

# Introduction to the C/C++11 concurrency model

Graduate seminar on weakly consistent concurrency

Ori Lahav



March 14, 2018

# The C11 memory model

- ▶ Introduced by the ISO C/C++ 2011 standards.
- ▶ Defines the semantics of **concurrent** memory accesses.
- ▶ Platform independent.
- ▶ Supported by popular compilers (GCC, LLVM, ...).

We will introduce a *simplified* version of (a fragment of) the C11 model as a *prototypical example* of a real world weak concurrency semantics.

Two types of accesses

**Ordinary  
(Non-Atomic)**

Races are **errors**

**Atomic**

Welcome to the  
**expert mode**

Two types of accesses

**Ordinary  
(Non-Atomic)**

Races are **errors**

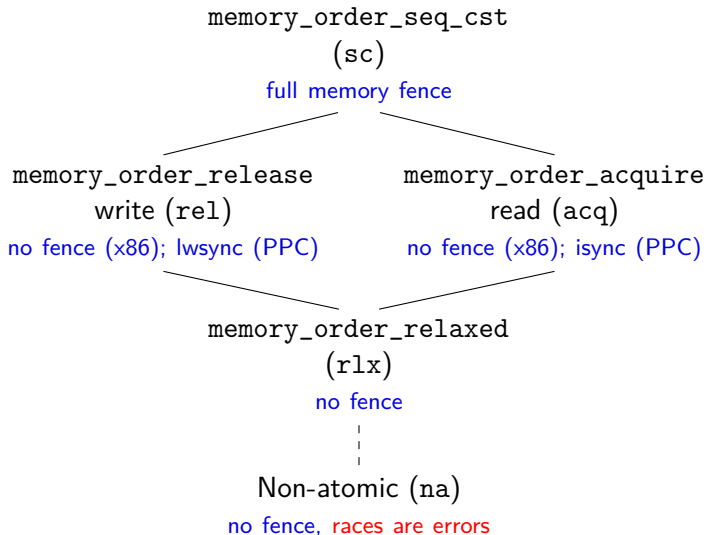
**Atomic**

Welcome to the  
**expert mode**

**DRF (data race freedom) guarantee**

no data races  
under SC  $\implies$  only  
SC behaviors

# A spectrum of access modes



+ Explicit primitives for fences

# C11: a declarative memory model

**Declarative semantics** abstracts away from implementation details.

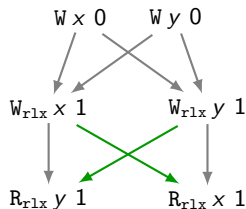
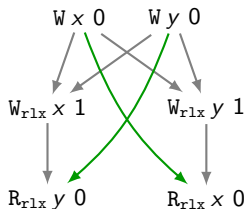
1. a program  $\rightsquigarrow$  a set of directed graphs (called: *execution graphs*)
2. The memory model defines what executions are *consistent*.
3. The semantics of a program is the set of its *consistent executions*.

## Exception: “catch-fire” semantics

- ▶ Existence of at least one “bad” consistent execution (with a forbidden data races) implies undefined behavior.

# Execution graphs

## Store buffering (SB)

$$\begin{array}{l} x = y = 0 \\ x :=_{rlx} 1 \quad \parallel \quad y :=_{rlx} 1 \\ a :=_{rly} x \quad \parallel \quad b :=_{rlx} y \end{array}$$


## Relations

- ▶ Program order,  $po$
- ▶ Reads-from,  $rf$  (we require that *every read reads from some write*).

$$\begin{aligned}
[-] &: \text{CExp} \rightarrow \mathbb{P}((res : \text{Val} \cup \{\perp\}, \mathcal{A} : \mathbb{P}(\text{AName}), \text{lab} : \mathcal{A} \rightarrow \text{Act}, \text{sb} : \mathbb{P}(\mathcal{A} \times \mathcal{A}), \text{fst} : \mathcal{A}, \text{lst} : \mathcal{A})) \\
[v] &\stackrel{\text{def}}{=} \{\langle v, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{skip}\} \\
[\text{alloc}()] &\stackrel{\text{def}}{=} \{\langle \ell, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \ell \in \text{Loc} \wedge \text{lab}(a) = \Lambda(\ell) \} \\
[[v]z ::= v'] &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{Wz}(v, v') \} \\
[[v]z &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge \text{lab}(a) = \text{Rz}(v, v') \} \\
[\text{CAS}_{X,Y}(v, v_0, v_n)] &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge v' \neq v_0 \wedge \text{lab}(a) = \text{Ry}(v, v') \} \\
&\cup \{\langle v_0, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{RMW}_X(v, v_0, v_n) \} \\
[\text{let } x = E_1 \text{ in } E_2] &\stackrel{\text{def}}{=} \{\langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \mid \langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \} \\
&\cup \{\langle res_2, \mathcal{A}_1 \uplus \mathcal{A}_2, \text{lab}_1 \cup \text{lab}_2, \text{sb}_1 \cup \text{sb}_2 \cup \{(lst_1, \text{fst}_2)\}, \text{fst}_1, \text{fst}_2, \text{lst}_2 \rangle \mid \\
&\quad \langle res_1, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \wedge \langle res_2, \mathcal{A}_2, \text{lab}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in [E_2[v_1/x]] \} \\
[\text{repeat } E \text{ end}] &\stackrel{\text{def}}{=} \{\langle res_N, \bigcup_{i \in [1..N]} \mathcal{A}_i, \bigcup_{i \in [1..N]} \text{lab}_i, \bigcup_{i \in [1..N]} \text{sb}_i \cup \{(lst_N, \text{fst}_2), \dots, (lst_{N-1}, \text{fst}_N)\}, \text{fst}_1, \text{lst}_N \rangle \mid \\
&\quad \forall i. \langle res_i, \mathcal{A}_i, \text{lab}_i, \text{sb}_i, \text{fst}_i, \text{lst}_i \rangle \in [E] \wedge (i \neq N \implies res_i = 0) \wedge res_N \neq 0 \} \\
[[E_1] \| E_2] &\stackrel{\text{def}}{=} \{\langle \text{combine}(res_1, res_2), \mathcal{A}_1 \uplus \mathcal{A}_2 \uplus \{a_{\text{fork}}, a_{\text{join}}\}, \text{lab}_1 \cup \text{lab}_2 \cup \{a_{\text{fork}} \mapsto \text{skip}, a_{\text{join}} \mapsto \text{skip}\}, \\
&\quad \text{sb}_1 \cup \text{sb}_2 \cup \{(a_{\text{fork}}, \text{fst}_1), (a_{\text{fork}}, \text{fst}_2), (lst_1, a_{\text{join}}), (lst_2, a_{\text{join}}), a_{\text{fork}}, a_{\text{join}}\} \mid \\
&\quad \langle res_1, \mathcal{A}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \wedge \langle res_2, \mathcal{A}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in [E_2] \wedge a_{\text{fork}}, a_{\text{join}} \in \text{AName} \}
\end{aligned}$$

Figure 2. Semantics of closed program expressions.

$$\begin{aligned}
&\not\exists x. \text{hb}(x, x) && \text{(IrreflexiveHB)} \\
&\forall \ell. \text{totalorder}(\{a \in \mathcal{A} \mid \text{iswrite}_\ell(a)\}, \text{mo}) \wedge \text{hb}_\ell \subseteq \text{mo} && \text{(ConsistentMO)} \\
&\text{totalorder}(\{a \in \mathcal{A} \mid \text{isSeqCst}(a)\}, \text{sc}) \wedge \text{hb}_{\text{SeqCat}} \subseteq \text{sc} \wedge \text{mo}_{\text{SeqCat}} \subseteq \text{sc} && \text{(ConsistentSC)} \\
&\forall b. \text{rf}(b) \neq \perp \iff \exists \ell. a. \text{iswrite}_\ell(a) \wedge \text{isread}_\ell(b) \wedge \text{hb}(a, b) && \text{(ConsistentRFdom)} \\
&\forall a, b. \text{rf}(b) = a \implies \exists \ell. v. \text{iswrite}_{\ell,v}(a) \wedge \text{isread}_{\ell,v}(b) \wedge \neg \text{hb}(a, b) && \text{(ConsistentRF)} \\
&\forall a, b. \text{rf}(b) = a \wedge (\text{mode}(a) = \text{na} \vee \text{mode}(b) = \text{na}) \implies \text{hb}(a, b) && \text{(ConsistentRFna)} \\
&\forall a, b. \text{rf}(b) = a \wedge \text{isSeqCst}(b) \implies \text{isc}(a, b) \vee \neg \text{isSeqCst}(a) \wedge (\forall x. \text{isc}(x, b) \implies \neg \text{hb}(a, x)) && \text{(RestrSCReads)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), \text{rf}(a)) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRRR)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), a) \wedge \text{iswrite}(a) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentWR)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(b, \text{rf}(a)) \wedge \text{iswrite}(b) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRW)} \\
&\forall a. \text{isrmw}(a) \wedge \text{rf}(a) \neq \perp \implies \text{mo}(\text{rf}(a), a) \wedge \not\exists c. \text{mo}(\text{rf}(a), c) \wedge \text{mo}(c, a) && \text{(AtomicRMW)} \\
&\quad \forall a, b, \ell. \text{lab}(a) = \text{lab}(b) = \Lambda(\ell) \implies a = b && \text{(ConsistentAlloc)}
\end{aligned}$$

where  $\text{iswrite}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{old}}. \text{lab}(a) \in \{W_X(\ell, v), \text{RMW}_X(\ell, v_{\text{old}}, v)\}$   $\text{iswrite}_\ell(a) \stackrel{\text{def}}{=} \exists v. \text{iswrite}_{\ell,v}(a)$

$\text{isread}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{new}}. \text{lab}(a) \in \{R_X(\ell, v), \text{RMW}_X(\ell, v, v_{\text{new}})\}$   $\text{etc.}$

$\text{rsElem}(a, b) \stackrel{\text{def}}{=} \text{sameThread}(a, b) \vee \text{isrmw}(b)$

$\text{rseq}(a) \stackrel{\text{def}}{=} \{a\} \cup \{b \mid \text{rsElem}(a, b) \wedge \text{mo}(a, b) \wedge (\forall c. \text{mo}(a, c) \wedge \text{mo}(c, b) \implies \text{rsElem}(a, c))\}$

$\text{sw} \stackrel{\text{def}}{=} \{(a, b) \mid \text{mode}(a) \in \{\text{rel}, \text{rel\_acq}, \text{sc}\} \wedge \text{mode}(b) \in \{\text{acq}, \text{rel\_acq}, \text{sc}\} \wedge \text{rf}(b) \in \text{rseq}(a)\}$

$\text{hb} \stackrel{\text{def}}{=} (\text{sb} \uplus \text{sw})^+$

$\text{hb}_\ell \stackrel{\text{def}}{=} \{(a, b) \in \text{hb} \mid \text{iswrite}_\ell(a) \wedge \text{iswrite}_\ell(b)\}$

$X_{\text{SeqCst}} \stackrel{\text{def}}{=} \{(a, b) \in X \mid \text{isSeqCst}(a) \wedge \text{isSeqCst}(b)\}$

$\text{isc}(a, b) \stackrel{\text{def}}{=} \text{iswrite}_{\text{locs}(b)}(a) \wedge \text{sc}(a, b) \wedge \not\exists c. \text{sc}(a, c) \wedge \text{sc}(c, b) \wedge \text{iswrite}_{\text{locs}(b)}(c)$

Figure 3. Axioms satisfied by consistent C11 executions, Consistent( $\mathcal{A}$ , lab, sb, rf, mo, sc).

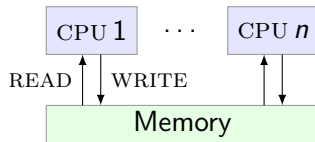
$c : W(\ell, 1) \xrightarrow{\text{rf}} a : R(\ell, 1)$ $\uparrow \text{mo}$ $d : W(\ell, 2) \xrightarrow{\text{rf}} b : R(\ell, 2)$ violates CoherentRR	$c : W(\ell, 2) \xrightarrow{\text{mo}} a : W(\ell, 1)$ $\text{rf} \searrow$ $b : R(\ell, 2)$ violates CoherentWR	$c : W(\ell, 1) \xrightarrow{\text{mo}} a : R(\ell, 1)$ $\text{mo} \searrow$ $b : W(\ell, 2)$ violates CoherentRW	$a \xrightarrow{\text{rf}} b$ means $a = \text{rf}(b)$ $a \xrightarrow{\text{mo}} b$ means $\text{mo}(a, b)$ $a \xrightarrow{\text{hb}} b$ means $\text{hb}(a, b)$
---	--	--	--

Figure 4. Sample executions violating coherency conditions (Batty et al. 2011).



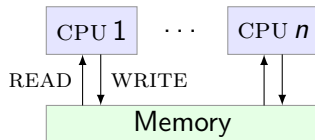
# Sequential consistency

Before C/C++11, how would you define SC as a declarative memory model?



# Sequential consistency

Before C/C++11, how would you define SC as a declarative memory model?



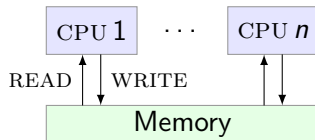
## Definition (Lamport '79)

$G$  is *SC-consistent* if there exists a relation  $sc$  such that:

- ▶  $sc$  is a total order on the events of  $G$ .
- ▶  $po \cup rf \subseteq sc$ .
- ▶ If  $\langle a, b \rangle \in rf$  then there does not exist  $c \in W_{loc(a)}$  such that  $\langle a, c \rangle \in sc$  and  $\langle c, b \rangle \in sc$ .

# Sequential consistency

Before C/C++11, how would you define SC as a declarative memory model?



## Definition (Lamport '79)

$G$  is *SC-consistent* if there exists a relation  $sc$  such that:

- ▶  $sc$  is a total order on the events of  $G$ .
  - ▶  $po \cup rf \subseteq sc$ .
  - ▶ If  $\langle a, b \rangle \in rf$  then there does not exist  $c \in W_{loc(a)}$  such that  $\langle a, c \rangle \in sc$  and  $\langle c, b \rangle \in sc$ .
- 
- ▶ We will later see alternative more “economical” definitions that do not require a total order on all events of  $G$ .

# Ingredients of execution graph consistency in C/C++11

1. SC-per-location (a.k.a. coherence)
2. Release/acquire synchronization
3. Global conditions on SC accesses

# Ingredients of execution graph consistency in C/C++11

1. SC-per-location (a.k.a. coherence)
2. Release/acquire synchronization
3. Global conditions on SC accesses

# SC-per-location

## Definition (Declarative definition of SC)

$G$  is *SC-consistent* if there exists a relation  $sc$  such that:

- ▶  $sc$  is a total order on the events of  $G$ .
- ▶  $po \cup rf \subseteq sc$ .
- ▶ If  $\langle a, b \rangle \in rf$  then there does not exist  $c \in W_{loc(a)}$  such that  $\langle a, c \rangle \in sc$  and  $\langle c, b \rangle \in sc$ .

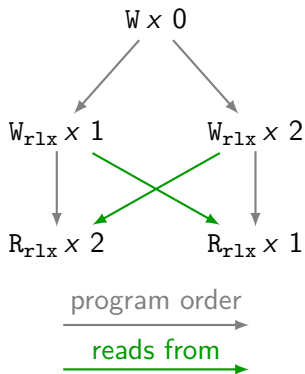
## Definition (SC-per-location)

$G$  satisfies *SC-per-location* if for every location  $x$ , there exists a relation  $sc_x$  such that:

- ▶  $sc_x$  is a total order on the events of  $G$  that access  $x$ .
- ▶  $(po \cup rf) \cap \{\langle a, b \rangle \mid loc(a) = loc(b) = x\} \subseteq sc_x$ .
- ▶ If  $\langle a, b \rangle \in rf$  then there does not exist  $c \in W_x$  such that  $\langle a, c \rangle \in sc_x$  and  $\langle c, b \rangle \in sc_x$ .

# SC-per-location: Example

$x = 0$   
 $x :=_{r1x} 1 \parallel x :=_{r1x} 2$   
 $a := x_{r1x} \parallel b := x_{r1x}$



*inconsistent!*

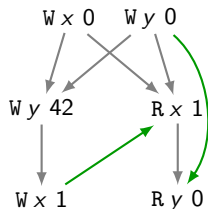
# Release/acquire synchronization

SC-per-location is often too weak:

- ▶ It does not support the message passing idiom:

## Message passing (MP)

```
y := 42; || a := x; // 1  
x := 1;  || b := y; // 0
```





## Implementing locks?

Can we implement a lock under SC-per-location?

## Implementing locks?

Can we implement a lock under SC-per-location?

### Spinlock implementation

lock(*l*) :

*r* := 0;

**while**  $\neg r$  **do**

*r* := **CAS**(*l*, 0, 1)

unlock(*l*) :

*l* := 0

# Implementing locks?

Can we implement a lock under SC-per-location?

## Spinlock implementation

<u><b>lock</b>(<math>l</math>) :</u>	<u><b>unlock</b>(<math>l</math>) :</u>
$r := 0;$	$l := 0$
<b>while</b> $\neg r$ <b>do</b>	
$r := \mathbf{CAS}(l, 0, 1)$	

## Lock example

<b>lock</b> ( $l$ )		<b>lock</b> ( $l$ )
$x := 1$		$y := 1$
$a := y \ // \ 0$		$b := x \ // \ 0$
<b>unlock</b> ( $l$ )		<b>unlock</b> ( $l$ )

Under SC-per-location, the spinlock implementation does not guarantee mutual exclusion.

## Synchronization in C/C++11 through examples

```
    int y = 0;
    int x = 0;
y = 42; || if(x == 1){
x = 1; ||     print(y);
      || }
```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||     print(y);
      || }
```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;  ← race print(y);
        || }

```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;  ||     print(y);
        || }
        ||
        ||
```

*race*

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||     print(y);
          || }
          ||
```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ← race print(y);
        || }

```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ← race print(y);
        || }

```



# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||     print(y);
         ||     }
         ||
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||     print(y);
         ||     }
         ||
```

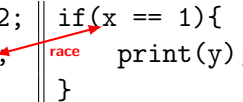
3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||     print(y);
         ||     }
         ||
```

# Synchronization in C/C++11 through examples

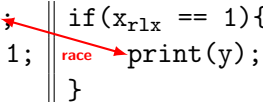
1

```
int y = 0;
int x = 0;
y = 42;
x = 1;
if(x == 1){
    print(y);
}
```



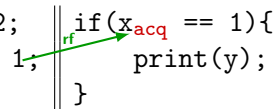
2

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rlx 1;
if(x_rlx == 1){
    print(y);
}
```



3

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rel 1;
if(x_acq == 1){
    print(y);
}
```



# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42;
x = 1;
if(x == 1){
    print(y);
}
```

A vertical double line separates the two threads. A red arrow labeled "race" points from the assignment `x = 1;` in the left thread to the condition `if(x == 1){` in the right thread.

2

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_r1x 1;
if(x_r1x == 1){
    print(y);
}
```

A vertical double line separates the two threads. A red arrow labeled "race" points from the assignment `x =_r1x 1;` in the left thread to the condition `if(x_r1x == 1){` in the right thread.

3

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rel 1;
if(x_acq == 1){
    print(y);
}
```

A vertical double line separates the two threads. A green arrow labeled "rf" points from the assignment `x =_rel 1;` in the left thread to the condition `if(x_acq == 1){` in the right thread. A green arrow labeled "sw" points from the condition back to the assignment, indicating a write-read dependency.

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||   print(y);
         ||   }
         ||
         ||
         ||
         ||
         ||
```

A red arrow labeled "race" points from the `x == 1` condition to the `x = 1` assignment.

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xr1x == 1){
x =r1x 1; ||   print(y);
         ||   }
         ||
```

A red arrow labeled "race" points from the `xr1x == 1` condition to the `x =r1x 1` assignment.

3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
         ||   }
         ||
         ||
         ||
         ||
         ||
```

A green arrow labeled "rf" points from the `xacq == 1` condition to the `x =rel 1` assignment. A green arrow labeled "sw" points from the `x =rel 1` assignment to the `xacq == 1` condition.

4

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xr1x == 1){
fencerel; ||   fenceacq;
x =r1x 1; ||   print(y);
         ||   }
         ||
```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||   print(y);
      ||   }
      ||
      ||   race
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||   print(y);
      ||   }
      ||
      ||   race
```

3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
      ||   }
      ||
      ||   rf
      ||   sw
```

4

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
fencerel; ||   fenceacq;
x =rlx 1; ||   print(y);
      ||   }
      ||
      ||   rf
```

# Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42;
x = 1;
if(x == 1){
    print(y);
}
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_r1x 1;
if(x_r1x == 1){
    print(y);
}
```

3

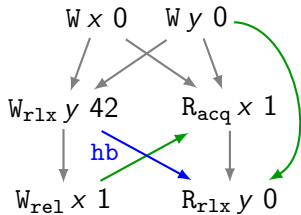
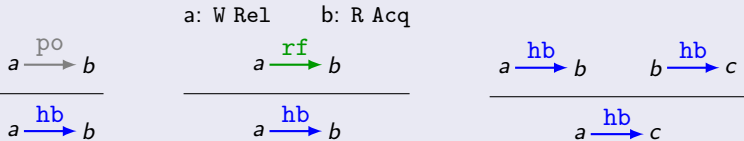
```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rel 1;
if(x_acq == 1){
    print(y);
}
```

4

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_r1x 1;
fence_rel;
if(x_r1x == 1){
    fence_acq;
    print(y);
}
```

# The “happens-before” relation

## Definition (happens-before)



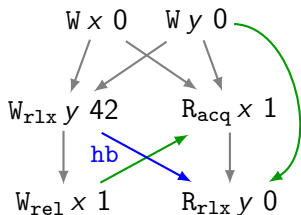
# C11 coherence

- ▶ Require (for every location  $x$ ) that :

$$\text{hb} \cap \{\langle a, b \rangle \mid \text{loc}(a) = \text{loc}(b) = x\} \subseteq \text{sc}_x$$

(instead of  $(\text{po} \cup \text{rf}) \cap \{\langle a, b \rangle \mid \text{loc}(a) = \text{loc}(b) = x\} \subseteq \text{sc}_x$ ).

- ▶ Using acquire CAS's and release writes, we can implement locks.





Using **hb**, we can formally define a race in this model.

## Definition

Two events  $a, b$  form a *race* in an execution  $G$  if the following hold:

- ▶  $\text{loc}(a) = \text{loc}(b)$ .
- ▶  $a \in W$  or  $b \in W$  ( $W$  denotes the set of writes and RMWs).
- ▶  $\langle a, b \rangle \notin \text{hb}$  and  $\langle b, a \rangle \notin \text{hb}$ .

## SC accesses and fences

Store buffer

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

How to guarantee only SC behaviors (*i.e.*,  $a = 1 \vee b = 1$ )?

# SC accesses and fences

## Store buffer

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

How to guarantee only SC behaviors (*i.e.*,  $a = 1 \vee b = 1$ )?

$$\begin{array}{l} x :=_{\text{sc}} 1; \\ a := y_{\text{sc}}; \end{array} \parallel \parallel \begin{array}{l} y :=_{\text{sc}} 1; \\ b := x_{\text{sc}}; \end{array} \quad \approx \quad \begin{array}{l} x :=_{\text{rlx}} 1; \\ \mathbf{fence}_{\text{sc}}; \\ a := y_{\text{rlx}}; \end{array} \parallel \parallel \begin{array}{l} y :=_{\text{rlx}} 1; \\ \mathbf{fence}_{\text{sc}}; \\ b := x_{\text{rlx}}; \end{array}$$

# Semantics of SC accesses

Basic idea: Require an **sc**-order on SC accesses:

- ▶ **sc** is a total order on all SC accesses in  $G$ .
- ▶  $\text{hb} \cap \{\langle a, b \rangle \mid \text{mod}(a) = \text{mod}(b) = \text{sc}\} \subseteq \text{sc}$ .
- ▶ If  $\langle a, b \rangle \in \text{rf}$  then there does not exist  $c \in W_{\text{loc}(a)}$  such that  $\langle a, c \rangle \in \text{sc}$  and  $\langle c, b \rangle \in \text{sc}$ .
  
- ▶ Semantics of SC fences?
- ▶ Accesses to the same location in both SC and non-SC modes?

This is a complicated part of the model:

[Repairing Sequential Consistency in C/C++11 PLDI'17]

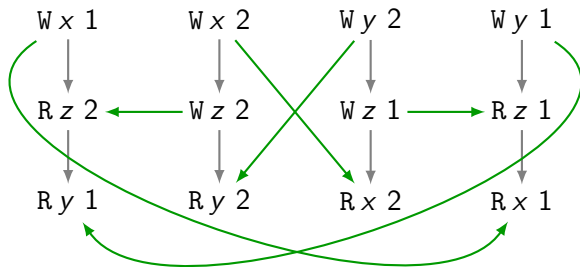
# Revisiting sequential consistency

## Definition

$G$  is *SC-consistent* if there exists a relation  $sc$  such that:

- ▶  $sc$  is a total order on the events of  $G$ .
- ▶  $po \cup rf \subseteq sc$ .
- ▶ If  $\langle a, b \rangle \in rf$  then there does not exist  $c \in W_{loc}(a)$  such that  $\langle a, c \rangle \in sc$  and  $\langle c, b \rangle \in sc$ .

Can we avoid ordering *all* events?



# Alternative definition of SC

## Definition (Modification order (a.k.a. coherence order))

$mo$  is called a *modification order* for an execution graph  $G$  if

$mo = \bigcup_{x \in Loc} mo_x$  where each  $mo_x$  is a total order on  $W_x$ .

## Definition (Alternative SC definition)

An execution graph  $G$  is called *SC-consistent* if the following holds:

- ▶ There exists a modification order  $mo$  for  $G$  such that  $po \cup rf \cup mo \cup fr$  is acyclic where:
  - ▶  $fr \stackrel{\text{def}}{=} rf^{-1}; mo \setminus id$  (from-reads / reads-before)

# Equivalence

## Theorem

*The two SC definitions are equivalent.*

## Proof (sketch).

### Original SC $\implies$ alternative SC:

- ▶ Take  $mo_x$  to be the restriction of  $sc$  on  $W_x$ .
- ▶ Then,  $po \cup rf \cup mo \cup fr \subseteq sc$ .

### Alternative SC $\implies$ original SC:

- ▶ Take  $sc$  to be any total order extending  $po \cup rf \cup mo \cup fr$ .  $\square$



# Alternative definition of SC-per-location

## Definition (SC-per-location)

$G$  satisfies *SC-per-location* if for every location  $x$ , there exists a relation  $\mathbf{sc}_x$  such that:

- ▶  $\mathbf{sc}_x$  is a total order on the events of  $G$  that access  $x$ .
- ▶  $(\mathbf{po} \cup \mathbf{rf}) \cap \{\langle a, b \rangle \mid \mathbf{loc}(a) = \mathbf{loc}(b) = x\} \subseteq \mathbf{sc}_x$ .
- ▶ If  $\langle a, b \rangle \in \mathbf{rf}$  then there does not exist  $c \in W_x$  such that  $\langle a, c \rangle \in \mathbf{sc}_x$  and  $\langle c, b \rangle \in \mathbf{sc}_x$ .

## Definition (Alternative SC-per-location definition)

An execution graph  $G$  satisfies *SC-per-location* if the following holds:

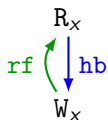
- ▶ There exists a modification order  $\mathbf{mo}$  for  $G$  such that  $\mathbf{po}|_{\mathbf{loc}} \cup \mathbf{rf} \cup \mathbf{mo} \cup \mathbf{fr}$  is acyclic where:
  - ▶  $\mathbf{fr} \stackrel{\text{def}}{=} \mathbf{rf}^{-1}; \mathbf{mo} \setminus \mathbf{id}$  (from-reads / reads-before)
  - ▶  $\mathbf{po}|_{\mathbf{loc}} = \mathbf{po} \cap (\langle a, b \rangle \mid \mathbf{loc}(a) = \mathbf{loc}(b))$

- ▶ Similar “optimization” is done in the full C11 model, requiring that:

$hb|_{loc} \cup rf \cup mo \cup fr$  is acyclic

- ▶ This acyclicity condition can be simplified to several “bad patterns”.

# “Bad patterns” I



$a := x // 1$

$x := 1$

no-future-read



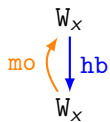
$r := \mathbf{CAS}(x, 1, 1) // 1$

rmw-1

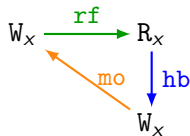
## Notations:

- ▶  $W$  is either a write or an RMW.
- ▶  $R$  is either a read or an RMW.

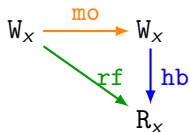
# “Bad patterns” II

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


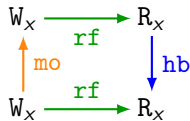
coherence-ww



coherence-rw



coherence-wr

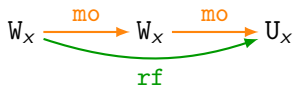


coherence-rr

# “Bad patterns” III



rmw-2



atomicity

# Alternative definition via “Bad patterns”

## Theorem

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

$hb|_{loc} \cup rf \cup mo \cup fr$  is acyclic iff the following hold:

- ▶  $rf; hb$  is irreflexive. (no-future-read)
- ▶  $mo; hb$  is irreflexive. (coherence-ww)
- ▶  $mo; rf; hb$  is irreflexive. (coherence-rw)
- ▶  $fr; hb$  is irreflexive. (coherence-wr)
- ▶  $fr; rf; hb$  is irreflexive. (coherence-rr)
- ▶  $rf$  is irreflexive. (rmw-1)
- ▶  $mo; rf$  is irreflexive. (rmw-2)
- ▶  $fr; mo$  is irreflexive. (rmw-atomicity)

Side note: Java plain accesses do not guarantee (coherence-rr).

# Takeaways

- ▶ C/C++11 concurrency semantics is specified *declaratively* using formal constraints on execution graphs.
- ▶ This is a *highly flexible* framework that abstracts away all execution details.
- ▶ Reasoning about program behaviors requires analyzing *possible cycles* in execution graphs.
- ▶ A declarative approach to concurrency semantics is prominent in other systems as well:
  - ▶ *Multicore architectures:*  
Jade Alglave, Luc Maranget, Michael Tautschnig.  
Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory.
  - ▶ *Distributed systems:*  
Sebastian Burckhardt.  
Principles of Eventual Consistency.