

# Multicore Programming: C++0x

Mark Batty

University of Cambridge

in collaboration with

Scott Owens, Susmit Sarkar, Peter Sewell, Tjark Weber

November, 2010

# C++0x: the next C++

Specified by the C++ Standards Committee

Defined in The Standard, a 1300 page prose document

The design is a detailed compromise:

- performance, optimisations and hardware
- usability
- compatibility with the next C, C1X
- legacy code

# C++0x: the next C++

Our mathematical model is faithful to the intent of, and has influenced The Standard

The model:

- syntactically separates out expert features
- has a weak memory
- defines a happens-before relation
- requires non-atomic reads and writes to be DRF
- provides atomic reads and writes for racy programs

# The syntactic divide

## An example of the syntax

```
// for regular programmers:  
atomic_int x = 0;  
x.store(1);  
y = x.load();
```

```
// for experts:  
x.store(2, memory_order);  
y = x.load(memory_order);  
atomic_thread_fence(memory_order);
```

## With a choice of *memory\_order*

mo_seq_cst	mo_release	mo_acquire
mo_acq_rel	mo_consume	mo_relaxed

# A model of two parts

An operational semantics:

Processes programs, identifying *memory actions*

Constructs candidate executions,  $E_{\text{opsem}}$

An axiomatic memory model:

Judges  $E_{\text{opsem}}$  paired with a memory ordering,  $X_{\text{witness}}$

Searches the consistent executions for races and unconstrained reads

# Judgement of the axiomatic model

```
cpp_memory_model opsem (p : program) =  
  let pre_executions = {(Eopsem, Xwitness).  
    opsem p Eopsem ∧  
    consistent_execution (Eopsem, Xwitness)} in  
  if ∃X ∈ pre_executions.  
    (indeterminate_reads X ≠ {}) ∨  
    (unsequenced_races X ≠ {}) ∨  
    (data_races X ≠ {})  
  then NONE  
  else SOME pre_executions
```

# The relations of a pre-execution

An  $E_{\text{opsem}}$  part containing:

sb — *sequenced before*, program order

asw — *additional synchronizes with*, inter-thread ordering

dd — *data-dependence*

An  $X_{\text{witness}}$  part containing:


rf — relates a write to any reads that take its value

sc — a total order over  $\text{mo\_seq\_cst}$  and mutex actions

mo — *modification order*, per location total order of writes

# A single threaded program

```
int main() {  
    int x = 2;  
    int y = 0;  
    y = (x == x);  
    return 0; }  
}
```



../examples/t1.c



# Memory actions

action ::=

a:R <sub>na</sub> x=v	non-atomic read
a:W <sub>na</sub> x=v	non-atomic write
a:R <sub>mo</sub> x=v	atomic read
a:W <sub>mo</sub> x=v	atomic write
a:RMW <sub>mo</sub> x=v1/v2	atomic read-modify-write
a:L x	lock
a:U x	unlock
a:F <sub>mo</sub>	fence

# Memory orders

Memory orders are shown as follows:

```
mo ::=  
    SC    memory_order_seq_cst  
|  RLX   memory_order_relaxed  
|  REL   memory_order_release  
|  ACQ   memory_order_acquire  
|  CON   memory_order_consume  
|  A/R   memory_order_acq_rel
```

# Location kinds

```
location_kind =  
    MUTEX  
    | NON_ATOMIC  
    | ATOMIC
```

```
actions_respect_location_kinds =
```

```
   $\forall a.$ 
```

```
    case location  $a$  of SOME  $l \rightarrow$ 
```

```
      (case location-kind  $l$  of
```

```
        MUTEX  $\rightarrow$  is_lock_or_unlock  $a$ 
```


```
        || NON_ATOMIC  $\rightarrow$  is_load_or_store  $a$ 
```

```
        || ATOMIC  $\rightarrow$  is_load_or_store  $a \vee$  is_atomic_action  $a$ )
```

```
        || NONE  $\rightarrow$  T
```

# That single threaded program again

```
int main() {  
    int x = 2;  
    int y = 0;  
    y = (x == x);  
    return 0; }  
}
```




../examples/t1.c

# Unsequenced race

$$\begin{aligned} \text{unsequenced\_races} = & \{(a, b). \\ & \text{is\_load\_or\_store } a \wedge \text{is\_load\_or\_store } b \wedge \\ & (a \neq b) \wedge \text{same\_location } a \ b \wedge (\text{is\_write } a \vee \text{is\_write } b) \wedge \\ & \text{same\_thread } a \ b \wedge \\ & \neg(a \xrightarrow{\text{sequenced-before}} b \vee b \xrightarrow{\text{sequenced-before}} a)\} \end{aligned}$$

# An unsequenced race

```
int main() {  
    int x = 2;  
    int y = 0;  
    y = (x == (x=3));  
    return 0; }  
}
```

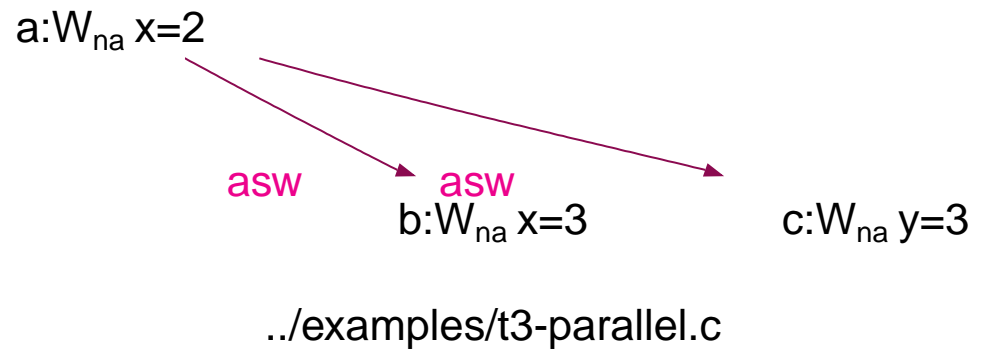


# A multi-threaded program

```
void foo(int* p) {*p=3;}
int main() {
    int x = 2;
    int y;
    thread t1(foo, &x);
    y = 3;
    t1.join();
    return 0; }
```

becomes:

```
int main() {
    int x = 2;
    int y;
    {{{ x = 3;
      ||| y = 3;
      }}}
    return 0; }
```



# Synchronizes-with and happens-before

The parent thread has synchronization edges, labeled  $asw$ , to its child threads. There are other ways to synchronize.

We will define the happens-before relation later. It contains the transitive closure of all synchronization edges and all sequenced before edges (amongst other things).



# Data race

$\text{data\_races} = \{(a, b).$

$(a \neq b) \wedge \text{same\_location } a \ b \wedge (\text{is\_write } a \vee \text{is\_write } b) \wedge$

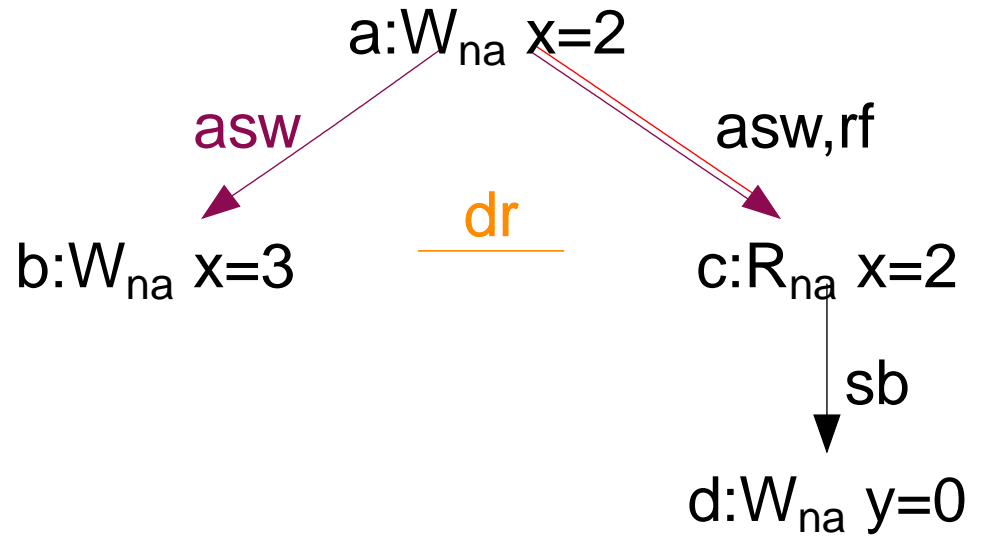
$\neg \text{same\_thread } a \ b \wedge$

$\neg(\text{is\_atomic\_action } a \wedge \text{is\_atomic\_action } b) \wedge$

$\neg(a \xrightarrow{\text{happens-before}} b \vee b \xrightarrow{\text{happens-before}} a)\}$

# A data race

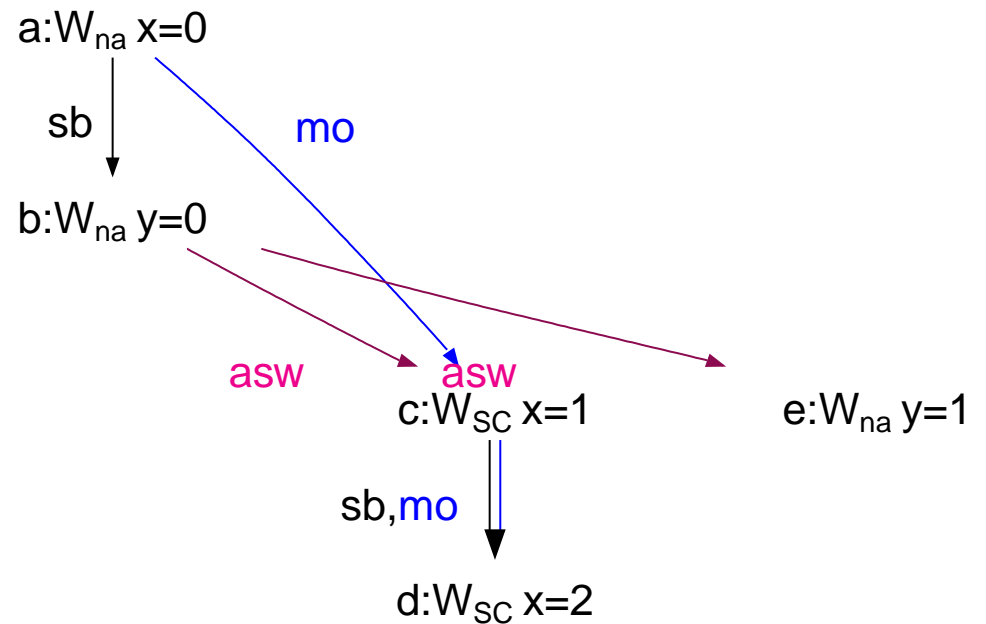
```
int main() {  
    int x = 2;  
    int y;  
    {{{ x=3;  
      ||| y=(x==3);  
      }}};  
    return 0; }
```



# Modification order

A total order of the writes at each atomic location, similar to coherence order on Power

```
int main() {
    atomic_int x = 0;
    int y = 0;
    {{{ { x.store(1);
          x.store(2); }
      ||| { y = 1; }
    }}}
    return 0; }
```



../examples/t70-na-mo.c

# SC order

There is a total order over all sequentially consistent atomic actions. SC atomics read the last prior write in SC order (or a non SC write).

consistent\_sc\_order =

**let** sc\_happens\_before =  $\xrightarrow{\text{happens-before}}|_{\text{all\_sc\_actions}}$  **in**

**let** sc\_mod\_order =  $\xrightarrow{\text{modification-order}}|_{\text{all\_sc\_actions}}$  **in**

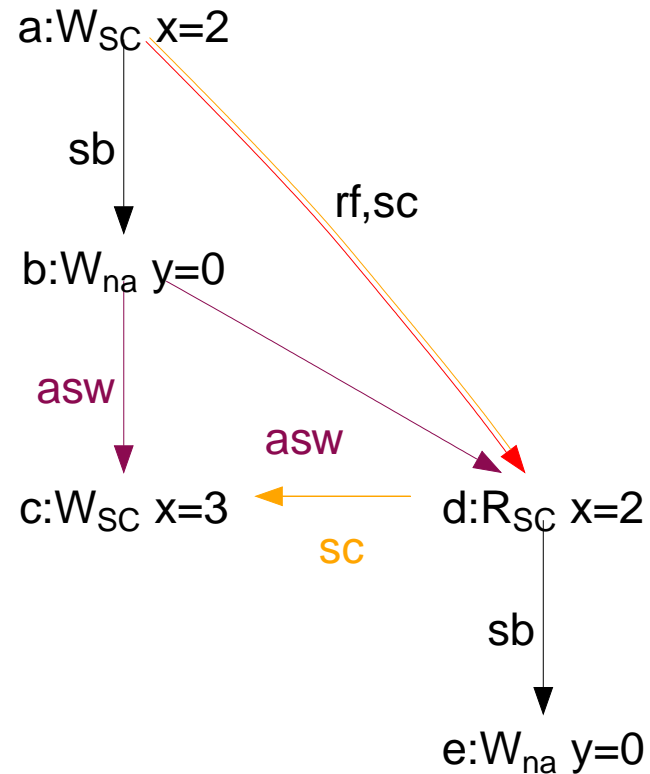
strict\_total\_order\_over all\_sc\_actions ( $\xrightarrow{\text{SC}}$ )  $\wedge$

$\xrightarrow{\text{sc\_happens\_before}} \subseteq \xrightarrow{\text{SC}} \wedge$

$\xrightarrow{\text{sc\_mod\_order}} \subseteq \xrightarrow{\text{SC}}$


# Atomic actions do not race

```
int main() {  
    atomic_int x;  
    x.store(2, mo_seq_cst);  
    int y = 0;  
    {{{ x.store(3);  
        ||| y = ((x.load()) == 3);  
        }}};  
    return 0; }
```




# The release-acquire idiom

```
// sender      | // receiver  
x = ...       | while (0 == y);  
y = 1;        | r = x;
```



../examples/t15.c

# Release-acquire synchronization



../examples/t8a.c

# The release sequence


The release sequence is a sub-sequence of the the modification order following a release

$$\text{rs\_element } rs\_head\ a = \\ \text{same\_thread } a\ rs\_head \vee \text{is\_atomic\_rmw } a$$

$$a_{rel} \xrightarrow{\text{release-sequence}} b = \\ \text{is\_at\_atomic\_location } b \wedge \\ \text{is\_release } a_{rel} \wedge ( \\ (b = a_{rel}) \vee \\ (\text{rs\_element } a_{rel}\ b \wedge a_{rel} \xrightarrow{\text{modification-order}} b \wedge \\ (\forall c. a_{rel} \xrightarrow{\text{modification-order}} c \xrightarrow{\text{modification-order}} b \implies \\ \text{rs\_element } a_{rel}\ c)))$$



# An execution with a release sequence



../examples/t8a-no-sw.c

# Synchronizes-with

$$a \xrightarrow{\text{synchronizes-with}} b =$$

(\* – additional synchronization, from thread create etc. – \*)

$$a \xrightarrow{\text{additional-synchronized-with}} b \vee$$

(same\_location  $a\ b \wedge a \in \text{actions} \wedge b \in \text{actions} \wedge$

(\* – mutex synchronization – \*)

$$(\text{is\_unlock } a \wedge \text{is\_lock } b \wedge a \xrightarrow{\text{sc}} b) \vee$$

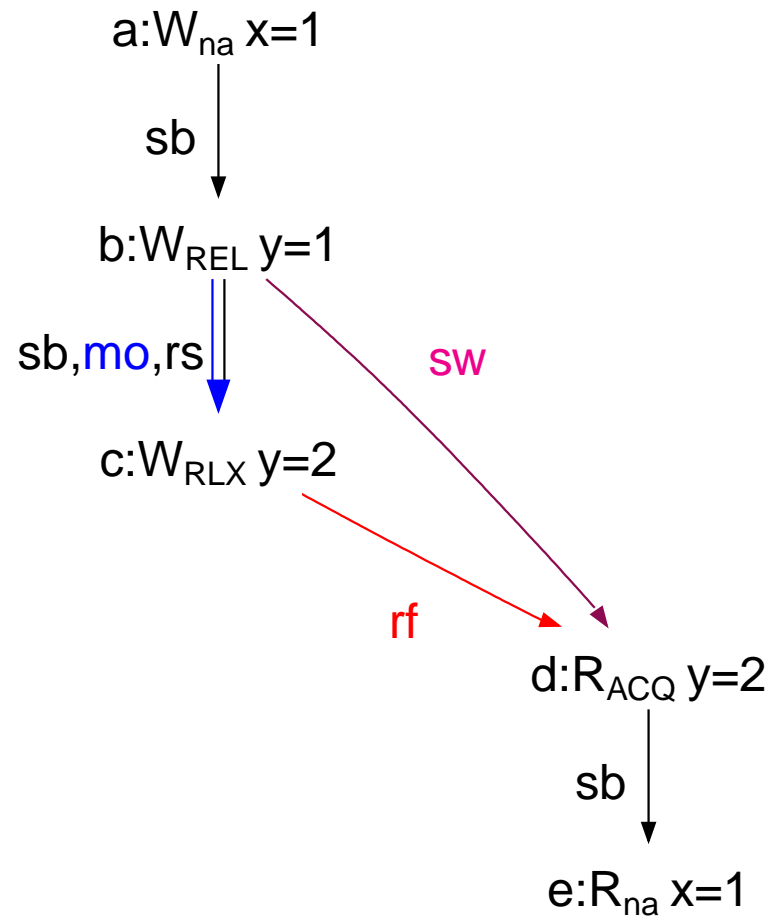
(\* – release/acquire synchronization – \*)

(is\_release  $a \wedge \text{is\_acquire } b \wedge \neg \text{same\_thread } a\ b \wedge$

$$(\exists c. a \xrightarrow{\text{release-sequence}} c \xrightarrow{\text{rf}} b)) \vee$$

[...]))

# Release-acquire synchronization



../examples/t8a.c

# Happens-before (without consume)

$$\frac{\textit{simple\_happens\_before}}{\longrightarrow} = \left( \frac{\textit{sequenced\_before}}{\longrightarrow} \cup \frac{\textit{synchronizes\_with}}{\longrightarrow} \right)^+$$

$$\textit{consistent\_simple\_happens\_before} = \textit{irreflexive} \left( \frac{\textit{simple\_happens\_before}}{\longrightarrow} \right)$$

# Happens-before

$$\begin{aligned}
 \xrightarrow{\text{inter-thread-happens-before}} &= \\
 \text{let } r &= \xrightarrow{\text{synchronizes-with}} \cup \\
 &\quad \xrightarrow{\text{dependency-ordered-before}} \cup \\
 &\quad \left( \xrightarrow{\text{synchronizes-with}} \circ \xrightarrow{\text{sequenced-before}} \right) \text{ in} \\
 \left( \xrightarrow{r} \cup \left( \xrightarrow{\text{sequenced-before}} \circ \xrightarrow{r} \right) \right)^+
 \end{aligned}$$

$$\begin{aligned}
 \text{consistent\_inter\_thread\_happens\_before} &= \\
 \text{irreflexive} \left( \xrightarrow{\text{inter-thread-happens-before}} \right)
 \end{aligned}$$

$$\begin{aligned}
 \xrightarrow{\text{happens-before}} &= \\
 \xrightarrow{\text{sequenced-before}} \cup &\quad \xrightarrow{\text{inter-thread-happens-before}}
 \end{aligned}$$

# Visible side effect

Non-atomic reads read from one of their visible side effects


$$\begin{aligned} a \xrightarrow{\text{visible-side-effect}} b = & \\ a \xrightarrow{\text{happens-before}} b \wedge & \\ \text{is\_write } a \wedge \text{is\_read } b \wedge \text{same\_location } a \ b \wedge & \\ \neg(\exists c. (c \neq a) \wedge (c \neq b) \wedge & \\ \text{is\_write } c \wedge \text{same\_location } c \ b \wedge & \\ a \xrightarrow{\text{happens-before}} c \xrightarrow{\text{happens-before}} b) & \end{aligned}$$

# Visible sequence of side effects

Atomic reads read from a write in one of their visible sequences of side effects.

$$\begin{aligned} \text{visible\_sequence\_of\_side\_effects\_tail } vsse\_head \ b = \\ \{ & c. vsse\_head \xrightarrow{\text{modification-order}} c \wedge \\ & \neg(b \xrightarrow{\text{happens-before}} c) \wedge \\ & (\forall a. vsse\_head \xrightarrow{\text{modification-order}} a \xrightarrow{\text{modification-order}} c \\ & \implies \neg(b \xrightarrow{\text{happens-before}} a)) \} \end{aligned}$$

# An atomic read



../examples/t8a.c




# Consistent reads-from mapping

```


consistent_reads_from_mapping =
  (∀b. (is_read b ∧ is_at_non_atomic_location b) ⇒
    (if (∃avse. avse  $\xrightarrow{\text{visible-side-effect}}$  b)
      then (∃avse. avse  $\xrightarrow{\text{visible-side-effect}}$  b ∧ avse  $\xrightarrow{\text{rf}}$  b)
      else ¬(∃a. a  $\xrightarrow{\text{rf}}$  b))) ∧
  (∀b. (is_read b ∧ is_at_atomic_location b) ⇒
    (if (∃(b', vsse) ∈ visible-sequences-of-side-effects. (b' = b))
      then (∃(b', vsse) ∈ visible-sequences-of-side-effects.
        (b' = b) ∧ (∃c ∈ vsse. c  $\xrightarrow{\text{rf}}$  b))
      else ¬(∃a. a  $\xrightarrow{\text{rf}}$  b))) ∧
  (∀(x, a) ∈  $\xrightarrow{\text{rf}}$ .
    ∀(y, b) ∈  $\xrightarrow{\text{rf}}$ .
      a  $\xrightarrow{\text{happens-before}}$  b ∧
      same_location a b ∧ is_at_atomic_location b
      ⇒ (x = y) ∨ x  $\xrightarrow{\text{modification-order}}$  y) ∧
  (∀(a, b) ∈  $\xrightarrow{\text{rf}}$ . is_atomic_rmw b
    ⇒ a  $\mid$   $\xrightarrow{\text{modification-order}}$  b) ∧
  (∀(a, b) ∈  $\xrightarrow{\text{rf}}$ . is_seq_cst b
    ⇒ ¬is_seq_cst a ∨
      a  $\xrightarrow{\text{sc}}_{\lambda c. \text{is\_write } c \wedge \text{same\_location } b \ c} b$ ) ∧
  [...]
  
```

# Coherence


Coherence is defined as absence of four execution fragments:




../examples/coherence-axiom-1.exc



../examples/coherence-axiom-2.exc



../examples/coherence-axiom-4.exc



../examples/coherence-axiom-3.exc

# Concurrency examples that can be observed

The model allows the following non-SC behaviour:

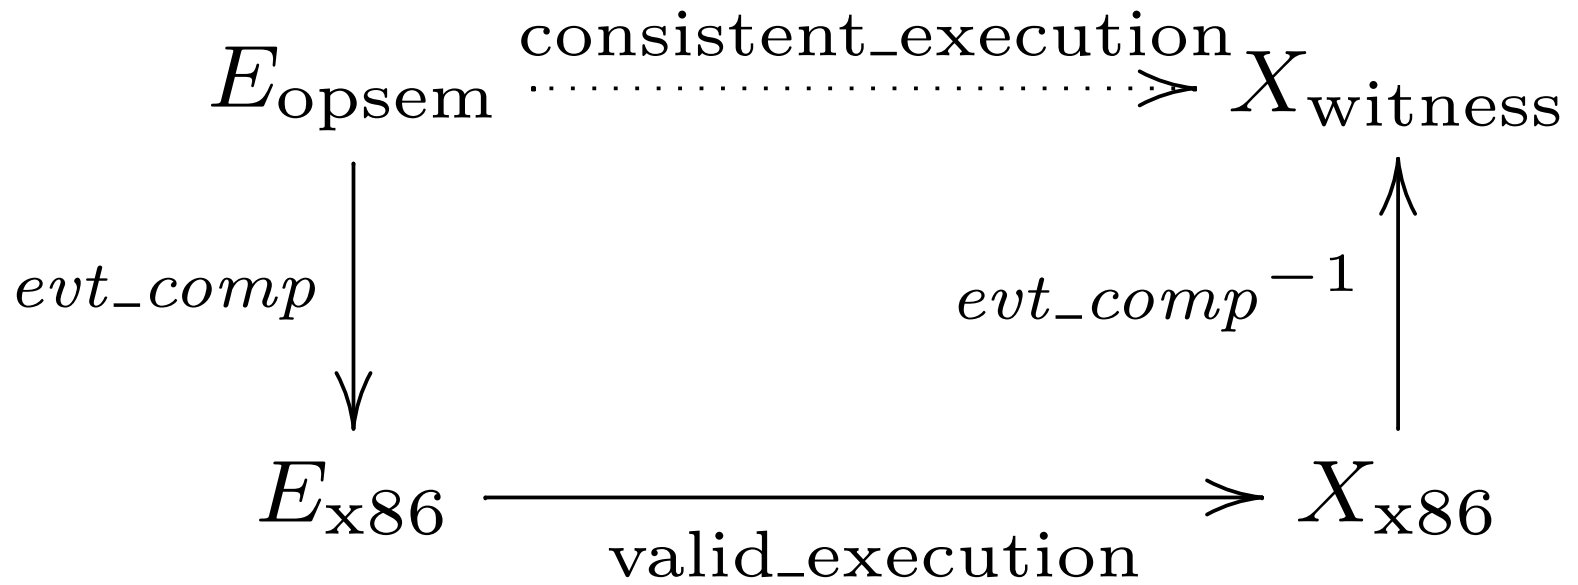
- message passing (RLX, REL-CON)
- store buffering (REL-ACQ, RLX, REL-CON)
- load buffering (RLX, CON)
- write-to-read causality (RLX, CON)
- IRIW (REL-ACQ, RLX, REL-CON)

...but DRF programs that use only the `memory_order_seq_cst` atomics should be sequentially consistent

# An execution compiler

Operation	x86 Implementation	
Load non-SC	mov	
Load Seq_cst	lock xadd(0)	OR: mfence, mov
Store non-SC	mov	
Store Seq_cst	lock xchg	OR: mov , mfence
Fence non-SC	no-op	
Fence Seq_cst	mfence	

# Theorem



# Conclusion

C++0x offers a simple model to normal programmers while experts get a highly configurable language that abstracts the hardware memory model

we have arrived just in time to point out a few bugs, and many changes have been made as a result of our work

the intricacy of such models makes tools important, CPPMEM helps in exploring and understanding the model

formal models provide an opportunity to provide guarantees about programs based on the specification, like our compiler correctness result