

Correspondence between operational and declarative concurrency semantics

Ori Lahav Viktor Vafeiadis

29 August 2017

Two alternative definitions of SC

Definition (Operational SC)

An outcome O is allowed for a program P under SC if there exists M such that $P, S_0, M_0 \Rightarrow^* \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, M$.

Definition (Declarative SC)

An outcome O is allowed for a program P under SC if there exists an SC-consistent execution graph of P with outcome O .

How do we show that
the two definitions
are **equivalent**?

State:

$$\langle P, S, G, \mathbf{sc} \rangle \in \text{Program} \times (\text{Tid} \rightarrow \text{Store}) \times \text{ExecutionGraph} \times \mathcal{P}(\text{Event} \times \text{Event})$$

- ▶ Initial stores: $S_0 \triangleq \lambda i. s_0$
- ▶ Initial execution: G_0 consisting only of the initialization events
- ▶ Initial **sc**-relation: \mathbf{sc}_0 is an arbitrary total order on $G_0.E$

$$\begin{array}{c}
 \text{NON-SILENT} \\
 P, S \xrightarrow{\text{tid}(a):\text{lab}(a)} P', S' \\
 G' \in \text{Add}(G, a) \\
 \mathbf{sc}' = \mathbf{sc} \cup (G.E \times \{a\}) \\
 G' \text{ is SC-consistent wrt } \mathbf{sc}' \\
 \hline
 P, S, G, \mathbf{sc} \Rightarrow P', S', G', \mathbf{sc}'
 \end{array}$$

$$\begin{array}{c}
 \text{SILENT} \\
 P, S \xrightarrow{i:\varepsilon} P', S' \\
 \hline
 P, S, G, \mathbf{sc} \Rightarrow P', S', G, \mathbf{sc}
 \end{array}$$

where $\text{Add}(G, a)$ is the set of all complete graphs G' satisfying:

- ▶ $G'.E = G.E \uplus \{a\}$
- ▶ $G'.\text{po} = G.\text{po} \cup ((E_0 \cup G.E^{\text{tid}(a)}) \times \{a\})$
- ▶ $G.\text{rf} \subseteq G'.\text{rf}$

Definition (Operational-declarative SC)

An outcome O is allowed for a program P under SC if there exist G, \mathbf{sc} such that $P, S_0, G_0, \mathbf{sc}_0 \Rightarrow^* \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, G, \mathbf{sc}$.

Establish correspondence between operational SC and declarative SC in two steps:

1. operational SC = intermediate SC
2. declarative SC = intermediate SC

We will use forward weak *simulation*. Consider two labeled state transition systems $M_1 = \langle Q_1, q_1^0, \rightarrow_1 \rangle$ and $M_2 = \langle Q_2, q_2^0, \rightarrow_2 \rangle$.

- ▶ $\mathcal{R} \subseteq Q_1 \times Q_2$ is a *simulation relation* from M_1 to M_2 if:
 - ▶ $q_1^0 \mathcal{R} q_2^0$, and
 - ▶ whenever $q_1 \mathcal{R} q_2$ and $q_1 \rightarrow_1 q'_1$, then there exists some $q'_2 \in Q_2$ such that $q_2 \rightarrow_2^* q'_2$ and $q'_1 \mathcal{R} q'_2$.
- ▶ $\mathcal{R} \subseteq Q_1 \times Q_2$ is called a *bisimulation relation* if it is a simulation relation from M_1 to M_2 and \mathcal{R}^{-1} is a simulation relation from M_2 to M_1 .

Lemma

If a simulation relation exists then for every state $q_1 \in Q_1$ that is reachable from q_1^0 in M_1 , there exists some $q_2 \in Q_2$ that is reachable from q_2^0 in M_2 and satisfies $q_1 \mathcal{R} q_2$.

Our bisimulation relation:

$\langle P, S, M \rangle \sim \langle P', S', G, \text{sc} \rangle$ if the following hold:

- ▶ $P = P'$
 - ▶ $S = S'$
 - ▶ $M = \lambda x. \text{val}_w(\max_{\text{sc}} G.W_x)$
 - ▶ G is complete and SC-consistent wrt sc .
-
- ▶ Show that \sim is a bisimulation relation.
 - ▶ Deduce that operational SC and intermediate SC have the same outcomes for any given program.

Declarative SC = intermediate SC

Two directions:

- ⊆ Every outcome allowed for P according to declarative SC is allowed according to intermediate SC
- ⊇ Every outcome allowed for P according to intermediate SC is allowed according to declarative SC

Reminders:

Definition

G is an execution graph of a program P with an outcome O if G^i is an execution of $P(i)$ with final store $O(i)$ for every $i \in \text{Tid}$.

Definition (Declarative SC)

An outcome O is allowed for a program P under SC if there exists an SC-consistent execution graph of P with outcome O .

Lemma (Execution generation)

Let G be an execution of a program P_0 with outcome O . Let a_1, \dots, a_n be an enumeration of $G.E \setminus E_0$ that respects $G.po$. Then, there exist $\langle P_1, S_1 \rangle, \dots, \langle P_n, S_n \rangle$ such that:

- ▶ $P_n = \mathbf{skip} \parallel \dots \parallel \mathbf{skip}$ and $S_n = O$
- ▶ For every $1 \leq j \leq n$, we have:

$$P_{j-1}, S_{j-1} \xrightarrow{\text{tid}(a_j):\varepsilon}^* \xrightarrow{\text{tid}(a_j):\text{lab}(a_j)} \xrightarrow{\text{tid}(a_j):\varepsilon}^* P_j, S_j$$

Declarative SC \subseteq intermediate SC

- ▶ Let G be an SC-consistent execution graph of P_0 with outcome O .
- ▶ Let \mathbf{sc} be a total order on $G.E$ such that G is SC-consistent wrt \mathbf{sc} .
- ▶ We show that $P, S_0, G_0, \mathbf{sc}_0 \Rightarrow^* \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, G, \mathbf{sc}$.
- ▶ Let a_1, \dots, a_n be an enumeration of $G.E \setminus E_0$ following \mathbf{sc} .
- ▶ Since $G.po \subseteq \mathbf{sc}$, by the previous lemma, there exist $\langle P_1, S_1 \rangle, \dots, \langle P_n, S_n \rangle$ such that:
 - ▶ $P_n = \mathbf{skip} \parallel \dots \parallel \mathbf{skip}$ and $S_n = O$
 - ▶ For every $1 \leq j \leq n$, we have:

$$P_{j-1}, S_{j-1} \xrightarrow{\text{tid}(a_j):\varepsilon}^* \xrightarrow{\text{tid}(a_j):\text{lab}(a_j)} \xrightarrow{\text{tid}(a_j):\varepsilon}^* P_j, S_j$$

- ▶ For every $0 \leq j \leq n$, let
 - ▶ G_j - the restriction of G to $E_0 \cup \{a_1, \dots, a_j\}$
 - ▶ \mathbf{sc}_j - the restriction of \mathbf{sc} to $E_0 \cup \{a_1, \dots, a_j\}$
- ▶ Then, for every $1 \leq j \leq n$, we have:

$$P_{j-1}, S_{j-1}, G_{j-1}, \mathbf{sc}_{j-1} \Rightarrow^* P_j, S_j, G_j, \mathbf{sc}_j$$

- ▶ Suppose that $P_0, S_0, G_0, \mathbf{sc}_0 \Rightarrow^* \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, G, \mathbf{sc}$.
- ▶ By definition, G is SC-consistent. It remains to show that each G^i is an execution of $P_0(i)$ with final store $O(i)$.
- ▶ We know: $P_0, S_0, G_0, \mathbf{sc}_0 \Rightarrow P_1, S_1, G_1, \mathbf{sc}_1 \Rightarrow \dots \Rightarrow P_n, S_n, G_n, \mathbf{sc}_n$ where $P_n, S_n, G_n, \mathbf{sc}_n = \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, G, \mathbf{sc}$.
- ▶ The sequence above induces the following sequence of transitions:

$$P_0, S_0 \xrightarrow{i_1:l_1} P_1, S_1 \xrightarrow{i_2:l_2} P_2, S_2 \xrightarrow{i_3:l_3} \dots \xrightarrow{i_n:l_n} P_n, S_n$$

- ▶ In turn, by filtering only the transitions of thread i we obtain:

$$P_0(i), S_0(i) \xrightarrow{l_{k_1}} P_{k_1}(i), S_{k_1}(i) \xrightarrow{l_{k_2}} \dots \xrightarrow{l_{k_{n_i}}} P_{k_{n_i}}(i), S_{k_{n_i}}(i) = \mathbf{skip}, O(i)$$

- ▶ It follows that $P_0(i), s_0, G_\emptyset \Rightarrow^* \mathbf{skip}, O(i), G^i$, and so G^i is an execution of $P_0(i)$ with final store $O(i)$.

Operational semantics for COH

Recall the following litmus tests:

Store buffering

$$\begin{array}{l} x = y = 0 \\ x := 1 \quad \parallel \quad y := 1 \\ a := y \text{ // } 0 \quad \parallel \quad b := x \text{ // } 0 \end{array}$$

Coherence test

$$\begin{array}{l} x = 0 \\ x := 1 \quad \parallel \quad x := 2 \\ a := x \text{ // } 2 \quad \parallel \quad b := x \text{ // } 1 \end{array}$$


Two approaches:

- ▶ Out-of-order execution with SC memory.
- ▶ In-order execution with non-standard memory:
 - ▶ Allow threads to observe different subsets of writes.
 - ▶ Use timestamps to order writes to the same location.

Store buffering

	$x = y = 0$	
$x := 1;$		$y := 1;$
$a := y; //0$		$b := x; //0$

Operational semantics for coherence



Memory

$\langle x : 0 @ 0 \rangle$
 $\langle y : 0 @ 0 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      |      y := 1;
▶ a := y; //0 |      ▶ b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      y := 1;
▶ a := y; //0  ▶ b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	1
	1

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      y := 1;
a := y; //0  ▶ b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      y := 1;
a := y; //0  b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Operational semantics for coherence

Store buffering

```
      x = y = 0
x := 1;  ||  y := 1;
a := y; //0 || b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

Coherence test

```
      x = 0
x := 1;  ||  x := 2;
a := x; //2 || b := x; //1
```

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      |      y := 1;
a := y; //0  |      b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	1
	1

Coherence test

```
x = 0
▶ x := 1;      |      ▶ x := 2;
a := x; //2    |      b := x; //1
```

Memory

```
⟨x : 0@0⟩
```

T_1 's view

x
0

T_2 's view

x
0

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      ||      y := 1;
a := y; //0  ||      b := x; //0
```

Memory

```
⟨x: 0@0⟩
⟨y: 0@0⟩
⟨x: 1@1⟩
⟨y: 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

Coherence test

```
x = 0
x := 1;      ||      x := 2;
a := x; //2  ||      b := x; //1
```

Memory

```
⟨x: 0@0⟩
⟨x: 1@1⟩
```

T_1 's view

x
0
1

T_2 's view

x
0

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      ||      y := 1;
a := y; //0  ||      b := x; //0
```

Memory

```
⟨x: 0@0⟩
⟨y: 0@0⟩
⟨x: 1@1⟩
⟨y: 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

Coherence test

```
x = 0
x := 1;      ||      x := 2;
a := x; //2  ||      b := x; //1
```

Memory

```
⟨x: 0@0⟩
⟨x: 1@1⟩
⟨x: 2@2⟩
```

T_1 's view

x
0
1

T_2 's view

x
0
2

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      ||      y := 1;
a := y; //0  ||      b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

Coherence test

```
x = 0
x := 1;      ||      x := 2;
a := x; //2  ||      b := x; //1
```

Memory

```
⟨x : 0@0⟩
⟨x : 1@1⟩
⟨x : 2@2⟩
```

T_1 's view

x
0
1
2

T_2 's view

x
0
2

Operational semantics for coherence

Store buffering

```
x = y = 0
x := 1;      ||      y := 1;
a := y; //0  ||      b := x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0
	1

Coherence test

```
x = 0
x := 1;      ||      x := 2;
a := x; //2  ||      b := x; //1
```

Memory

```
⟨x : 0@0⟩
⟨x : 1@1⟩
⟨x : 2@2⟩
```

T_1 's view

x
0
1
2

T_2 's view

x
0
2

2+2W

```
      x = y = 0
x := 1;  ||  y := 1;
y := 2;  ||  x := 2;
a := y //1 || b := x //1
```

Supporting write-write reordering

2+2W

$x = y = 0$		
▶ $x := 1;$ $y := 2;$ $a := y$ //1		▶ $y := 1;$ $x := 2;$ $b := x$ //1

Memory

$\langle x : 0@0 \rangle$
 $\langle y : 0@0 \rangle$

T_1 's view

x	y

0	0

T_2 's view

x	y

0	0

Supporting write-write reordering

2+2W

```
x = y = 0
x := 1;      ||      ▶ y := 1;
▶ y := 2;    ||      x := 2;
a := y //1  ||      b := x //1
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
```

T_1 's view

x	y
0	0
1	

T_2 's view

x	y
0	0

Supporting write-write reordering

2+2W

$x = y = 0$	
$x := 1;$	$\blacktriangleright y := 1;$
$y := 2;$	$x := 2;$
$\blacktriangleright a := y \ //1$	$b := x \ //1$

Memory

$\langle x : 0@0 \rangle$
 $\langle y : 0@0 \rangle$
 $\langle x : 1@1 \rangle$
 $\langle y : 2@1 \rangle$

T_1 's view

x	y
0	0
1	1

T_2 's view

x	y
0	0

Supporting write-write reordering

2+2W

```
x = y = 0
x := 1;      ||      y := 1;
y := 2;      ||      ▶ x := 2;
▶ a := y //1 ||      b := x //1
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 2@1⟩
⟨y : 1@2⟩
```

T_1 's view

x	y
0	0
1	1

T_2 's view

x	y
0	0
	2

Supporting write-write reordering

2+2W

```
x = y = 0
x := 1;      ||      y := 1;
y := 2;      ||      x := 2;
▶ a := y //1  ||  ▶ b := x //1
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@1⟩
⟨y : 2@1⟩
⟨y : 1@2⟩
⟨x : 2@0.5⟩
```

T₁'s view

x	y
0	0
1	1

T₂'s view

x	y
0	0
0.5	2

- Writes choose timestamp *greater than the thread's view*, not necessarily the globally greatest one.

Load buffering (LB)

$$\begin{array}{c} x = y = 0 \\ a := x \ //1 \quad \parallel \quad b := y \ //1 \\ y := 1 \quad \parallel \quad x := 1 \end{array}$$

- ▶ COH allows this outcome.
- ▶ But, the suggested operational semantics disallows it!
- ▶ We will see later an approach to fix this mismatch (using out-of-order execution).
- ▶ For now, we will strengthen the declarative semantics.

Definition (Strong coherence)

An execution G is *strongly coherent* if the following hold:

- ▶ G is complete.
- ▶ G is coherent wrt some modification order mo for G .
- ▶ $G.po \cup G.rf$ is acyclic.

A note about the implementability of StrongCOH

Some hardware implementations (e.g., ARM) allow $po \cup rf$ cycles involving only plain loads and stores. To implement StrongCOH on those architectures, a syntactic dependency or a fence has to be introduced between every load and subsequent store.

- ▶ Time $\triangleq \{t \in \mathbb{Q} \mid t \geq 0\}$ is the set of timestamps.
- ▶ A message is a triple $\langle x : v @ t \rangle$ where $x \in \text{Loc}$, $v \in \text{Val}$, and $t \in \text{Time}$.
- ▶ A memory is a finite set of messages.
- ▶ A view is a function $\text{view} : \text{Loc} \rightarrow \text{Time}$.
- ▶ A thread view function is a function $V : \text{Tid} \rightarrow (\text{Loc} \rightarrow \text{Time})$ assigning a view to every thread.

The state consists of

- ▶ a program P
- ▶ a store function S
- ▶ a memory M
- ▶ a thread view function V

Initial state $\langle P, S_0, M_0, V_0 \rangle$ where

- ▶ $S_0 = \lambda i. s_0 = \lambda i. \lambda r. 0$
- ▶ $M_0 = \{\langle x : 0 @ 0 \rangle \mid x \in \text{Loc}\}$
- ▶ $V = \lambda i. \text{view}_0 = \lambda i. \lambda x. 0.$

SILENT-THREAD

$$\frac{P, S \xrightarrow{i:\varepsilon} P', S'}{P, S, M, V \Rightarrow P', S', M, V}$$

READ

$$\frac{\begin{array}{l} P, S \xrightarrow{i:l} P', S' \quad l = R(x, v) \\ \langle x : v@t \rangle \in M \quad V(i)(x) \leq t \\ \text{view}' = V(i)[x \mapsto t] \end{array}}{P, S, M, V \Rightarrow P', S', M, V[i \mapsto \text{view}']}$$

WRITE

$$\frac{\begin{array}{l} P, S \xrightarrow{i:l} P', S' \quad l = W(x, v) \\ V(i)(x) < t \quad \forall v'. \langle x : v'@t \rangle \notin M \\ M' = M \cup \{ \langle x : v@t \rangle \} \quad \text{view}' = V(i)[x \mapsto t] \end{array}}{P, S, M, V \Rightarrow P', S', M', V[i \mapsto \text{view}']}$$

Definition (Operational StrongCOH)

An outcome O is allowed for a program P under StrongCOH if there exist M, V such that $P, S_0, M_0, V_0 \Rightarrow^* \text{skip} \parallel \dots \parallel \text{skip}, O, M, V$.

As for SC, we will introduce an “intermediate” semantics for StrongCOH.

Establish correspondence between operational StrongCOH and declarative StrongCOH in two steps:

1. operational StrongCOH = intermediate StrongCOH
2. declarative StrongCOH = intermediate StrongCOH

Operational version of StrongCOH declarative semantics

State $\langle P, S, G, \text{mo} \rangle$ **where** $P \in \text{Program}$, $S \in (\text{Tid} \rightarrow \text{Store})$,
 $G \in \text{ExecutionGraph}$, $\text{mo} \subseteq G.E \times G.E$.

- ▶ Initial stores: $S_0 \triangleq \lambda i. s_0$
- ▶ Initial execution: G_0 consisting only of the initialization events
- ▶ Initial modification order: $\text{mo}_0 = \emptyset$

SILENT

$$\frac{P, S \xrightarrow{i:\varepsilon} P', S'}{P, S, G, \text{mo} \Rightarrow P', S', G, \text{mo}}$$

NON-SILENT

$$\frac{\begin{array}{l} P, S \xrightarrow{i:l} P', S' \quad l \neq \varepsilon \\ G' \in \text{Add}(G, \langle n, i, l \rangle, i) \quad \text{mo} \subseteq \text{mo}' \\ \text{mo}' \text{ is a modification order for } G' \\ G' \text{ is COH-consistent wrt } \text{mo}' \end{array}}{P, S, G, \text{mo} \Rightarrow P', S', G', \text{mo}'}$$

Definition (Operational-declarative StrongCOH)

An outcome O is allowed for a program P under StrongCOH if there exist G, mo such that $P, S_0, G_0, \text{mo}_0 \Rightarrow^* \text{skip} \parallel \dots \parallel \text{skip}, O, G, \text{mo}$.

Our bisimulation relation:

$P, S, M, V \sim P', S', G, \text{mo}$ if the following hold:

- ▶ $P = P'$
- ▶ $S = S'$
- ▶ there exists a function $ts : G.W \rightarrow \text{Time}$ such that:
 - ▶ $ts(w_1) < ts(w_2)$ whenever $\langle w_1, w_2 \rangle \in \text{mo}$
 - ▶ $M = \{ \langle \text{loc}(w) : \text{val}_w(w) @ ts(w) \rangle \mid w \in G.W \}$
 - ▶ $V = \lambda i x. \max \{ ts(w) \mid w \in \text{dom}([G.W_x]; G.\text{rf}^?; [G.E^i]) \}$
- ▶ G is strongly coherent (wrt mo).

Exercise

- ▶ Show that \sim is a bisimulation relation.
- ▶ Hence deduce that the operational StrongCOH model and the intermediate StrongCOH model have the same outcomes for any given program.

Declarative StrongCOH \subseteq intermediate StrongCOH

- ▶ Let G be an StrongCOH-consistent execution graph of P_0 with outcome O .
- ▶ Let mo be a modification order for G such that G is COH-consistent wrt mo .
- ▶ We show that $P_0, S_0, G_0, \text{mo}_0 \Rightarrow^* \text{skip} \parallel \dots \parallel \text{skip}, O, G, \text{mo}$.
- ▶ Let a_1, \dots, a_n be an enumeration of $G.E \setminus E_0$ following $G.\text{po} \cup G.\text{rf}$.
- ▶ By the “execution generation” lemma, there exist $\langle P_1, S_1 \rangle, \dots, \langle P_n, S_n \rangle$ such that:
 - ▶ $P_n = \text{skip} \parallel \dots \parallel \text{skip}$ and $S_n = O$
 - ▶ For every $1 \leq j \leq n$, we have:

$$P_{j-1}, S_{j-1} \xrightarrow{\text{tid}(a_j):\varepsilon}^* \xrightarrow{\text{tid}(a_j):\text{lab}(a_j)} \xrightarrow{\text{tid}(a_j):\varepsilon}^* P_j, S_j$$

- ▶ For every $0 \leq j \leq n$, let
 - ▶ G_j - the restriction of G to $E_0 \cup \{a_1, \dots, a_j\}$
 - ▶ mo_j - the restriction of mo to $E_0 \cup \{a_1, \dots, a_j\}$
- ▶ Then, for every $1 \leq j \leq n$, we have: (why?)

$$P_{j-1}, S_{j-1}, G_{j-1}, \text{mo}_{j-1} \Rightarrow^* P_j, S_j, G_j, \text{mo}_j$$

- ▶ Suppose that $P_0, S_0, G_0, \text{mo}_0 \Rightarrow^* \text{skip} \parallel \dots \parallel \text{skip}, O, G, \text{mo}$.
- ▶ We show that G is a StrongCOH-consistent execution graph of P with outcome O .
- ▶ We know:

$$P_0, S_0, G_0, \text{mo}_0 \Rightarrow P_1, S_1, G_1, \text{mo}_1 \Rightarrow \dots \Rightarrow P_n, S_n, G_n, \text{mo}_n$$

where $P_n, S_n, G_n, \text{mo}_n = \text{skip} \parallel \dots \parallel \text{skip}, O, G, \text{mo}$.

- ▶ By definition, G is COH-consistent.
- ▶ Using induction on the length of the sequence, we also have that $G.\text{po} \cup G.\text{rf}$ is acyclic.
- ▶ It remains to show that each G^i is an execution of $P_0(i)$ with final store $O(i)$. (This is done exactly as for SC.)

Operational semantics for RA

Can we extend the operational semantics to support message passing (*i.e.*, release-acquire synchronization)?

Message passing (MP)

$$\begin{array}{l} x = y = 0 \\ x := 42; \quad \parallel \quad a := y; \text{ // } 1 \\ y := 1 \quad \parallel \quad b := x \text{ // } 0 \end{array}$$

Double message passing

$$\begin{array}{l} x = y = 0 \\ x := 42; \quad \parallel \quad a := y; \text{ // } 1 \quad \parallel \quad b := z; \text{ // } 1 \\ y := 1 \quad \parallel \quad z := 1 \quad \parallel \quad c := x \text{ // } 0 \end{array}$$

Desired semantics

When reading a message the thread becomes aware of all messages that the writer of the message was aware of when the message was written.

We implement this using *message views*:

- ▶ Each message m will carry a view: the view of the thread who wrote m when m was written.
- ▶ When reading a message m , the thread will update its view to include at least the view contained in m .

Operational semantics for RA

- ▶ A message is a tuple $\langle x : v@t \text{ view} \rangle$ where $x \in \text{Loc}$, $v \in \text{Val}$, $t \in \text{Time}$ and $\text{view} : \text{Loc} \rightarrow \text{Time}$
- ▶ Initially, $M_0 \triangleq \{ \langle x : 0@0 \perp \rangle \mid x \in \text{Loc} \}$
- ▶ Bottom view: $\perp \triangleq \lambda x. 0$
- ▶ Joining views: $\text{view}_1 \sqcup \text{view}_2 \triangleq \lambda x. \max\{\text{view}_1(x), \text{view}_2(x)\}$

$$\begin{array}{c}
 \text{SILENT-THREAD} \\
 \frac{P, S \xrightarrow{i:\varepsilon} P', S'}{P, S, M, V \Rightarrow P', S', M, V}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{READ} \\
 \frac{
 \begin{array}{l}
 P, S \xrightarrow{i:l} P', S' \quad l = R(x, v) \\
 \langle x : v@t \text{ view} \rangle \in M \quad V(i)(x) \leq t \\
 \text{view}' = V(i) \sqcup \text{view}
 \end{array}
 }{P, S, M, V \Rightarrow P', S', M, V[i \mapsto \text{view}']}
 \end{array}$$

WRITE

$$\frac{
 \begin{array}{l}
 P, S \xrightarrow{i:l} P', S' \quad l = W(x, v) \\
 V(i)(x) < t \quad \forall v', \text{view}. \langle x : v'@t \text{ view} \rangle \notin M \\
 \text{view}' = V(i)[x \mapsto t] \quad M' = M \cup \{ \langle x : v@t \text{ view}' \rangle \}
 \end{array}
 }{P, S, M, V \Rightarrow P', S', M', V[i \mapsto \text{view}']}$$

Definition (Operational RA)

An outcome O is allowed for a program P under RA if there exist M, V such that $P, S_0, M_0, V_0 \Rightarrow^* \mathbf{skip} \parallel \dots \parallel \mathbf{skip}, O, M, V$.

Exercise

Prove the correspondence between the declarative and the operational definitions of RA.

Exercise: Strong release/acquire semantics

Suppose we change the write step in the operational semantics of RA as follows:

WRITE

$$\frac{\begin{array}{l} P, S \xrightarrow{i:t} P', S' \quad I = W(x, v) \\ \forall t', v', \text{view}. \langle x : v' @ t' \text{ view} \rangle \in M \Rightarrow t' < t \\ \text{view}' = V(i)[x \mapsto t] \quad M' = M \cup \{ \langle x : v @ t \text{ view}' \rangle \} \end{array}}{P, S, M, V \Rightarrow P', S', M', V[i \mapsto \text{view}']}$$

Here, when writing a message, the thread may only choose a timestamp larger than *all* timestamps that were used for the given location.

- ▶ Show an example which differentiates this model from RA.
- ▶ What will be the corresponding declarative semantics?

Exercise: Plain accesses in Java

Recall the following alternative definition of coherence:

Let mo be a modification order for an execution graph G .

G is *coherent wrt* mo iff the following hold:

- ▶ $rf; po$ is irreflexive. (no-future-read)
- ▶ $mo; po$ is irreflexive. (coherence-ww)
- ▶ $mo; rf; po$ is irreflexive. (coherence-rw)
- ▶ $rf^{-1}; mo; po$ is irreflexive. (coherence-wr)
- ▶ $rf^{-1}; mo; rf; po$ is irreflexive. (coherence-rr)

Plain accesses in the Java memory model do not provide full coherence. In particular, they do not ensure “coherence-rr”.

- ▶ Adapt the StrongCOH timestamp machine to match this weaker variant.

- ▶ **Taming release-acquire consistency.** Ori Lahav, Nick Giannarakis, Viktor Vafeiadis. POPL 2016: 649-662
- ▶ **A promising semantics for relaxed-memory concurrency.** Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, Derek Dreyer. POPL 2017: 175-189