

An Asynchronous Soundness Theorem for Concurrent Separation Logic

Paul-André Melliès

Léo Stefanescu

Institut de Recherche en Informatique Fondamentale (IRIF)
CNRS & Université Paris Diderot

CRECOGI+ELICA+GDRILL Plenary Meeting, 2018, Paris

Concurrent Separation Logic (CSL)

Concurrent Separation Logic is a logic for

concurrent programs with shared memory and locks.

Hoare triples

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

are proved using **derivation trees**

$$\Gamma \vdash \{P\} C \{Q\} \quad \vdots \pi$$

Soundness theorem

Soundness theorem A

Starting from a state $s \models P$,

no execution of C will crash

and if $(C, s) \Downarrow s'$, then $s' \models Q$.

Soundness theorem B

Starting from $s \models P$,

no execution of C will produce any data races.

Data races

A **data race** occurs when two instructions are executed **in parallel** and either:

- one **writes** to a location the other is **reading**,
- or both instructions **write** to the **same** location.

$$x := 89 \quad || \quad y := x + z$$

Soundness theorem

Soundness theorem A

Starting from a state $s \models P$,

no execution of C will crash

and if $(C, s) \Downarrow s'$, then $s' \models Q$.

- depends on **individual** instructions

Soundness theorem B

Starting from $s \models P$,

no execution of C will produce any data races.

- depends on **pairs** of instructions

Soundness theorem

Soundness theorem A

Starting from a state $s \models P$,

no execution of C will crash

and if $(C, s) \Downarrow s'$, then $s' \models Q$.

- depends on **individual** instructions
- will be seen as a **1-dimensional** property

Soundness theorem B

Starting from $s \models P$,

no execution of C will produce any data races.

- depends on **pairs** of instructions
- will be seen as a **2-dimensional** property

A Concurrent Shared Memory Language

resource r do

while $x > 0$ do

$$\left\{ \begin{array}{l} x := x - 1; \\ \text{with } r \text{ do } \quad \parallel \quad \text{with } r \text{ do} \\ y := y + 1 \quad \parallel \quad y := y + 1 \end{array} \right\}$$

Static Semantics of Imperative Languages

- For non-concurrent languages, a good abstraction for a program C is that of a **state transformer**:

$$\begin{aligned} \llbracket C \rrbracket : \quad & \mathbf{States} \longrightarrow \mathbf{States} \\ & \text{initial state} \longmapsto \text{final state} \end{aligned}$$

This is enough for sequential composition.

- For concurrent languages, **more information is needed**:
the final value of y in

$$x := 0; y := x$$

can be **any value**, depending on what the **environment** does:

$$x := 0; y := x \quad \parallel \quad x := 77$$

Stateful traces

The **traces** contain the **state** at each step of the execution

$$\llbracket C \rrbracket \ni s_1 \xrightarrow{\text{env}} s_2 \xrightarrow{m_1} s_3 \xrightarrow{\text{env}} s_4 \xrightarrow{m_2} s_5 \xrightarrow{\text{env}} s_6$$

where:

- the m_i are **elementary instructions**,

$$\text{eg. } s \xrightarrow{x := y + 4z} s', \quad s \xrightarrow{P(r)} s';$$

- the **env** transitions are played by the **Environment**.

Example:

$$\llbracket x := y \rrbracket = \left\{ s_1 \xrightarrow{\text{env}} s_2 \xrightarrow{x := y} s_2[x := s_1(y)] \xrightarrow{\text{env}} s_3 \mid \forall s_1, s_2, s_3 \right\}$$

Stateless traces

The **traces** are sequences of **events**

$$[[C]] \ni a_1 a_2 a_3 \dots a_{n-1} a_n$$

where:

- the events a_i are **independent** from the state,
- the **Environment** implies **non sequentially consistent traces**.

eg. $Wr(x, 89) \cdot Rd(x, 70)$

Our approach in this talk

Interpret programs as a **2-dimensional asynchronous graphs**.

Each program C will have **two related semantics**, corresponding to the **stateful** and the **stateless** trace semantics:

$$\boxed{\llbracket C \rrbracket_s \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_L}$$

Our approach in this talk

Interpret programs as a **2-dimensional asynchronous graphs**.

Each program C will have **two related semantics**, corresponding to the **stateful** and the **stateless** trace semantics:

$$\boxed{\llbracket C \rrbracket_S \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_L}$$

Interpret **derivation trees** of CSL as **asynchronous graphs**:

$$\boxed{\left[\begin{array}{c} \vdots \pi \\ \Gamma \vdash \{P\} C \{Q\} \end{array} \right]_{Sep} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_S}$$

Our approach in this talk

Interpret programs as a **2-dimensional asynchronous graphs**.

Each program C will have **two related semantics**, corresponding to the **stateful** and the **stateless** trace semantics:

$$\boxed{\llbracket C \rrbracket_S \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_L}$$

Interpret **derivation trees** of CSL as **asynchronous graphs**:

$$\boxed{\left[\begin{array}{c} \vdots \pi \\ \Gamma \vdash \{P\} C \{Q\} \end{array} \right]_{Sep} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_S}$$

Soundness = properties of these maps of **asynchronous graphs**.

Asynchronous graphs

Definition

An **asynchronous graph** is

- a **graph** $G = (V, E, s, t)$, with $s, t : E \rightarrow V$,
- a **relation** \diamond between paths $f, g : x \twoheadrightarrow y$ of length 2, with the same source and target nodes.

Two related paths form a **tile**.

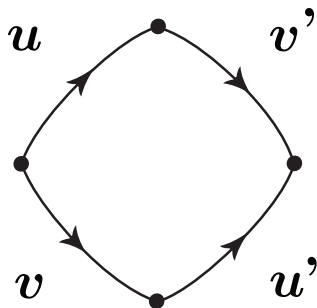
Asynchronous graphs

Definition

An **asynchronous graph** is

- a **graph** $G = (V, E, s, t)$, with $s, t : E \rightarrow V$,
- a **relation** \diamond between paths $f, g : x \rightarrow y$ of length 2, with the same source and target nodes.

Two related paths form a **tile**.



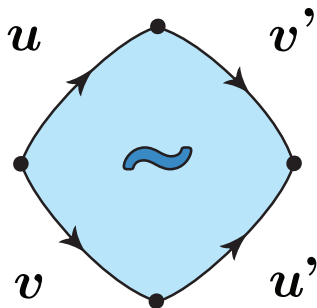
Asynchronous graphs

Definition

An **asynchronous graph** is

- a **graph** $G = (V, E, s, t)$, with $s, t : E \rightarrow V$,
- a **relation** \diamond between paths $f, g : x \rightarrow y$ of length 2, with the same source and target nodes.

Two related paths form a **tile**.



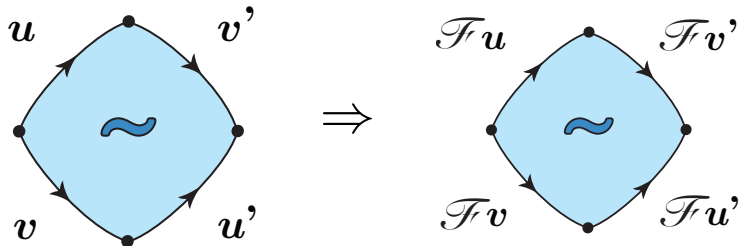
Maps between asynchronous graphs

Definition

A **map of asynchronous graphs**

$$\mathcal{F} : (G, \diamond) \longrightarrow (G', \diamond')$$

is a **graph homomorphism** $\mathcal{F} : G \rightarrow G'$ that maps **tiles into tiles**.



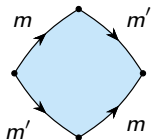
Can be seen as a **labeling** for G .

The stateful machine model

The **stateful machine model** \mathfrak{M}_S is defined as follows:

- its **nodes** are machine states $s \in (\mathbf{Loc} \rightarrow_{fin} \mathbf{Val}) \times \mathcal{P}(\mathit{Locks})$,
- there is an **edge** $s \xrightarrow{m} s'$ whenever $\llbracket m \rrbracket(s) = s'$,

- a square is a **tile** when



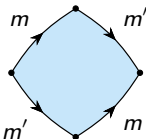
1. m and m' **do not synchronize**: $\mathit{lock}(m) \cap \mathit{lock}(m') = \emptyset$
2. and do **not** induce a **data race**:

$$\begin{aligned}(\mathit{rd}(m) \cup \mathit{wr}(m)) \cap \mathit{wr}(m') &= \emptyset \\(\mathit{rd}(m') \cup \mathit{wr}(m')) \cap \mathit{wr}(m) &= \emptyset\end{aligned}$$

The stateless machine model

The **stateless machine model** \mathfrak{A}_L is defined as follows:

- its **nodes** are lock states $L \subseteq Locks$,
- there is an **edge** $L \xrightarrow{m} L'$ whenever $\llbracket m \rrbracket(L) = L'$,

- a square  is a **tile** when

m and m' **do not synchronize**: $\text{lock}(m) \cap \text{lock}(m') = \emptyset$.

Machine Models

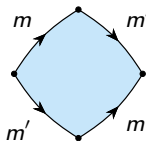
 \mathfrak{s}

Nodes machine states $\mathfrak{s} = (\mu, L)$ with

- memory state
 $\mu : \mathbf{Loc} \rightarrow_{fin} \mathbf{Val}$
- lock state $L \subseteq Locks$

Edges $\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$ for each $\llbracket m \rrbracket(\mathfrak{s}) = \mathfrak{s}'$

Tiles data race-freedom



is a tile when:

$$\text{lock}(m) \cap \text{lock}(m') = \emptyset$$

$$(\text{rd}(m) \cup \text{wr}(m)) \cap \text{wr}(m') = \emptyset$$

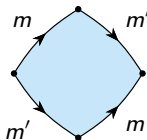
$$(\text{rd}(m') \cup \text{wr}(m')) \cap \text{wr}(m) = \emptyset$$

 \mathfrak{L}

lock states $L \subseteq Locks$

$L \xrightarrow{m} L'$ for $\llbracket m \rrbracket(L) = L'$

parallelism



is a tile when:

$$\text{lock}(m) \cap \text{lock}(m') = \emptyset$$

Asynchronous Transition Systems (ATS)

Definition (ATS)

An **ATS** over a **machine model** \mathcal{M} is an **asynchronous graph** G together with a **map of asynchronous graphs**

$$\lambda : G \longrightarrow \mathcal{M}$$

with a partition on G 's edges, for the **Code** and the **Environment**

- **nodes** are labeled by **states**,
- **edges** are labeled by **instructions** m ,
- **labels** are consistent with the semantics of **instructions**:

$$x \xrightarrow{m} y \quad \Longrightarrow \quad \llbracket m \rrbracket(\lambda(x)) = \lambda(y)$$

Semantics of the Code

The semantics of a **program** C is a pair of **related ATs**.

$$\llbracket C \rrbracket_S : G_S(C) \longrightarrow \mathcal{A}_S$$

$$\llbracket C \rrbracket_L : G_L(C) \longrightarrow \mathcal{A}_L$$

Definition

The semantics is defined by **induction of the structure of the Code**:

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

where $\llbracket \cdot \rrbracket$ is either stateful $\llbracket \cdot \rrbracket_S$ or stateless $\llbracket \cdot \rrbracket_L$

Semantics of the Code

The semantics of a **program** C is a pair of **related ATs**.

$$\begin{array}{ccc} \llbracket C \rrbracket_S : G_S(C) & \longrightarrow & \mathfrak{A}_S \\ \mathcal{L} \downarrow & & \downarrow (\mu, L) \mapsto L \\ \llbracket C \rrbracket_L : G_L(C) & \longrightarrow & \mathfrak{A}_L \end{array}$$

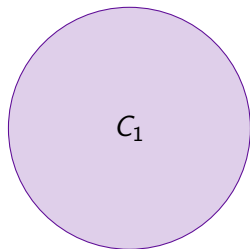
Definition

The semantics is defined by **induction of the structure of the Code**:

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

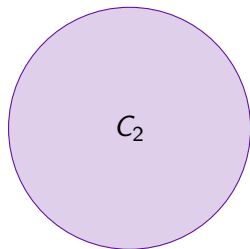
where $\llbracket \cdot \rrbracket$ is either stateful $\llbracket \cdot \rrbracket_S$ or stateless $\llbracket \cdot \rrbracket_L$

The parallel product



C_1

C_1

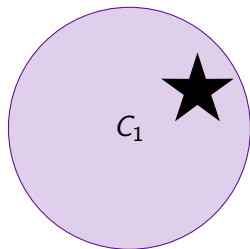


C_2

C_2

$C_1 \parallel C_2$

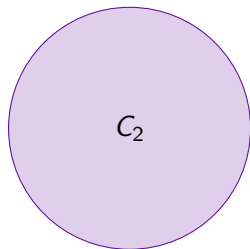
The parallel product



C_1

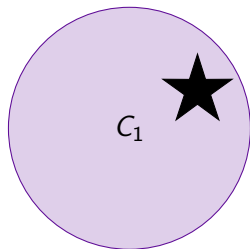
$$s \xrightarrow{m} s'$$

$C_1 \parallel C_2$



C_2

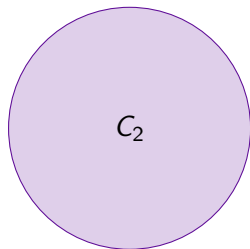
The parallel product



C_1

$$s \xrightarrow{m} s'$$

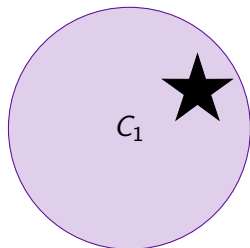
$C_1 \parallel C_2$



C_2

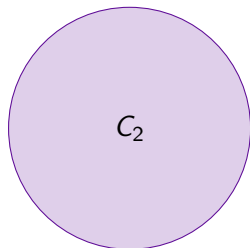
$$s \xrightarrow{m} s'$$

The parallel product



C_1

$$s \xrightarrow{m} s'$$



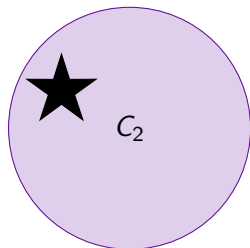
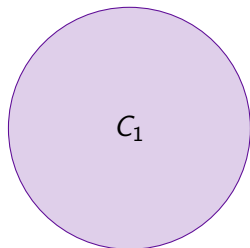
C_2

$$s \xrightarrow{m} s'$$

$C_1 \parallel C_2$

$$s \xrightarrow{m} s'$$

The parallel product



C_1

$$s \xrightarrow{m} s'$$

$$s \xrightarrow{m} s'$$

$C_1 \parallel C_2$

$$s \xrightarrow{m} s'$$

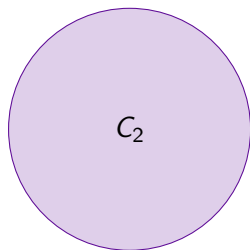
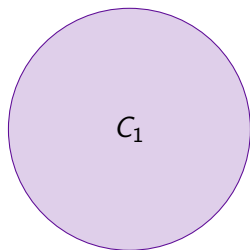
$$s \xrightarrow{m} s'$$

C_2

$$s \xrightarrow{m} s'$$

$$s \xrightarrow{m} s'$$

The parallel product



C_1

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$C_1 \parallel C_2$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

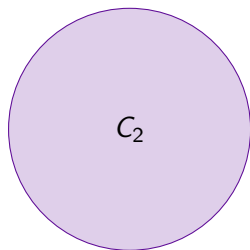
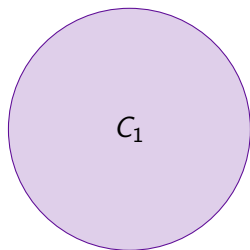
C_2

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

The parallel product



C_1

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$C_1 \parallel C_2$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

C_2

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}'$$

The parallel product

$$G_1 \xrightarrow{\lambda_1} \mathfrak{A}, G_2 \xrightarrow{\lambda_2} \mathfrak{A} \quad \mapsto \quad G_1 \parallel G_2 \xrightarrow{\lambda_{1\parallel 2}} \mathfrak{A}$$

Nodes of $G_1 \parallel G_2$: $x_1|x_2 \in G_1 \times G_2$, such that $\lambda_1(x_1) = \lambda_2(x_2)$

$$\lambda_{1\parallel 2}(x_1|x_2) := \lambda_1(x_1) = \lambda_2(x_2)$$

Three types of transitions, e.g. $x_1|x_2 \xrightarrow{m|m} x'_1|x'_2$ is a pair:

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

This transition is mapped by $\lambda_{1\parallel 2}$ to $\lambda_{1\parallel 2}(x_1|x_2) \xrightarrow{m} \lambda_{1\parallel 2}(x'_1|x'_2)$

Tiles of $G_1 \parallel G_2$ are tiles of G_1 , of G_2 , or made of instructions of the form

$$m_1|m_1 \quad \text{and} \quad m_2|m_2$$

Data races in the stateless semantics $\llbracket C \rrbracket_L$

A **data race** needs two **unsynchronized** instructions in some trace



Data races in the stateless semantics $\llbracket C \rrbracket_L$

A **data race** needs two **unsynchronized** instructions in some trace



thus coming with **another schedule** of m_1 and m_2 :



Data races in the stateless semantics $\llbracket C \rrbracket_L$

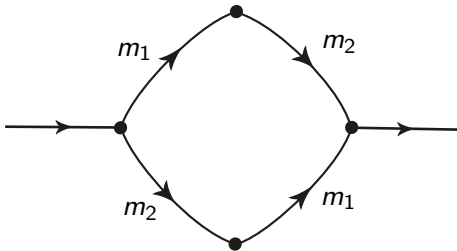
A **data race** needs two **unsynchronized** instructions in some trace



thus coming with **another schedule** of m_1 and m_2 :



In our approach, we turn it into a square



Data races in the stateless semantics $\llbracket C \rrbracket_L$

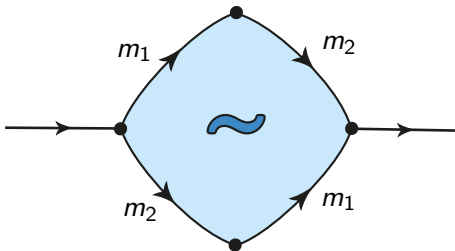
A **data race** needs two **unsynchronized** instructions in some trace



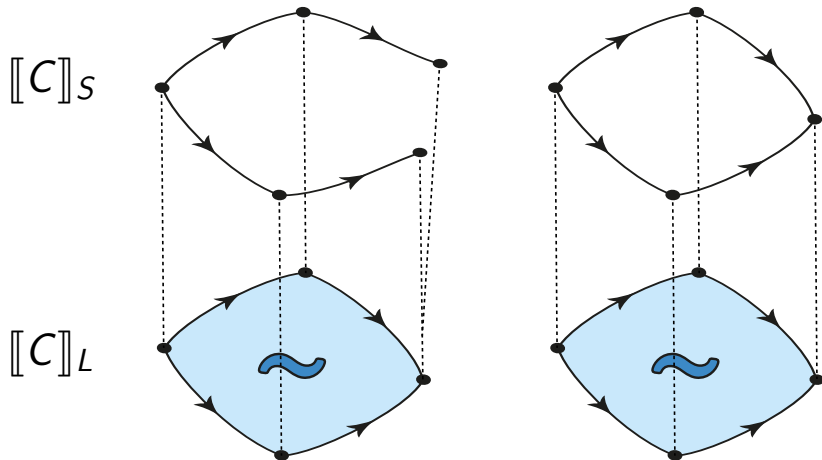
thus coming with **another schedule** of m_1 and m_2 :



In our approach, we turn it into a **tile**



A topological account of data races



Data race = stateful hole above a **stateless tile**

Separation Logic (sequential)

Hoare logic with extended predicates:

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

Meaning of triple: $\forall \sigma, \sigma', (\sigma \models P \wedge C, \sigma \Downarrow \sigma') \Rightarrow \sigma' \models Q$

Predicates on the memory

$P, Q, J ::= \top \mid \perp \mid P \vee Q \mid P \wedge Q \mid \forall v. P \mid \exists v. P \mid P * Q \mid \mathbf{emp} \mid v \mapsto w$

$$\sigma \models P \wedge Q \iff \sigma \models P \text{ and } \sigma \models Q$$

$$\sigma \models P * Q \iff \exists \sigma_1 \sigma_2, \sigma = \sigma_1 \uplus \sigma_2 \text{ and } \sigma_1 \models P \text{ and } \sigma_2 \models Q$$

$$\sigma \models \mathbf{emp} \iff \sigma = \emptyset$$

$$\sigma \models v \mapsto w \iff \sigma = [v \mapsto w]$$

The Frame Rule

$$\frac{\vdash \{P\} C \{Q\}}{\vdash \{P * R\} C \{Q * R\}}$$

C depends only on the **resource** P

The Frame Rule

$$\frac{\vdash \{P\} C \{Q\}}{\vdash \{P * R\} C \{Q * R\}}$$

C depends only on the **resource** P

$$\frac{\vdash \{P_1\} C_1 \{Q_1\} \quad \vdash \{P_2\} C_2 \{Q_2\}}{\vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}}$$

The Frame Rule

$$\frac{\vdash \{P\} C \{Q\}}{\vdash \{P * R\} C \{Q * R\}}$$

C depends only on the **resource** P

$$\frac{\vdash \{P_1\} C_1 \{Q_1\} \quad \vdash \{P_2\} C_2 \{Q_2\}}{\vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}}$$

But C_1 and C_2 **cannot communicate!**

Concurrent Separation Logic

Associate an **invariant** J to each **lock** r in a context:

$$r_1:J_1, \dots, r_n:J_n \vdash \{P\} C \{Q\}$$

Inference rules:

$$\frac{\Gamma \vdash \{P * J\} C \{Q * J\}}{\Gamma, r : J \vdash \{P\} \text{ with } r \text{ do } C \{Q\}}$$

$$\frac{\Gamma, r : J \vdash \{P\} C \{Q\}}{\Gamma \vdash \{P * J\} \text{ resource } r \text{ do } C \{Q * J\}}$$

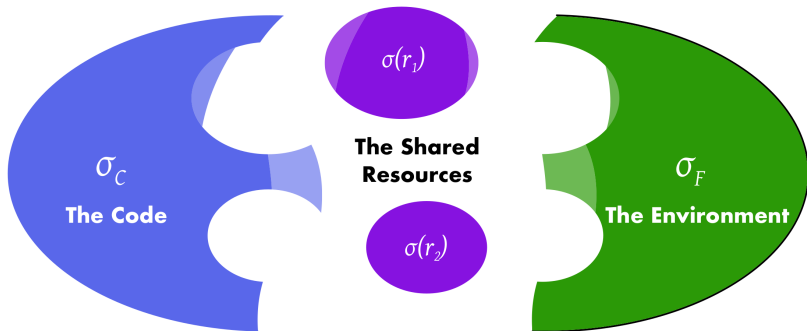
Separated States

A **separated state** is a triple

$$(\sigma_C, \sigma, \sigma_F)$$

\cap

$$\mathbf{States} \times (\mathit{Locks} \rightarrow \mathbf{States} + \{C, F\}) \times \mathbf{States}$$



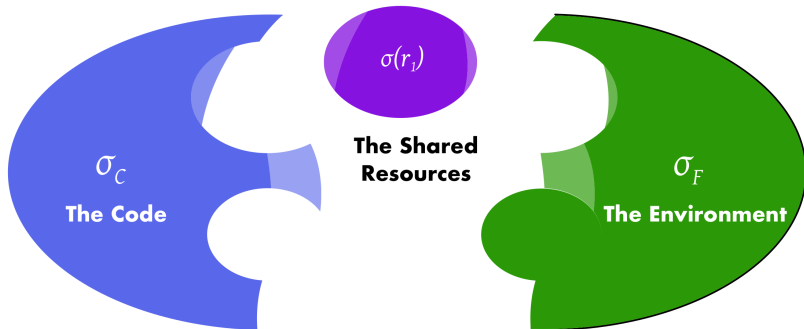
Separated States

A **separated state** is a triple

$$(\sigma_C, \sigma, \sigma_F)$$

$$\cap$$

$$\mathbf{States} \times (\mathbf{Locks} \rightarrow \mathbf{States} + \{C, F\}) \times \mathbf{States}$$



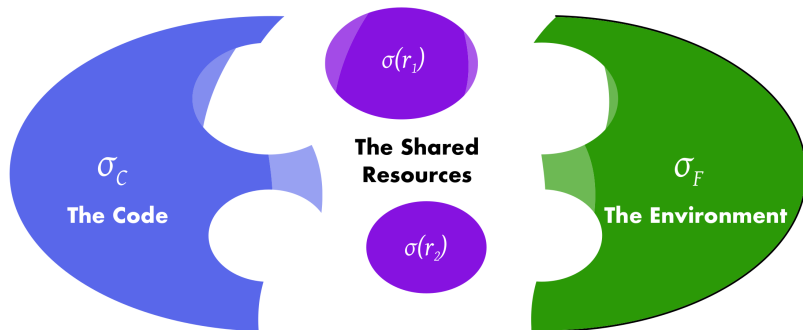
Separated States

A **separated state** is a triple

$$(\sigma_C, \boldsymbol{\sigma}, \sigma_F)$$

$$\cap$$

$$\mathbf{States} \times (\mathit{Locks} \rightarrow \mathbf{States} + \{C, F\}) \times \mathbf{States}$$



The separated machine model

The **separated machine model** \mathfrak{A}_{Sep} is defined in a similar fashion.

Its nodes are **separated states** $(\sigma_C, \sigma, \sigma_F)$,

and there are two kinds of **transitions**

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

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

There is a **map of asynchronous graphs**

$$\mathfrak{A}_{Sep} \longrightarrow \mathfrak{A}_S$$

Semantics of the derivation trees

Every **CSL derivation tree** $\Gamma \vdash \{P\} C \{Q\} \stackrel{\pi}{\vdots}$ is interpreted as an **ATS**:

$$\llbracket \pi \rrbracket_{Sep} : G_{Sep}(\pi) \longrightarrow \mathfrak{A}_{Sep}$$

$$\begin{array}{ccccc} \llbracket C \rrbracket_S : & G_S(C) & \longrightarrow & \mathfrak{A}_S & \\ \mathcal{L} \downarrow & \mathcal{L}_G \downarrow & & \downarrow (\mu, L) \mapsto L & \\ \llbracket C \rrbracket_L : & G_L(C) & \longrightarrow & \mathfrak{A}_L & \end{array}$$

Semantics of the derivation trees

Every **CSL derivation tree** $\Gamma \vdash \{P\} C \{Q\} \stackrel{\vdots \pi}{}$ is interpreted as an **ATS**:

$$\begin{array}{ccccc}
 \llbracket \pi \rrbracket_{Sep} : & G_{Sep}(\pi) & \longrightarrow & \mathfrak{A}_{Sep} & \\
 \downarrow \mathcal{S} & \downarrow \mathcal{S}_G & & \downarrow & \\
 \llbracket C \rrbracket_S : & G_S(C) & \longrightarrow & \mathfrak{A}_S & \\
 \downarrow \mathcal{L} & \downarrow \mathcal{L}_G & & \downarrow (\mu, L) \mapsto L & \\
 \llbracket C \rrbracket_L : & G_L(C) & \longrightarrow & \mathfrak{A}_L &
 \end{array}$$

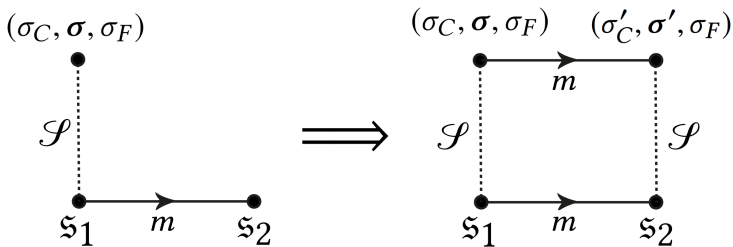
An asynchronous soundness theorem

Theorem 1

The map of asynchronous graphs

$$\left[\left[\Gamma \vdash \{P\} C \{Q\} \right]_{Sep} \right] \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_S$$

is a **1-dimensional fibration on Code transitions**.



An asynchronous soundness theorem

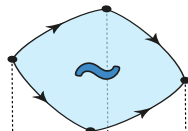
Theorem 2

The map of asynchronous graphs

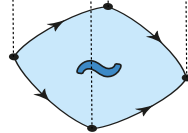
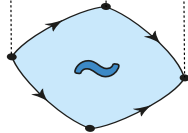
$$\left[\left[\Gamma \vdash \{P\} C \{Q\} \right]_{Sep} \right]_{\pi} \xrightarrow{\mathcal{I}} \llbracket C \rrbracket_S \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_L$$

is a **2-dimensional fibration**.

$\llbracket \pi \rrbracket_{Sep}$



$\llbracket C \rrbracket_L$



Conclusion & Future work

A topological account of **data races**

A **truly concurrent** semantics for Concurrent Separation Logic derivation trees gives a **strategy** for **memory management**

Extension to more sophisticated logics and programming languages
eg. higher order, more general concurrency primitives

Understand our semantics in a general categorical framework

Conclusion & Future work

A topological account of **data races**

A **truly concurrent** semantics for Concurrent Separation Logic derivation trees gives a **strategy** for **memory management**

Extension to more sophisticated logics and programming languages
eg. higher order, more general concurrency primitives

Understand our semantics in a general categorical framework

Thank you!