{ do } C\{Q\}} \text{WITH} \\
\frac{\Gamma \vdash \{P_1\}C_1\{Q_1\} \quad \Gamma \vdash \{P_2\}C_2\{Q_2\}}{\Gamma \vdash \{P_1 * P_2\}C_1 \parallel C_2\{Q_1 * Q_2\}} \text{PAR} \qquad \frac{\Gamma \vdash \{P\}C\{Q\}}{\Gamma \vdash \{P * R\}C\{Q * R\}} \text{FRAME}
\end{array}$$

Figure 2. Inference rules of Concurrent Separation Logic

A Code transition in $\llbracket C_1 \parallel C_2 \rrbracket_S$ is (without loss of generality) a pair of compatible transitions: one Code transition from $\llbracket C_1 \rrbracket_S$ and one Frame transition from $\llbracket C_2 \rrbracket_S$. We need to lift this transition into a move in $\llbracket \pi \rrbracket_{\text{Sep}} = \llbracket \pi_1 \rrbracket_{\text{Sep}} \parallel \llbracket \pi_2 \rrbracket_{\text{Sep}}$. Using the induction hypothesis, we can show that Eve can lift the former into $\llbracket \pi_1 \rrbracket_{\text{Sep}}$, and we can lift the latter because λ_{π_2} is a Frame 1-fibration. Therefore, we can lift the Code transition from $\llbracket C_1 \parallel C_2 \rrbracket_S$ into $\llbracket \pi_1 \parallel \pi_2 \rrbracket_{\text{Sep}}$.

Proof. Consider a node x in $\llbracket \pi \rrbracket_{\text{Sep}}$ whose label is $(\sigma_C, \sigma, \sigma_F)$, and let $s = \otimes(\sigma_C, \sigma, \sigma_F)$. Consider a Code transition in $\llbracket C_1 \parallel C_2 \rrbracket_S$ of the form (writing the images of the nodes under λ)

$$s \xrightarrow{m:C} s' \quad (17)$$

whose starting node is $\mathcal{S}(x)$. By definition of the parallel product there exists a three-party separated state

$$(\sigma_1, \sigma_2, \sigma, \sigma_F)$$

whose projection through **proj** is $(\sigma_C, \sigma, \sigma_F)$. By definition of $\llbracket C_1 \parallel C_2 \rrbracket_S$, the transition (17) is of the form

$$a_1 | a_2 \xrightarrow{m:C} b_1 | b_2$$

with $\lambda_1(a_1) = \lambda_2(a_2) = s$ and $\lambda_1(b_1) = \lambda_2(b_2) = s'$. Our goal is to find two moves, one in $\llbracket \pi_1 \rrbracket_{\text{Sep}}$ of the form

$$(\sigma_1, \sigma, \sigma_F * \sigma_2) \xrightarrow{m:C} (\sigma'_1, \sigma', \sigma_F * \sigma_2)$$

that is mapped under \mathcal{S}_1 to the transition:

$$a_1 \xrightarrow{m} a_2$$

and one in $\llbracket \pi_2 \rrbracket_{\text{Sep}}$ of the form

$$(\sigma_2, \sigma, \sigma_F * \sigma_1) \xrightarrow{m:F} (\sigma_2, \sigma', \sigma_F * \sigma'_1)$$

that is mapped under \mathcal{S}_2 to the transition:

$$b_1 \xrightarrow{m:F} b_2.$$

(Note that we omit to write the change of perspective on the σ .) The first exists according to the hypothesis on \mathcal{S}_1 , and the second because λ_2 is a 1-fibration on Frame transitions. \square

B.2 Sequential composition

The case of sequential composition is easy, since a Code transition of $C_1; C_2$ is either a transition from C_1 or a transition from C_2 , and both cases follow immediately from the induction hypotheses.

Lemma B.2. *Suppose that π is the derivation tree*

$$\begin{array}{c}
\vdots \pi_1 \qquad \vdots \pi_2 \\
\frac{\Gamma \vdash \{P\}C_1\{Q\} \quad \Gamma \vdash \{Q\}C_2\{R\}}{\Gamma \vdash \{P\}C_1; C_2\{R\}} \text{SEQ}
\end{array}$$

and that the interpretation

$$\mathcal{S}_i : \llbracket \pi_i \rrbracket_{\text{Sep}} \rightarrow \llbracket C_i \rrbracket_S$$

is a 1-fibration on Code transitions, for $i = 1, 2$. In that case, the asynchronous graph morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C_1; C_2 \rrbracket_S$$

is also a Code 1-fibration.

Proof. A Code transition in $\llbracket C_1; C_2 \rrbracket_S$ is either a transition in $\llbracket C_1 \rrbracket_S$ or a transition in $\llbracket C_2 \rrbracket_S$. The result follows from the hypotheses. \square

B.3 The frame rule

Lemma B.3. *Suppose that π is the derivation tree*

$$\begin{array}{c}
\vdots \pi' \\
\frac{\Gamma \vdash \{P\}C\{Q\}}{\Gamma \vdash \{P * R\}C\{Q * R\}} \text{FRAME}
\end{array}$$

and that the interpretation

$$\mathcal{S} : \llbracket \pi' \rrbracket_{\text{Sep}} \rightarrow \llbracket C \rrbracket_S$$

is a 1-fibration on Code transitions. In that case, the asynchronous graph morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C \rrbracket_S$$

is also a Code 1-fibration.

Proof. Let us consider a node x in $\llbracket \pi \rrbracket_{\text{Sep}}$ and a Code transition

$$a \xrightarrow{m:C} b \in \llbracket C \rrbracket_S$$

where $\mathcal{S}(x) = a$. Recall that the frame rule is interpreted as

$$\text{frame}[R](\llbracket \pi' \rrbracket_{\text{Sep}}) = \bigcup_{\sigma_R \models R} \text{frame}[\sigma_R](\llbracket \pi' \rrbracket_{\text{Sep}})$$

Therefore, the node x belongs to one of the copies of $\text{frame}[\sigma_R](\llbracket \pi' \rrbracket_{\text{Sep}})$, for some σ_R , and there is a node x' above a in $\llbracket \pi' \rrbracket_{\text{Sep}}$ such that

$$\lambda(x) = (\sigma_C * \sigma_R, \sigma, \sigma_F) \quad \text{and} \quad \lambda'(x') = (\sigma_C, \sigma, \sigma_F * \sigma_R)$$

with $\sigma_R \models R$. By the hypothesis on \mathcal{S}' , there is a node $y' \in \llbracket \pi' \rrbracket_{\text{Sep}}$ and a move

$$x' \xrightarrow{m:C} y'$$

above the Code transition above. Hence, $\lambda'(y')$ is of the form

$$\lambda'(y') = (\sigma'_C, \sigma', \sigma_F * \sigma_R)$$

which implies that there is a transition

$$x \xrightarrow{m:C} y$$

above the Code transition in frame $[\sigma_R](\llbracket \pi' \rrbracket_{Sep}) \subseteq \llbracket \pi \rrbracket_{Sep}$ with

$$\lambda(y) = (\sigma'_C * \sigma_R, \sigma', \sigma_F)$$

□

B.4 Resource introduction

The rule for resource introduction is interpreted using the $\text{hide}[r]$ construction, which hides the new resource. The proof consists basically in showing that if some Code transition t can be lifted into a move T , then $\text{hide}[r](t)$ can be lifted into $\text{hide}[r](T)$.

Lemma B.4. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi' \\ \Gamma, r : J \vdash \{P\} C' \{Q\} \end{array}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C' \{Q * J\}} \text{RES}$$

and that the interpretation

$$\mathcal{S}' : \llbracket \pi' \rrbracket_{Sep} \rightarrow \llbracket C' \rrbracket_S$$

is a 1-fibration on Code transitions. In that case, the asynchronous morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{Sep} \rightarrow \llbracket \text{resource } r \text{ do } C' \rrbracket_S$$

is also a 1-fibration on Code transitions.

Proof. Let us consider a node x in $\llbracket \pi \rrbracket_{Sep}$ and a Code transition t

$$a \xrightarrow{m:C} b$$

where $\mathcal{S}(x) = a$. By definition of $\llbracket \text{resource } r \text{ do } C' \rrbracket_S$, t is the image under $\text{hide}[r]$ of a transition t' :

$$a \xrightarrow{m:C} b \in \llbracket C' \rrbracket_S$$

Case 1. Suppose \underline{m} is neither $P(r)$ nor $V(r)$. Then $m = \underline{m}$ and the situation is:

$$\lambda(x) = (\sigma_C * \sigma, \sigma, \sigma_F) \quad \text{and} \quad \lambda'(x) = (\sigma_C, \sigma \uplus [r \mapsto \sigma], \sigma_F)$$

By the hypothesis on \mathcal{S}' , there is a node $y \in \llbracket \pi' \rrbracket_{Sep}$ and a move

$$x \xrightarrow{m:C} y$$

above the Code transition t' and such that

$$\lambda'(y) = (\sigma'_C, \sigma' \uplus [r \mapsto \sigma], \sigma_F)$$

Hence, there is a lifting of t starting from x .

Case 2. Suppose, for example, that $\underline{m} = P(r)$. In that case, $m = \text{nop}$, and

$$\lambda(x) = (\sigma_C * \sigma, \sigma, \sigma_F) \quad \text{and} \quad \lambda'(x) = (\sigma_C, \sigma \uplus [r \mapsto \sigma], \sigma_F)$$

By the hypothesis on \mathcal{S}' , there is a node $y \in \llbracket \pi' \rrbracket_{Sep}$ and a move

$$x \xrightarrow{m:C} y$$

above the Code transition T' and such that

$$\lambda'(y) = (\sigma'_C * \sigma, \sigma' \uplus [r \mapsto C], \sigma_F)$$

This means that there is a move above T of the form (with labels instead of nodes):

$$(\sigma_C * \sigma, \sigma, \sigma_F) \xrightarrow{\text{nop}:C} (\sigma_C * \sigma, \sigma, \sigma_F)$$

□

B.5 Critical sections

Lemma B.5. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi' \\ \Gamma \vdash \{(P * J) \wedge B\} C' \{Q * J\} \end{array}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ when } B \text{ do } C' \{Q\}} \text{WITH}$$

and that the interpretation

$$\mathcal{S}' : \llbracket \pi' \rrbracket_{Sep} \rightarrow \llbracket C' \rrbracket_S$$

is a 1-fibration on Code transitions. In that case, the asynchronous morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{Sep} \rightarrow \llbracket \text{with } r \text{ when } B \text{ do } C' \rrbracket_S$$

is also a 1-fibration on Code transitions.

Proof. Recall that, in that case, the semantics of the Code and of the derivation trees are defined as

$$\begin{aligned} \llbracket C \rrbracket_{Sep} &= \text{whentruel}[B](\llbracket P(r) \rrbracket_S; \text{when}[r](\llbracket C' \rrbracket_S); \llbracket V(r) \rrbracket_S) \\ &\quad \oplus \text{whenabort}[B] \end{aligned}$$

$$\llbracket \pi \rrbracket_{Sep} = \text{whentruel}[B](\text{acquire}[r]; \text{when}[r](\llbracket \pi' \rrbracket_{Sep}); \text{release}[r])$$

Consider a Code transition t of $\llbracket C \rrbracket_S$

$$a \xrightarrow{m:C} b$$

and a node x in $\llbracket C \rrbracket_{Sep}$ above a . According to the side-condition $P \Rightarrow \text{def}(B)$, we know that this transition is not in $\text{whenabort}[B]$, because all the variables in B are necessarily well-defined. There are now three possibilities:

Case 1: $t \in \text{when}[r](\llbracket C' \rrbracket_S)$. In that case the image of x under λ is of the form

$$\lambda(x) = (\sigma_C, \sigma \uplus [r \mapsto C], \sigma_F)$$

and t is also a transition in $\llbracket C' \rrbracket_{Sep}$, whose image under λ' is:

$$(\mu, L \uplus \{r\}) \xrightarrow{m:C} (\mu', L' \uplus \{r\})$$

such that the image of t under λ is:

$$(\mu, L) \xrightarrow{m:C} (\mu', L').$$

Moreover, by hypothesis, we can lift the transition t from $\llbracket C' \rrbracket_S$ to a move whose labels are of the form

$$(\sigma_C, \sigma, \sigma_F) \xrightarrow{m:C} (\sigma'_C, \sigma', \sigma_F)$$

Therefore, the same move is a lifting of t in $\llbracket C \rrbracket_{Sep}$, and its labels are

$$(\sigma_C, \sigma \uplus [r \mapsto C], \sigma_F) \xrightarrow{m:C} (\sigma'_C, \sigma' \uplus [r \mapsto C], \sigma_F)$$

Cases 2 and 3: t is in $\text{whentrue}[B](\llbracket P(r) \rrbracket_S)$ or in $\llbracket V(r) \rrbracket_S$. Follows from the definitions. \square

B.6 The conjunction

In this section, we suppose that Γ is *precise* and that π is the following derivation tree:

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\}C\{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\}C\{Q_2\} \end{array}}{\Gamma \vdash \{P_1 \wedge P_2\}C\{Q_1 \wedge Q_2\}} \text{CONJ}$$

This rule is not sound when the context Γ is not precise, see [7]. In our setting, precision implies that Eve has at most one way of lifting a given transition from the Code. This is very useful, because, on the one hand, the induction hypothesis tells us that there are two ways of lifting the same Code transition. And, on the other hand, precision implies that they are actually the same, and therefore define a move in the interpretation of CONJ, which is defined using a synchronous product.

Lemma B.6. *Suppose the interpretation*

$$\mathcal{S}_i : \llbracket \pi_i \rrbracket_{Sep} \rightarrow \llbracket C \rrbracket_S$$

is a 1-fibration on Code transitions, for $i = 1, 2$. In that case, the asynchronous graph morphism

$$\mathcal{S} : \llbracket \pi \rrbracket_{Sep} \rightarrow \llbracket C \rrbracket_S$$

is also a Code 1-fibration.

Proof. Consider a Code transition of $\llbracket C \rrbracket_S$

$$a \xrightarrow{m:C} b$$

and a node x in $\llbracket \pi \rrbracket_{Sep}$ above a . By definition of the semantics of π , there exist two nodes $x_1 \in \llbracket \pi_1 \rrbracket_{Sep}$ and $x_2 \in \llbracket \pi_2 \rrbracket_{Sep}$ each above a and such that

$$\lambda_1(x_1) = \lambda_2(x_2) = \lambda(x).$$

Therefore, according to the hypotheses there exist two moves in $\llbracket \pi_1 \rrbracket_{Sep}$ and in $\llbracket \pi_2 \rrbracket_{Sep}$ each of the form

$$x_i \xrightarrow{m:C} y_i.$$

It suffices to show that $\lambda(y_1) = \lambda(y_2)$, since it implies that there exists y in $\llbracket \pi \rrbracket_{Sep}$ such that

$$x \xrightarrow{m:C} y$$

is above the transition on $\llbracket C \rrbracket_S$ we considered.

Let us show, then, that $\lambda(y_1) = \lambda(y_2)$. Since all the permissions are preserved, by definition of the strategies, the first and the last component of the two separated states coincide. Moreover, precision tells us that there is at most one sub-logical state that satisfies the invariants, which means that the middle components coincide as well the middle components coincide as well. \square

C Proof of the 2-soundness theorem (Thm. 8.2)

In this section, we establish the 2-soundness theorem (Theorem 8.2) by induction on the proof π of a Hoare triple of the form $\Gamma \vdash \{P\}C\{Q\}$. First, we begin with the case of the PAR rule, which we find the most interesting.

C.1 The parallel rule

Lemma C.1. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\}C_1\{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\}C_2\{Q_2\} \end{array}}{\Gamma \vdash \{P_1 * P_2\}C_1 \parallel C_2\{Q_1 * Q_2\}} \text{PAR}$$

and that the interpretation

$$\mathcal{L}_i : \llbracket \pi_i \rrbracket_{Sep} \rightarrow \llbracket C_i \rrbracket_L$$

is a 2-fibration, for $i = 1, 2$. In that case, the asynchronous graph morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{Sep} \rightarrow \llbracket C_1 \parallel C_2 \rrbracket_L$$

is also a 2-fibration.

Proof. Write $C = C_1 \parallel C_2$. Let us consider a tile in $\llbracket C \rrbracket_L$:

$$\begin{array}{ccc} a_1|a_2 & \begin{array}{c} \xrightarrow{m} b_1|b_2 \\ \xrightarrow{m'} b'_1|b'_2 \end{array} & \begin{array}{c} \xrightarrow{m'} c_1|c_2 \\ \xrightarrow{m} c_1|c_2 \end{array} \\ & \approx & \end{array} \quad (18)$$

such that there exist two transitions in $\llbracket C \rrbracket_S$

$$\begin{array}{ccc} & \xrightarrow{m} y_1|y_2 & \xrightarrow{m'} z_1|z_2 \\ x_1|x_2 & & \end{array} \quad (19)$$

that are sent through \mathcal{L} onto the upper path in (18).

By definition of the parallel product of ATSS, since (18) is a tile, its two projections are tiles as well:

$$\begin{array}{ccc} u_{11} \rightarrow b_1 \rightarrow v'_{11} & & u_{12} \rightarrow b_2 \rightarrow v'_{12} \\ a_1 \rightarrow \approx \rightarrow c_1 & & a_2 \rightarrow \approx \rightarrow c_2 \\ v_{11} \rightarrow b'_1 \rightarrow u'_{11} & & v_{12} \rightarrow b'_2 \rightarrow u'_{12} \end{array}$$

By hypothesis, the asynchronous morphisms \mathcal{L}_1 and \mathcal{L}_2 are both 2-fibrations. This means that the following two squares above them in $\llbracket C_1 \rrbracket_S$ and in $\llbracket C_2 \rrbracket_S$ are tiles as well:

$$\begin{array}{ccc} u_{11} \rightarrow y_1 \rightarrow v'_{11} & & u_{12} \rightarrow y_2 \rightarrow v'_{12} \\ x_1 \rightarrow \sim \rightarrow z_1 & & x_2 \rightarrow \sim \rightarrow z_2 \\ s_1 \rightarrow y'_1 \rightarrow t_1 & & s_2 \rightarrow y'_2 \rightarrow t_2 \end{array}$$

for some nodes y'_1 and y'_2 in $\llbracket C_1 \rrbracket_S$ and $\llbracket C_2 \rrbracket_S$ respectively. By definition of tiles in $\llbracket C_1 \parallel C_2 \rrbracket_S$, to show that there exists a tile completing (19) above (18), it suffices to show that the two states $\lambda_1(y'_1)$ and $\lambda_2(y'_2)$ are *compatible*, in the sense of Definition 7.1. This follows from the following lemma. \square

Lemma C.2. *Suppose given two ATSS over separated states G_1 and G_2 , and two tiles:*

$$\begin{array}{ccc} u_{11} \rightarrow y_1 \rightarrow v'_{11} & & u_{12} \rightarrow y_2 \rightarrow v'_{12} \\ x_1 \rightarrow \sim \rightarrow z_1 & & x_2 \rightarrow \sim \rightarrow z_2 \\ s_1 \rightarrow y'_1 \rightarrow t_1 & & s_2 \rightarrow y'_2 \rightarrow t_2 \end{array}$$

where the upper paths are compatible. Then y'_1 and y'_2 , and the paths $s_1; t_1$ and $s_2; t_2$ define a path in $G_1 \parallel G_2$.

Proof. By case analysis on the polarities of the transitions. \square

C.2 Sequential composition

Lemma C.3. Suppose that π is the derivation tree

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P\}C_1\{Q\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{Q\}C_2\{R\} \end{array}}{\Gamma \vdash \{P\}C_1; C_2\{R\}} \text{SEQ}$$

and that the interpretation

$$\mathcal{L}_i : \llbracket \pi_i \rrbracket_{\text{Sep}} \rightarrow \llbracket C_i \rrbracket_L$$

is a 2-fibration, for $i = 1, 2$. In that case, the asynchronous graph morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C_1; C_2 \rrbracket_L$$

is also a 2-fibration.

Proof. Recall that the semantics of sequential composition is defined by:

$$\begin{aligned} \llbracket C_1; C_2 \rrbracket_L &= \llbracket C_1 \rrbracket_L; \llbracket C_2 \rrbracket_L \\ \llbracket \pi_1; \pi_2 \rrbracket_{\text{Sep}} &= \llbracket \pi_1 \rrbracket_{\text{Sep}}; \llbracket \pi_2 \rrbracket_{\text{Sep}} \end{aligned}$$

This means that a tile in $\llbracket C_1; C_2 \rrbracket_L$, is either a tile in $\llbracket C_1 \rrbracket_L$ or a tile in $\llbracket C_2 \rrbracket_L$. By the hypothesis on \mathcal{L}_i , it is clear that in either case we can lift the tile in either $\llbracket \pi_1 \rrbracket_{\text{Sep}}$ or in $\llbracket \pi_2 \rrbracket_{\text{Sep}}$, and thus in $\llbracket \pi_1; \pi_2 \rrbracket_{\text{Sep}}$. \square

C.3 Resource introduction

Lemma C.4. Suppose that π is the derivation tree

$$\frac{\begin{array}{c} \vdots \pi' \\ \Gamma, r : J \vdash \{P\}C'\{Q\} \end{array}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C'\{Q * J\}} \text{RES}$$

and that the interpretation

$$\mathcal{L}' : \llbracket \pi' \rrbracket_{\text{Sep}} \rightarrow \llbracket C' \rrbracket_L$$

is a 2-fibration. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket \text{resource } r \text{ do } C' \rrbracket_L$$

is also a 2-fibration.

Proof. Write $C = \text{resource } r \text{ do } C'$. Suppose there is a tile in $\llbracket C \rrbracket_L$ of the form

$$\begin{array}{ccc} a & \xrightarrow{m} & b \\ & \searrow m' & \nearrow m' \\ & & c \end{array} \approx \begin{array}{ccc} & & b \\ & \nearrow m' & \searrow m' \\ a & & b' \\ & \searrow m' & \nearrow m \\ & & c \end{array}$$

and that its upper path is the image of the following tile by \mathcal{L}

$$\begin{array}{ccc} & & (\sigma'_C, \sigma_2, \sigma_F) \\ & \nearrow m & \searrow m' \\ (\sigma_C, \sigma_1, \sigma_F) & & (\sigma''_C, \sigma_3, \sigma'_F) \end{array} \quad (20)$$

Recall that the semantics of resource introduction is given by

$$\begin{aligned} \llbracket \text{resource } r \text{ do } C' \rrbracket_L &:= \text{hide}[r](\llbracket C' \rrbracket_L) \\ \llbracket \text{RES}(\pi') \rrbracket_{\text{Sep}} &:= \text{hide}[r](\llbracket \pi' \rrbracket_{\text{Sep}}) \end{aligned}$$

By the definition of the semantics, it is the image under $\text{hide}[r]$ of some tile

$$\begin{array}{ccc} a & \xrightarrow{m} & b \\ & \searrow m' & \nearrow m' \\ & & c \end{array} \approx \begin{array}{ccc} & & b \\ & \nearrow m' & \searrow m' \\ a & & b' \\ & \searrow m' & \nearrow m \\ & & c \end{array}$$

Suppose, first, that \underline{m} touches r ; say $\underline{m} = P(r)$ for example. Since the square above is a tile, we know that, in that case, \underline{m}' is neither $P(r)$ nor $V(r)$. This implies that the path (20) is the image by $\text{hide}[r]$ of a path:

$$\begin{array}{ccc} & & (\underline{\sigma}_C * \sigma, \sigma \uplus [r \mapsto C], \sigma_F) \\ & \nearrow P(r) & \searrow m' \\ (\underline{\sigma}_C, \sigma \uplus [r \mapsto \sigma], \sigma_F) & & (\underline{\sigma}''_C, \sigma_3 \uplus [r \mapsto C], \sigma'_F) \end{array}$$

with $\sigma_C = \underline{\sigma}_C * \sigma$. By induction, this path can be completed into a tile in $\llbracket \pi' \rrbracket_{\text{Sep}}$ of the form

$$\begin{array}{ccc} & & (\underline{\sigma}_C * \sigma, \sigma \uplus [r \mapsto C], \sigma_F) \\ & \nearrow P(r) & \searrow m' \\ (\underline{\sigma}_C, \sigma \uplus [r \mapsto \sigma], \sigma_F) & \sim & (\underline{\sigma}''_C, \sigma_3 \uplus [r \mapsto C], \sigma'_F) \\ & \searrow m' & \nearrow P(r) \\ & & (\underline{\sigma}'''_C, \sigma \uplus [r \mapsto \sigma], \sigma_F) \end{array}$$

which implies that $\underline{\sigma}''_C = \underline{\sigma}'''_C * \sigma$. Finally, this implies that (20) can be completed into a tile:

$$\begin{array}{ccc} & & (\sigma_C, \sigma, \sigma_F) \\ & \nearrow \text{nop} & \searrow m' \\ (\sigma_C, \sigma, \sigma_F) & \sim & (\sigma''_C, \sigma_3, \sigma'_F) \\ & \searrow m' & \nearrow \text{nop} \\ & & (\sigma''_C, \sigma, \sigma'_F) \end{array}$$

Note that this reasoning holds when m' is replaced with a Frame move. (indeed, in the proof above, we accepted that m' change σ_F). The case where neither \underline{m} nor \underline{m}' touch r is similar. \square

C.4 Critical sections

Before we prove the 2-fibration lemma for the rule WITH, we analyze the structure of $\text{when}[r]$. Recall that this ATS is built in two steps (see §4.2): first, the nodes and transitions of G are lifted by adding the new resource r into the sets of locked resources of all the states (and in the case of separated states, we add that r is locked by the Code), and, second, we add Environment transitions to make it an ATS. Call the first kind of transition natural, and the second artificial. The artificial transitions correspond to the case where the Environment touches the lock while it is held by the Code. This is, of course, a highly incorrect behavior. Thankfully, the constraints on the Frame moves in the Machine model of Separated States, and hence in the interpretation of proofs, rule these transitions out, in the following sense. First, natural transitions are stable by homotopy.

Lemma C.5. Given a tile T in $\text{when}[r](G)$, if the upper path is made of two natural transitions, then so does the lower path.

Proof sketch. In a tile, opposite transitions have the same footprints, and therefore the same behavior on locks. \square

In the interpretations of proofs, all Code transitions are natural.

Lemma C.6. If G is an ATS over the machine model of separated states, then all the Code transitions of $\text{when}[r]$ are natural.

Proof sketch. Since the lock r is held by the Code, the Frame cannot touch it. \square

Lemma C.7. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi' \\ \Gamma \vdash \{(P * J) \wedge B\} C' \{Q * J\} \end{array}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ when } B \text{ do } C' \{Q\}} \text{WITH}$$

and that the interpretation

$$\mathcal{L}' : \llbracket \pi' \rrbracket_{\text{Sep}} \longrightarrow \llbracket C' \rrbracket_L$$

is a 2-fibration. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{\text{Sep}} \longrightarrow \llbracket \text{with } r \text{ when } B \text{ do } C' \rrbracket_L$$

is also a 2-fibration.

Proof. Write $C = \text{with } r \text{ when } B \text{ do } C'$. Recall that the semantics of C is defined as:

$$\llbracket C' \rrbracket_{\text{Sep}} = \text{whentrue}[B](\llbracket P(r) \rrbracket_{\text{Sep}}; \text{when}[r](\llbracket C \rrbracket_{\text{Sep}}); \llbracket V(r) \rrbracket_{\text{Sep}}) \oplus \text{whenabort}[B].$$

By definition of sequential composition of ATSS, a tile in $\llbracket C \rrbracket_L$ contains a Code transition from $\llbracket P(r) \rrbracket_L$ (or $\llbracket V(r) \rrbracket_L$) only if it is a Code/Frame tile. This case is easy because $P(r)$ and the Adam move touch distinct components of the separated states (since it is a tile, Adam cannot touch the r component of σ), and λ is a Code-Frame 2-fibration by definition of ATSS.

Consider a path of the following form in $\llbracket \pi \rrbracket_{\text{Sep}}$:

$$\begin{array}{ccc} & (\sigma_C, \sigma_2 \uplus [r \mapsto C], \sigma_F) & \\ m \nearrow & & \searrow m' \\ (\sigma_C, \sigma \uplus [r \mapsto C], \sigma_F) & & (\sigma''_C, \sigma_3 \uplus [r \mapsto C], \sigma'_F) \end{array} \quad (21)$$

and consider a tile in $\llbracket C \rrbracket_L$ whose upper path is the image of the above path under \mathcal{L} .

We can suppose that this tile is in $\text{when}[r](G)$. According to the two lemmas above, it is of the form (where we write, instead of the nodes themselves, their images under λ')

$$\begin{array}{ccc} L_1 \uplus \{r\} & \xrightarrow{m} & L_2 \uplus \{r\} & \xrightarrow{m'} & L_3 \uplus \{r\} \\ & \searrow m' & \approx & \nearrow m & \\ & & L'_2 \uplus \{r\} & & \end{array} \quad (22)$$

and, moreover, there is a tile in $\llbracket C' \rrbracket_L$ (on the same nodes, since the map $\text{when}[r]$ is defined pointwise) of the form:

$$\begin{array}{ccc} L_1 & \xrightarrow{m} & L_2 & \xrightarrow{m'} & L_3 \\ & \searrow m' & \approx & \nearrow m & \\ & & L'_2 & & \end{array} \quad (23)$$

By hypothesis, and by the definition of the semantics of the rule WITH, there exists a tile in $\llbracket \pi' \rrbracket_{\text{Sep}}$ above the tile (23), which is of

the form:

$$\begin{array}{ccc} & (\sigma_C, \sigma_2, \sigma_F) & \\ m \nearrow & & \searrow m' \\ (\sigma_C, \sigma, \sigma_F) & \sim & (\sigma''_C, \sigma_3, \sigma'_F) \\ & \searrow m' & \nearrow m \\ & (\sigma''_C, \sigma'_2, \sigma'_F) & \end{array}$$

where the domain of the σ 's does not contain r . This finally means that we can complete the path (21) into the following tile above (22) in $\llbracket \pi \rrbracket_{\text{Sep}}$:

$$\begin{array}{ccc} & (\sigma_C, \sigma_2 \uplus [r \mapsto C], \sigma_F) & \\ m \nearrow & & \searrow m' \\ (\sigma_C, \sigma \uplus [r \mapsto C], \sigma_F) & \sim & (\sigma''_C, \sigma_3 \uplus [r \mapsto C], \sigma'_F) \\ & \searrow m' & \nearrow m \\ & (\sigma''_C, \sigma'_2 \uplus [r \mapsto C], \sigma'_F) & \end{array}$$

\square

C.5 Conjunction

Lemma C.8. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\} C \{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\} C \{Q_2\} \end{array}}{\Gamma \vdash \{P_1 \wedge P_2\} C \{Q_1 \wedge Q_2\}} \text{CONJ}$$

where Γ is precise, and that the interpretation

$$\mathcal{L}_i : \llbracket \pi_i \rrbracket_{\text{Sep}} \longrightarrow \llbracket C \rrbracket_L$$

is a 2-fibration, for $i = 1, 2$. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{\text{Sep}} \longrightarrow \llbracket C \rrbracket_L$$

is also a 2-fibration.

Proof. Consider a tile T in $\llbracket C \rrbracket_L$ such that its upper path can be lifted to a path p in $\llbracket \pi \rrbracket_{\text{Sep}}$. By definition of $\llbracket \pi \rrbracket_{\text{Sep}} := \llbracket \pi_1 \rrbracket_{\text{Sep}} \cap \llbracket \pi_2 \rrbracket_{\text{Sep}}$, the path p corresponds to a path p_1 in $\llbracket \pi_1 \rrbracket_{\text{Sep}}$ and p_2 in $\llbracket \pi_2 \rrbracket_{\text{Sep}}$, which all have the same image under λ , λ_1 and λ_2 respectively. Moreover, by hypothesis, the tile T can be lifted to tiles T_1 and T_2 in $\llbracket \pi_1 \rrbracket_{\text{Sep}}$ and in $\llbracket \pi_2 \rrbracket_{\text{Sep}}$ respectively, such that their upper paths are p_1 and p_2 , respectively. Since there is at most one tile which has a given upper path (Axiom 2 of asynchronous graphs), the two tiles have equal images under λ_1 and λ_2 . Therefore T_1 and T_2 define a tile in $\llbracket \pi \rrbracket_{\text{Sep}}$ extending the path p . \square

C.6 Conditionals

Lemma C.9. *Suppose π is the following derivation tree:*

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P \wedge B\} C_1 \{Q\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P \wedge \neg B\} C_2 \{Q\} \end{array}}{\Gamma \vdash \{P\} C \{Q\}} \text{IF}$$

where $C := \text{if } B \text{ then } C_1 \text{ else } C_2$, and that the interpretation

$$\mathcal{L}_i : \llbracket \pi_i \rrbracket_{\text{Sep}} \longrightarrow \llbracket C_i \rrbracket_L$$

is a 2-fibration, for $i = 1, 2$. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{\text{Sep}} \longrightarrow \llbracket C \rrbracket_L$$

is also a 2-fibration.

Proof. Recall that the semantics of if statements is defined as

$$\begin{aligned} \llbracket \text{if } B \text{ then } C_1 \text{ else } C_2 \rrbracket_{Sep} &= \text{whentrue}[B](\llbracket \text{nop} \rrbracket_{Sep}); \llbracket C_1 \rrbracket_{Sep} \\ &\oplus \text{whenfalse}[B](\llbracket \text{nop} \rrbracket_{Sep}); \llbracket C_2 \rrbracket_{Sep} \\ &\oplus \text{whenabort}[B] \end{aligned}$$

All non trivial tiles are either in $\llbracket C_1 \rrbracket_L$ or in $\llbracket C_2 \rrbracket_L$. In either case, the hypothesis on \mathcal{L}_i tells us that this tile can be lifted to $\llbracket \pi_i \rrbracket_{Sep}$. \square

C.7 Disjunction

Lemma C.10. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\}C\{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\}C\{Q_2\} \end{array}}{\Gamma \vdash \{P_1 \vee P_2\}C\{Q_1 \vee Q_2\}} \text{DISJ}$$

and that the interpretation

$$\mathcal{L}_i : \llbracket \pi_i \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_L$$

is a 2-fibration, for $i = 1, 2$. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_L$$

is also a 2-fibration.

Proof. Similarly to the previous case, if T is a tile in $\llbracket C \rrbracket_L$, there we can either lift it to $\llbracket \pi_1 \rrbracket_{Sep}$ or to $\llbracket \pi_2 \rrbracket_{Sep}$, so in any case we can lift it to $\llbracket \pi \rrbracket_{Sep}$. \square

C.8 The frame rule

Lemma C.11. *Suppose that π is the derivation tree*

$$\frac{\begin{array}{c} \vdots \pi' \\ \Gamma \vdash \{P\}C\{Q\} \end{array}}{\Gamma \vdash \{P * R\}C\{Q * R\}} \text{FRAME}$$

and that the interpretation

$$\mathcal{L}' : \llbracket \pi' \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_L$$

is a 2-fibration. In that case, the asynchronous morphism

$$\mathcal{L} : \llbracket \pi \rrbracket_{Sep} \longrightarrow \llbracket C \rrbracket_L$$

is also a 2-fibration.

Proof. Consider a path in $\llbracket \pi \rrbracket_{Sep}$ of the following form (it is in one of the $\text{frame}[\sigma_R](\llbracket \pi' \rrbracket_{Sep})$)

$$\begin{array}{ccc} & (\sigma'_C * \sigma_R, \sigma_2, \sigma_F) & \\ m \nearrow & & \searrow m' \\ (\sigma_C * \sigma_R, \sigma_1, \sigma_F) & & (\sigma'_C * \sigma_R, \sigma_3, \sigma_F) \end{array} \quad (24)$$

and a tile in $\llbracket C \rrbracket_L$ whose upper path is the image of (24) under \mathcal{L}_π :

$$\begin{array}{ccc} a & \xrightarrow{m} & b \\ & \searrow m' & \nearrow m' \\ & & c \end{array} \quad \approx \quad \begin{array}{ccc} & b & \\ & \nearrow m' & \searrow m \\ a & \xrightarrow{m} & c \end{array} \quad (25)$$

By definition of $\text{frame}[\sigma_R]$, and according to the hypothesis on $\mathcal{L}_{\pi'}$, the path (24) in $\text{frame}[\sigma_R](\llbracket \pi' \rrbracket_{Sep})$ corresponds to a path in

$\llbracket \pi' \rrbracket_{Sep}$ that is the upper path a tile of the form:

$$\begin{array}{ccc} & (\sigma'_C, \sigma_2, \sigma_F * \sigma_R) & \\ m \nearrow & & \searrow m' \\ (\sigma_C, \sigma_1, \sigma_F * \sigma_R) & \sim & (\sigma_C, \sigma_3, \sigma_F * \sigma_R) \\ & \searrow m' & \nearrow m \\ & (\sigma_C^\dagger, \sigma'_3, \sigma_F * \sigma_R) & \end{array}$$

which means that there is a tile that extends (24) above (25):

$$\begin{array}{ccc} & (\sigma'_C * \sigma_R, \sigma_2, \sigma_F) & \\ m \nearrow & & \searrow m' \\ (\sigma_C * \sigma_R, \sigma_1, \sigma_F) & \sim & (\sigma_C * \sigma_R, \sigma_3, \sigma_F) \\ & \searrow m' & \nearrow m \\ & (\sigma_C^\dagger * \sigma_R, \sigma'_3, \sigma_F) & \end{array}$$

\square