



Shared memory: an elusive abstraction

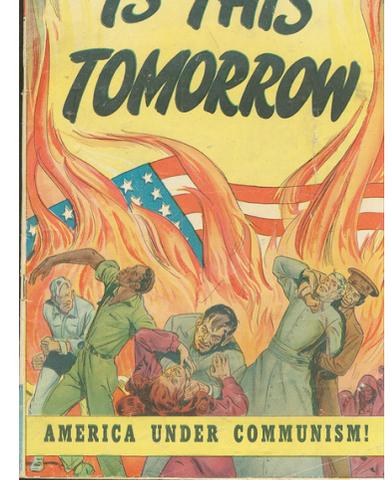
Francesco Zappa Nardelli

Inria Paris

<http://www.di.ens.fr/~zappa/projects/weakmemory>

Based on work done by or with

Peter Sewell, Jaroslav Ševčík, Susmit Sarkar, Tom Ridge, Scott Owens, Viktor Vafeiadis, Magnus O. Myreen, Kayvan Memarian, Luc Maranget, Derek Williams, Pankaj Pawan, Thomas Braibant, Mark Batty, Jade Alglave.



High-level languages, compilers, multiprocessors... an elusive mix?

Francesco Zappa Nardelli

Inria Paris

<http://www.di.ens.fr/~zappa/projects/weakmemory>

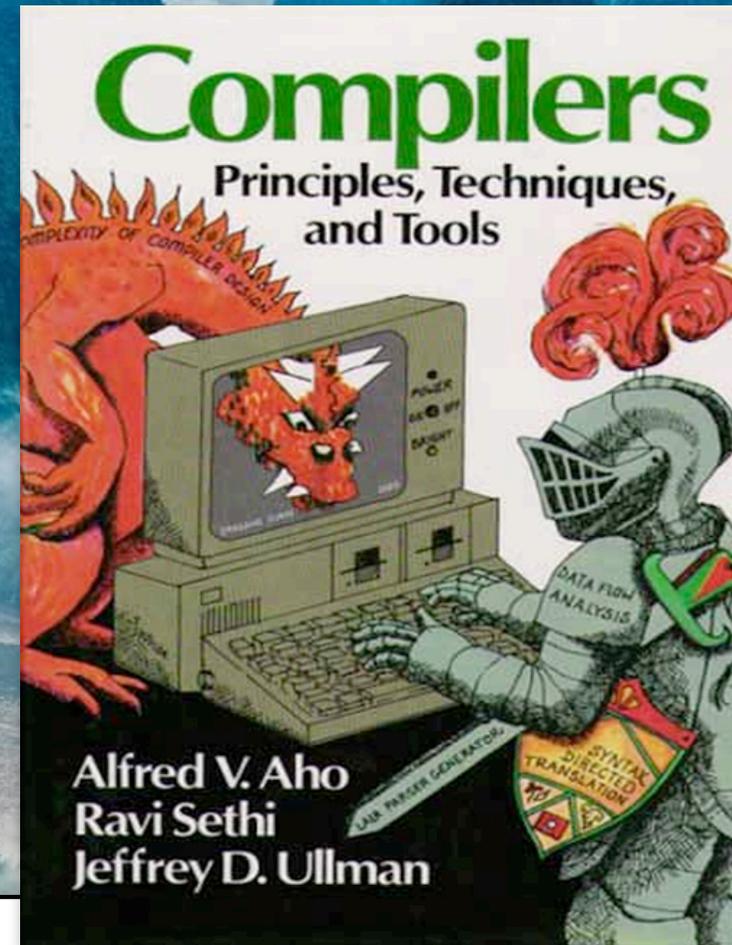
Based on work done by or with

Peter Sewell, Jaroslav Ševčík, Susmit Sarkar, Tom Ridge, Scott Owens,
Viktor Vafeiadis, Magnus O. Myreen, Kayvan Memarian, Luc Maranget,
Derek Williams, Pankaj Pawan, Thomas Braibant, Mark Batty, Jade Alglave.

Imagine an ideal world



Imagine an ideal world



Programmers and compilers cooperate
to make great software

Constant propagation

A simple, and *innocuous*, optimisation:

Source code

$$\begin{aligned}x &= 14 \\ y &= 7 - x / 2\end{aligned}$$


Optimised code

$$\begin{aligned}x &= 14 \\ y &= 7 - 14 / 2\end{aligned} \quad \longrightarrow \quad \begin{aligned}x &= 14 \\ y &= 0\end{aligned}$$

Shared memory concurrency

Shared memory

`x = y = 0`

Thread 1

```
x = 1
if (y == 1)
    print x
```

```
if (x == 1) {
    x = 0
    y = 1 }
```

Thread 2

Shared memory concurrency

Shared memory

`x = y = 0`

Thread 1

```
x = 1
if (y == 1)
    print x
```



```
if (x == 1) {
    x = 0
    y = 1 }
```

Thread 2

Intuitively this program always prints 0

Shared memory concurrency

But if the compiler propagates the *constant* $x = 1$...

$x = y = 0$

Thread 1

```
x = 1
if (y == 1)
    print x
```

Shared memory concurrency

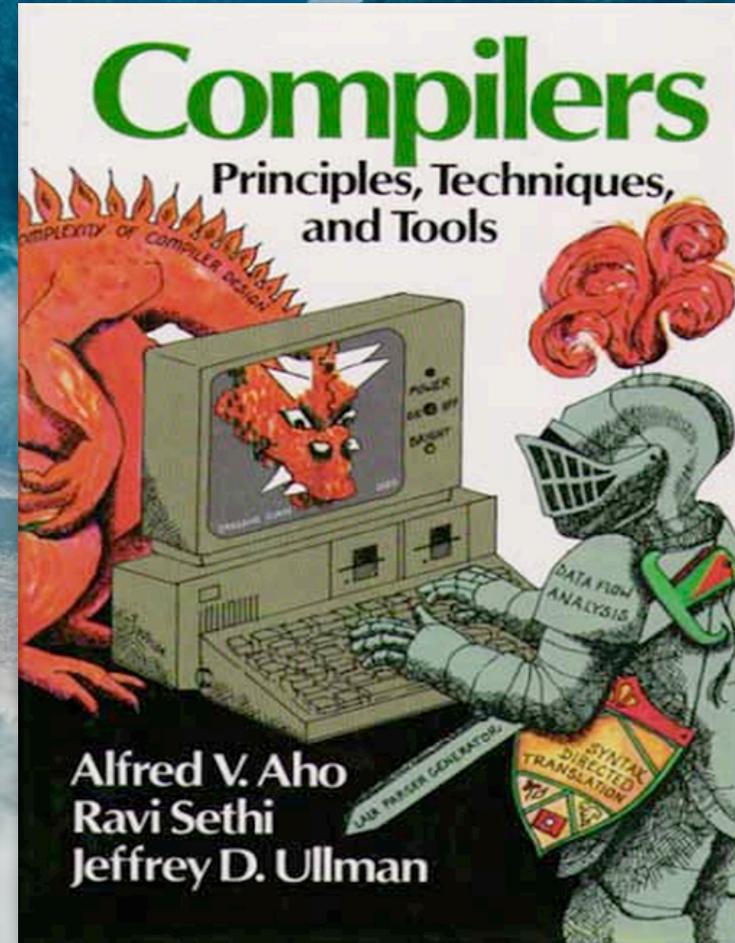
But if the compiler propagates the *constant* $x = 1$...

$x = y = 0$

	<pre>x = 1</pre>		<pre>if (x == 1) {</pre>	
Thread 1	<pre> if (y == 1)</pre>		<pre> x = 0</pre>	Thread 2
	<pre> print x</pre>		<pre> y = 1 }</pre>	
	<pre> print 1</pre>			

...the program always writes 1 rather than 0.

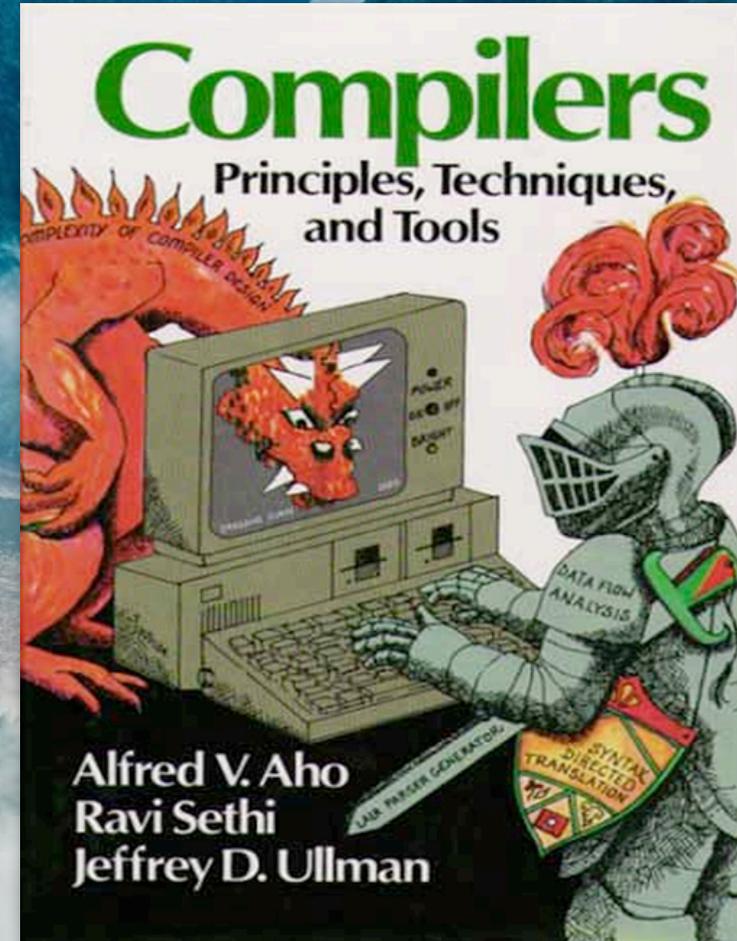
The fundamental problem



The fundamental problem



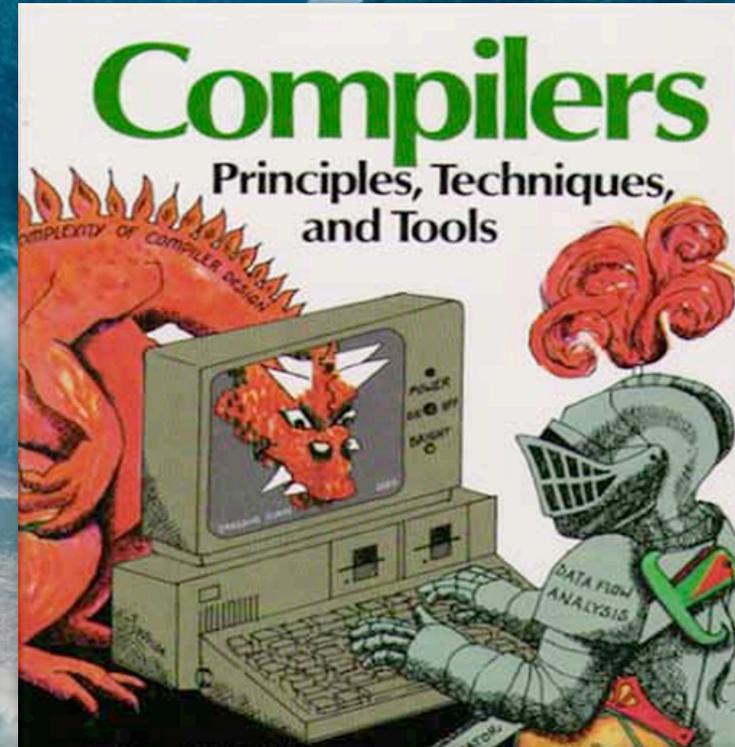
The programmer wants to understand the code he writes



The fundamental problem



The programmer wants to understand the code he writes



The compiler (and the hardware) try hard to optimise it

Background: lock and unlock

- Suppose that two threads increment a shared memory location:

$x = 0$

```
tmp1 = *x;  
*x = tmp1 + 1;
```

```
tmp2 = *x;  
*x = tmp2 + 1;
```

- If both threads read 0, (even in an ideal world) $x == 1$ is possible:

```
tmp1 = *x; tmp2 = *x; *x = tmp1 + 1; *x = tmp2 + 1
```

Background: lock and unlock

- **Lock** and **unlock** are primitives that prevent the two threads from interleaving their actions.

`x = 0`

<pre>lock(); tmp1 = *x; *x = tmp1 + 1; unlock();</pre>	<pre>lock(); tmp2 = *x; *x = tmp2 + 1; unlock();</pre>
--	--

- In this case, the interleaving below is forbidden, and we are guaranteed that `x == 2` at the end of the execution.

FORBIDDEN

```
tmp1 = *x;
```

```
tmp2 = *x;
```

```
*x = tmp1 + 1;
```

```
*x = tmp2 + 1
```

Lazy initialisation (an unoptimising compiler breaks your program)

Deferring an object's initialisation until first use: a big win if an object is never used (e.g. device drivers code). Compare:

```
int x = computeInitValue();    // eager initialization
...                            // clients refer to x
```

with:

```
int xValue() {
    static int x = computeInitValue(); // lazy initialization
    return x;
} ...                                // clients refer to xValue()
```

The singleton pattern

Lazy initialisation is a pattern commonly used. In C++ you would write:

```
class Singleton {
public:
    static Singleton *instance (void) {
        if (instance_ == NULL)
            instance_ = new Singleton;
        return instance_;
    }
... // other methods omitted
private:
    static Singleton *instance_; // other fields omitted
};

...
Singleton::instance () -> method ();
```

But this code is not thread safe! Why?

Making the singleton pattern thread safe

A simple thread safe version:

```
class Singleton {
public:
    static Singleton *instance (void) {
        Guard<Mutex> guard (lock_); // only one thread at a time
        if (instance_ == NULL)
            instance_ = new Singleton;
        return instance_;
    }
private:
    static Mutex lock_;
    static Singleton *instance_;
};
```

Every call to instance must acquire and release the lock: excessive overhead.

Obvious (broken) optimisation

```
class Singleton {
public:
    static Singleton *instance (void) {
        if (instance_ == NULL) {
            Guard<Mutex> guard (lock_); // lock only if unitialised
            instance_ = new Singleton; }
        return instance_;
    }

private:
    static Mutex lock_;
    static Singleton *instance_;
};
```

Exercise: why is it broken?

Clever programmers use double-check locking

```
class Singleton {
public:
    static Singleton *instance (void) {
        // First check
        if (instance_ == NULL) {
            // Ensure serialization
            Guard<Mutex> guard (lock_);
            // Double check
            if (instance_ == NULL)
                instance_ = new Singleton;
        }
        return instance_;
    }
private: [..]
};
```

Idea: re-check that the Singleton has not been created after acquiring the lock.

Double-check locking: clever but broken

The instruction

```
instance_ = new Singleton;
```

does three things:

- 1) allocate memory
- 2) construct the object
- 3) assign to `instance_` the address of the memory

Not necessarily in this order! For example:

```
instance_ = // 3  
    operator new(sizeof(Singleton)); // 1  
new (instance_) Singleton // 2
```

If this code is generated, the order is 1,3,2.

Broken...

```
if (instance_ == NULL) { // Line 1
    Guard<Mutex> guard (lock_);
    if (instance_ == NULL) {
        instance_ =
            operator new(sizeof(Singleton)); // Line 2
        new (instance_) Singleton; }}
```

Thread 1:

executes through Line 2 and is suspended; at this point, `instance_` is non-NULL, but no singleton has been constructed.

Thread 2:

executes Line 1, sees `instance_` as non-NULL, returns, and dereferences the pointer returned by `Singleton` (i.e., `instance_`).

Thread 2 attempts to reference an object that is not there yet!

The fundamental problem

Problem: You need a way to specify that step 3 come after steps 1 and 2.

There is no way to specify this in C++

Similar examples can be built for any programming language...

That pesky hardware (1)

Consider misaligned 4-byte accesses

```
int32_t a = 0
```

<pre>a = 0x44332211</pre>	<pre>if (a == 0x00002211) print "error"</pre>
---------------------------	---

(Disclaimer: compiler will normally ensure alignment)

Intel SDM x86 atomic accesses:

- n -bytes on an n -byte boundary ($n = 1, 2, 4, 16$)
- P6 or later: ... or if unaligned but within a cache line

Question: what about multi-word high-level language values?

That pesky hardware (1)

Consider misaligned 4-byte accesses

```
int32_t a = 0
```

<pre>a = 0x44332211</pre>	<pre>if (a == 0x00002211) print "error"</pre>
---------------------------	---

(Disclaimer)

Intel SDM

- *n*-byte

- P6 or I

Question: what about multi-word high-level language values?

This is called a *out-of-thin air read*:
the program reads a value
that the programmer never wrote.

That pesky hardware (2)

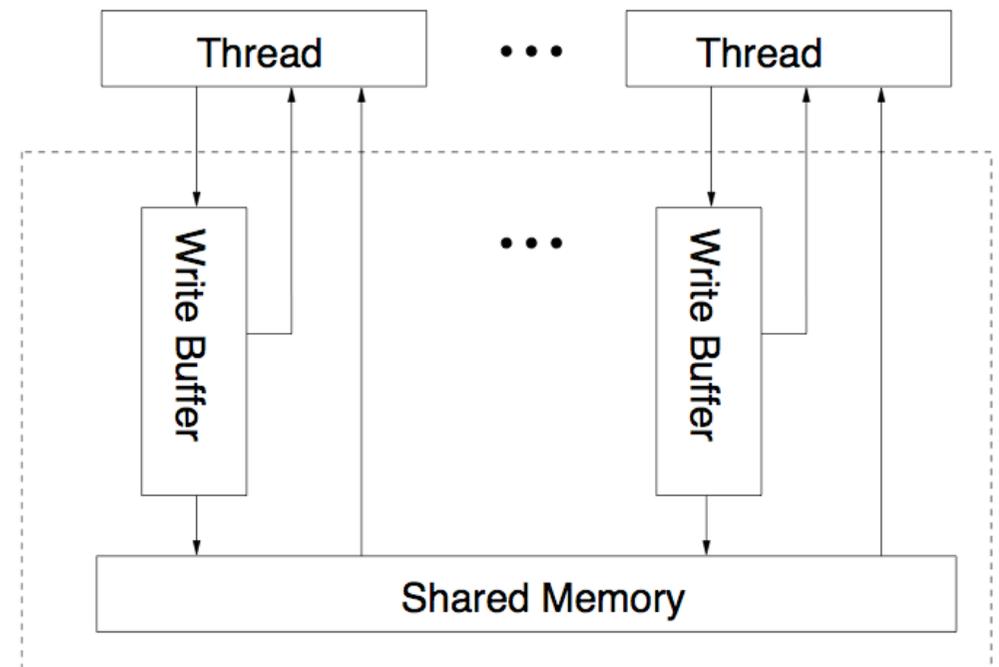
Hardware optimisations can be observed by concurrent code:

Thread 0	Thread 1
<code>x = 1</code> <code>print y</code>	<code>y = 1</code> <code>print x</code>

At the end of some executions:

0 0

is printed on the screen,
both on x86 and Power/ARM).



That pesky hardware (2)

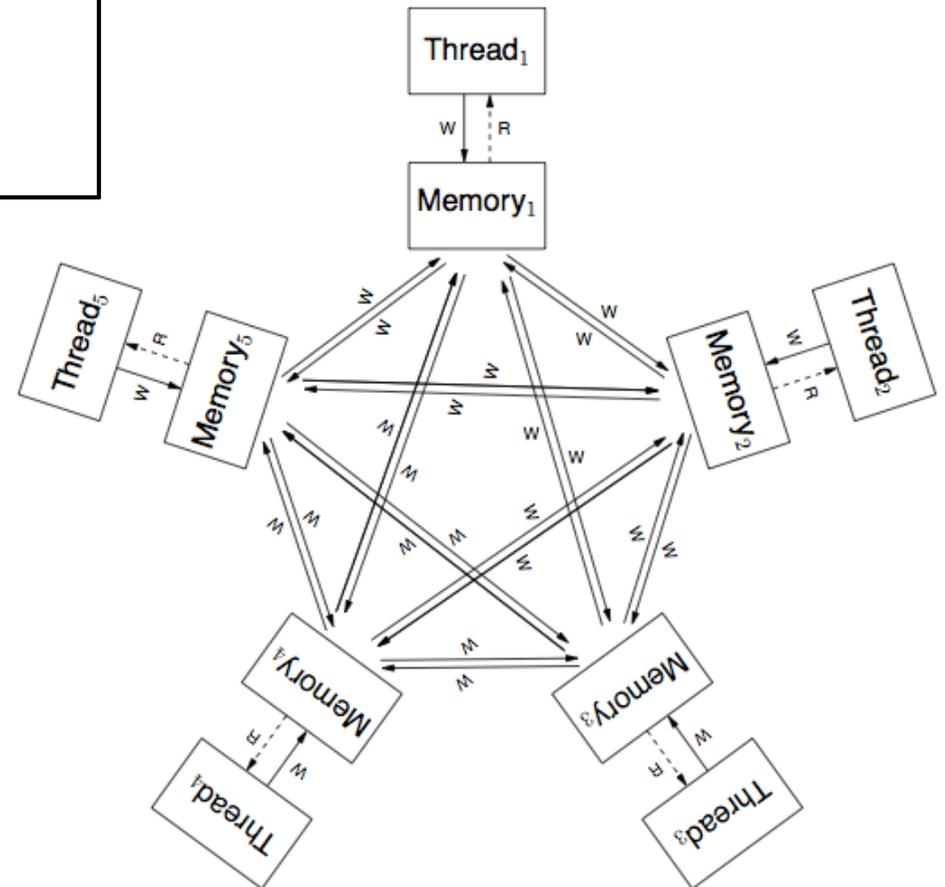
...and differ between architectures...

Thread 0	Thread 1
<code>x = 1</code>	<code>print y</code>
<code>y = 1</code>	<code>print x</code>

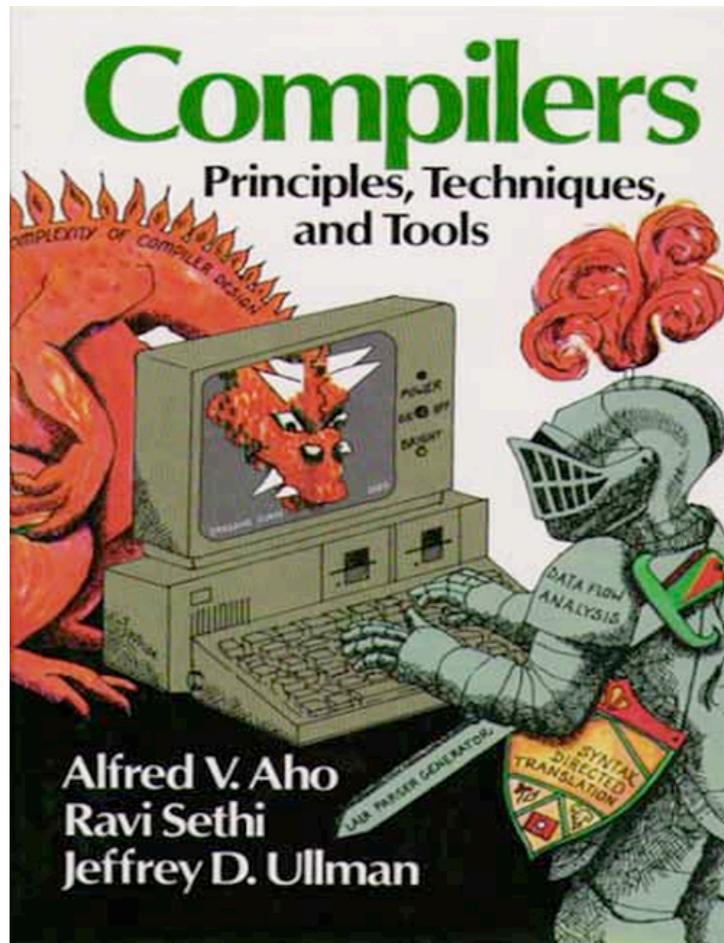
At the end of some executions:

1 0

is printed on the screen on Power/ARM
but not on x86.



Compilers vs. programmers



Compilers vs. programmers

Tension:

- the programmer wants to understand the code he writes
- the compiler and the hardware want to optimise it.

Which are the valid optimisations that the compiler or the hardware can perform without breaking the expected semantics of a concurrent program?

Which is the semantics of a concurrent program?

This lecture

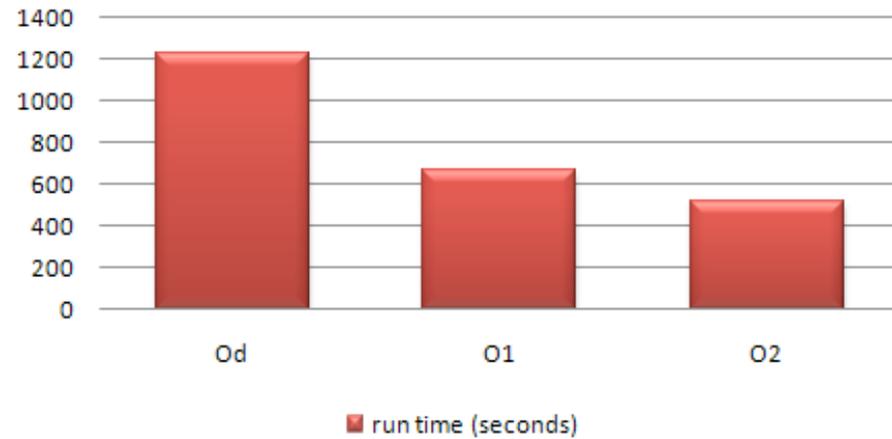
Programming language models

- 1) defining the semantics of a concurrent programming language
- 2) data-race freedom
- 3) soundness of compiler optimisations

Previous lecture: hardware models

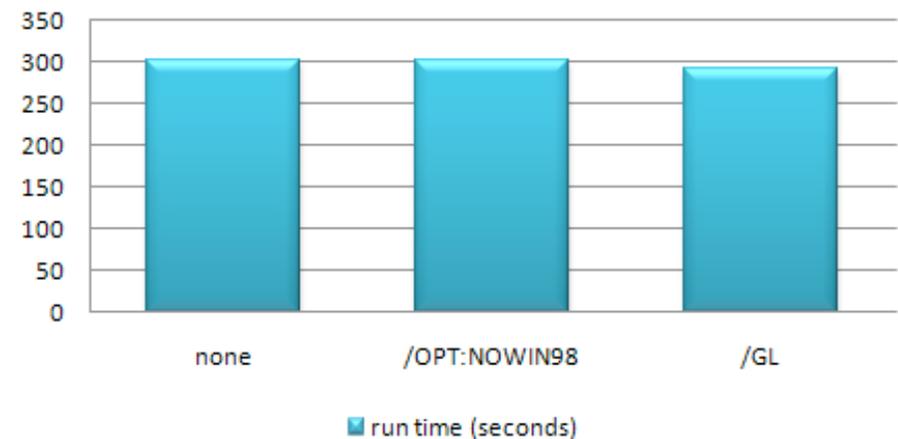
- 1) why are industrial specs so often flawed?
focus on x86, with a glimpse of Power/ARM
- 2) usable models: x86-TSO, PowerARM

effect of VS2005 compiler optimisations on speed



A brief tour of compiler optimisations

effect of additional VS2005 optimisations on speed



World of optimisations

A typical compiler performs many optimisations.

[gcc 4.4.1](#) with `-O2` option goes through [147](#) compilation passes.

computed using `-fdump-tree-all` and `-fdump-rtl-all`

[Sun Hotspot Server JVM](#) has 18 high-level passes with each pass composed of one or more smaller passes.

<http://www.azulsystems.com/blog/cliff-click/2009-04-14-odds-ends>

World of optimisations

A typical compiler performs many optimisations.

- Common subexpression elimination
(copy propagation, partial redundancy elimination, value numbering)
- (conditional) constant propagation
- dead code elimination
- loop optimisations
(loop invariant code motion, loop splitting/peeling, loop unrolling, etc.)
- vectorisation
- peephole optimisations
- tail duplication removal
- building graph representations/graph linearisation
- register allocation
- call inlining
- local memory to registers promotion
- spilling
- instruction scheduling

World of optimisations

However only some optimisations change shared-memory traces:

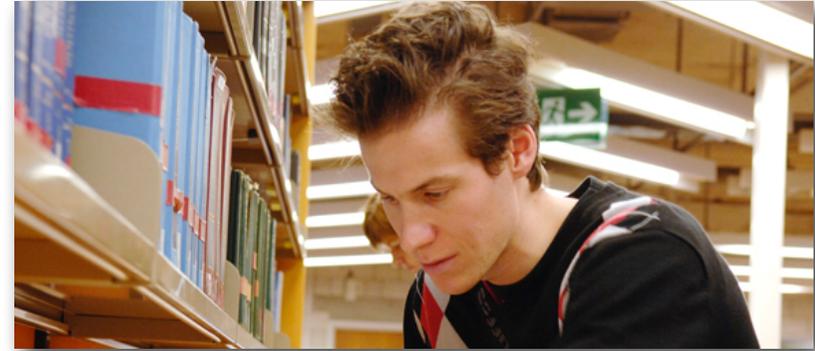
- *Common subexpression elimination*
(*copy propagation, partial redundancy elimination, value numbering*)
- (conditional) constant propagation
- dead code elimination
- loop optimisations
(*loop invariant code motion*, loop splitting/peeling, loop unrolling, etc.)
- vectorisation
- *peephole optimisations*
- tail duplication removal
- building graph representations/graph linearisation
- register allocation
- call inlining
- *local memory to registers promotion*
- *spilling*
- instruction scheduling

What is an optimisation?

Compiler Writer



Semanticist



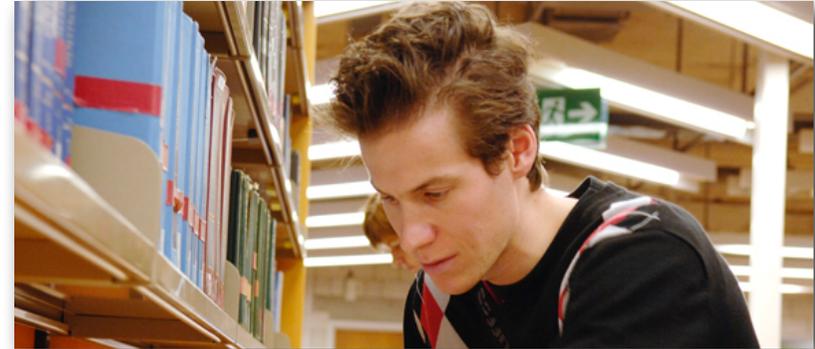
What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

Semanticist



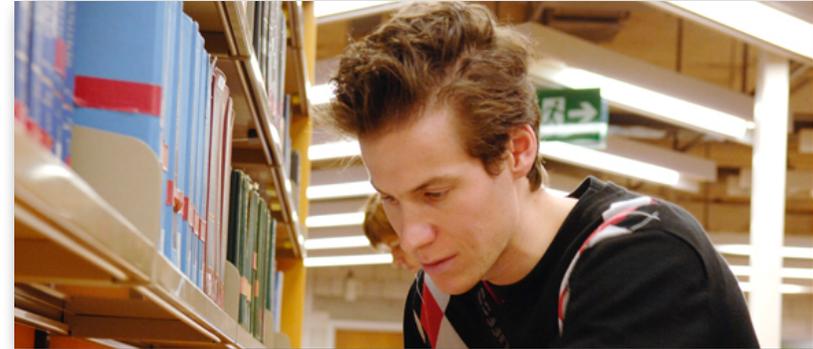
What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

Semanticist



```
for (int i=0; i<2; i++) {  
    z = i;  
    x[i] += y+1 ;  
}
```

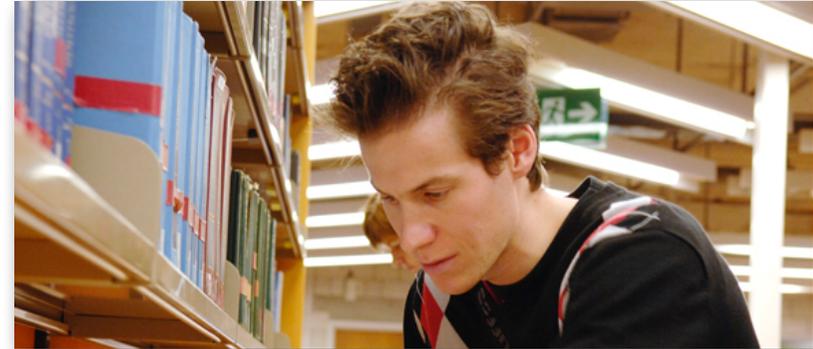
What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

Semanticist



```
tmp = y+1 ;  
for (int i=0; i<2; i++) {  
    z = i;  
    x[i] += tmp ;  
}
```

What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

Semanticist



Elimination of run-time events
Reordering of run-time events
Introduction of run-time events
Operations on sets of events

```
tmp = y+1 ;  
for (int i=0; i<2; i++) {  
    z = i;  
    x[i] += tmp ;  
}
```

What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

```
tmp = y+1 ;  
for (int i=0; i<2; i++) {  
    z = i;  
    x[i] += tmp ;  
}
```

Semanticist



Elimination of run-time events
Reordering of run-time events
Introduction of run-time events
Operations on sets of events

```
Store z 0  
Load y 42  
Store x[0] 43  
Store z 1  
Load y 42  
Store x[1] 43
```

What is an optimisation?

Compiler Writer



Sophisticated program analyses
Fancy algorithms
Source code or IR
Operations on AST

```
tmp = y+1 ;  
for (int i=0; i<2; i++) {  
    z = i;  
    x[i] += tmp ;  
}
```

Semanticist



Elimination of run-time events
Reordering of run-time events
Introduction of run-time events
Operations on sets of events

```
Load y 42  
Store z 0  
  
Store x[0] 43  
Store z 1  
  
Store x[1] 43
```

Eliminations

This includes common subexpression elimination, dead read elimination, overwritten write elimination, redundant write elimination.

Irrelevant read elimination:

$$r=*x; C \rightarrow C$$

where r is not free in C .

Redundant read after read elimination:

$$r1=*x; r2=*x \rightarrow r1=*x; r2=r1$$

Redundant read after write elimination:

$$*x=r1; r2=*x \rightarrow *x=r1; r2=r1$$

Reordering

Common subexpression elimination, some loop optimisations, code motion.

Normal memory access reordering:

```
r1=*x; r2=*y → r2=*y; r1=*x
*x=r1; *y=r2 → *y=r2; *x=r1
r1=*x; *y=r2 ⇔ *y=r2; r1=*x
```

Roach motel reordering:

```
memop; lock m → lock m; memop
unlock m; memop → memop; unlock m
where memop is *x=r1 or r1=*x
```

Memory access introduction

Can an optimisation introduce memory accesses?

Yes, but rarely:

```

i = 0;
...
while (i != 0) {
    j = *x + 1;
    i = i-1 }

```

→

```

i = 0;
...
tmp = *x;
while (i != 0) {
    j = tmp + 1;
    i = i-1 }

```

Note that the loop body is not executed.

Memory access introduction

Back to our question now:

Which is the semantics of a concurrent program?

Note that the loop body is not executed.

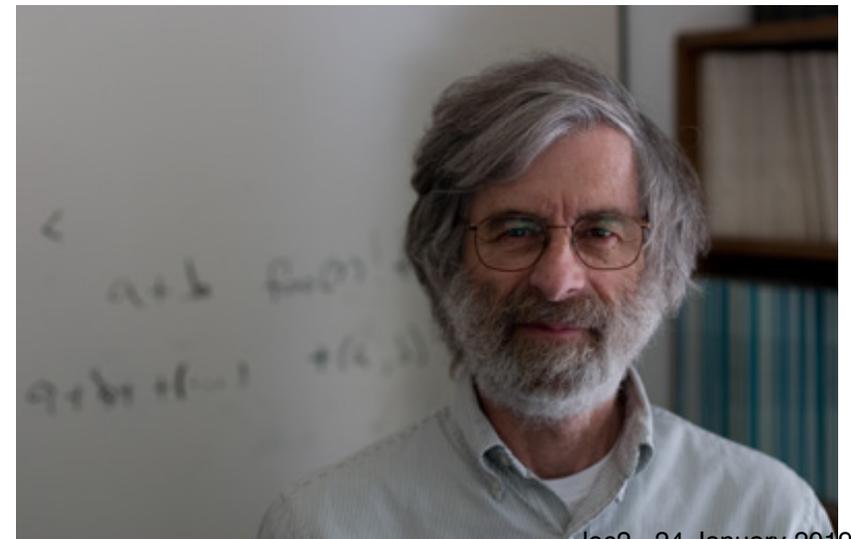
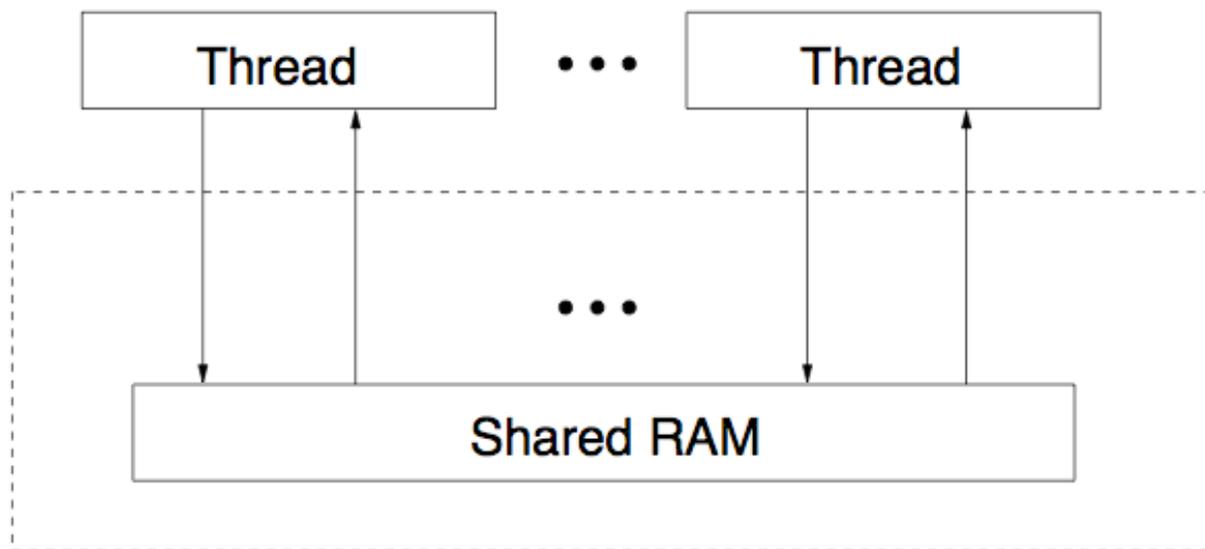
Naive answer: enforce sequential consistency

Sequential consistency

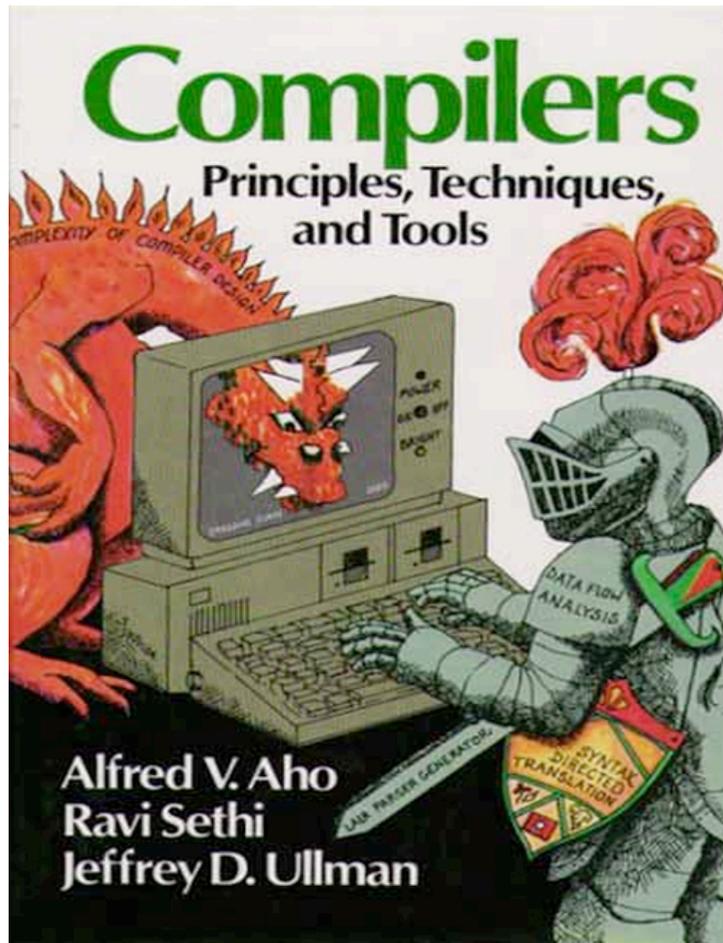
Multiprocessors have a *sequentially consistent* shared memory when:

...the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program...

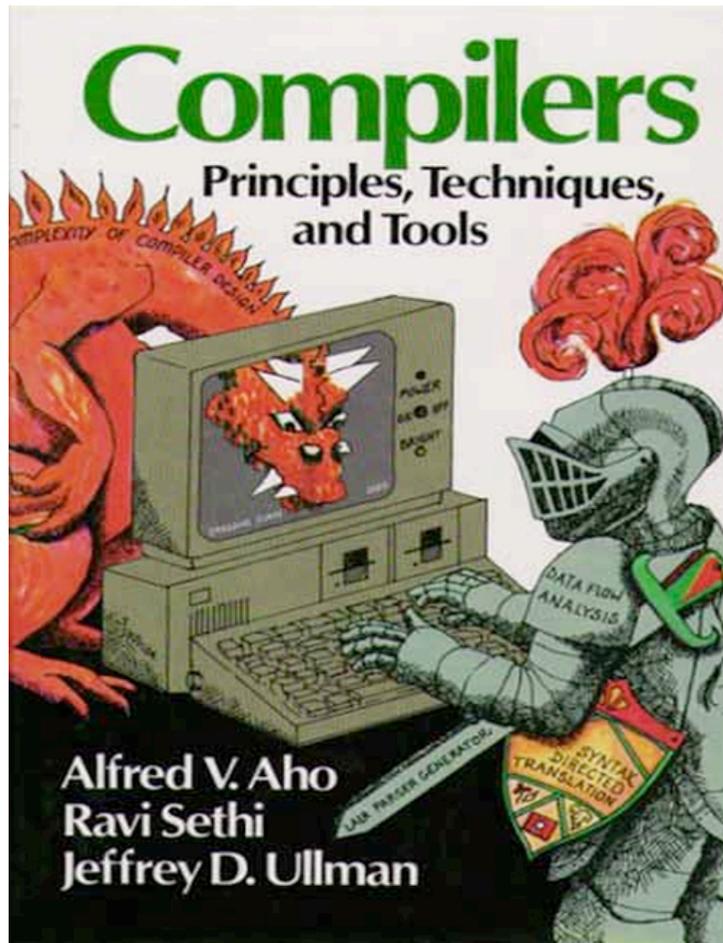
Lamport, 1979.



Compilers, programmers & sequential consistency

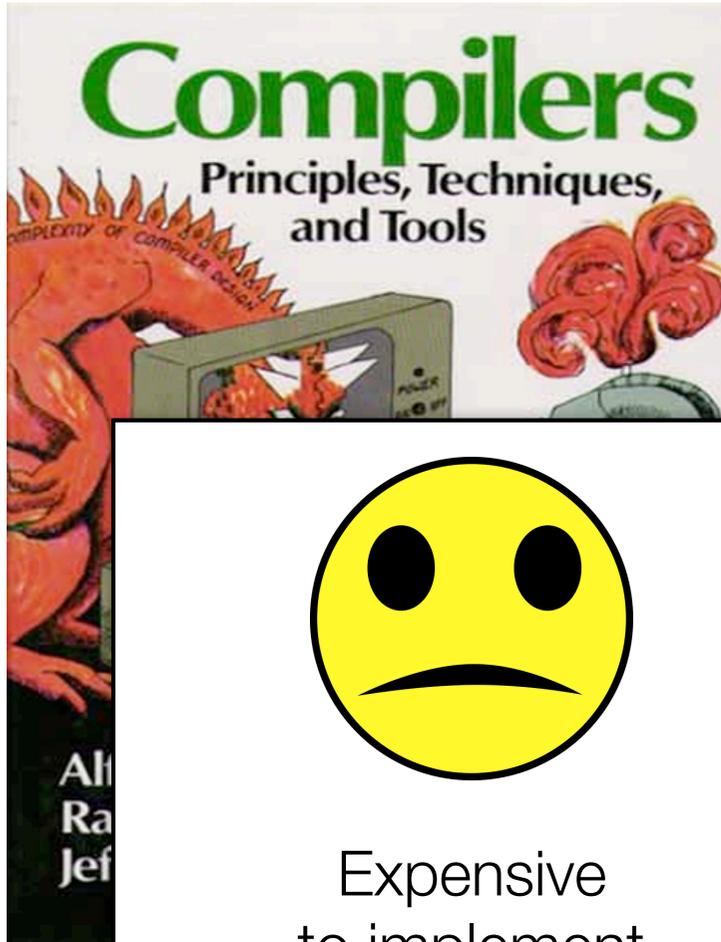


Compilers, programmers & sequential consistency



Simple and intuitive programming model

Compilers, programmers & sequential consistency



Expensive
to implement



Simple and intuitive
programming model

A Case for an SC-Preserving Compiler

Daniel Marino[†] Abhayendra Singh* Todd Millstein[†] Madanlal Musuvathi[‡] Satish Narayanasamy*

[†]University of California, Los Angeles

*University of Michigan, Ann Arbor

[‡]Microsoft Research, Redmond

An SC-preserving compiler, obtained by restricting the optimization phases in LLVM, a state-of-the-art C/C++ compiler, incurs an average slowdown of 3.8% and a maximum slowdown of 34% on a set of 30 programs from the SPLASH-2, PARSEC, and SPEC CINT2006 benchmark suites.

And this study supposes that the hardware is SC.



Expensive
to implement

SC and hardware

The compiler must insert enough synchronising instructions to prevent hardware reorderings. On x86 we have:

- **MFENCE**
flush the local write buffer
- **LOCK prefix (e.g. CMPXCHG)**
flush the local write buffer
globally lock the memory

Initial: $[x]=0 \wedge [y]=0$	
proc 0	proc 1
MOV $[x] \leftarrow \$1$	MOV $[y] \leftarrow \$1$
MFENCE	MFENCE
MOV $EAX \leftarrow [y]$	MOV $EBX \leftarrow [x]$
Forbid: $EAX=0 \wedge EBX=0$	

Initially, $[100] = 0$
At the end, $[100] = 2$

proc:0	proc:1
LOCK; INC $[100]$	LOCK; INC $[100]$

These consumes hundreds of cycles... ideally should be avoided.

Naively recovering SC on x86 incurs in a ~40% overhead.

A Case for an SC-Preserving Compiler

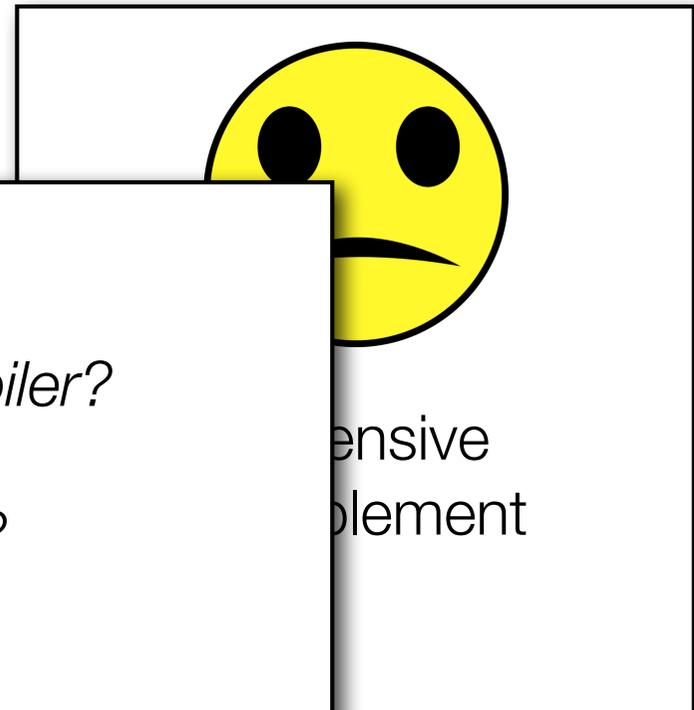
Daniel Marino[†] Abhayendra Singh* Todd Millstein[†] Madanlal Musuvathi[‡] Satish Narayanasamy*

[†]University of California, Los Angeles

^{*}University of Michigan, Ann Arbor

[‡]Microsoft Research, Redmond

An SC-preserving compiler, obtained by restricting the optimization phases in LLVM, a state-of-the-art C/C++ compiler, incurs an average slowdown of 3.8% and a maximum slowdown of 6.24% on a set of 6.20 programs. The average speedup is 1.01 and the maximum speedup is 1.01. The average SPE is 1.01 and the maximum SPE is 1.01.



What is an SC-preserving compiler?

When is a compiler correct?

And this st

When is a compiler correct?

A compiler is correct if any behaviour of the compiled program could be exhibited by the original program.

i.e. for any execution of the compiled program, there is an execution of the source program with the *same observable behaviour*.

Intuition: we represent programs as sets of memory action traces, where the trace is a sequence of memory actions of a single thread.

Intuition: the observable behaviour of an execution is the subtrace of external actions.

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Is the transformation from P1 to P2 correct (in an SC semantics)?

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Executions of P1:

$W_{t_1} x=1, R_{t_2} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 2$
 $R_{t_2} x=0, R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Executions of P1:

$W_{t_1} x=1, R_{t_2} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 2$
 $R_{t_2} x=0, R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Executions of P2:

$W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Example

$P_1 = *x = 1$	$\left \begin{array}{l} r1 = *x; r2 = *x; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$
$P_2 = *x = 1$	$\left \begin{array}{l} r1 = *x; r2 = r1; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$

Executions of P1:

$W_{t_1} x=1, R_{t_2} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 2$
 $R_{t_2} x=0, R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Executions of P2:

$W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Behaviours of P1: $[P_{t_2} 1], [P_{t_2} 2]$

Behaviours of P2: $[P_{t_2} 1]$

Example

$P_1 = *x = 1$	<pre>r1 = *x; r2 = *x; if r1=r2 then print 1 else print 2</pre>
$P_2 = *x = 1$	<pre>r1 = *x; r2 = r1; if r1=r2 then print 1 else print 2</pre>

Executions of P1:

Executions of P2:

W_{t_1}
 R_{t_2}
 R_{t_2}
 R_{t_2}

It is correct to rewrite P1 into P2, but not the opposite!

Behaviours of P1: $[P_{t_2} 1], [P_{t_2} 2]$

Behaviours of P2: $[P_{t_2} 1]$

General CSE incorrect in SC

<pre>*x = 1; *y = 1; if *y = 2 then print *x</pre>	<pre>if *x=1 then (*x = 2; *y = 2)</pre>
--	--

There is only one execution with a printing behaviour:

$W_{t_1} x=1, W_{t_1} y=1, R_{t_2} x=1, W_{t_2} x=2, W_{t_2} y=2, R_{t_1} y=2, R_{t_1} x=2, P_{t_1} 2$

General CSE incorrect in SC

<pre>*x = 1; *y = 1; if *y = 2 then print *x</pre>	<pre>if *x=1 then (*x = 2; *y = 2)</pre>
--	--

But a compiler would optimise to:

<pre>*x = 1; *y = 1; if *y = 2 then print 1</pre>	<pre>if *x=1 then (*x = 2; *y = 2)</pre>
---	--

General CSE incorrect in SC

<pre>*x = 1; *y = 1; if *y = 2 then print 1</pre>		<pre>if *x=1 then (*x = 2; *y = 2)</pre>
---	--	--

The only execution with a printing behaviour in the optimised code is:

$W_{t_1} x=1, W_{t_1} y=1, R_{t_2} x=1, W_{t_2} x=2, W_{t_2} y=2, R_{t_1} y=2, P_{t_1} 1$

So the optimisation is not correct.

General CSE incorrect in SC

<code>*x = 1;</code>	<code>r = *x;</code>
<code>*y = 1;</code>	<code>print r;</code>

Our first example highlighted that CSE is incorrect in SC.

Here is another example.

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

General CSE incorrect in SC

<code>*x = 1;</code>		<code>r = *x;</code>
<code>*y = 1;</code>		<code>print r;</code>
		<code>print *y;</code>
		<code>print *x;</code>

The observable behaviours are (note that 0 - 1 - 0 is not observable):

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

General CSE incorrect in SC

<code>*x = 1;</code>	<code>r = *x;</code>
<code>*y = 1;</code>	<code>print r;</code>
	<code>print *y;</code>
	<code>print *x;</code>

But a compiler would optimise as:

<code>*x = 1;</code>	<code>r = *x;</code>
<code>*y = 1;</code>	<code>print r;</code>
	<code>print *y;</code>
	<code>print r;</code>

General CSE incorrect in SC

```
*x = 1;
*y = 1;

r = *x;
print r;
print *y;
print *x;
```

```
*x = 1;
*y = 1;

r = *x;
print r;
print *y;
print r;
```

Let's compare the behaviours of the two programs:

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 0]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

General CSE incorrect in SC

<code>*x = 1;</code>	<code>r = *x;</code>	<code>*x = 1;</code>	<code>r = *x;</code>
----------------------	----------------------	----------------------	----------------------

*

The optimised program exhibits a new, unexpected, behaviour.

Let's compare the behaviours of the two programs.

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 0]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

Reordering incorrect

<code>*x = 1;</code>		<code>*y = 1;</code>		<code>r1 = *y</code>		<code>*y = 1;</code>
<code>r1 = *y</code>		<code>r2 = *x;</code>	\Rightarrow	<code>*x = 1;</code>		<code>r2 = *x;</code>
<code>print r1</code>		<code>print r2</code>		<code>print r1</code>		<code>print r2</code>

Again, the optimised program exhibits a new behaviour:

$[P_{t_1} 0, P_{t_2} 1]$
 $[P_{t_1} 1, P_{t_2} 0]$
 $[P_{t_1} 1, P_{t_2} 1]$

$[P_{t_1} 0, P_{t_2} 1]$
 $[P_{t_1} 1, P_{t_2} 0]$
 $[P_{t_1} 1, P_{t_2} 1]$
 $[P_{t_1} 0, P_{t_2} 0]$

Elimination of adjacent accesses

There are some correct optimisations under SC. For example it is correct to rewrite:

`r1 = *x; r2 = *x` → `r1 = *x; r2 = r1`

The basic idea: whenever we perform the read `r1 = *x` in the optimised program, we perform *both* reads in the source program.

(More on this later)

Elimination of adjacent accesses

There are some correct optimisations under SC. For example it is correct to rewrite:

```
r1 = *x; r2 = *x    →    r1 = *x; r2 = r1
```

Can we define a model that:

- 1) enables more optimisations than SC, and
- 2) retains the simplicity of SC?

(more on this later)



The layman solution *forbid data-races*



Data-race freedom

Our examples again:

Thread 0	Thread 1
<code>*y = 1</code>	<code>if *x == 1</code>
<code>*x = 1</code>	<code>then print *y</code>

Observable behaviour: 0

- the problematic transformations (e.g. swapping the two writes in thread 0) **do not change the meaning of single-threaded programs;**
- the problematic transformations **are detectable** only by code that allows two threads to **access the same data simultaneously in conflicting ways** (e.g. one thread writes the data read by the other).

Data-race freedom

Our exam

- the prok
(e.g. sw
thread (

- the prok

allows two threads to **access the same data simultaneously in conflicting ways** (e.g. one thread writes the datas read by the other).

...intuition...

Programming languages provide
synchronisation mechanisms

if these are used (and implemented) correctly,
we might avoid the issues above...

Thread 1

```
x == 1
```

```
print *y
```

viour: 0

programs;

de that

The basic solution

Prohibit *data races*

Thread 0	Thread 1
<code>*y = 1</code> <code>*x = 1</code>	<code>if *x == 1</code> <code>then print *y</code>

Observable behaviour: 0

Defined as follows:

- two memory operations **conflict** if they access the same memory location and at least one is a store operation;
- a SC execution (interleaving) contains a data race if **two conflicting operations corresponding to different threads are adjacent** (maybe executed concurrently).

Example: a data race in the example above:

$W_{t_1} y=1, W_{t_1} x=1, R_{t_2} x=1, R_{t_2} y=1, P_{t_2} 1$

The basic solution

Prohibit *data races*

Thread 0	Thread 1
<code>*y = 1</code>	<code>if *x == 1</code>
<code>*x = 1</code>	<code>then print *y</code>

Observable behaviour: 0

Defined as follows:

- two memory locations
- a SC execution of two conflicting operations (maybe executed concurrently).

The definition of data race quantifies only over the sequential consistent executions

Example: a data race in the example above:

$$W_{t_1} y=1, W_{t_1} x=1, R_{t_2} x=1, R_{t_2} y=1, P_{t_2} 1$$

How do we avoid data races? (focus on high-level languages)

- Locks

No `lock(l)` can appear in the interleaving unless prior `lock(l)` and `unlock(l)` calls from other threads balance.

- Atomic variables

Allow concurrent access “exempt” from data races. Called `volatile` in Java.

Example:

Thread 0	Thread 1
<pre>*y = 1 lock(); *x = 1 unlock();</pre>	<pre>lock(); tmp = *x; unlock(); if tmp = 1 then print *y</pre>

How do we avoid data races? (focus on high-level languages)

Thread 0	Thread 1
<pre>*y = 1 lock(); *x = 1 unlock();</pre>	<pre>lock(); tmp = *x; unlock(); if tmp = 1 then print *y</pre>

This program is data-race free:

```
*y = 1; lock();*x = 1;unlock(); lock();tmp = *x;unlock(); if tmp=1 then print *y
```

```
*y = 1; lock(); tmp = *x; unlock(); lock(); *x = 1; unlock(); if tmp=1
```

```
*y = 1; lock(); tmp = *x; unlock(); if tmp=1; lock(); *x = 1; unlock();
```

```
lock();tmp = *x;unlock(); *y = 1; lock(); *x = 1; unlock(); if tmp=1
```

```
lock(); tmp = *x; unlock(); if tmp=1; *y = 1; lock();*x = 1;unlock();
```

```
lock();tmp = *x;unlock(); *y = 1; if tmp=1; lock(); *x = 1; unlock();
```

How do we avoid data races? (focus on high-level languages)

- `lock()`, `unlock()` are opaque for the compiler: viewed as potentially modifying any location, memory operations cannot be moved past them
- `lock()`, `unlock()` contain "*sufficient fences*" to prevent hardware reordering across them and global ordering

```
*y = 1; lock(); *x = 1; unlock(); lock(); tmp = *x; unlock(); if tmp=1 then print *y
```

```
*y = 1; lock(); tmp = *x; unlock(); lock(); *x = 1; unlock(); if tmp=1
```

```
*y = 1; lock(); tmp = *x; unlock(); if tmp=1; lock(); *x = 1; unlock();
```

```
lock(); tmp = *x; unlock(); *y = 1; lock(); *x = 1; unlock(); if tmp=1
```

```
lock(); tmp = *x; unlock(); if tmp=1; *y = 1; lock(); *x = 1; unlock();
```

```
lock(); tmp = *x; unlock(); *y = 1; if tmp=1; lock(); *x = 1; unlock();
```

How do w

Compiler/hardware can continue to reorder accesses

Intuition:

compiler/hardware do not know about threads, but only racing threads can tell the difference!

- `lock()`, `unlock()` potentially moved past them
- `lock()`, `unlock()` contain "sufficient fences" to prevent hardware reordering across them and global ordering

```
*y = 1; lock(); *x = 1; unlock(); lock(); tmp = *x; unlock(); if tmp=1 then print *y
```

```
*y = 1; lock(); tmp = *x; unlock(); lock(); *x = 1; unlock(); if tmp=1
```

```
*y = 1; lock(); tmp = *x; unlock(); if tmp=1; lock(); *x = 1; unlock();
```

```
lock(); tmp = *x; unlock(); *y = 1; lock(); *x = 1; unlock(); if tmp=1
```

```
lock(); tmp = *x; unlock(); if tmp=1; *y = 1; lock(); *x = 1; unlock();
```

```
lock(); tmp = *x; unlock(); *y = 1; if tmp=1; lock(); *x = 1; unlock();
```

Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1
<pre>if *x == 1 then *y = 1</pre>	<pre>if *y == 1 then *x = 1</pre>

Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1
<pre>if *x == 1 then *y = 1</pre>	<pre>if *y == 1 then *x = 1</pre>

Answer: yes!

The writes cannot be executed in any SC execution, so they cannot participate in a data race.

Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1
<pre>if *x == 1 then *y = 1</pre>	<pre>if *y == 1 then *x = 1</pre>

Data-race freedom is not the ultimate panacea

Ans

The
par

- the absence of data-races is hard to verify / test (undecidable)
- imagine debugging: my program ended with a wrong result, then either my program has a bug OR it has a data-race

Validity of compiler optimisations, summary

Transformation	SC	DRF
Memory trace preserving transformations	✓	✓
Redundant read after read elimination	✓*	✓
Redundant read after write elimination	✓*	✓
Irrelevant read elimination	✓	✓
Redundant write before write elimination	✓*	✓
Redundant write after read elimination	✓*	✓
Irrelevant read introduction	✓	×
Normal memory accesses reordering	×	✓
Roach-motel reordering	× (✓ for locks)	✓
External action reordering	×	✓

* Optimisations legal only on adjacent statements.

Validity of compiler optimisations, summary

Transformation	SC
Memory trace preserving transformations	✓



Jaroslav Sevcik

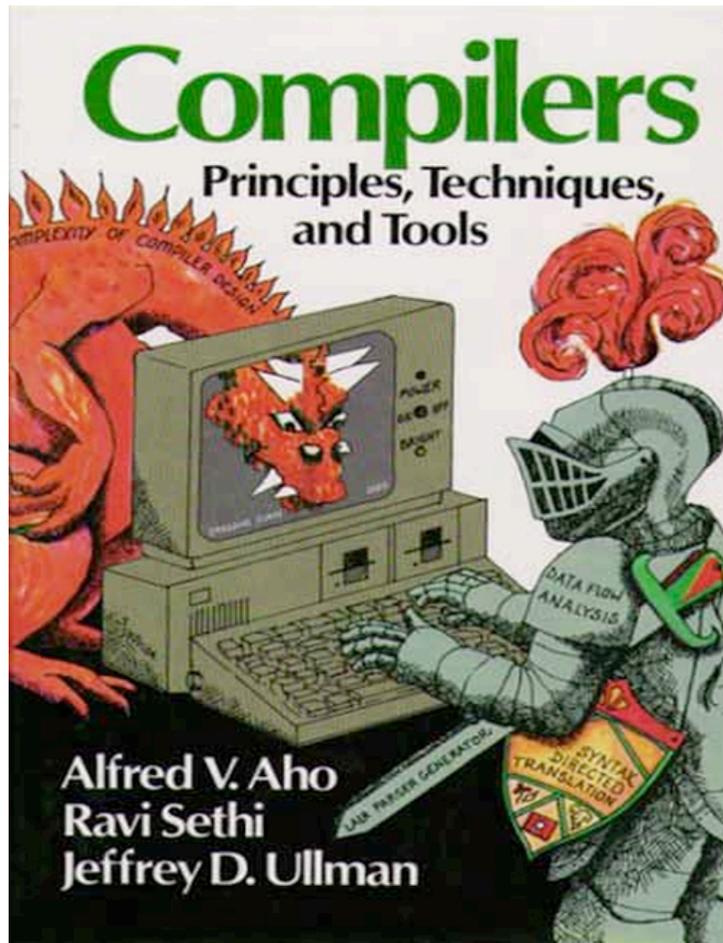
Safe Optimisations for Shared-Memory Concurrent Programs

PLDI 2011

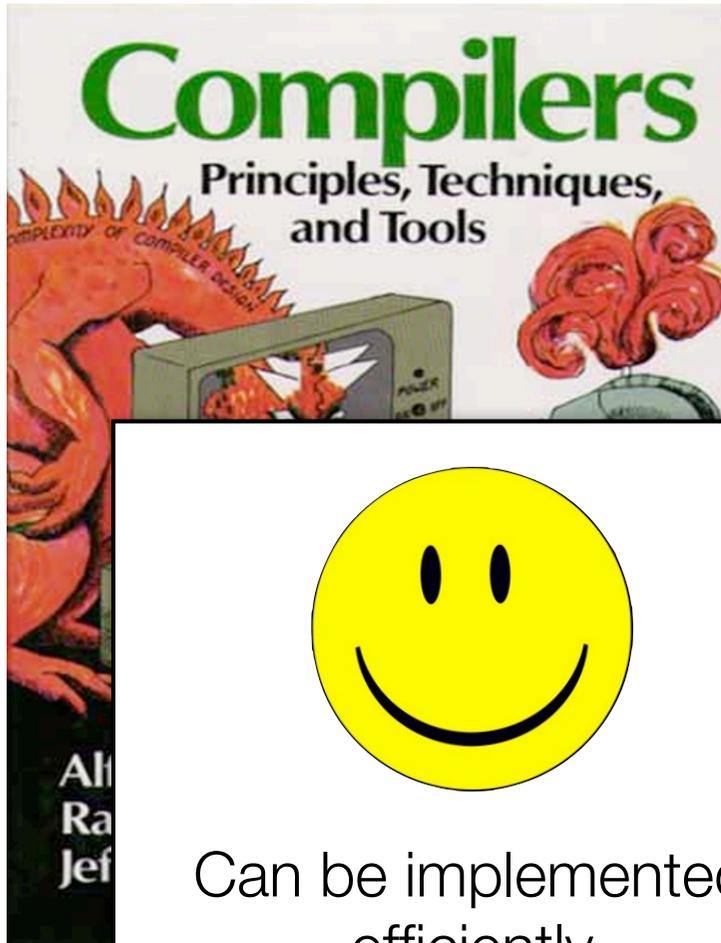
Roach-motel reordering	× (✓ for locks)	✓
External action reordering	×	✓

* Optimisations legal only on adjacent statements.

Compilers, programmers & data-race freedom



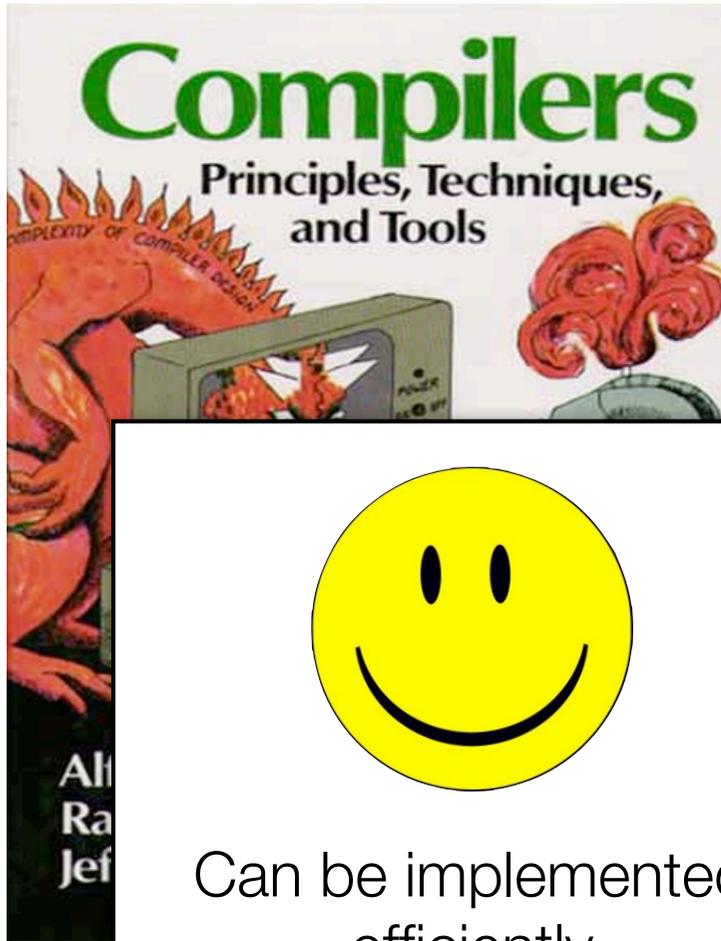
Compilers, programmers & data-race freedom



Can be implemented
efficiently



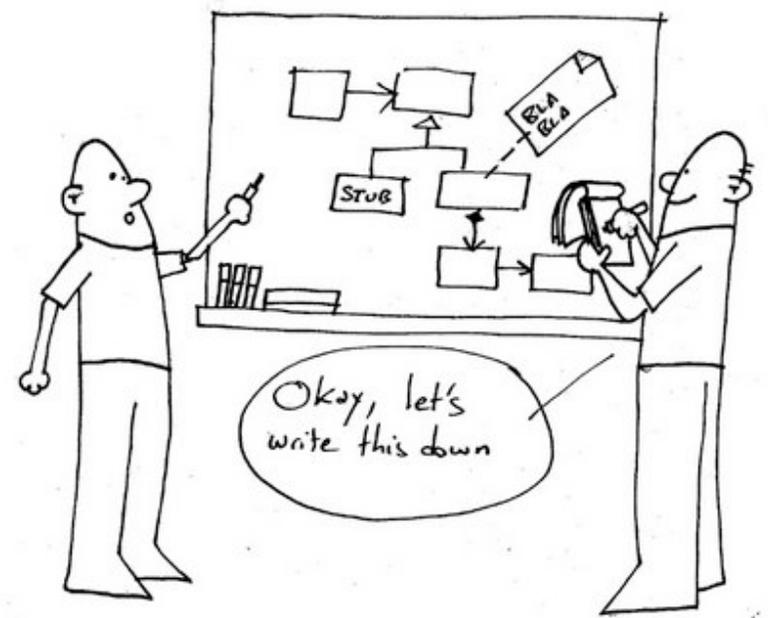
Compilers, programmers & data-race freedom



Can be implemented efficiently



Intuitive programming model (but detecting races is tricky!)



Data-race freedom, formalisation

A toy language: semantics

location, x shared memory location

register, r thread-local variable

integer, n integers

thread_id, t thread identifier

statement, s ::= statements

| *r := x* read from memory

| *x := r* write to memory

| *r := n* load constant into register

| *lock* lock

| *unlock* unlock

| *print r* output

program, p ::= s ; ... ; s a program is a sequence of statements

system ::= concurrent system

| *t₀ : p₀* | ... | *t_n : p_n* parallel composition of n threads

A toy language: semantics

location, x shared memory location
register, r thread-local variable

int
th

sta

We work with a toy language, but this approach scales to the full Java Memory Model or C11/C++11.

| lock lock
| unlock unlock
| print r output

program, p ::= s;...;s a program is a sequence of statements

system ::= concurrent system

| $t_0:p_0$ | ... | $t_n:p_n$ parallel composition of n threads

Traces and tracesets

Definition [trace]: a sequence of memory operations (read, write, thread start, I/O, synchronisation). Thread start is always the first action of thread. All actions in a trace belong to the same thread.

Definition [traceset]: a traceset is a prefix-closed set of traces.

Sample traceset:

Thread 0	Thread 1
$r1 := x$	$r2 := y$
$y := r1$	$x := 1$
	print r2

$$\begin{aligned} & \{[S(0), R[x=v], W[y=v]] \mid v \in V\} \\ \cup & \{[S(1), R[y=v], W[x=1], X(v)] \mid v \in V\} \end{aligned}$$

Remarks:

1. Reads can read arbitrary values from memory.
2. Tracesets should not be confused with interleavings.
3. Tracesets do not enforce receptiveness or determinism:

$$\{[S(0)], [S(0), R[x=1]], [S(0), W[y=1]]\}$$

is also a valid traceset for the example below.

Sample traceset:

Thread 0	Thread 1
<code>r1:=x</code>	<code>r2:=y</code>
<code>y:=r1</code>	<code>x:=1</code>
	<code>print r2</code>

$$\begin{aligned} & \{[S(0), R[x=v], W[y=v]] \mid v \in V\} \\ \cup & \{[S(1), R[y=v], W[x=1], X(v)] \mid v \in V\} \end{aligned}$$

Associate tracesets to toy language programs

$$\langle S, r := x; s \rangle \xrightarrow{R[x=v]} \langle S[r=v], s \rangle$$

$$\langle S, x := r; s \rangle \xrightarrow{W[x=S(r)]} \langle S, s \rangle$$

$$\langle S, r := n; s \rangle \xrightarrow{T} \langle S[r=n], s \rangle$$

$$\langle S, \text{lock}; s \rangle \xrightarrow{L} \langle S, s \rangle$$

$$\langle S, \text{unlock}; s \rangle \xrightarrow{U} \langle S, s \rangle$$

$$\langle S, \text{print } r; s \rangle \xrightarrow{X(S(r))} \langle S, s \rangle$$

$$\langle S, t_0:p_0 \mid \dots \mid t_n:p_n \rangle \xrightarrow{S(i)} \langle S, p_i \rangle$$

Tracesets and interleavings

Definition [interleaving]: an interleaving is a sequence of thread-identifier-action pairs.

Example: `y:=1; || r2:=v;print r2;`

$$I' = [\langle 0, S(0) \rangle, \langle 1, S(1) \rangle, \langle 0, W[y=1] \rangle, \langle 1, R[v=0] \rangle, \langle 1, X(0) \rangle]$$

Given an interleaving I , the trace of tid in I is the sequence of actions of thread tid in I , e.g.:

$$\text{trace } 1 \ I' = [S(1), R[v=0], X(0)].$$

Conversely, given a traceset, we can compute all the well-formed interleavings (called *executions*)... (next slide)

Tracesets and interleavings

An interleaving I is an *execution* of a traceset T if:

- for all tid , trace $tid \in T$ (traces belong to the traceset)
- $tids$ correspond to entry points $S(tid)$
- lock / unlock alternates correctly
- each read sees the most recent write to the same location (read/from).

(The last property enforces the sequentially consistent semantics for memory accesses).

Tracesets and interleavings

An interleaving I is an *execution* of a traceset T if:

- for

- *tids*

- loc

- each

(The

Remarks:

1. Interleavings order totally the actions, and do not keep track of which actions happen in parallel.

2. It is however possible to put more structure on interleavings, and recover informations about concurrency.

).

sses).

Happens-before

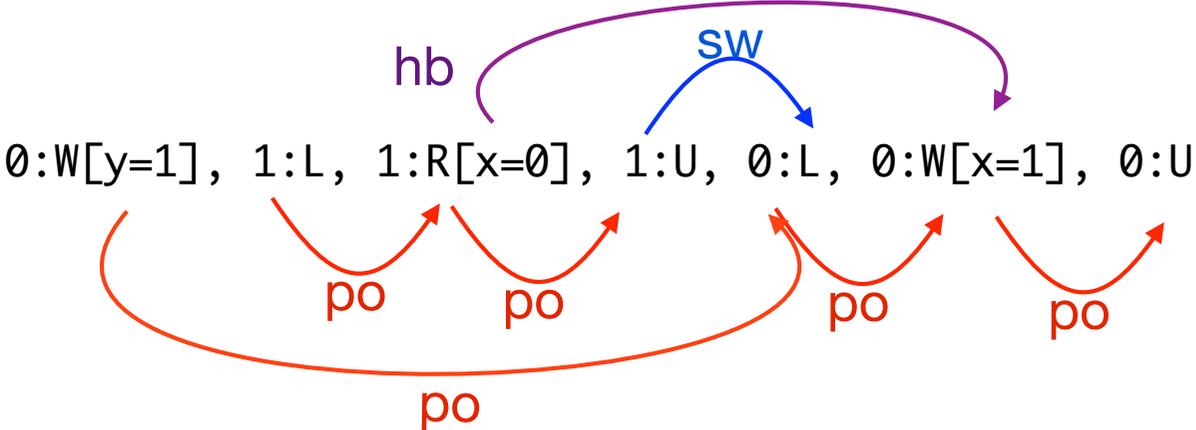
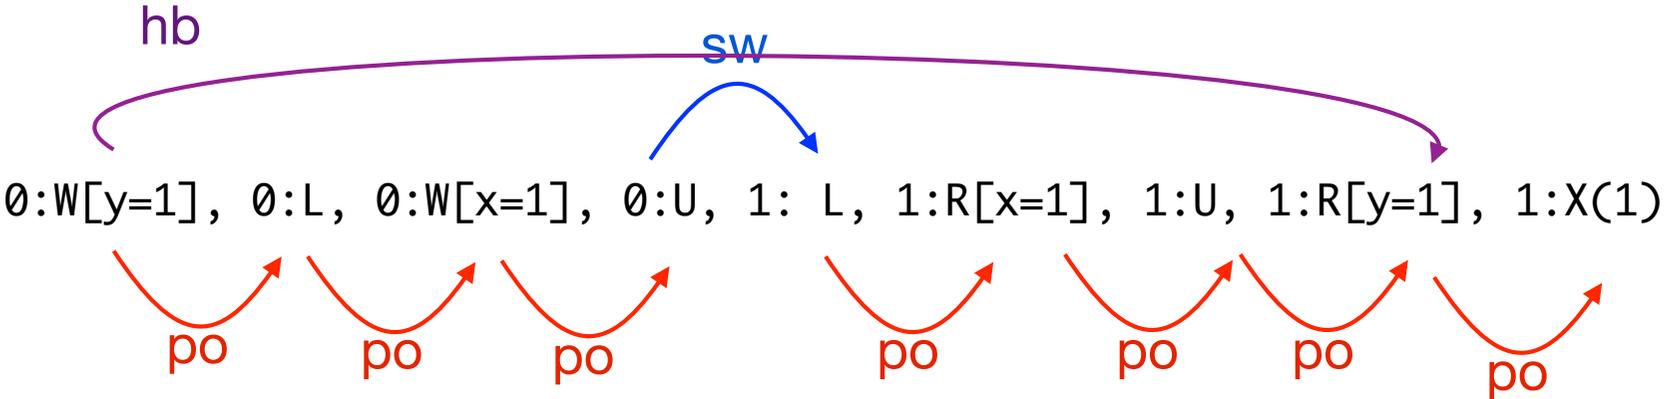
Definition [program order]: **program order**, $<_{po}$, is a total order over the actions of *the same thread in an interleaving*.

Definition [synchronises with]: in an interleaving I , index i **synchronises-with** index j , $i <_{sw} j$, if $i < j$ and $A(I_i) = U$ (unlock), $A(I_j) = L$ (lock).

Definition [happens-before]: **Happens-before** is the transitive closure of program order and synchronises with.

Examples of happens before

Thread 0	Thread 1
<pre>*y = 1 lock(); *x = 1 unlock();</pre>	<pre>lock(); tmp = *x; unlock(); if tmp = 1 then print *y</pre>



S(tid) actions omitted.

Data-race freedom

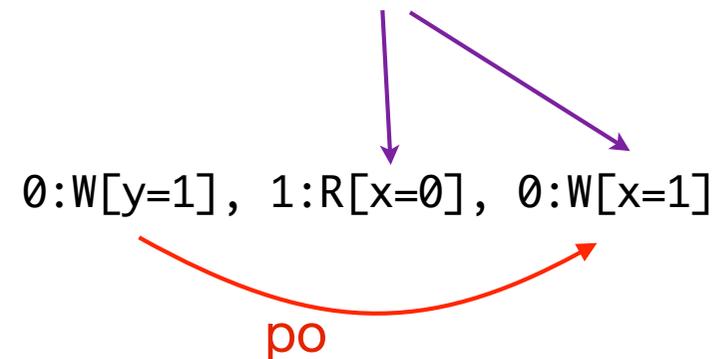
Definition [data-race-freedom]: A traceset is **data-race free** if none of its executions has two adjacent conflicting actions from different threads.

Equivalently, a traceset is data-race free if in all its executions all pairs of conflicting actions are ordered by happens-before.

A racy program

Thread 0	Thread 1
<code>*y = 1</code>	<code>if *x == 1</code>
<code>*x = 1</code>	<code>then print *y</code>

Two conflicting accesses
not related by happens before.



Data-race freedom: equivalence of definitions

Given an execution

$$\alpha \text{ ++ } [a] \text{ ++ } \beta \text{ ++ } [b]$$

of a traceset T where $[a]$ and $[b]$ are the first conflicting actions not related by happen-before, we build the interleaving

$$\alpha \text{ ++ } \beta' \text{ ++ } [a] \text{ ++ } [b]$$

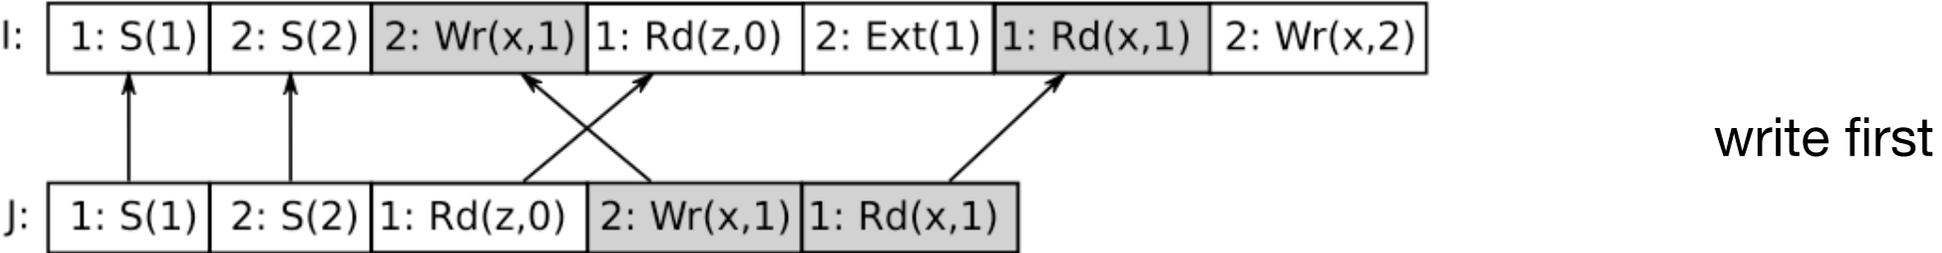
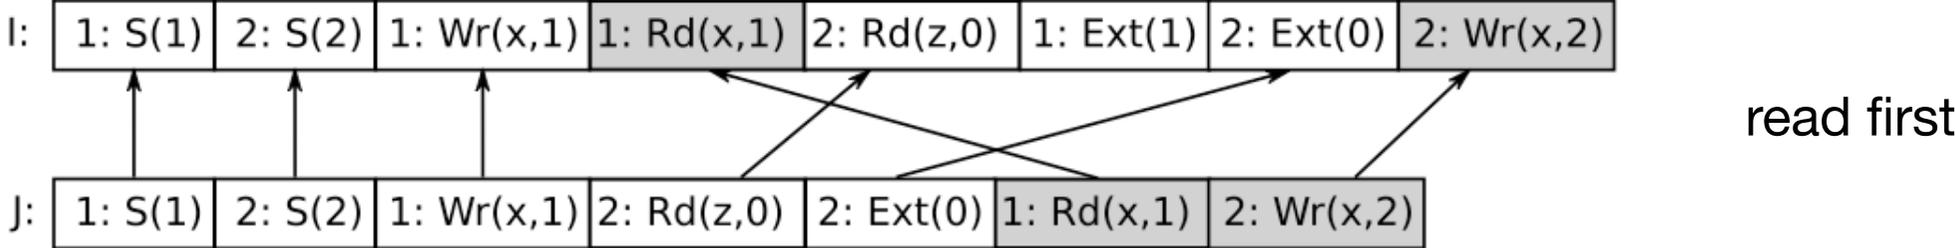
where β' are all the actions from β that strictly happen-before $[b]$.

It remains to show that $\alpha \text{ ++ } \beta' \text{ ++ } [a] \text{ ++ } [b]$ is an execution of T .

The formal proof is tedious and not easy (see Boyland 2008, Bohem & Adve 2008, Sevcik), here will give the intuitions of the construction on an example.

Data-race freedom: equivalence of definitions

```
Thread 1: x := 1; r1 := x; print r1;
Thread 2: r2 := z; print r2; x := 2;
```



Option 1

Don't.

No concurrency.

Implemented by highly-successful programming languages (**OCaml**)

Poor match for current trends

Option 2

Don't.

No shared memory

A good match for some problems (see Erlang, MPI, ...)

Option 3

Don't.

But language ensures data-race freedom

Possible:

- syntactically ensuring data accesses protected by associated locks
- fancy effect type systems

Not suitable for general purpose programming.

Option 4

Don't.

Leave it (sort of) up to the hardware

Example:

MLton, a high performance ML-to-x86 compiler with concurrency extensions

Accesses to ML refs exhibit the underlying x86-TSO behaviour (atomicity is guaranteed though)

Option 5

Do.

Use data race freedom as a definition

1. Programs that race-free have only sequentially consistent behaviours
2. Programs that have a race in some execution can behave in any way

Sarita Adve & Mark Hill, 1990



Option 5

Do.

Use data race freedom as a definition

Pro:

- simple
- strong guarantees for most code
- allows lots of freedom for compiler and hardware optimisations

Cons:

- undecidable premise
- can't write racy programs (escape mechanisms?)

Ada 83

[ANSI-STD-1815A-1983, 9.11] For the actions performed by a program that uses shared variables, the following assumptions can always be made:

- If between two synchronization points in a task, this task reads a shared variable whose type is a scalar or access type, then the variable is not updated by any other task at any time between these two points.
- If between two synchronization points in a task, this task updates a shared variable whose task type is a scalar or access type, then the variable is neither read nor updated by any other task at any time between these two points.

The execution of the program is erroneous if any of these assumptions is violated.

Data-races are errors

Posix Threads Specification

[IEEE 1003.1-2008, Base Definitions 4.11] Applications shall ensure that access to any memory location by more than one thread of control (threads or processes) is restricted such that no thread of control can read or modify a memory location while another thread of control may be modifying it.

Data-races are errors

C++ 2011

[C++ 2011 FDIS (WG21/N3290) 1.10p21] The execution of a program contains a *data race* if it contains two conflicting actions in different threads, at least one of which is not atomic, and neither happens before the other. Any such data race results in undefined behavior.

Data-races are errors

Data race freedom as a definition

- Core of the C11/C++11 standard.

Hans Boehm & Sarita Adve, PLDI 2008.



- Part of the JSR-133 standard.

Jeremy Manson & Bill Pugh & Sarita Adve, PLDI 2008.



Data race freedom as a definition

- Core of the C11/C++11 standard.

Hans Boehm & Sarita Adve, PLDI 2008.

with some escape mechanism called "low level atomics".

Mark Batty & al., POPL 2011.

- Part of the JSR-133 standard.

Jeremy Manson & Bill Pugh & Sarita Adve, PLDI 2008.

DRF gives no guarantees for untrusted code: a disaster for Java, which relies on unforgeable pointers for its security guarantees.

JSR-133 is **DRF + some out-of-thin-air guarantees** for all code.

*Escape lanes
for expert programmers*



Low-level atomics in C11/C++11

```
std::atomic<int> flag0(0),flag1(0),turn(0);
```

```
void lock(unsigned index) {  
    if (0 == index) {  
        flag0.store(1, std::memory_order_relaxed);  
        turn.exchange(1, std::memory_order_acq_rel);  
  
        while (flag1.load(std::memory_order_acquire)  
            && 1 == turn.load(std::memory_order_relaxed))  
            std::this_thread::yield();  
    } else {  
        flag1.store(1, std::memory_order_relaxed);  
        turn.exchange(0, std::memory_order_acq_rel);  
  
        while (flag0.load(std::memory_order_acquire)  
            && 0 == turn.load(std::memory_order_relaxed))  
            std::this_thread::yield();  
    }  
}
```

```
void unlock(unsigned index) {  
    if (0 == index) {  
        flag0.store(0, std::memory_order_release);  
    } else {  
        flag1.store(0, std::memory_order_release);  
    }  
}
```

Atomic variable declaration

New syntax
for memory accesses

Qualifier

Low level atomics

MO_SEQ_CST

MO_RELEASE / MO_ACQUIRE

MO_RELEASE / MO_CONSUME

MO_RELAXED

LESS RELAXED



MORE RELAXED

Low level atomics

MO_SEQ_CST

MO_RELEASE / MO_ACQUIRE

MO_RELEASE / MO_CONSUME

MO_RELAXED

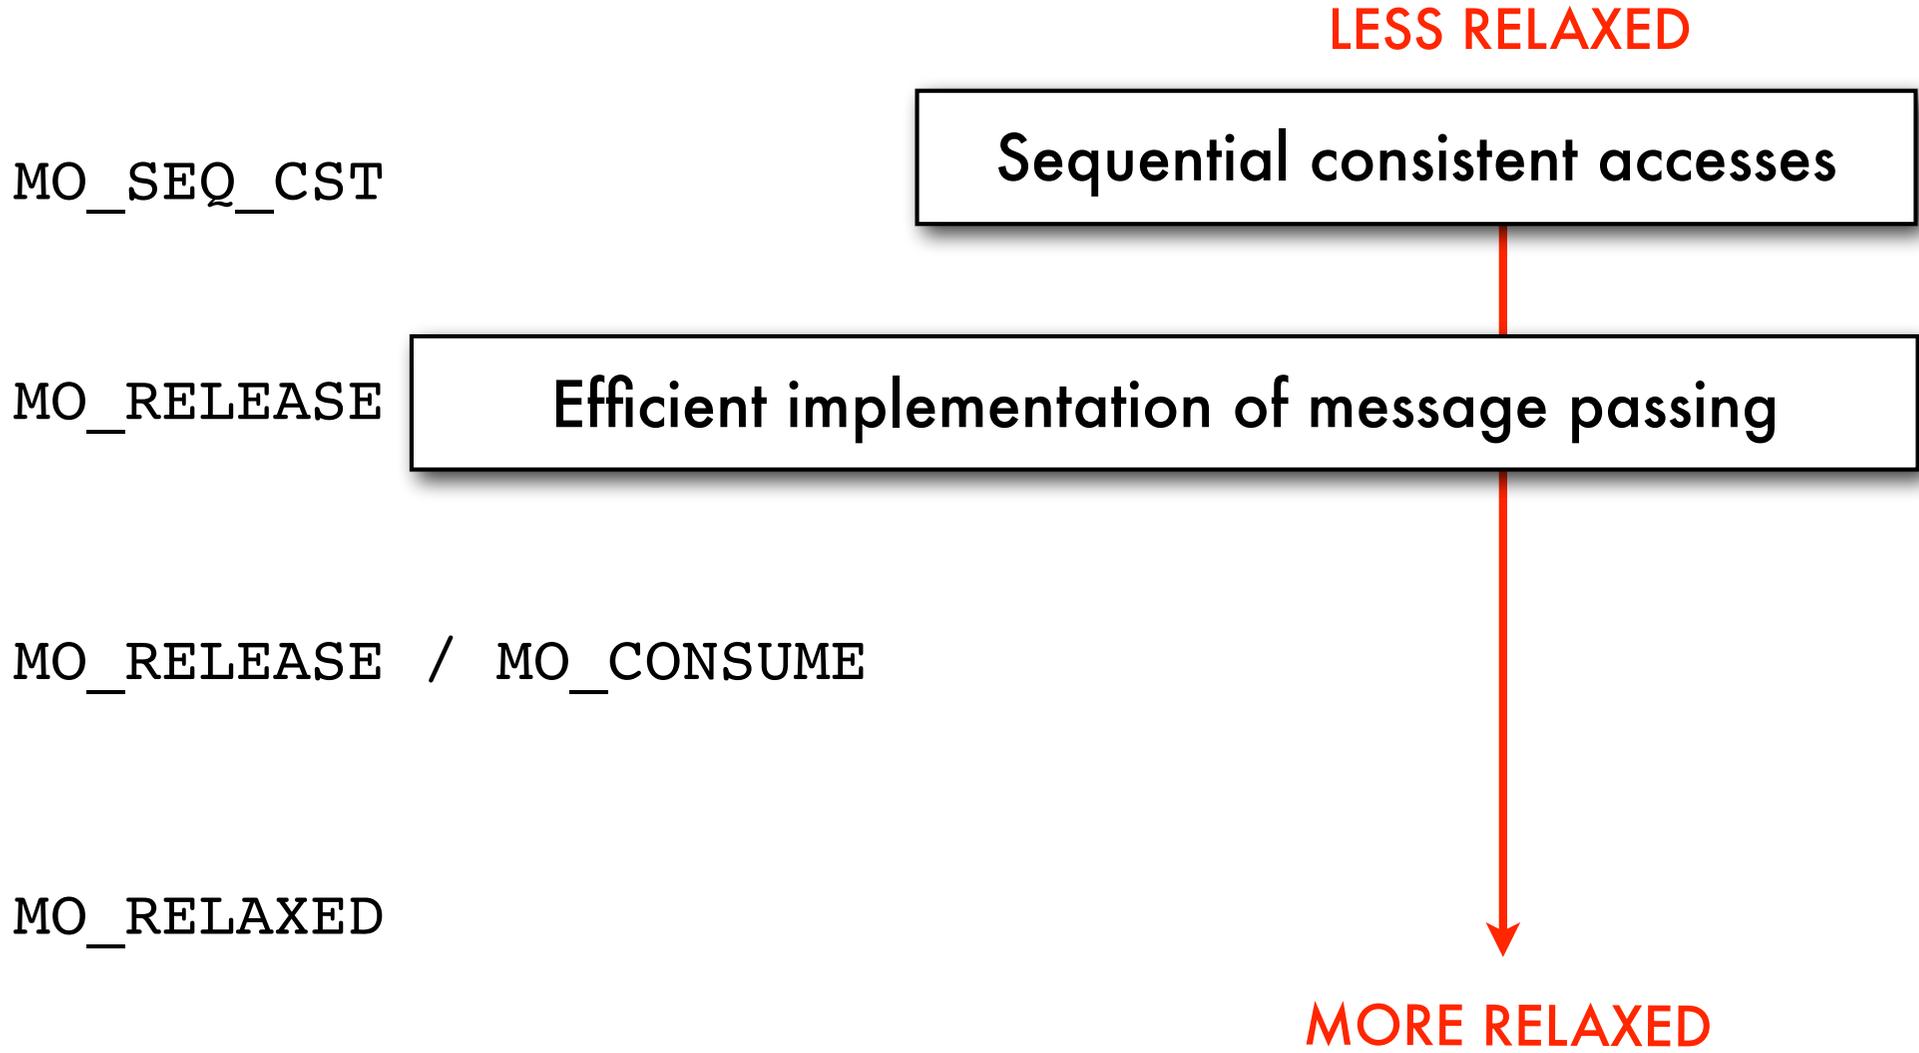
Sequential consistent accesses

LESS RELAXED

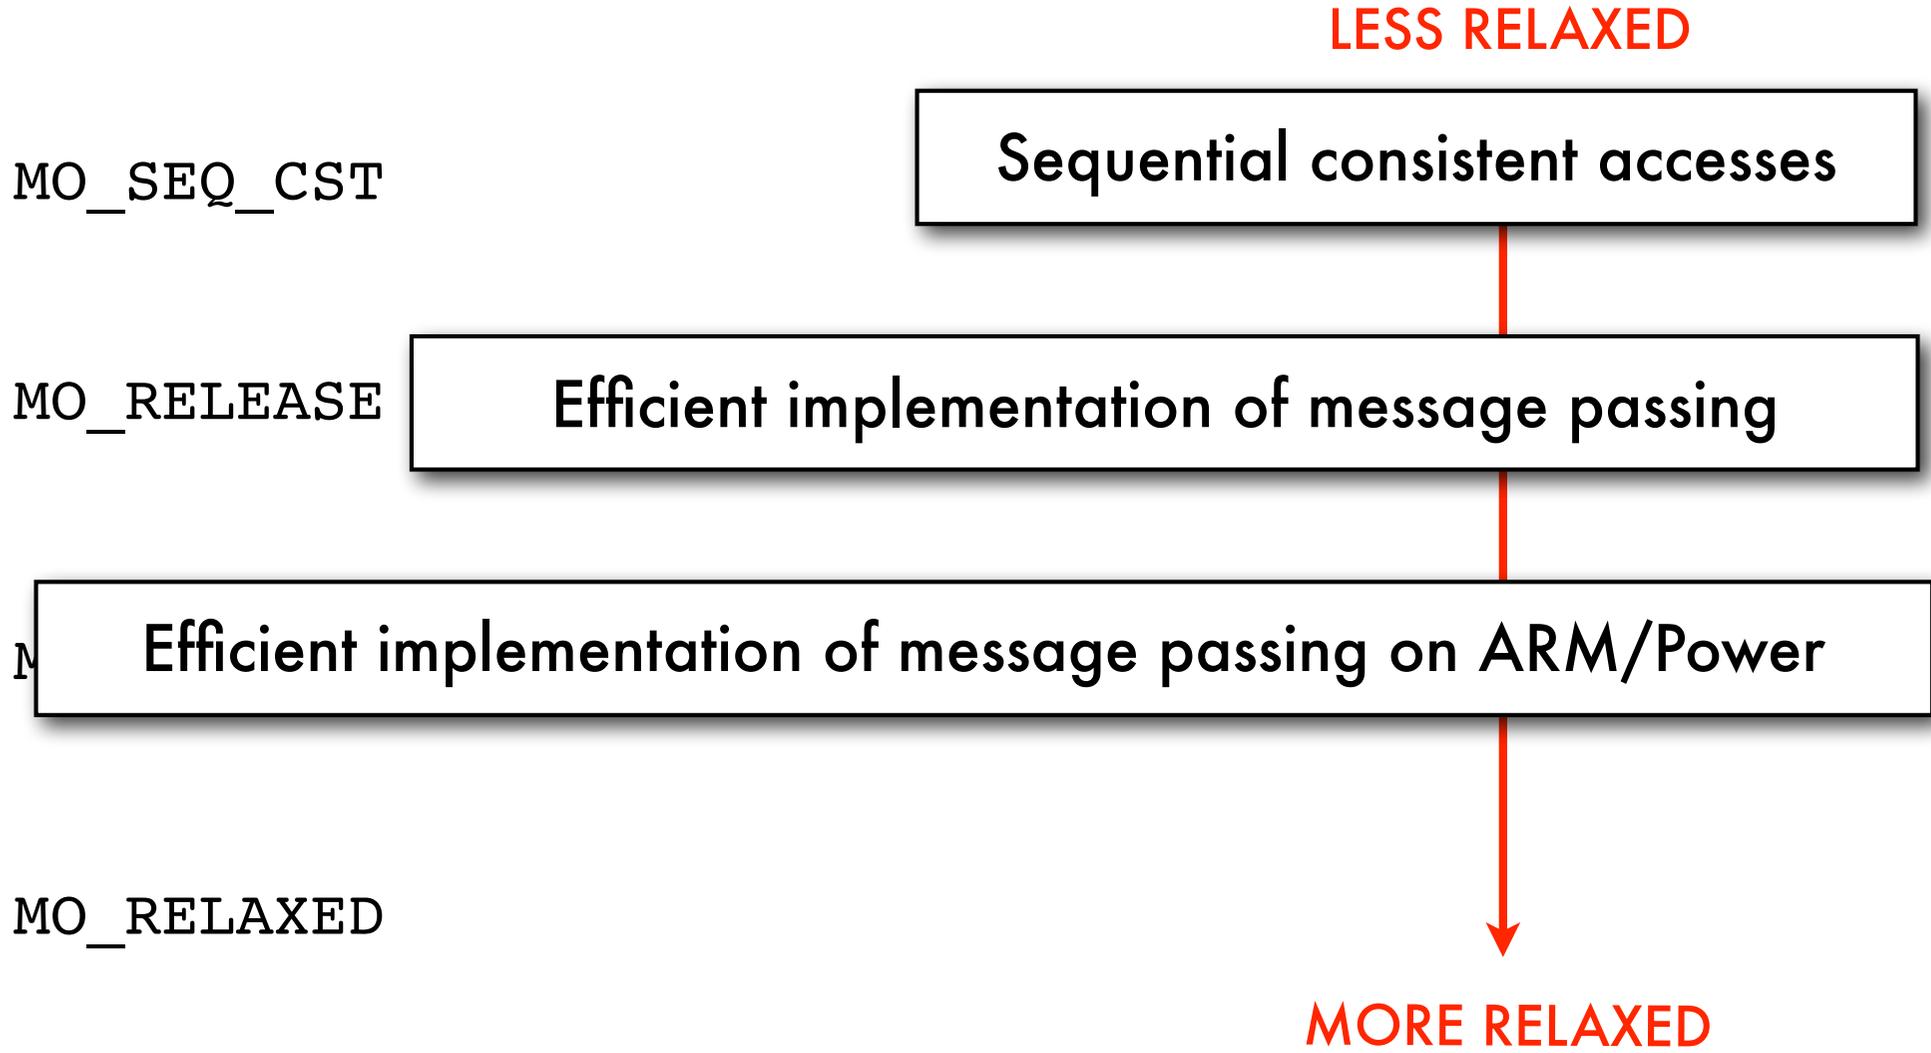


MORE RELAXED

Low level atomics



Low level atomics



Low level atomics

LESS RELAXED

MO_SEQ_CST

Sequential consistent accesses

MO_RELEASE

Efficient implementation of message passing

MO_RELAX

Efficient implementation of message passing on ARM/Power

MO_RELAX

No synchronisation; direct access to hardware

MORE RELAXED

MO_SEQ_CST

The compiler must ensure that `MO_SEQ_CST` accesses have sequentially consistent semantics.

Thread 0	Thread 1
<code>x.store(1,MO_SEQ_CST)</code> <code>r1 = y.load(MO_SEQ_CST)</code>	<code>y.store(1,MO_SEQ_CST)</code> <code>r2 = x.load(MO_SEQ_CST)</code>

The program above cannot end with `r1 = r2 = 0`.

Sample compilation on x86:

store: `MOV; MFENCE`
load: `MOV`

Sample compilation on Power:

store: `HWSYNC; ST`
load: `HWSYNC; LD; CMP; BC; ISYNC`

MO_RELEASE / MO_ACQUIRE

Supports a fast implementation of the message passing idiom:

Thread 0	Thread 1
<code>x.store(1,MO_RELAXED)</code>	<code>r1 = y.load(MO_ACQUIRE)</code>
<code>y.store(1,MO_RELEASE)</code>	<code>r2 = x.load(MO_RELAXED)</code>



The program above cannot end with `r1 = 1` and `r2 = 0`.

Accesses to the data structure can be reordered/optimised (`MO_RELAXED`).

Sample compilation on x86:

store: MOV
load: MOV

Sample compilation on Power:

store: LWSYNC; ST
load: LD; CMP; BC; ISYNC

MO_RELEASE / MO_CONSUME

Supports a fast implementation of the message passing idiom on Power:

Thread 0	Thread 1
<code>x.store(1,MO_RELAXED)</code> <code>y.store(&x,MO_RELEASE)</code>	<code>r1 = y.load(x,MO_CONSUME)</code> <code>r2 = (*x).load(MO_RELAXED)</code>

The program above cannot end with `r1 = 1` and `r2 = 0`.

The two loads have an address dependency, Power won't reorder them.

Sample compilation on x86:

store: MOV
load: MOV

Sample compilation on Power:

store: LWSYNC; ST
load: LD

Memory access synchronisation

`x = y = 0`

Thread 1

`y = 1`

`x.store(1, MO_RELEASE)`

Thread 2

`if (x.load(MO_ACQUIRE) == 1)`

`r2 = y`

Memory access synchronisation

`x = y = 0`

Thread 1

Thread 2

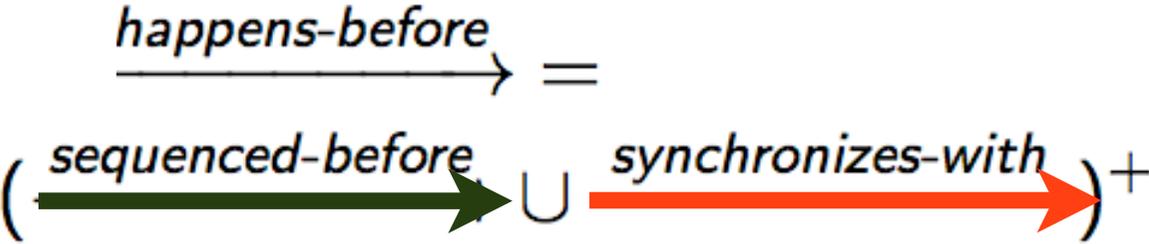
```

    y = 1
    x.store(1, MO_RELEASE)
  
```



```

    if (x.load(MO_ACQUIRE) == 1)
      r2 = y
  
```



Non-atomic loads must return the *most recent write* in the happens-before order (unique in a DRF program)

Understanding MO_RELAXED

`x = y = 0`

Thread 1

```
    y = 1  
x.store(1,MO_RELAXED)
```

Thread 2

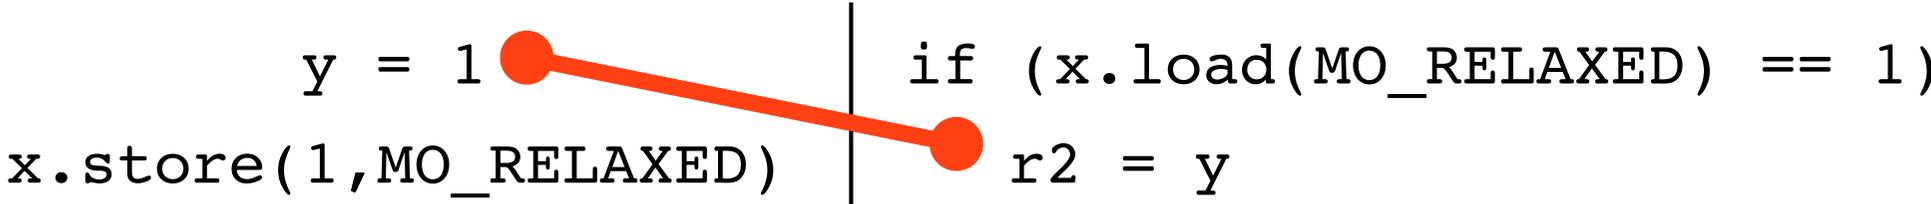
```
if (x.load(MO_RELAXED) == 1)  
    r2 = y
```

Understanding MO_RELAXED

`x = y = 0`

Thread 1

Thread 2



DATA RACE

Two conflicting accesses not related by happens-before

Understanding MO_RELAXED

$x = y = 0$

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

WELL DEFINED

but $r2 = 0$ is possible

Intuition

the compiler (or hardware) can reorder independent accesses

$x = y = 0$

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

WELL DEFINED

but $r2 = 0$ is possible

Intuition

the compiler (or hardware) can reorder independent accesses

$x = y = 0$

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

Allow a RELAXED load to see any store that:

- does not happen-after it
 - is not hidden by an intervening store hb-ordered between them
-

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return NULL;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Thread 2 is not affected by Thread 1 and vice-versa

This program is data-race free

This program must print 42

Shared memory

```
int a = 1;  
int b = 0;
```

This is a *compiler bug*

```
int s, b = 42,  
for (s=0; s!=4; s++) {      printf("%d\n", b);  
    if (a==1)  
        return NULL;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2 is not affected by Thread 1 and vice-versa

This program is data-race free

This program must print 42

Shared memory

```
int a = 1;  
int b = 0;
```

This is a *concurrency compiler bug*

```
int s, b = 42,  
for (s=0; s!=4; s++) {      printf("%d\n", b);  
    if (a==1)  
        return NULL;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2 is not affected by Thread 1 and vice-versa

This program is data-race free

This program must print 42

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 1 returns without modifying b

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;
int b = 0;
```

Thread 1

```
int s;
for (s=0; s!=4; s++) {
    if (a==1)
        return 0;
    for (b=0; b<=26; ++b)
        ;
}
```

Thread 2

```
b = 42;
printf("%d\n", b);
```

Thread 1 returns without modifying **b**

Thread 2 is not affected by Thread 1 and vice-versa
(this program is *data-race free*)

This program must always print **42**

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Typical system code!

Shared memory

```
int a = 1;
int b = 0;
```

Thread 1

```
int s;
for (s=0; s!=4; s++) {
    if (a==1)
        return 0;
    for (b=0; b<=26; ++b)
        ;
}
```

Thread 2

```
b = 42;
printf("%d\n", b);
```



...in some executions might print 0

gcc 4.7 -O2



```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)   # store ebx into b  
movl    $0, %eax        # return 0  
ret
```

gcc 4.7 -O2



```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

The outer loop can be (and is) compiled away

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)    # store ebx into b  
movl    $0, %eax         # return 0  
ret
```

gcc 4.7 -O2



```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)   # store ebx into b  
movl    $0, %eax        # return 0  
ret
```

gcc 4.7 -O2

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

Prefetch b, in case it comes handy later

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)    # store ebx into b  
movl    $0, %eax         # return 0  
ret
```

gcc 4.7 -O2



```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)    # store ebx into b  
movl    $0, %eax         # return 0  
ret
```

gcc 4.7 -O2



```
int s;
for (s=0; s!=4; s++) {
    if (a==1)
        return 0;
    for (b=0; b<=26; ++b)
        ;
}
```

Restore the prefetched value of b

```
movl  a(%rip), %eax    # load a into eax
movl  b(%rip), %ebx    # load b into ebx
testl %edx, %edx      # if a==1
jne   .L2              # jump to .L2
movl  $0, b(%rip)
ret
.L2:
movl  %ebx, b(%rip)   # store ebx into b
movl  $0, %eax        # return 0
ret
```

gcc 4.7 -O2



```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return 0;  
    for (b=0; b<=26; ++b)  
        ;  
}
```

```
movl    a(%rip), %eax    # load a into eax  
movl    b(%rip), %ebx    # load b into ebx  
testl   %edx, %edx      # if a==1  
jne     .L2              # jump to .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)    # store ebx into b  
movl    $0, %eax        # return 0  
ret
```

The compiler has introduced
the prefetch and restore of **b**

Surprising but correct in sequential executions

```
movl  a(%rip), %eax    # load a into eax
movl  b(%rip), %ebx    # load b into ebx
testl %edx, %edx      # if a==1
jne   .L2              # jump to .L2
movl  $0, b(%rip)
ret
.L2:
movl  %ebx, b(%rip)    # store ebx into b
movl  $0, %eax        # return 0
ret                    #
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl    a(%rip),%eax  
movl    b(%rip),%ebx  
testl   %eax, %eax  
jne     .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)  
movl    $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl a(%rip),%eax  
movl b(%rip),%ebx  
testl %eax, %eax  
jne .L2  
movl $0, b(%rip)  
ret  
.L2:  
movl %ebx, b(%rip)  
movl $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

- Read a (1) into eax

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl    a(%rip),%eax  
movl    b(%rip),%ebx  
testl   %eax, %eax  
jne     .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)  
movl    $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

- Read a (1) into eax
- Read b (0) into ebx

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl    a(%rip),%eax  
movl    b(%rip),%ebx  
testl   %eax, %eax  
jne     .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)  
movl    $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

- Read a (1) into eax
- Read b (0) into ebx
- Store 42 into b

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl    a(%rip),%eax  
movl    b(%rip),%ebx  
testl   %eax, %eax  
jne     .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)  
movl    $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

- Read a (1) into eax
- Read b (0) into ebx
- Store 42 into b
- Store ebx (0) into b

Shared memory

```
int a = 1;  
int b = 0;
```

Thread 1

```
movl    a(%rip),%eax  
movl    b(%rip),%ebx  
testl   %eax, %eax  
jne     .L2  
movl    $0, b(%rip)  
ret  
.L2:  
movl    %ebx, b(%rip)  
movl    $0, %eax  
ret
```

Thread 2

```
b = 42;  
printf("%d\n", b);
```

- Read a (1) into eax
- Read b (0) into ebx
- Store 42 into b
- Store ebx (0) into b
- Print b... 0 is printed

*Introduces unexpected behaviours
in some concurrent context*

```
testl %eax, %eax
jne .L2
movl $0, b(%rip)
ret
.L2:
movl %ebx, b(%rip)
movl $0, %eax
ret
```

- Read a (1) into eax
- Read b (0) into ebx
- Store 42 into b
- Store ebx (0) into b
- Print b... 0 is printed

*Introduces unexpected behaviours
in some concurrent context*

This is a concurrency compiler bug

```
movl    $0, b(%rip)
ret
.L2:
movl    %ebx, b(%rip)
movl    $0, %eax
ret
```

- Read a (1) into eax
- Read b (0) into ebx
- Store 42 into b
- Store ebx (0) into b
- Print b... 0 is printed

Introduces unexpected behaviours

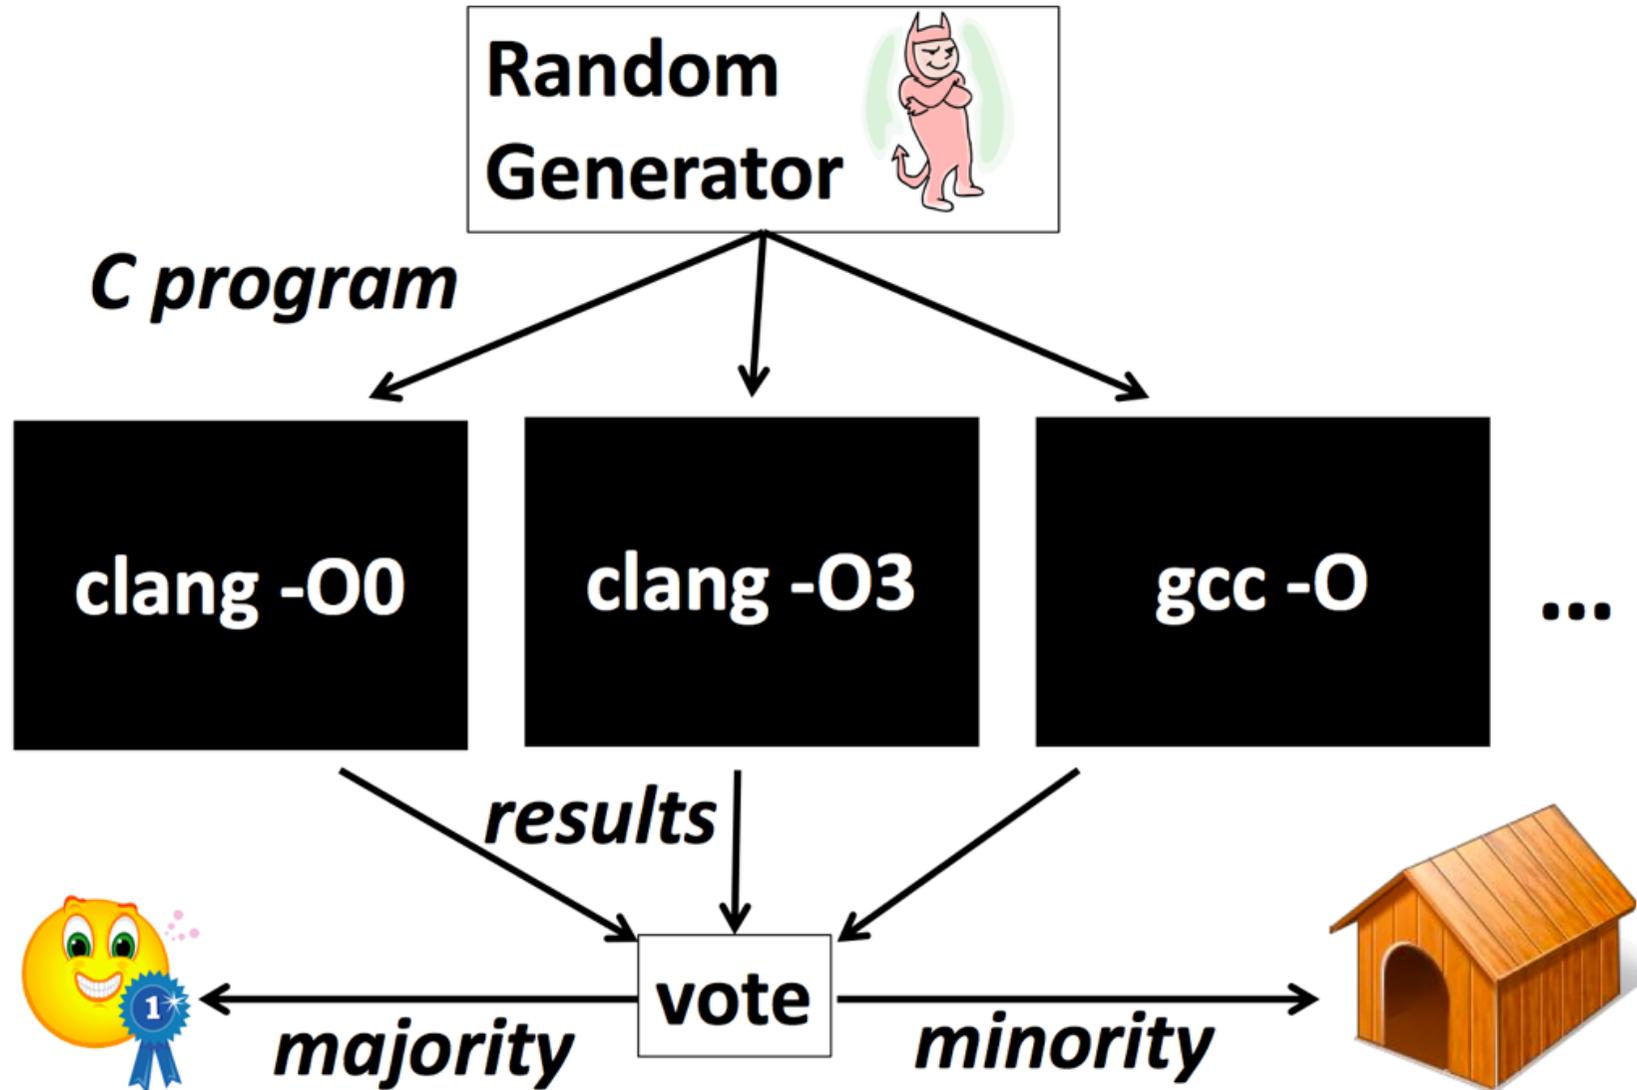
A bug report is not research.

*A technique to identify
concurrency compiler bugs
in existing compilers is!*

- Store ebx (0) into b
- Print b... 0 is printed

Compiler testing: state of the art

Yang, Chen, Eide, Regehr - PLDI 2011



Compiler testing: state of the art

Yang, Chen, Eide, Regehr - PLDI 2011

Random



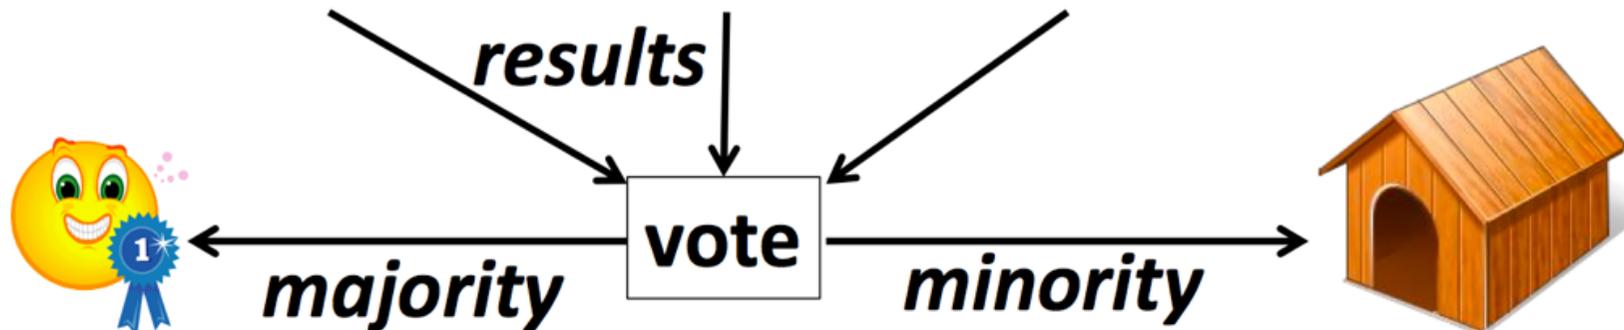
Reported hundreds of bugs
on various versions of gcc, clang and other compilers

clang -O0

clang -O5

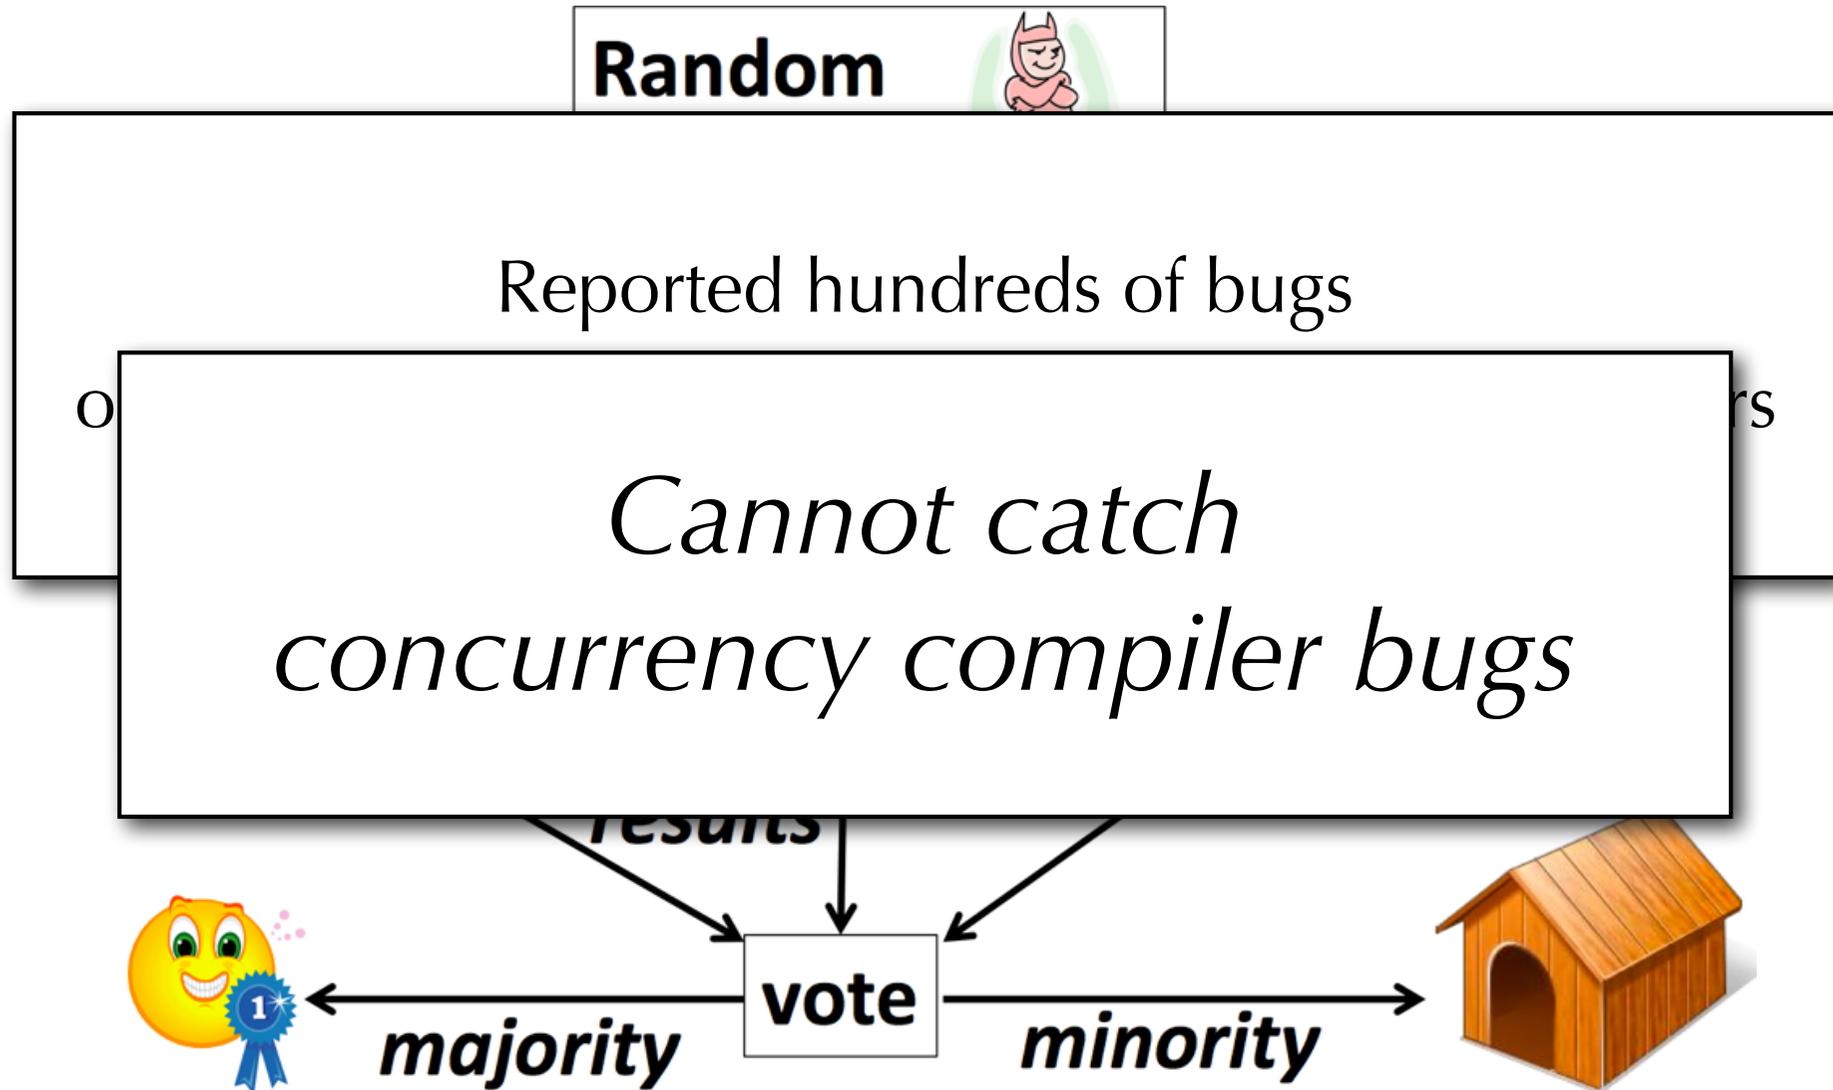
gcc -O

...



Compiler testing: state of the art

Yang, Chen, Eide, Regehr - PLDI 2011



Hunting concurrency compiler bugs?

How to deal with non-determinism?

How to generate non-racy interesting programs?

How to capture all the behaviours of concurrent programs?

A compiler can optimise away behaviours:

how to test for correctness?

limit case: two compilers generate correct code with disjoint final states

Idea

C/C++ compilers support separate compilation
Functions can be called in arbitrary non-racy concurrent contexts

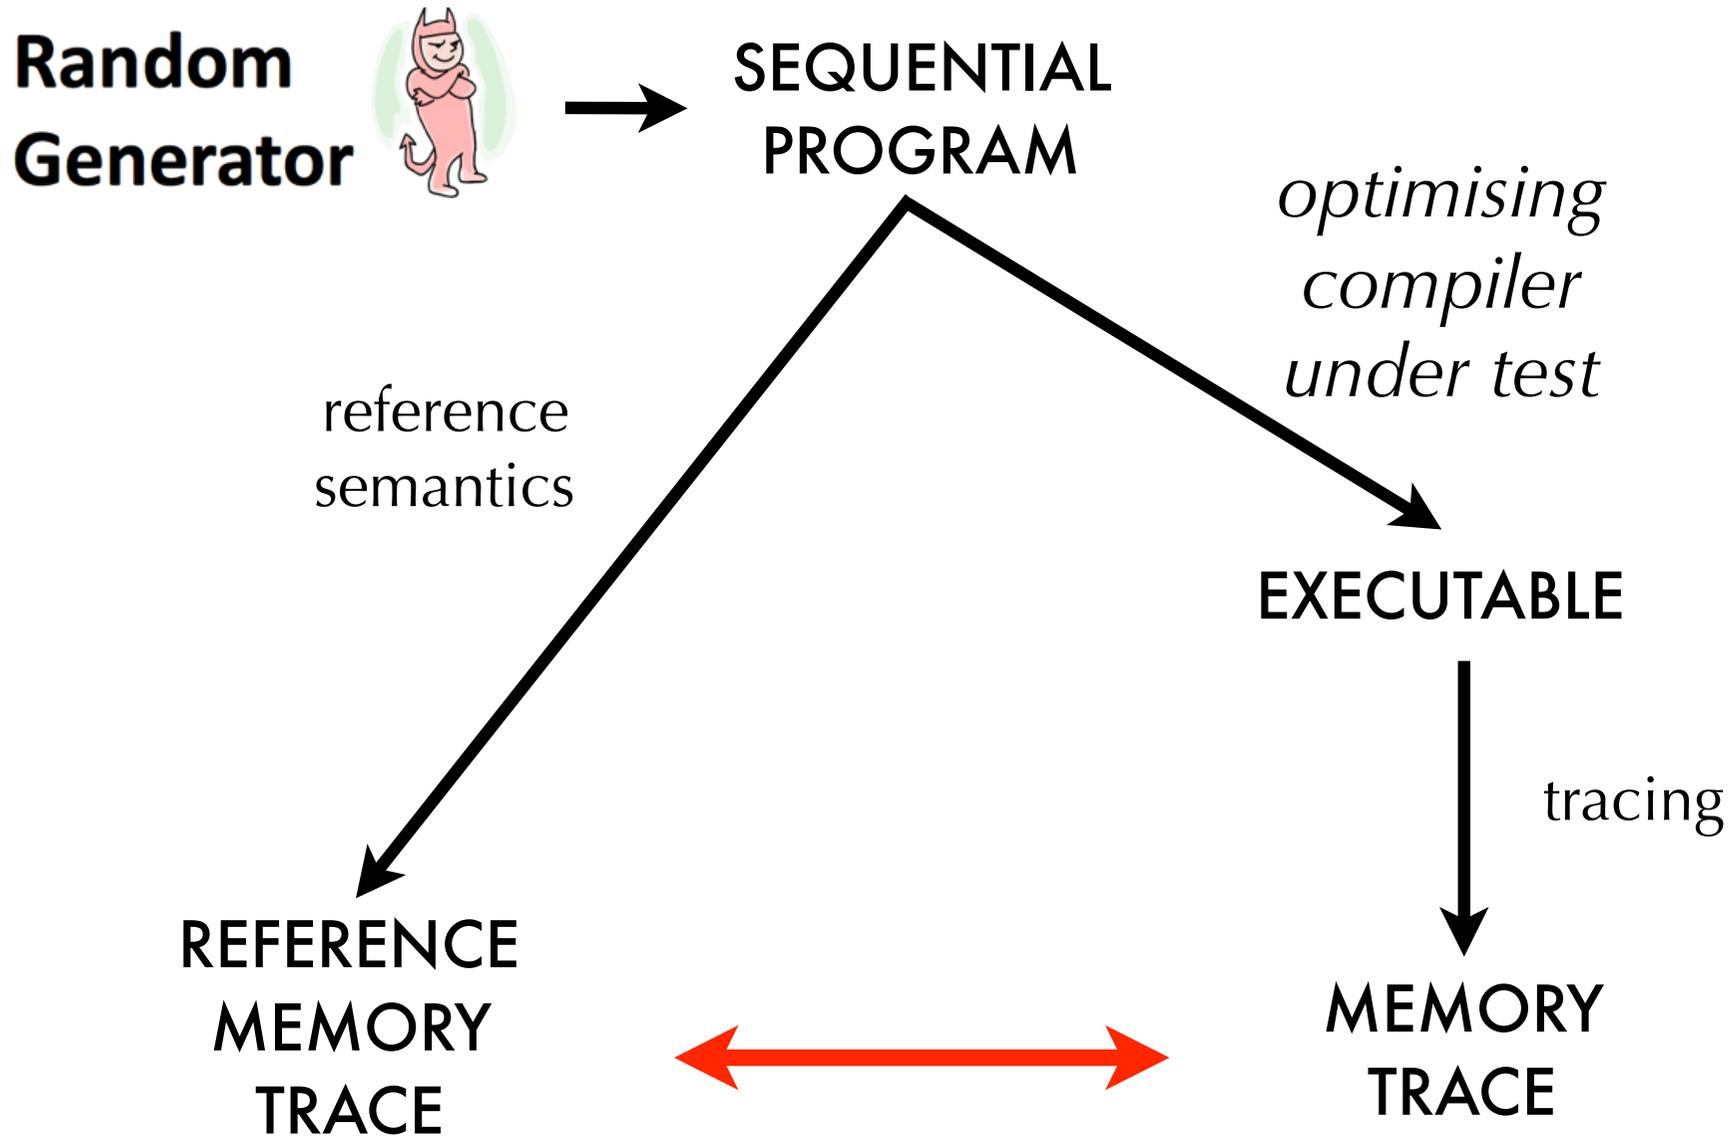


C/C++ compilers can only apply transformations sound
with respect to an arbitrary non-racy concurrent context

Hunt concurrency compiler bugs

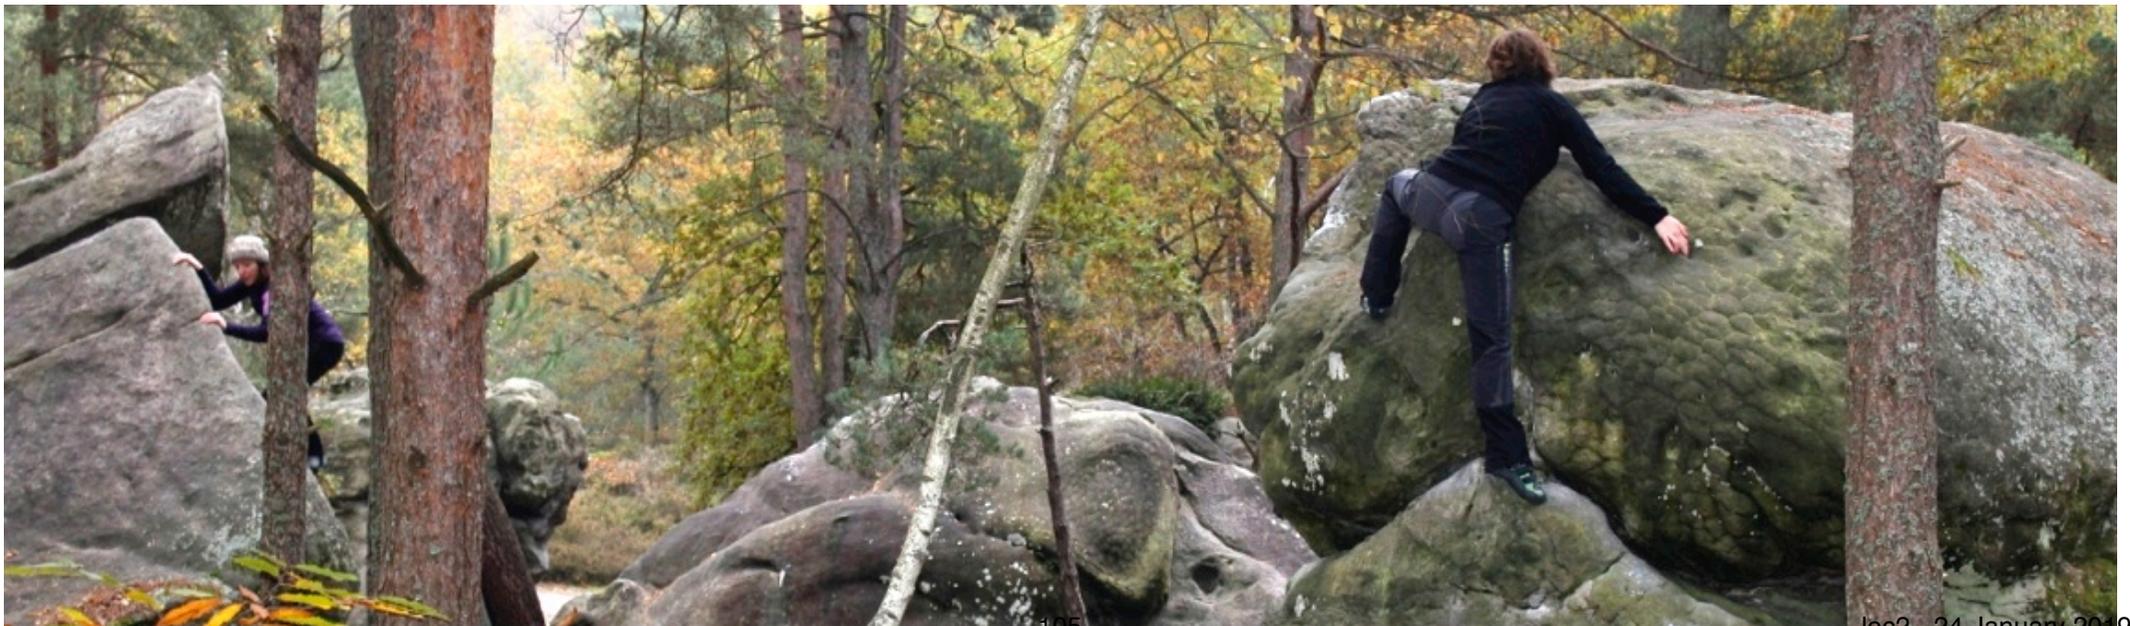
=

search for transformations of sequential code
not sound in an arbitrary non-racy context

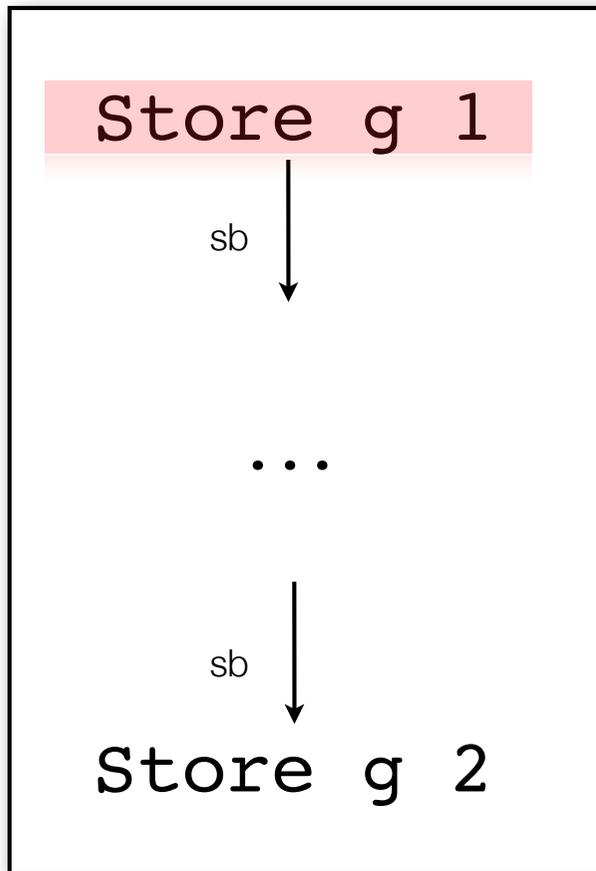


Check: only transformations sound in any concurrent non-racy context

Soundness of compiler optimisations in the C11/C++11 memory model



Elimination of *overwritten writes*



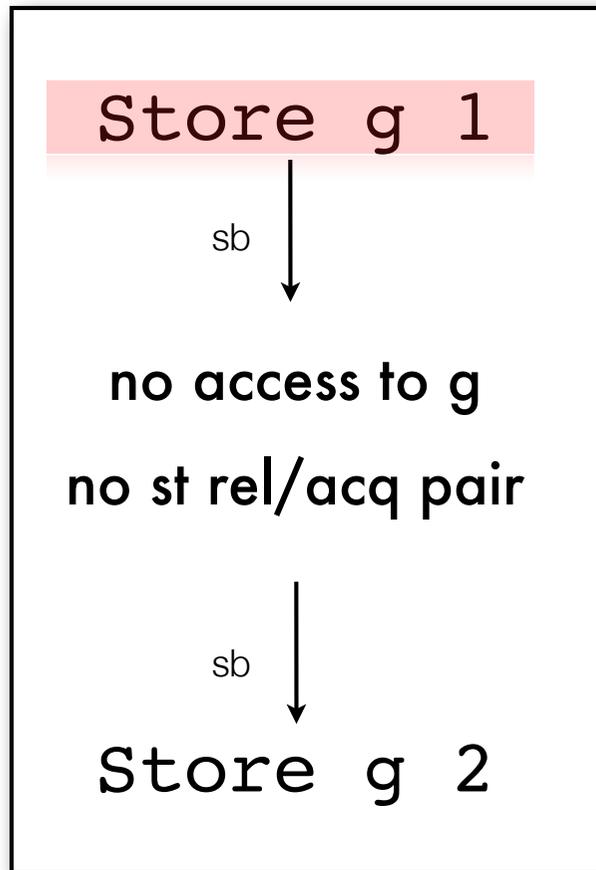
Under which conditions is it correct to eliminate the first store?

A **same-thread release-acquire pair** is a pair of a release action followed by an acquire action in program order.

An action is a *release* if it is a possible source of a synchronisation
unlock mutex, release or seq_cst atomic write

An action is an *acquire* if it is a possible target of a synchronisation
lock mutex, acquire or seq_cst atomic read

Elimination of *overwritten writes*



It is safe to eliminate the first store if there are:

1. no intervening accesses to **g**
2. no intervening same-thread release-acquire pair

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

candidate overwritten write

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;
```

```
f1.store(1, RELEASE);
```

```
while(f2.load(ACQUIRE) == 0);
```

```
g = 2;
```

candidate overwritten write



same-thread release-acquire pair



The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

Thread 2 is non-racy

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

Thread 2 is non-racy
The program should only print **1**

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;
```

```
f1.store(1, RELEASE);
```

```
while(f2.load(ACQUIRE)==0);
```

```
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE)==0);
```

```
printf("%d", g);
```

```
f2.store(1, RELEASE);
```

Thread 2 is non-racy

The program should only print 1

If we perform overwritten write elimination it prints 0

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);  
while(f2.load(ACQUIRE) == 0);  
g = 2;
```

sync

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);
```

```
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

sync



The soundness condition

Shared memory

```
g = 0; atomic f1 = f2 = 0;
```

Thread 1

```
g = 1;  
f1.store(1, RELEASE);
```

```
g = 2;
```

Thread 2

```
while(f1.load(ACQUIRE) == 0);  
printf("%d", g);  
f2.store(1, RELEASE);
```

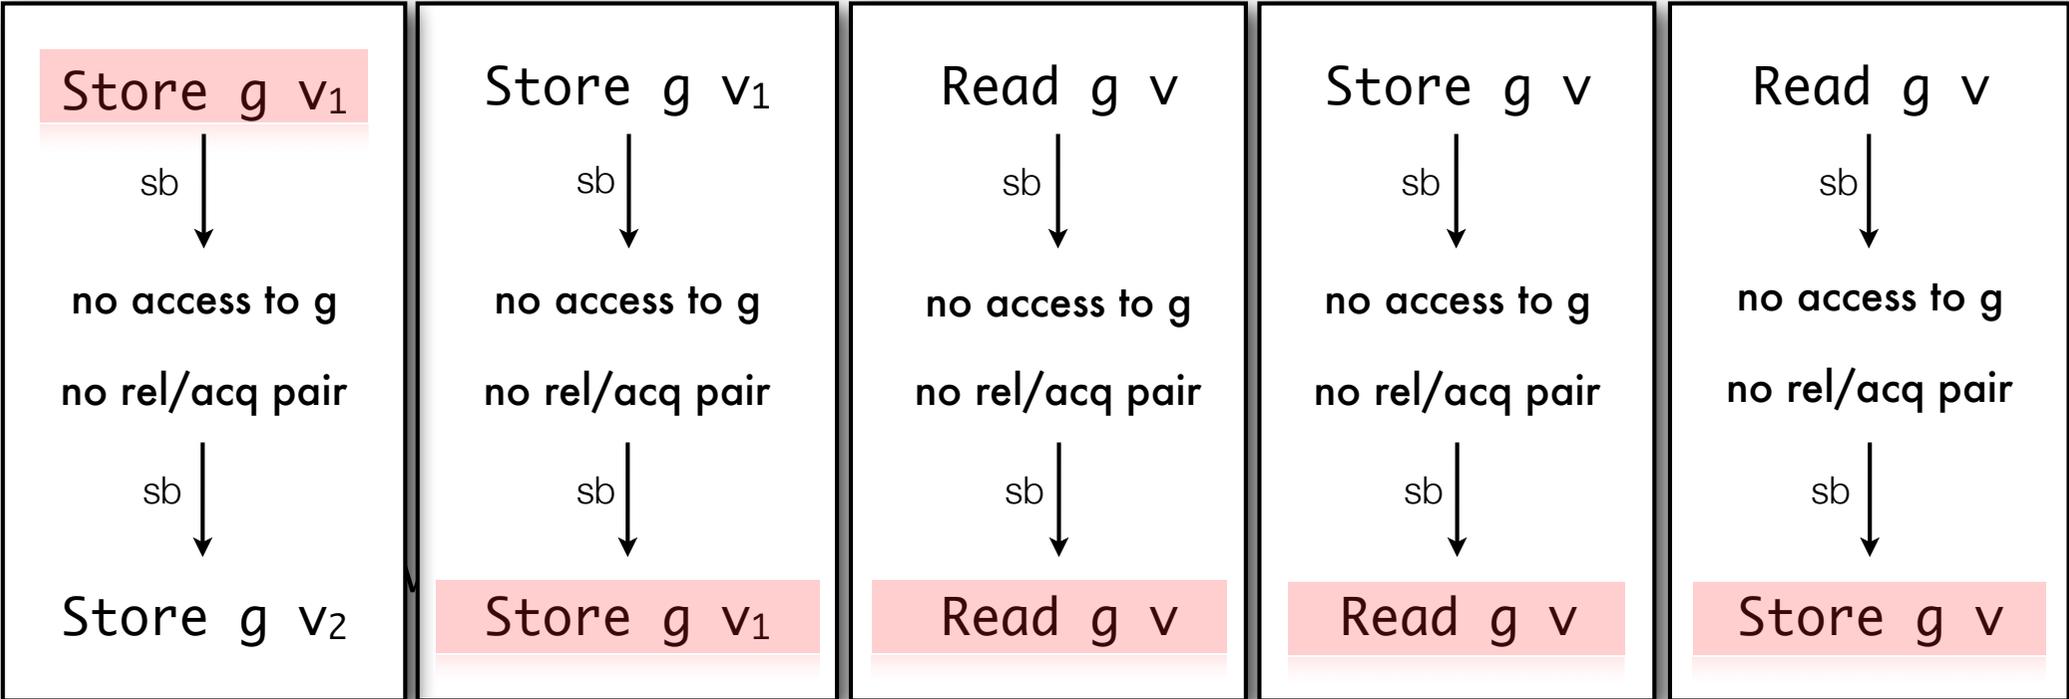
sync

data race

If only a release (or acquire) is present, then
all discriminating contexts *are racy*.

It is sound to optimise the overwritten write.

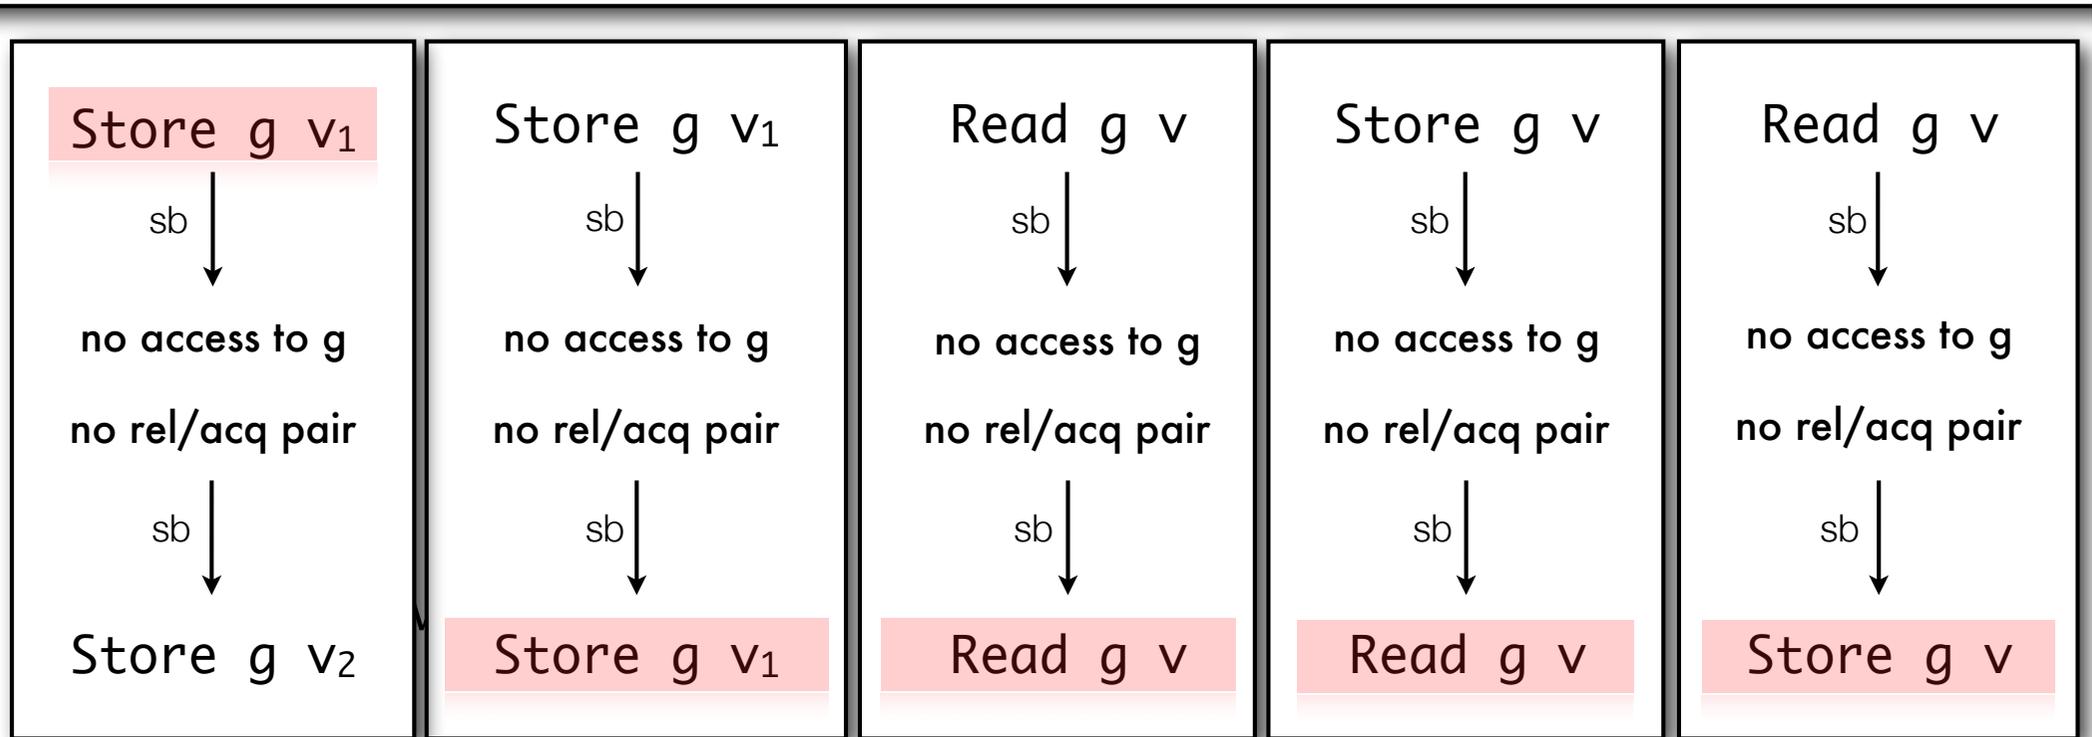
Eliminations: bestiary



Overwritten-Write Write-after-Write Read-after-Read Read-after-Write Write-after-Read

Reads which are not used (via data or control dependencies) to decide a write or synchronisation event are also eliminable (*irrelevant reads*).

Also correctness statements for reorderings, merging, and introductions of events.

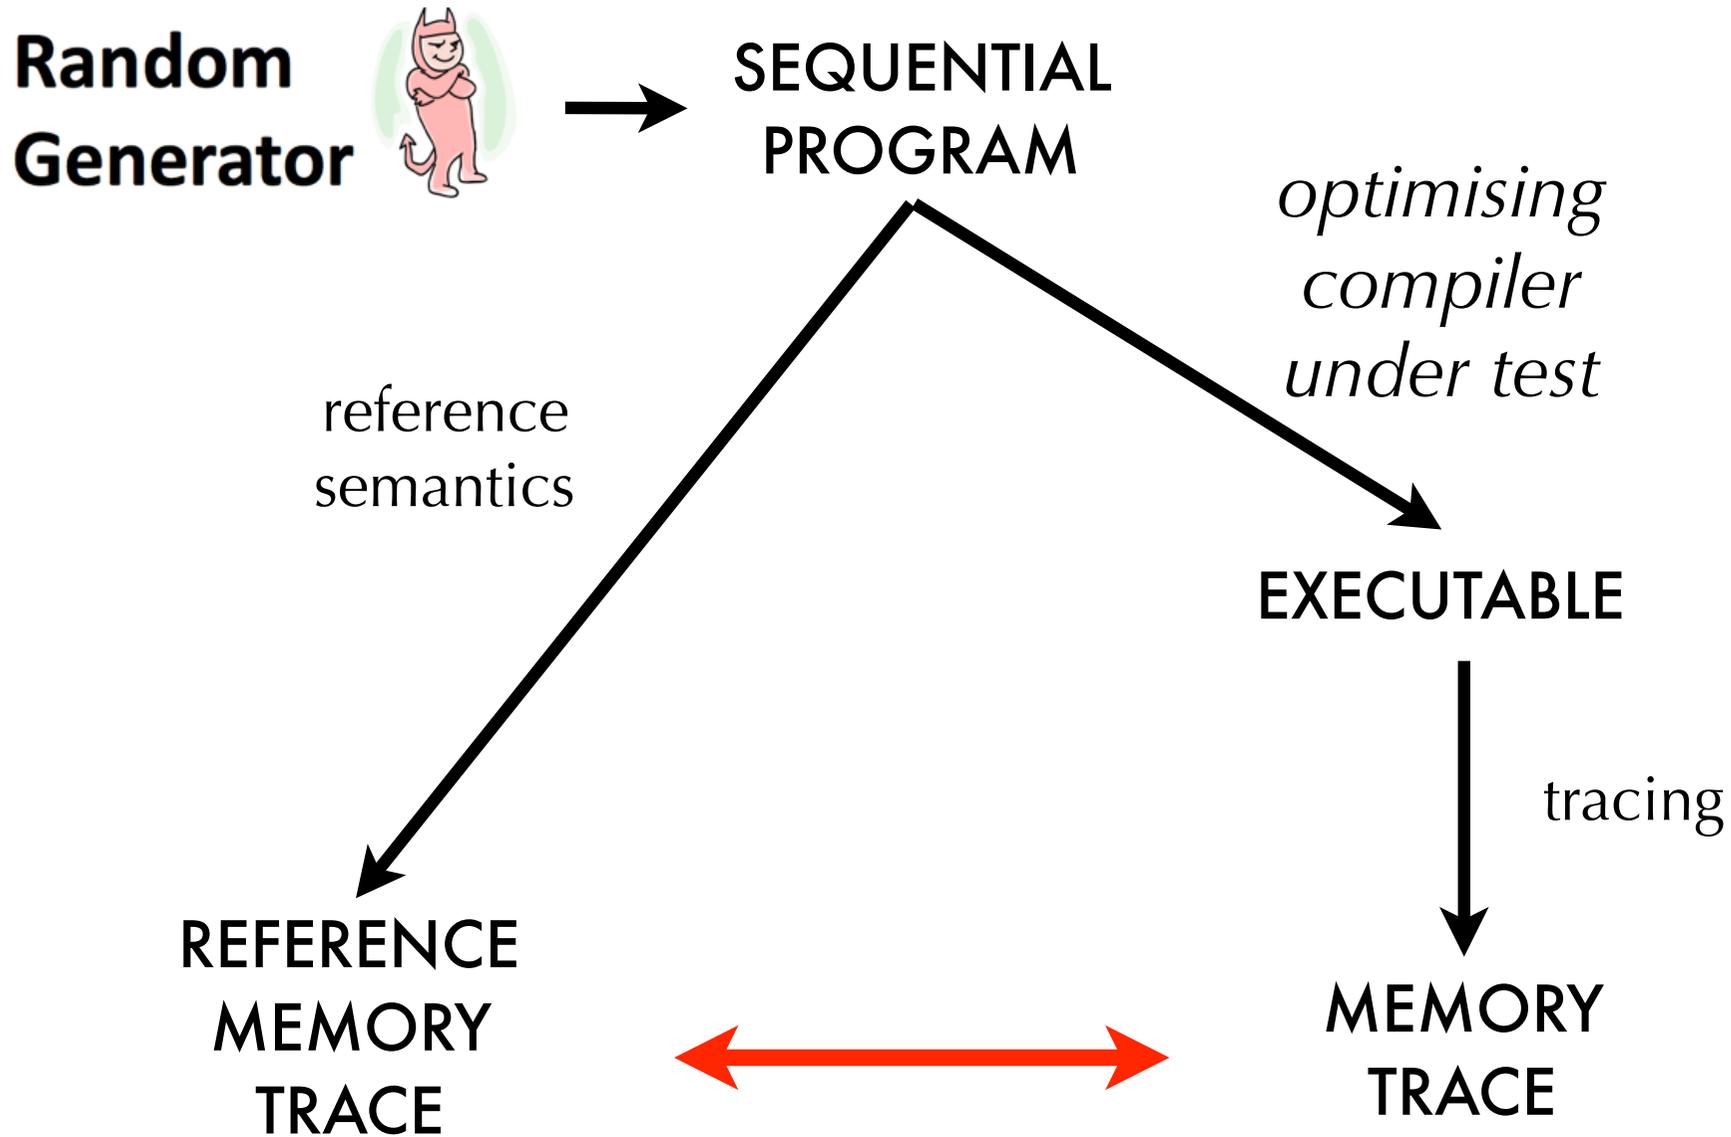


Overwritten-Write Write-after-Write Read-after-Read Read-after-Write Write-after-Read

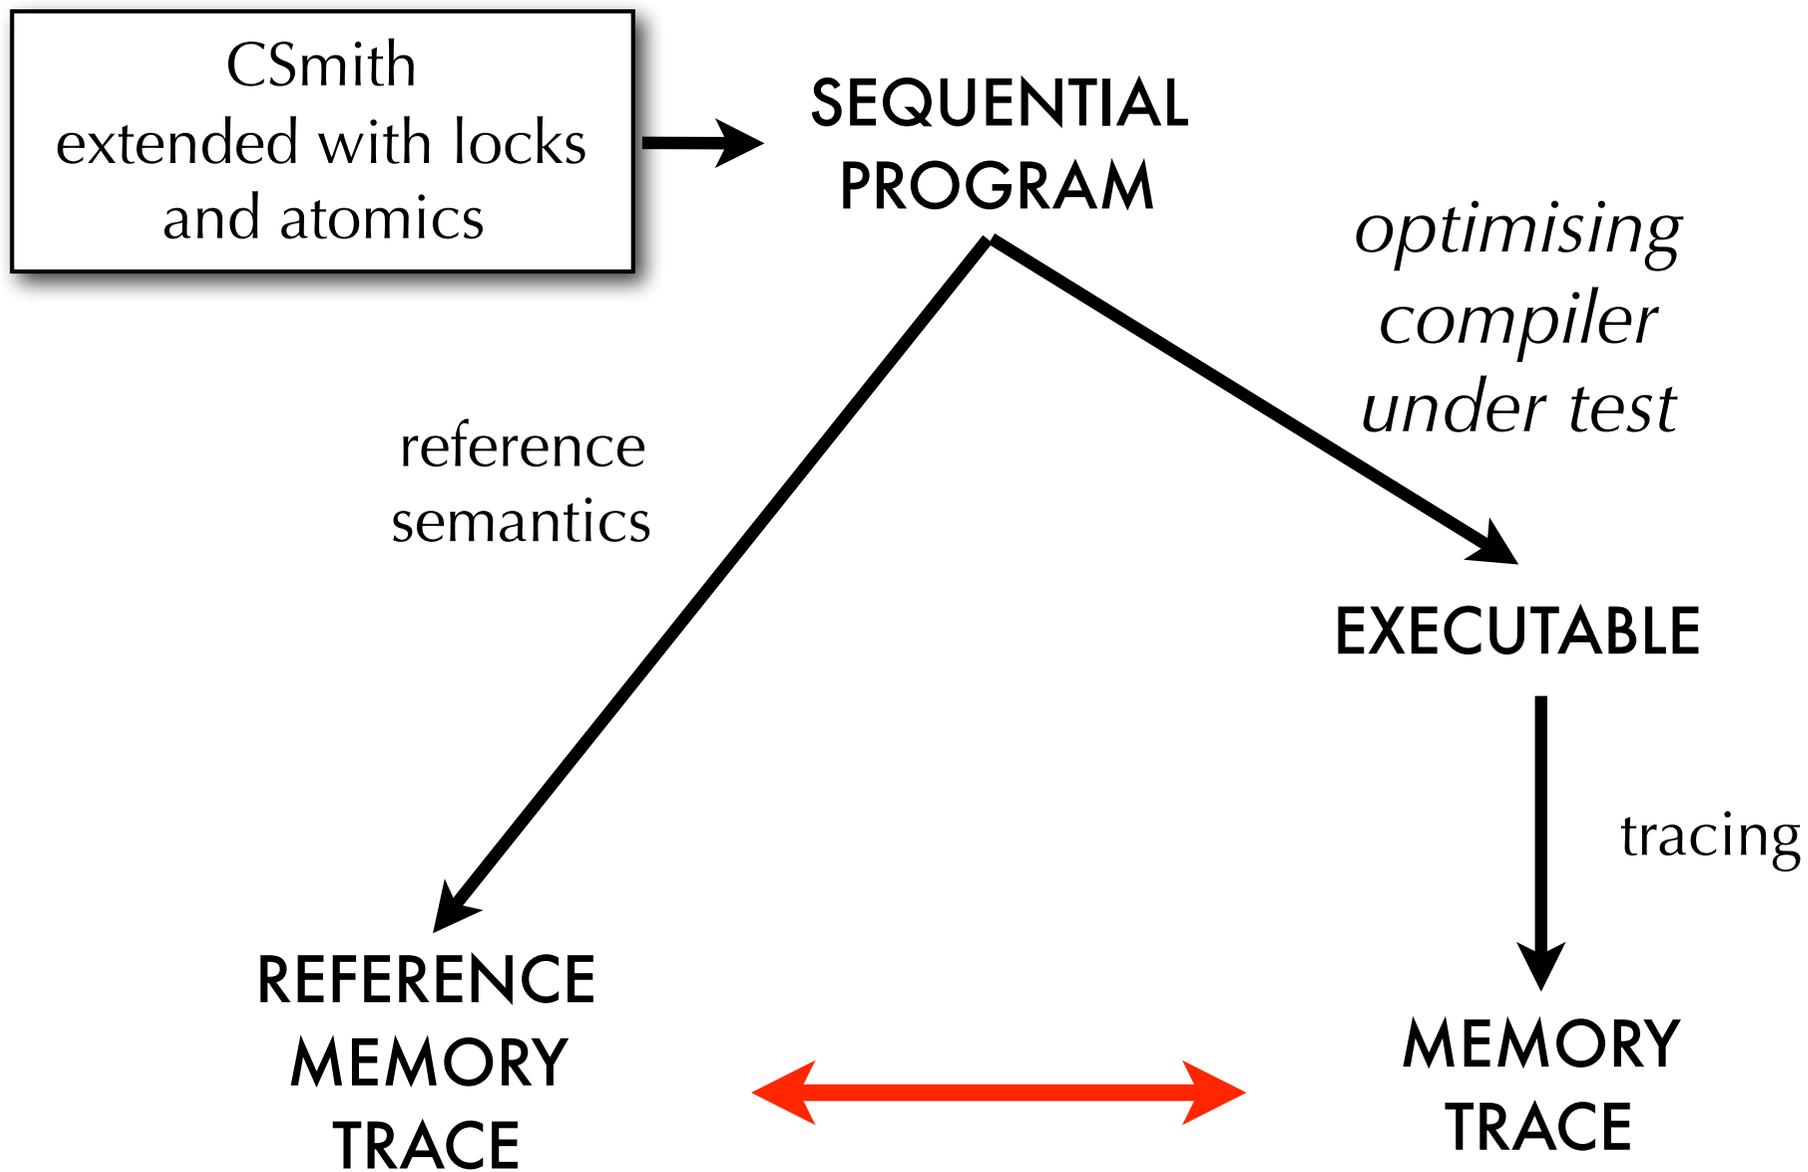
Reads which are not used (via data or control dependencies) to decide a write or synchronisation event are also eliminable (*irrelevant reads*).

From theory to the Cmmtest tool

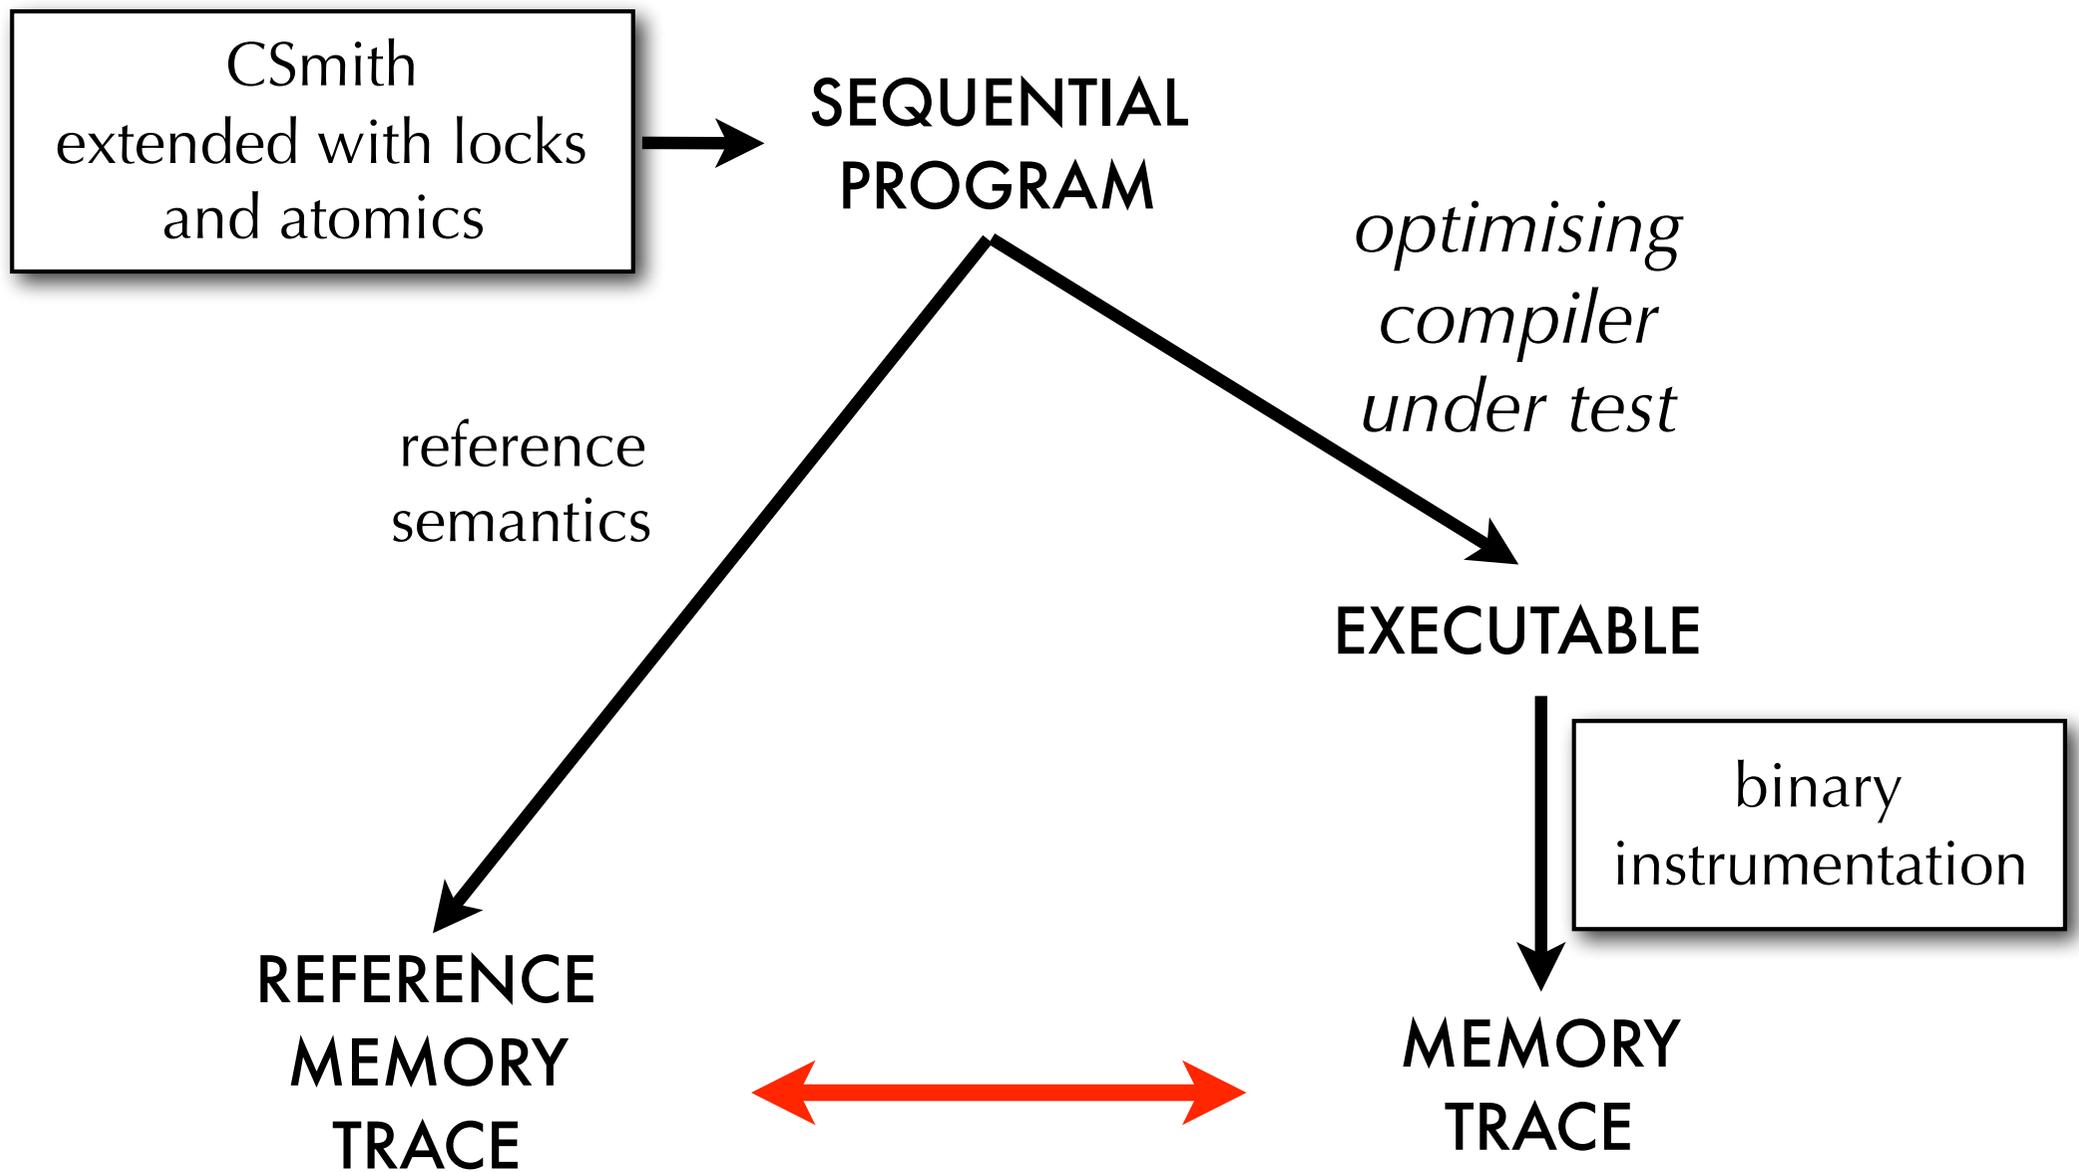




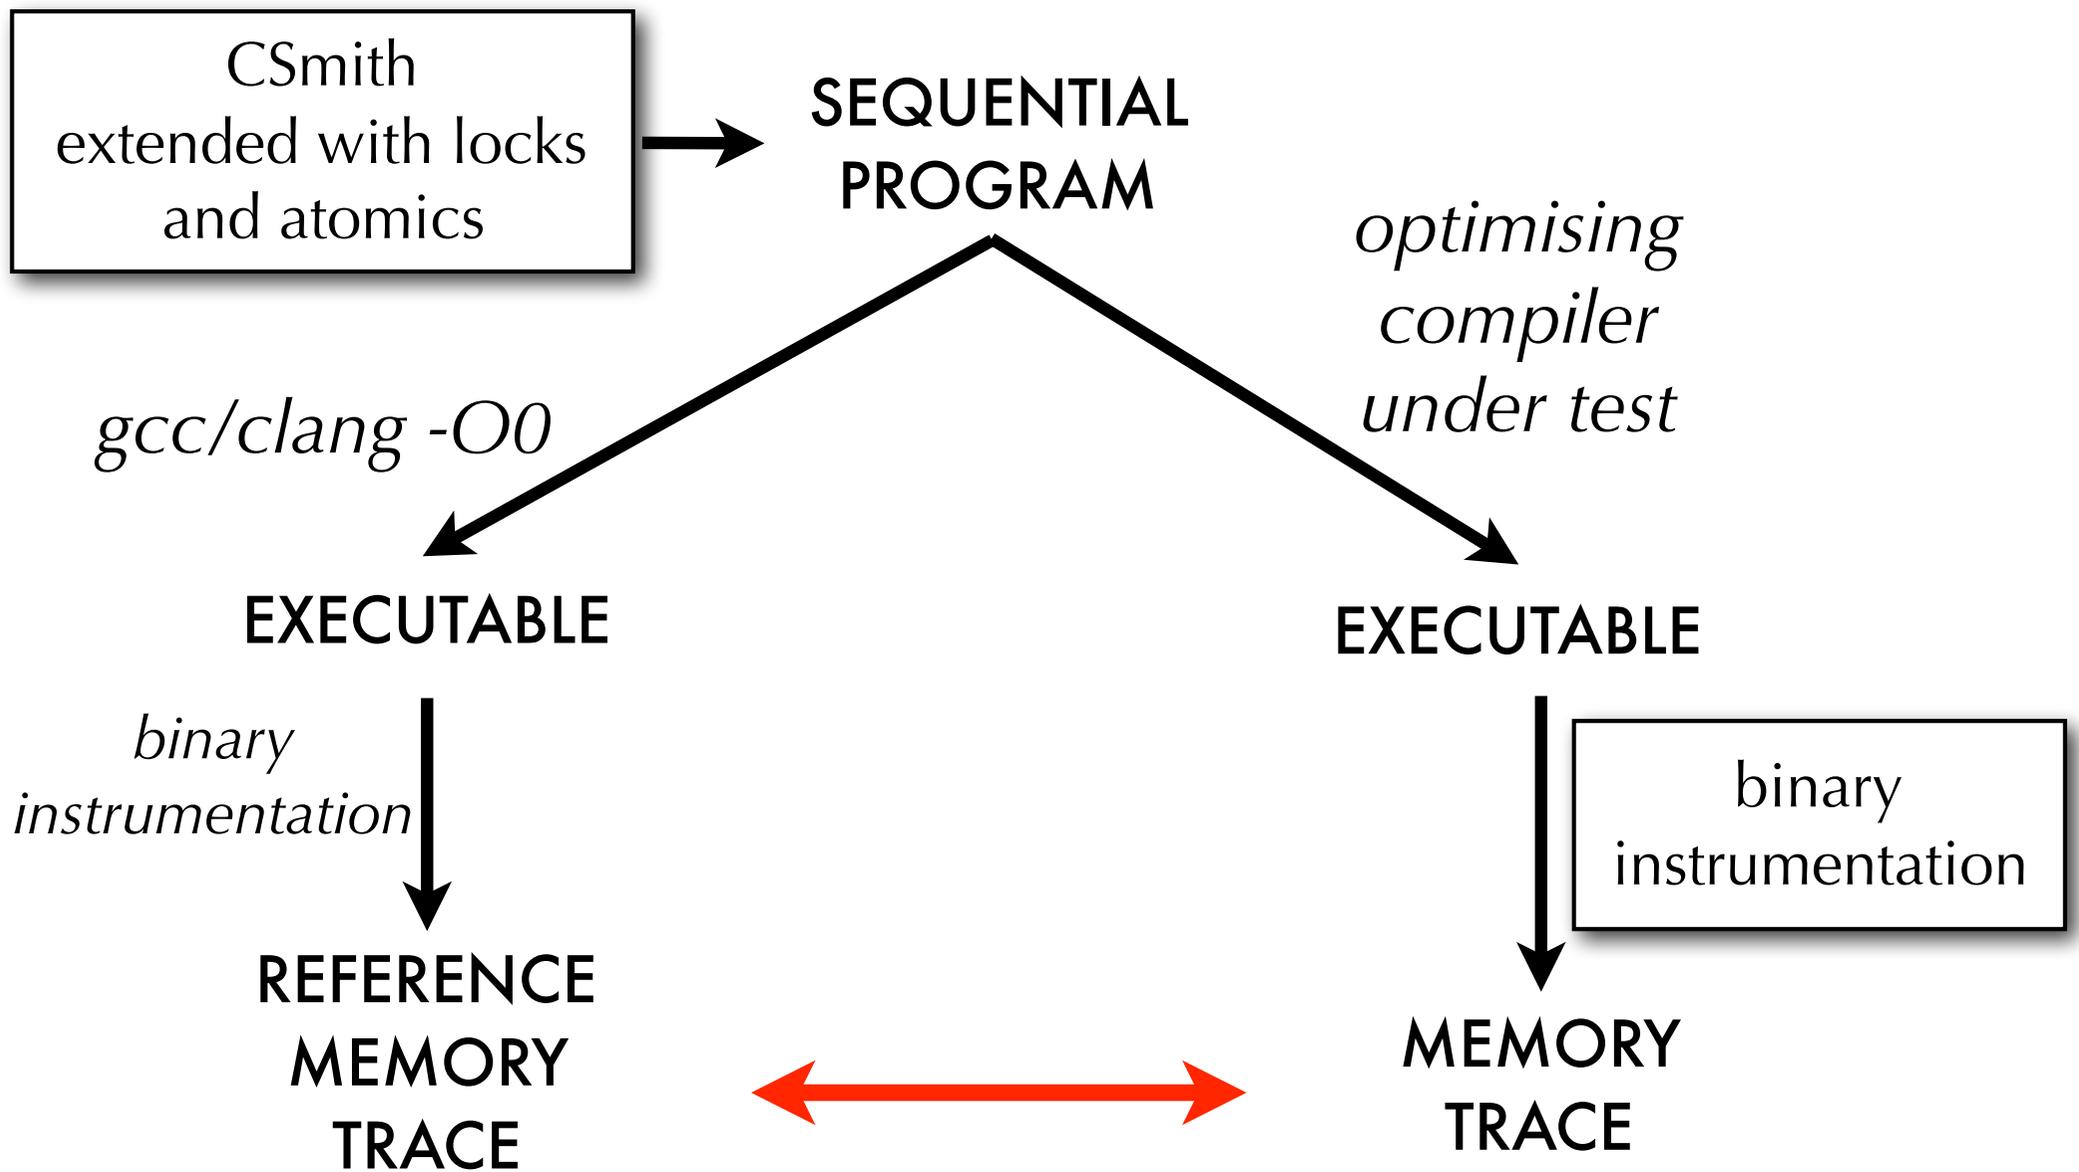
Check: only transformations sound in any concurrent non-racy context



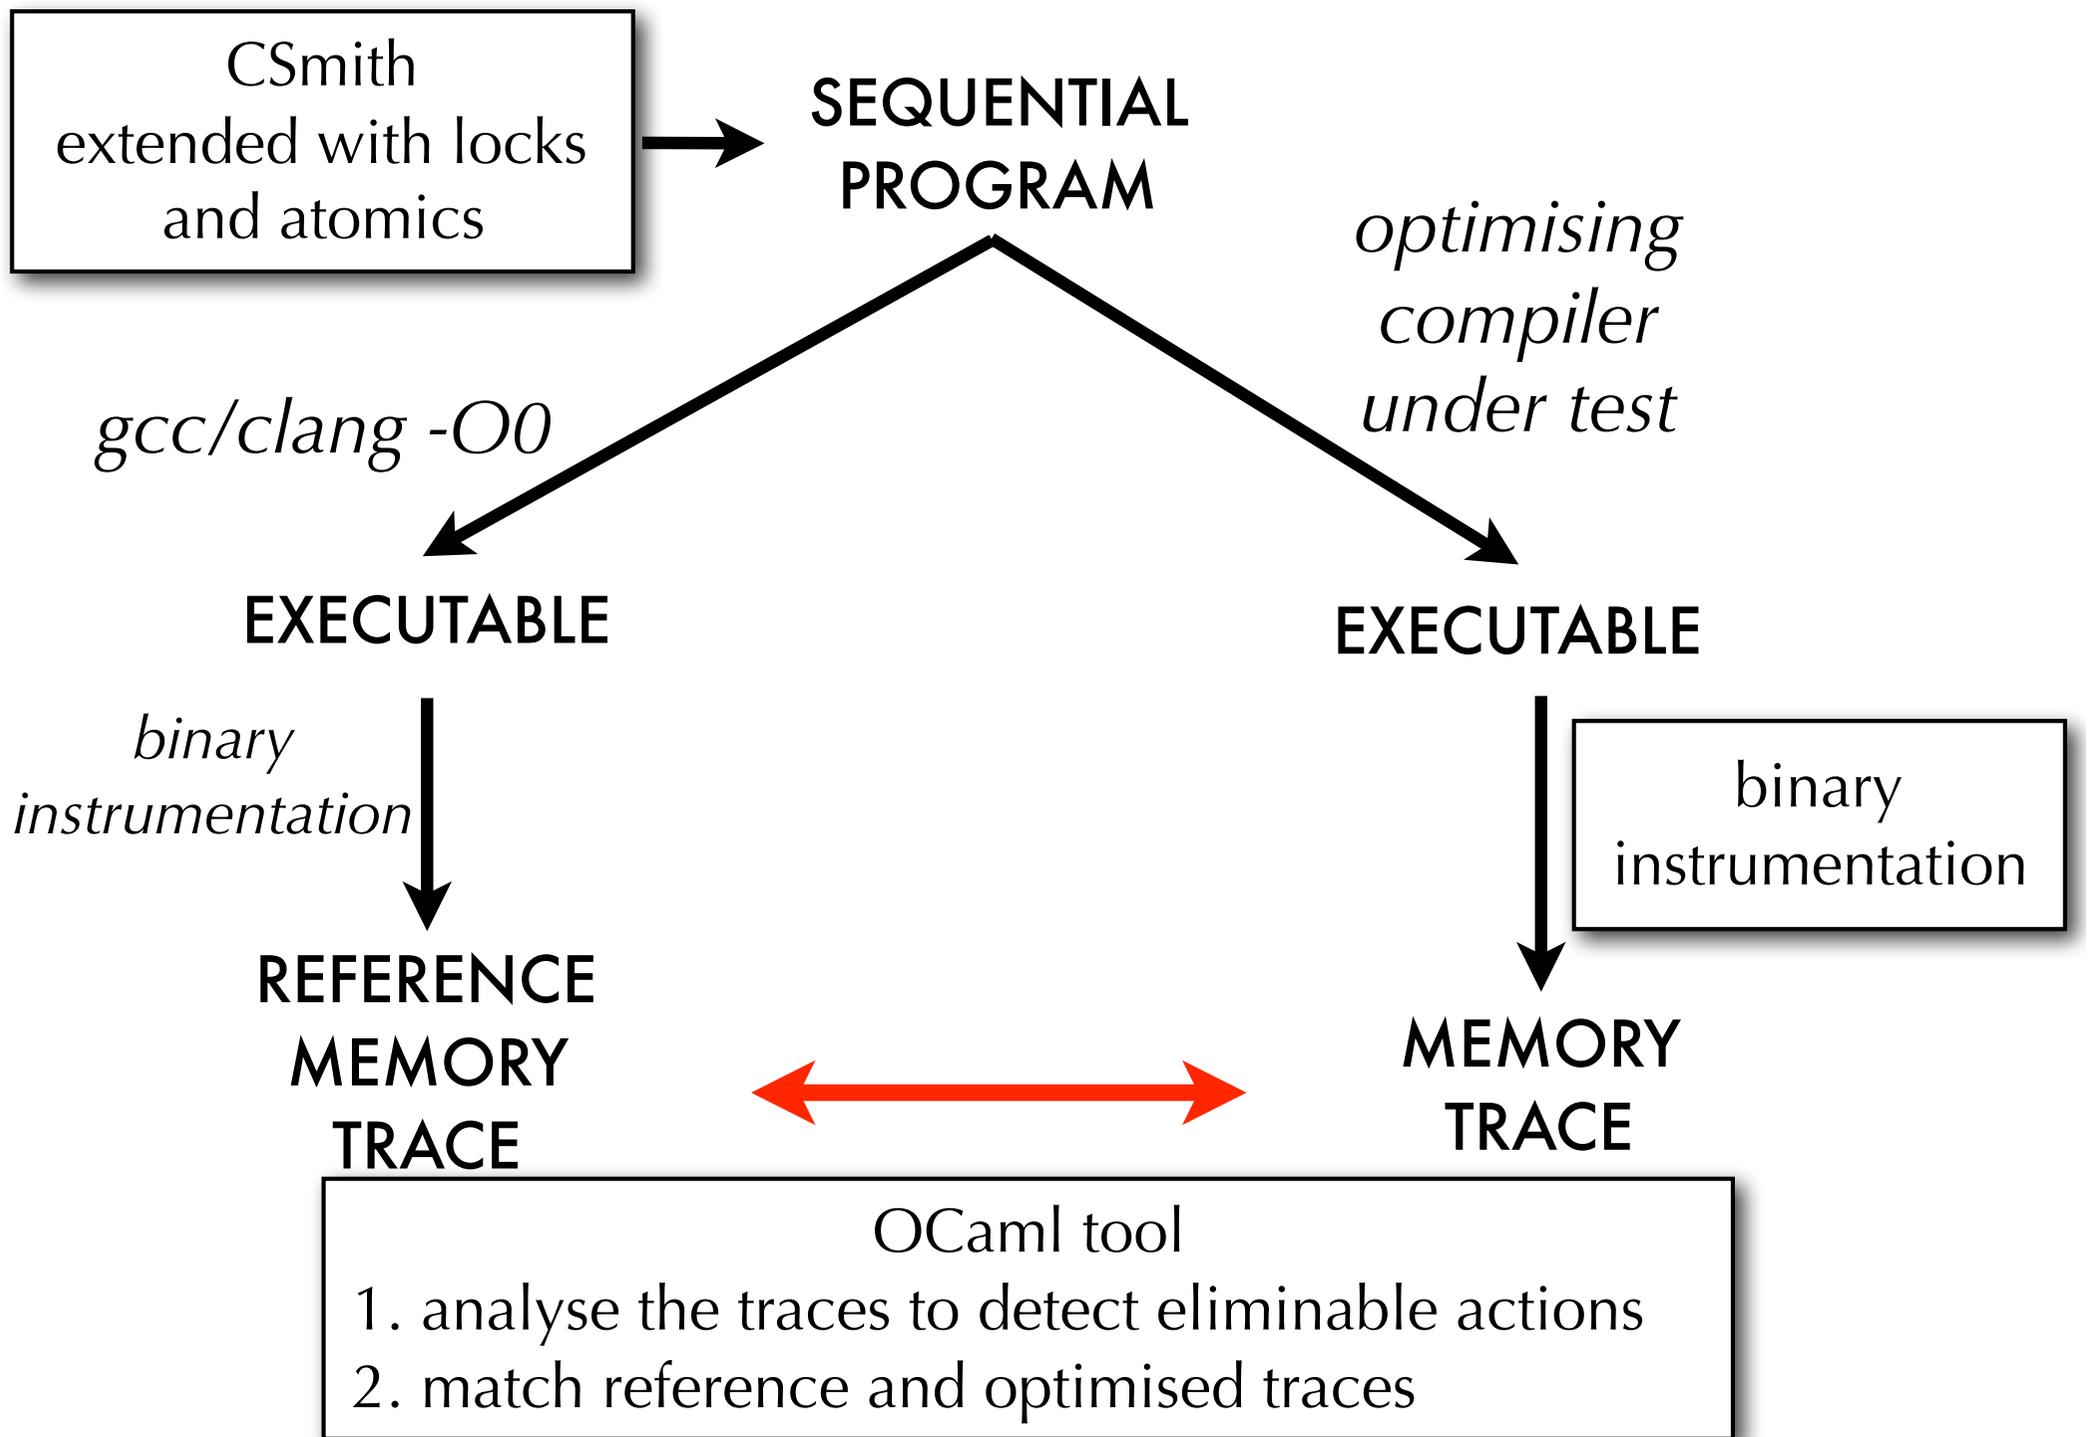
Check: only transformations sound in any concurrent non-racy context



Check: only transformations sound in any concurrent non-racy context



Check: only transformations sound in any concurrent non-racy context



```
const unsigned int g3 = 0UL;
long long g4 = 0x1;
int g6 = 6L;
volatile unsigned int g5 = 1UL;

void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102);
}
```

Start with a randomly generated well-defined program

```
const unsigned int g3 = 0UL; void func_1(void){
long long g4 = 0x1;          int *l8 = &g6;
int g6 = 6L;                int l36 = 0x5E9D070FL;
volatile unsigned int g5 = 1UL; unsigned int l107 = 0xAA37C3ACL;
                             g4 &= g3;
                             g5++;
                             int *l102 = &l36;
                             for (g6 = 4; g6 < (-3); g6 += 1);
                             l102 = &g6;
                             *l102 = ((*l8) && (l107 << 7))*(*l102));
                             }
}
```

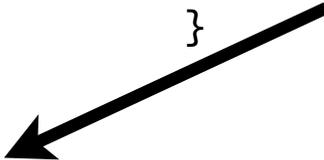
```
Init g3 0
Init g4 1
Init g5 1
Init g6 6
```

```
void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102));
}
```

```
Init g3 0
Init g4 1
Init g5 1
Init g6 6
```

```
void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102));
}
```

reference
semantics



```
Load g4 1
Store g4 0
Load g5 1
Store g5 2
Store g6 4
Load g6 4
Load g6 4
Load g6 4
Store g6 1
Load g4 0
```

```
Init g3 0
Init g4 1
Init g5 1
Init g6 6
```

```
void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102));
}
```

reference
semantics

gcc -O2 memory trace

```
Load g4 1
Store g4 0
Load g5 1
Store g5 2
Store g6 4
Load g6 4
Load g6 4
Load g6 4
Store g6 1
Load g4 0
```

```
Load g5 1
Store g4 0
Store g6 1
Store g5 2
Load g4 0
```

```

Init g3 0
Init g4 1
Init g5 1
Init g6 6

```

```

void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102));
}

```

reference semantics

gcc -O2 memory trace

```

RaW* Load g4 1
      Store g4 0
RaW* Load g5 1
      Store g5 2
OW* Store g6 4
RaW* Load g6 4
RaR* Load g6 4
RaR* Load g6 4
      Store g6 1
RaW* Load g4 0

```

```

Load g5 1
Store g4 0
Store g6 1
Store g5 2
Load g4 0

```



```

Init g3 0
Init g4 1
Init g5 1
Init g6 6

```

```

void func_1(void){
    int *l8 = &g6;
    int l36 = 0x5E9D070FL;
    unsigned int l107 = 0xAA37C3ACL;
    g4 &= g3;
    g5++;
    int *l102 = &l36;
    for (g6 = 4; g6 < (-3); g6 += 1);
    l102 = &g6;
    *l102 = ((*l8) && (l107 << 7))*(*l102));
}

```

reference semantics

gcc -O2 memory trace

```

RaW* Load g4 1
      Store g4 0
RaW* Load g5 1
      Store g5 2
OW* Store g6 4
RaW* Load g6 4
RaR* Load g6 4
RaR* Load g6 4
      Store g6 1
RaW* Load g4 0

```

```

Load g5 1
Store g4 0
Store g6 1
Store g5 2
Load g4 0

```

Init g3 0

```
void func_1(void){
int *l8 = &g6;
```

Can match applying
only correct eliminations and reorderings

```
*l102 = ((*l8) && (l107 << 7))*(*l102));
```

reference semantics

gcc -O2 memory trace

~~RaW*~~ ~~Load~~ ~~g4~~ ~~1~~

Store g4 0

RaW* Load g5 1

Store g5 2

~~OW*~~ ~~Store~~ ~~g6~~ ~~4~~

~~RaW*~~ ~~Load~~ ~~g6~~ ~~4~~

~~RaR*~~ ~~Load~~ ~~g6~~ ~~4~~

~~RaR*~~ ~~Load~~ ~~g6~~ ~~4~~

Store g6 1

RaW* Load g4 0

Load g5 1

Store g4 0

Store g6 1

Store g5 2

Load g4 0

```
int a = 1;
int b = 0;

int s;
for (s=0; s!=4; s++) {
    if (a==1)
        return NULL;
    for (b=0; b>=26; ++b)
        ;
}
```

If we focus on the miscompiled initial example...

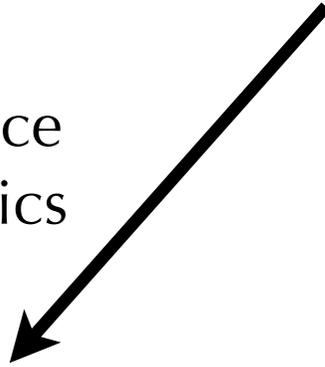
```
int a = 1;  
int b = 0;
```

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return NULL;  
    for (b=0; b>=26; ++b)  
        ;  
}
```

```
int a = 1;  
int b = 0;
```

```
int s;  
for (s=0; s!=4; s++) {  
    if (a==1)  
        return NULL;  
    for (b=0; b>=26; ++b)  
        ;  
}
```

reference
semantics

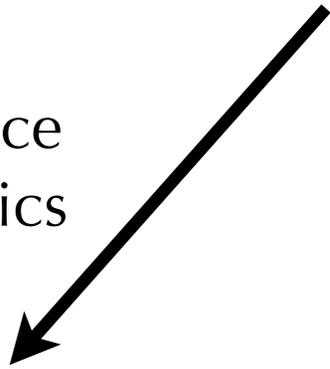


Load a 1

```
int a = 1;
int b = 0;
```

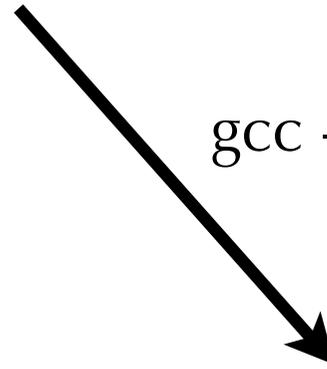
```
int s;
for (s=0; s!=4; s++) {
    if (a==1)
        return NULL;
    for (b=0; b<=26; ++b)
        ;
}
```

reference
semantics



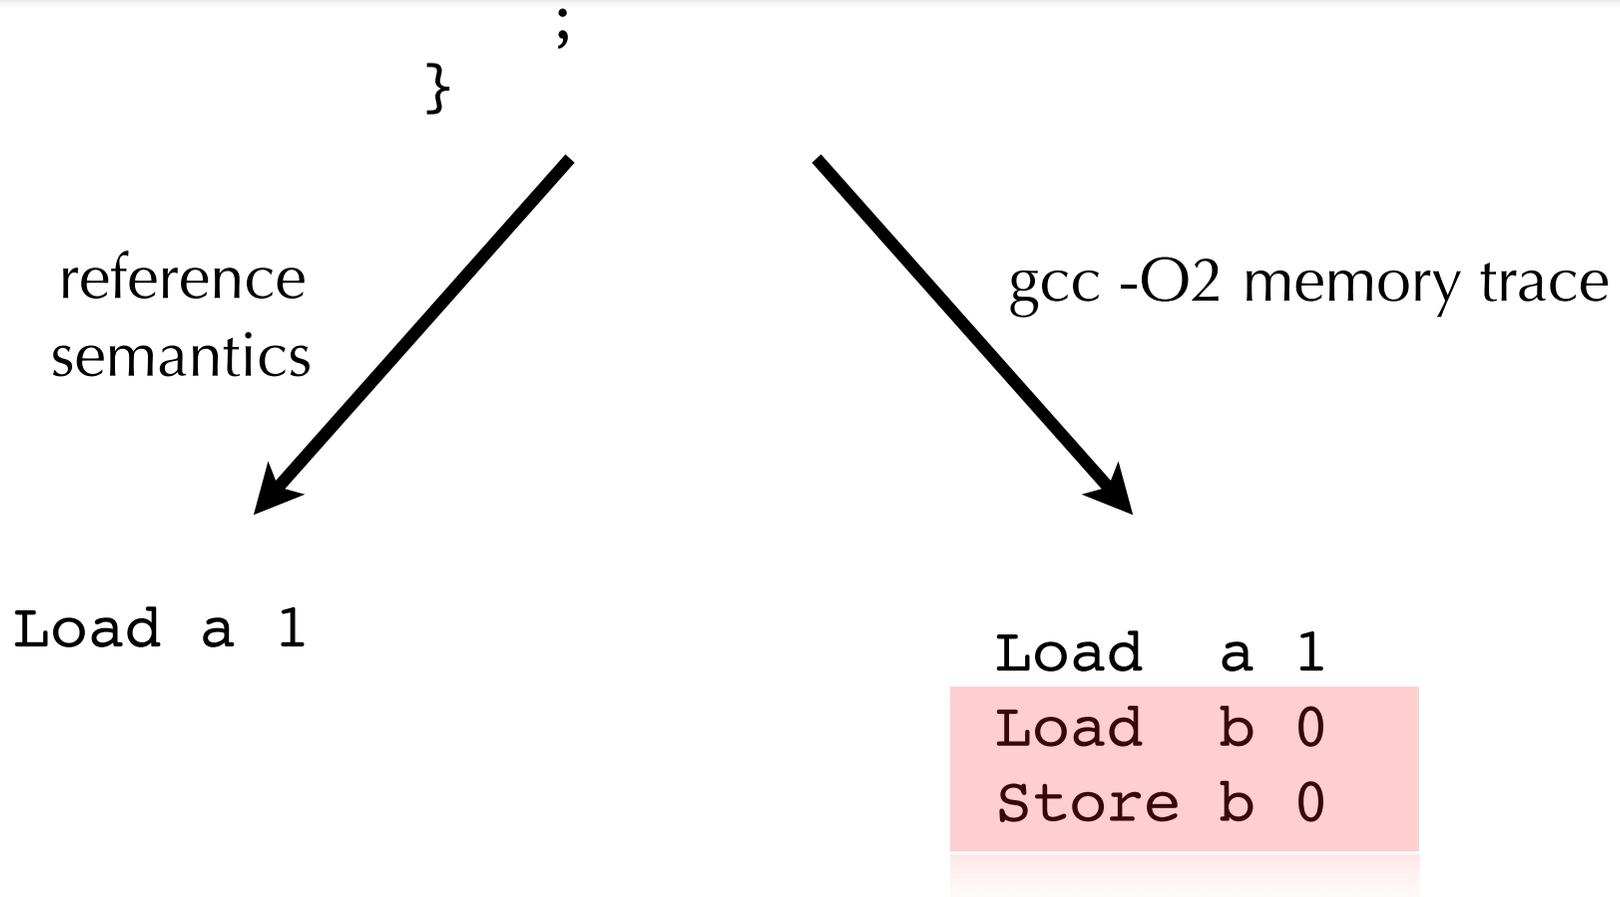
Load a 1

gcc -O2 memory trace



Load a 1
Load b 0
Store b 0

Cannot match some events → detect compiler bug



Applications

2013 - 2016



1. Testing C compilers (GCC, Clang, ICC)

Some concurrency compiler bugs found
in the latest version of GCC.

Store introductions performed by loop invariant motion or
if-conversion optimisations.

Remark: these bugs break the Posix thread model too.

All promptly fixed.

2. Checking compiler invariants

GCC internal invariant: never reorder with an atomic access

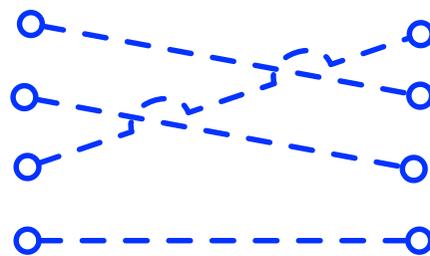
Baked this invariant into the tool and found a counterexample...

...not a bug, but fixed anyway

```
atomic_uint a;  
int32_t g1, g2;
```

```
int main (int, char *[]) {  
    a.load() & a.load ();  
    g2 = g1 != 0;  
}
```

```
ALoad  a  0  
ALoad  a  0  
Load   g1 0  
Store  g2 0
```



```
Load   g1 0  
ALoad  a  0  
ALoad  a  0  
Store  g2 0
```

3. Detecting unexpected behaviours

```
uint16_t g
```

```
for (; g==0; g--);
```



```
uint16_t g
```

```
g=0;
```

Correct or not?

3. Detecting unexpected behaviours

uint16_t g

for (; g==0; g--);

uint16_t g

g=0;



If `g` is initialised with `0`, a load gets replaced by a store:

Load g 0) — ? — (Store g 0

The introduced store cannot be observed by a non-racy context.

Still, arguable if a compiler should do this or not.

3. Detecting unexpected behaviours

uint16_t g

for (; g==0; g--);

uint16_t g

g=0;



If `g` is initialised with `0`, a load gets replaced by a store:

Load

g 0



Store

g 0

False positives in Thread Sanitizer

The formalisation of the C11 memory model enables compiler testing... what else?



Proving the correctness of mappings for atomics

<https://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html>

C/C++11 Operation	ARM implementation
Load Relaxed:	ldr
Load Consume:	ldr + preserve dependencies until next kill_dependency OR ldr; teq; beq; isb OR ldr; dmb
Load Acquire:	ldr; teq; beq; isb OR ldr; dmb
Load Seq Cst:	ldr; dmb
Store Relaxed:	str
Store Release:	dmb; str
Store Seq Cst:	dmb; str; dmb
Cmpxchg Relaxed (32 bit):	_loop: ldrex roldval, [rptr]; mov rres, 0; teq roldval, rold; strexeq rres, rnewval, [rptr]; teq rres, 0; bne _loop
Cmpxchg Acquire (32 bit):	_loop: ldrex roldval, [rptr]; mov rres, 0; teq roldval, rold; strexeq rres, rnewval, [rptr]; teq rres, 0; bne _loop; isb
Cmpxchg Release (32 bit):	dmb; _loop: ldrex roldval, [rptr]; mov rres, 0; teq roldval, rold; strexeq rres, rnewval, [rptr]; teq rres, 0; bne _loop;
Cmpxchg AcqRel (32 bit):	dmb; _loop: ldrex roldval, [rptr]; mov rres, 0; teq roldval, rold; strexeq rres, rnewval, [rptr]; teq rres, 0; bne _loop; isb
Cmpxchg SeqCst (32 bit):	dmb; _loop: ldrex roldval, [rptr]; mov rres, 0; teq roldval, rold; strexeq rres, rnewval, [rptr]; teq rres, 0; bne _loop; dmb
Acquire Fence:	dmb
Release Fence:	dmb
AcqRel Fence:	dmb
SeqCst Fence:	dmb

Inform new optimisations

e.g. the work by Robin Morisset on the Arm LLVM backend

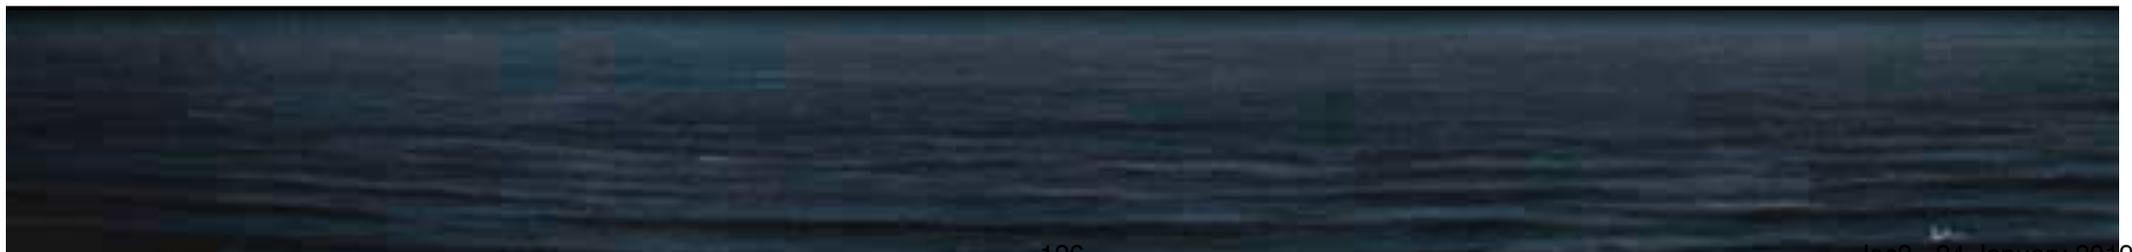
```
while (flag.load(acquire))  
{  
}
```

```
.loop  
ldr r0, [r1]  
dmb ish  
bnz .loop
```

```
.loop  
ldr r0, [r1]  
bnz .loop  
dmb ish
```



Out of thin-air reads



Memory access synchronisation

`x = y = 0`

Thread 1

```
    y = 1  
x.store(1, MO_RELEASE)
```

Thread 2

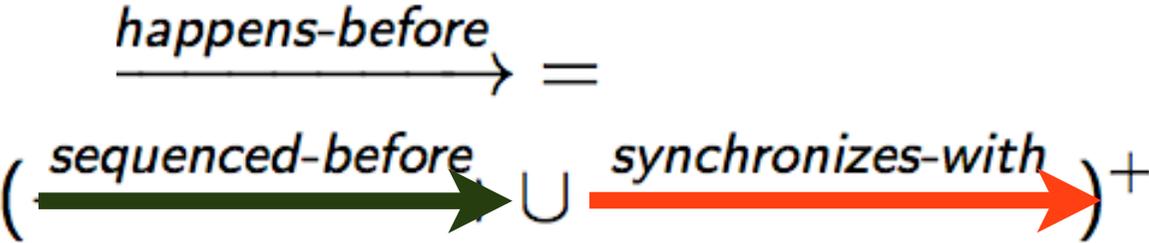
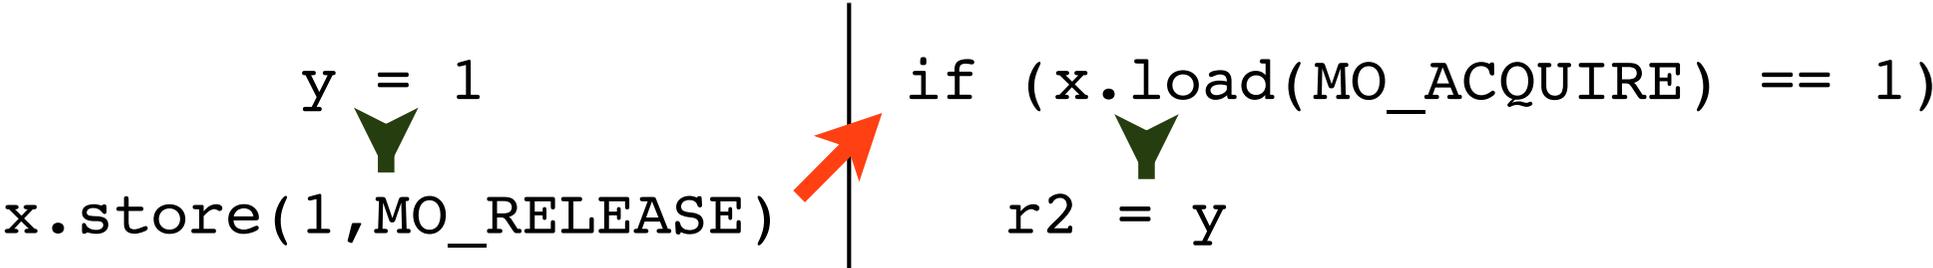
```
if (x.load(MO_ACQUIRE) == 1)  
    r2 = y
```

Memory access synchronisation

`x = y = 0`

Thread 1

Thread 2



Non-atomic loads must return the most recent write in the happens-before order

Understanding MO_RELAXED

`x = y = 0`

Thread 1

Thread 2

```
    y = 1  
x.store(1,MO_RELAXED)
```

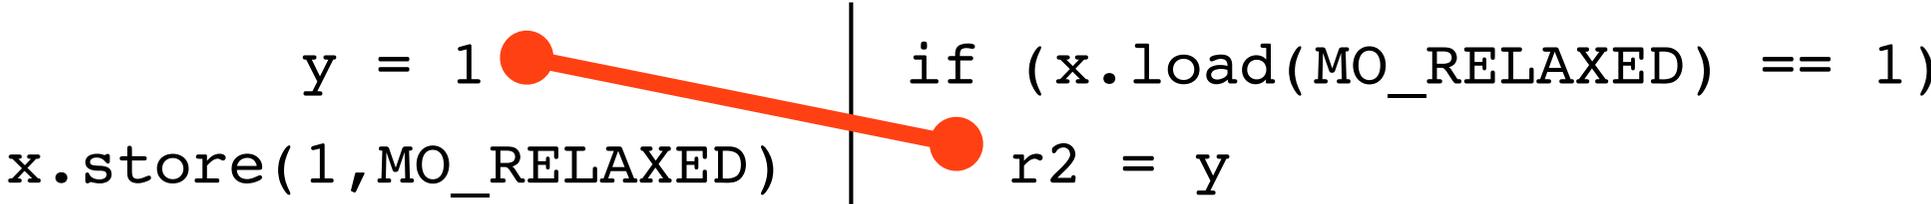
```
if (x.load(MO_RELAXED) == 1)  
    r2 = y
```

Understanding MO_RELAXED

`x = y = 0`

Thread 1

Thread 2



DATA RACE

Two conflicting accesses not related by happens-before

Understanding `MO_RELAXED`

`x = y = 0`

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

WELL DEFINED

but `r2 = 0` is possible

Understanding MO_RELAXED

$x = y = 0$

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

Allow a RELAXED load to see any store that:

- does not happens-after it
 - is not hidden by an intervening store hb-ordered between them
-

Intuition

the compiler (or hardware) can reorder independent accesses

$x = y = 0$

Thread 1

```
y.store(1,MO_RELAXED)
x.store(1,MO_RELAXED)
```

Thread 2

```
if (x.load(MO_RELAXED) == 1)
    r2 = y.load(MO_RELAXED)
```

Allow a RELAXED load to see any store that:

- does not happens-after it
 - is not hidden by an intervening store hb-ordered between them
-

Shorthand

from now on, all the memory accesses are
atomic with `MO_RELAXED` semantics

Out-of-thin-air

Thread 1

`r1 = x`

`y = r1`

`x = y = 0`

|

Thread 2

`r2 = y`

`x = 42`

Out-of-thin-air

$x = y = 0$

Thread 1

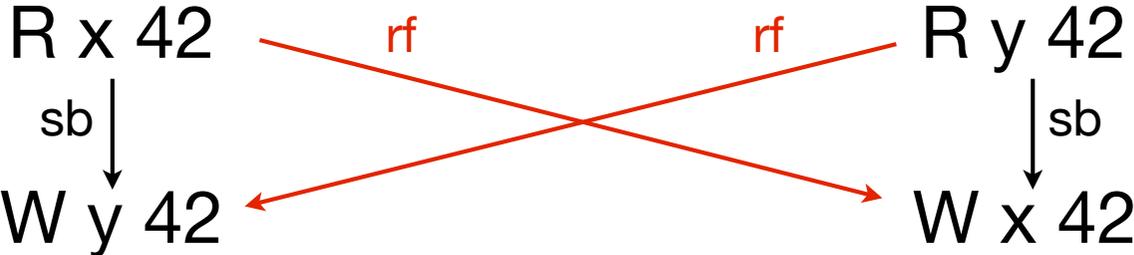
$r1 = x$
 $y = r1$

Thread 2

$r2 = y$
 $x = 42$

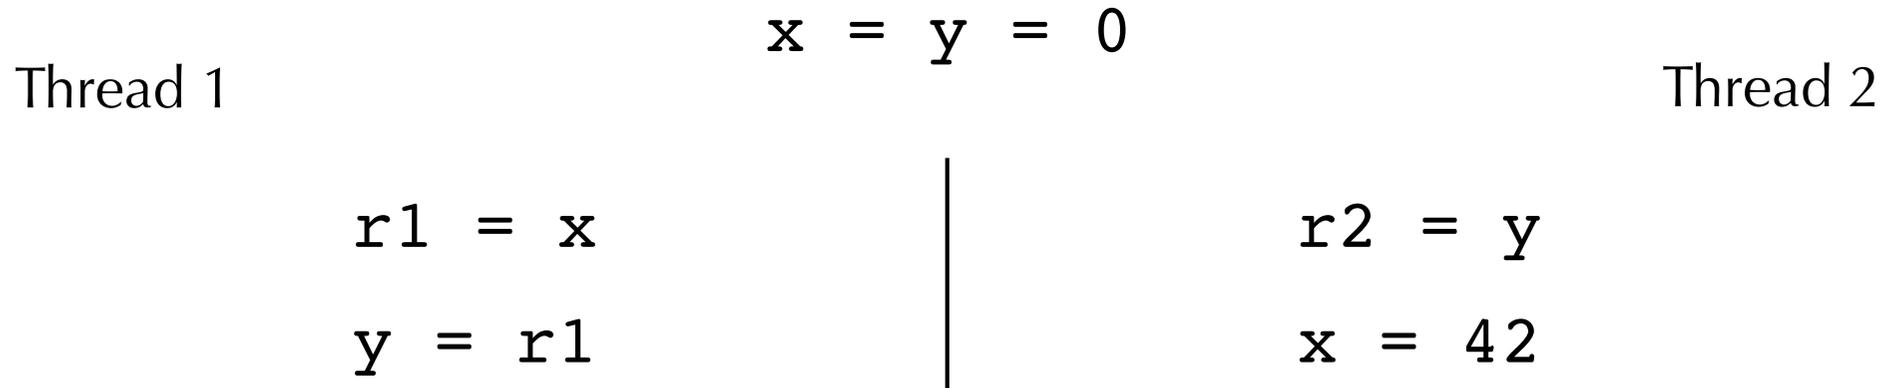
$r1 = r2 = 42$

is a valid execution.

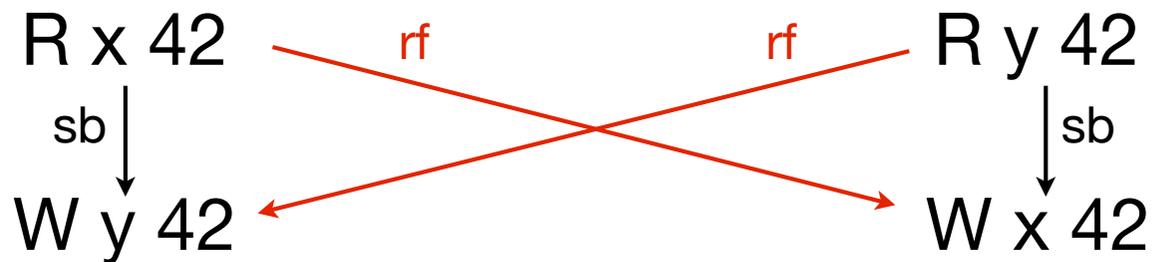


Intuition

the compiler (or hardware) can reorder independent accesses



$r1 = r2 = 42$
is a valid execution.



Out-of-thin-air reads

Thread 1

```
r1 = x  
y = r1
```

$x = y = 0$

|

Thread 2

```
r2 = y  
x = r2
```

Out-of-thin-air reads

Thread 1

```
r1 = x
y = r1
```

```
x = y = 0
```

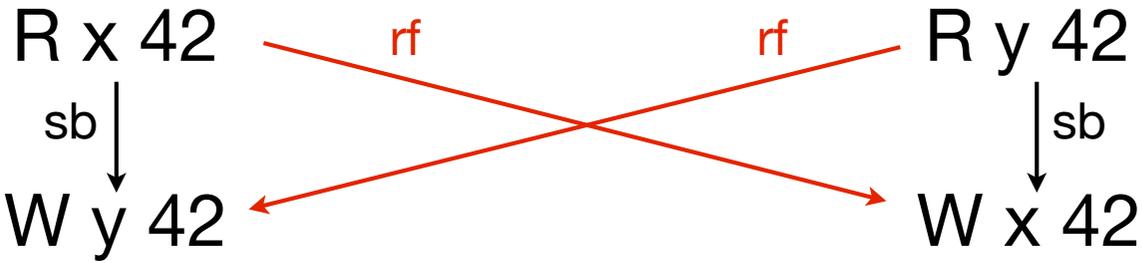


Thread 2

```
r2 = y
x = r2
```

$$r1 = r2 = 42$$

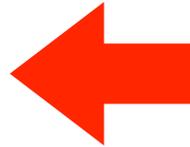
is also an allowed execution



Speculation can justify out-of-thin-air reads

If the compiler states that x is likely to hold 42...

```
y := 42
r1 := x
if (r1 != 42) y := r1;
print r1
```

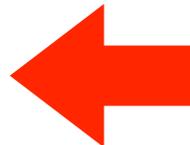


initially $x = y = 0$	
r1 := x	r2 := y
y := r1	x := r2
print r1	print r2

Speculation can justify out-of-thin-air reads

If the compiler states that x is likely to hold 42...

```
y := 42
r1 := x
if (r1 != 42) y := r1;
print r1
```

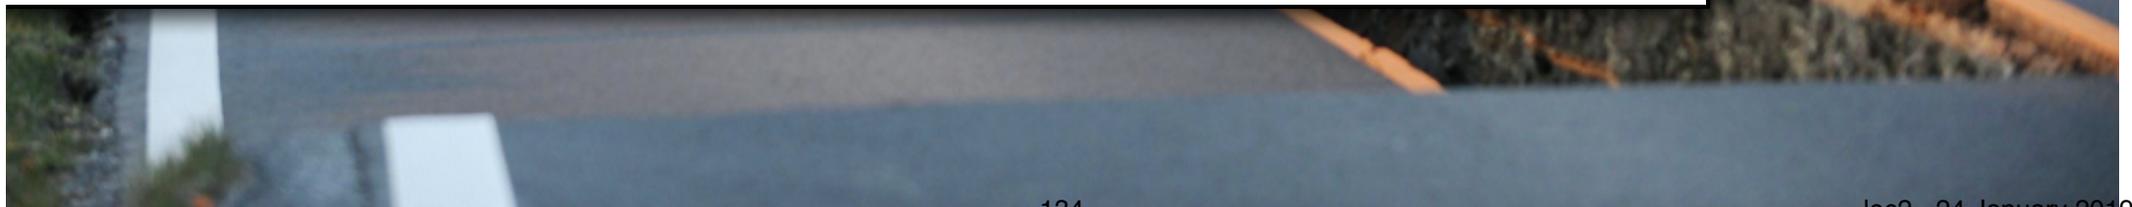


initially $x = y = 0$	
r1 := x	r2 := y
y := r1	x := r2
print r1	print r2

It does not happen in practice... even if it might!



Consequences of out-of-thin-air reads

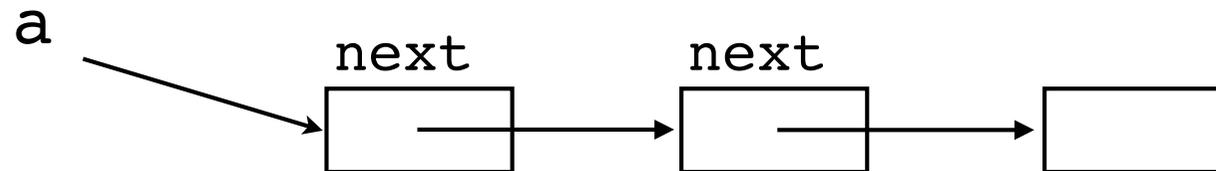


```
struct foo {  
    atomic<struct foo *> next;  
}  
struct foo *a;
```



Thread 1

```
r1 = a->next  
r1->next = a
```

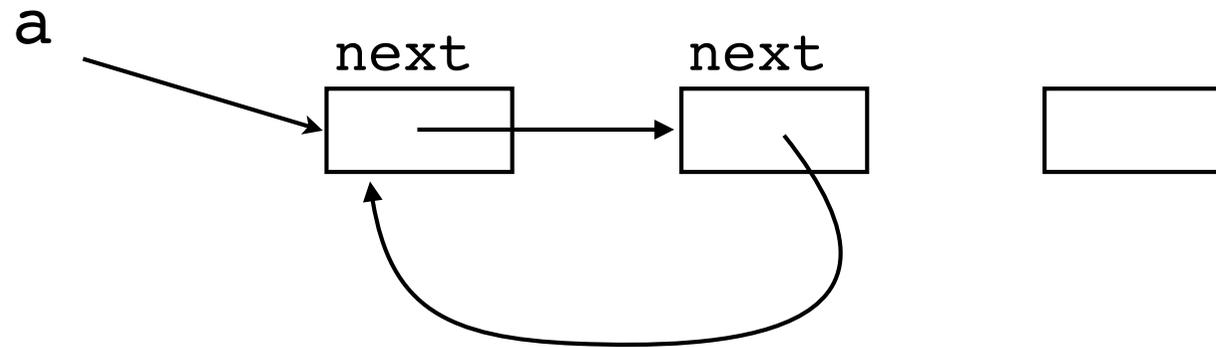


```
struct foo {  
    atomic<struct foo *> next;  
}  
struct foo *a;
```



Thread 1

```
r1 = a->next  
r1->next = a
```



```
struct foo {  
    atomic<struct foo *> next;  
}  
struct foo *a, *b;
```



Thread 1

```
r1 = a->next  
r1->next = a
```

Thread 2

```
r2 = b->next  
r2->next = b
```

```
struct foo {
    atomic<struct foo *> next;
}
struct foo *a, *b;
```



Thread 1

```
r1 = a->next
r1->next = a
```

Thread 2

```
r2 = b->next
r2->next = b
```

If **a** and **b** initially reference disjoint data-structures
we expect **a** and **b** to remain disjoint

```
struct foo {  
    atomic<struct foo *> next;  
}  
struct foo *a, *b;
```

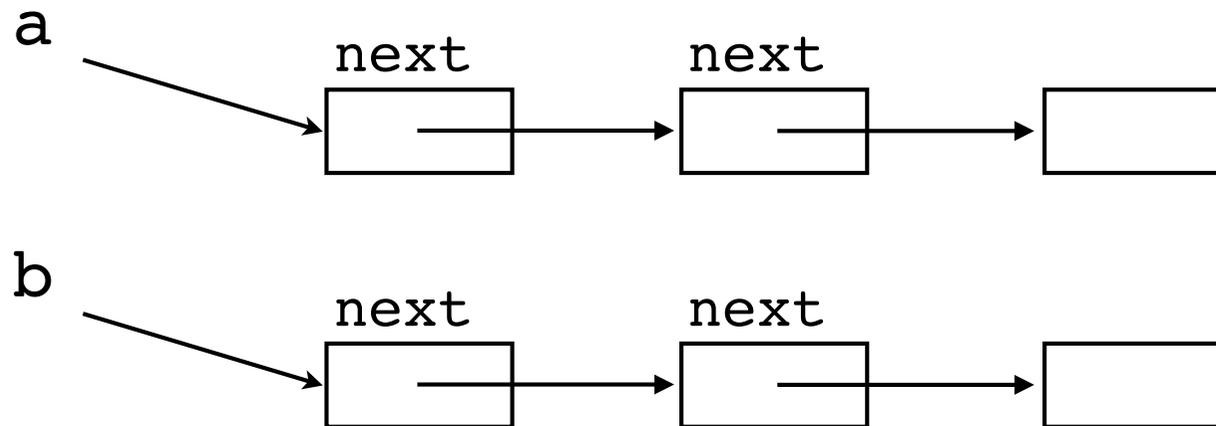


Thread 1

```
r1 = a->next  
r1->next = a
```

Thread 2

```
r2 = b->next  
r2->next = b
```



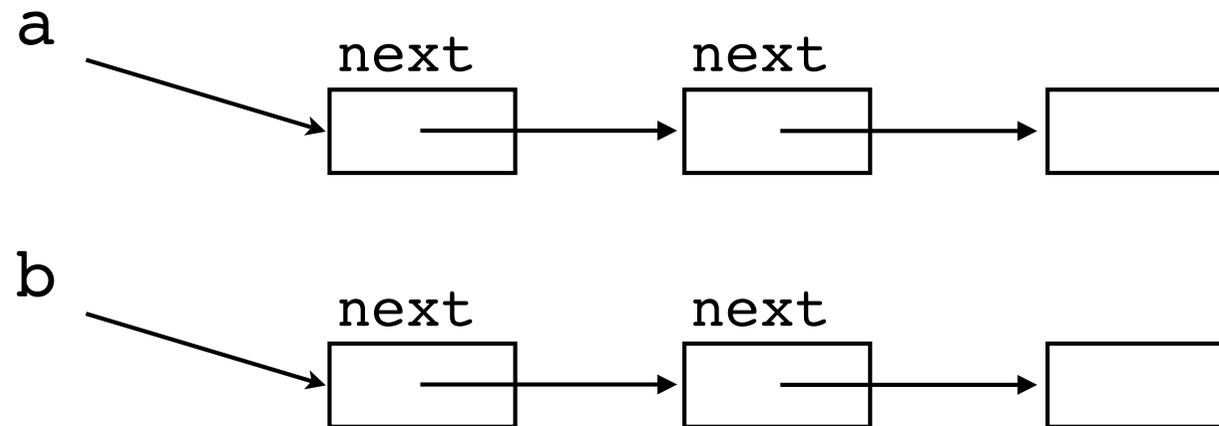
If the compiler speculates $r1=b$ and $r2=a$, then
the store $r1 \rightarrow next = a$ justifies $r2 = b \rightarrow next$ assigning $r2 = a$
(and symmetrically to justify $r1 = b$)

Thread 1

```
r1 = a->next  
r1->next = a
```

Thread 2

```
r2 = b->next  
r2->next = b
```



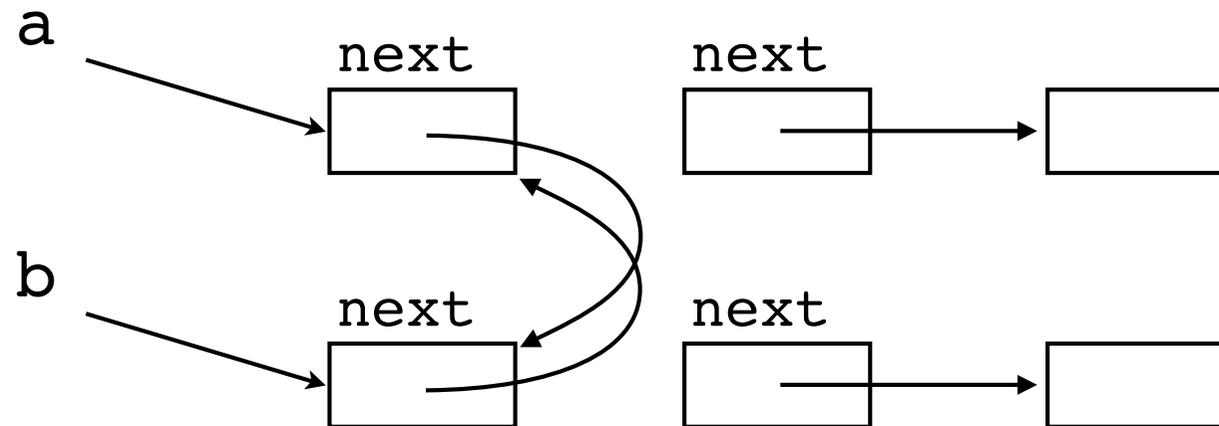
If the compiler speculates $r1=b$ and $r2=a$, then
the store $r1 \rightarrow next = a$ justifies $r2 = b \rightarrow next$ assigning $r2 = a$
(and symmetrically to justify $r1 = b$)

Thread 1

```
r1 = a->next  
r1->next = a
```

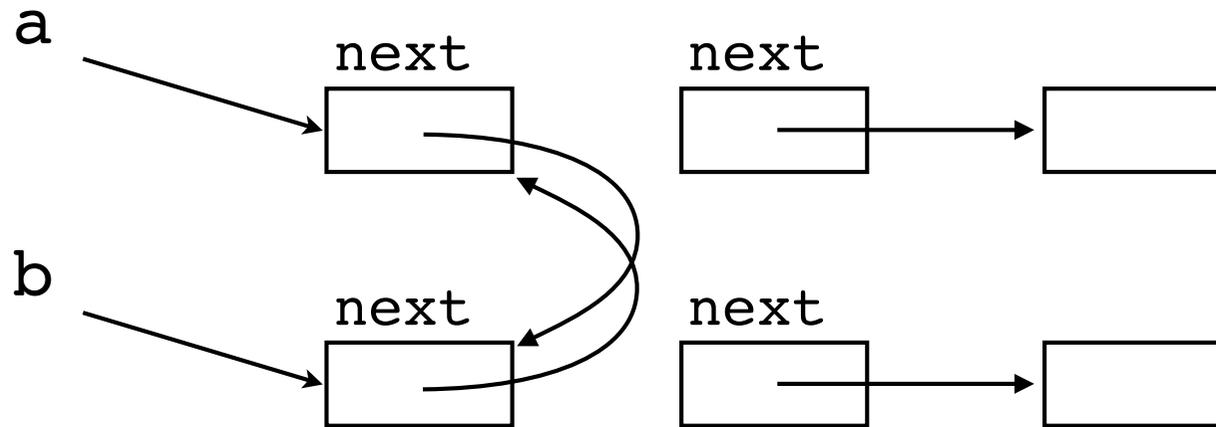
Thread 2

```
r2 = b->next  
r2->next = b
```



If the compiler speculates $r1=b$ and $r2=a$, then
the store $r1 \rightarrow next = a$ justifies $r2 = b \rightarrow next$ assigning $r2 = a$
(and symmetrically to justify $r1 = b$)

Break our basic intuitions about memory and sharing!





CN

Common compiler optimisations are unsound in C11

XINHUA

`x = y = a = 0`

```
if (x.load(rlx)==42)
    y.write(42,rlx)
```

```
if (y.load(rlx)==42)
    if (a==1)
        x.write(42,rlx)
```

`a = 1`

`x = y = a = 0`

<code>if (x.load(rlx)==42)</code>		<code>if (y.load(rlx)==42)</code>		<code>a = 1</code>
<code> y.write(42,rlx)</code>		<code> if (a==1)</code>		
		<code> x.write(42,rlx)</code>		

Remark 1

This code is not racy!

There is no consistent execution in which the read of `a` occurs.

$x = y = a = 0$

<pre>if (x.load(rlx)==42) y.write(42,rlx)</pre>		<pre>if (y.load(rlx)==42) if (a==1) x.write(42,rlx)</pre>		<pre>a = 1</pre>
---	--	---	--	------------------

Remark 2

$a = 1 \wedge x = y = 0$

is the only possible final state

`x = y = a = 0`

```
if (x.load(rlx)==42)
    y.write(42,rlx)
```

```
if (y.load(rlx)==42)
    if (a==1)
        x.write(42,rlx)
```

`a = 1`

Consider sequentialisation:

`C || D \implies C ; D`

(ought to be correct)

`x = y = a = 0`

```
if (x.load(rlx)==42)
  y.write(42,rlx)
```

```
if (y.load(rlx)==42)
  if (a==1)
    x.write(42,rlx)
```

`a = 1`



```
if (x.load(rlx)==42)
  y.write(42,rlx)
```

```
a = 1
if (y.load(rlx)==42)
  if (a==1)
    x.write(42,rlx)
```

`x = y = a = 0`

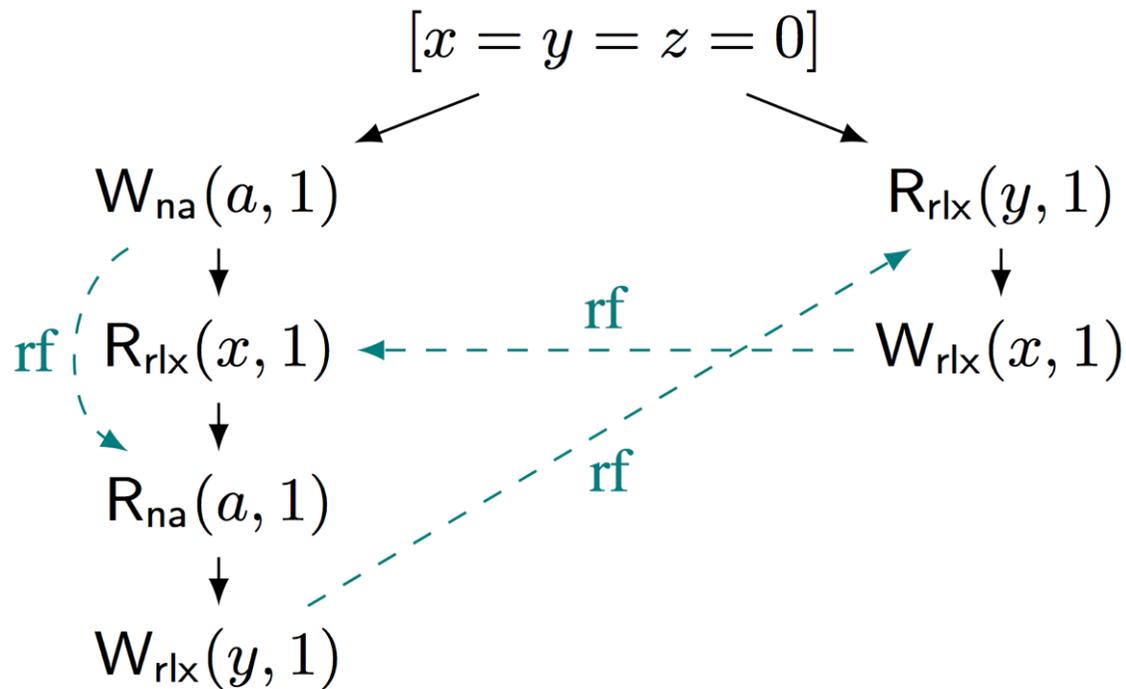
```
if (x.load(rlx)==42)
    y.write(42,rlx)
```

```
a = 1
if (y.load(rlx)==42)
    if (a==1)
        x.write(42,rlx)
```

$x = y = a = 0$

```
if (x.load(rlx)==42)  
    y.write(42,rlx)
```

```
a = 1  
if (y.load(rlx)==42)  
    if (a==1)  
        x.write(42,rlx)
```



$a = 1$
 $x = y = 42$
is also possible

`x = y = a = 0`

```
if (x.load(r1x)==42)
    y.write(42,r1x)
```

```
a = 1
if (y.load(r1x)==42)
    if (a==1)
        x.write(42,r1x)
```

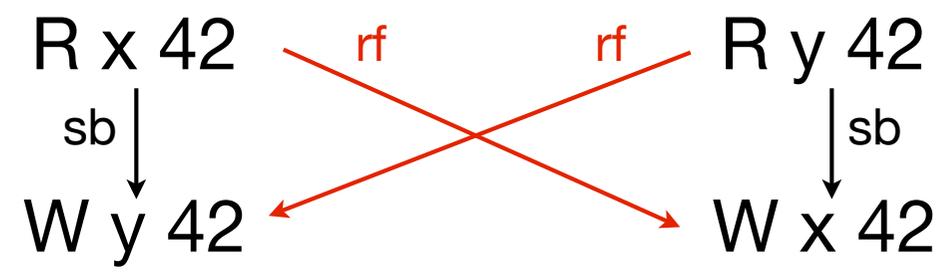
Break common source-to-source (or LLVM IR - to - LLVM IR) compiler optimisations

including *expression linearisation* and *roach-motel reorderings*

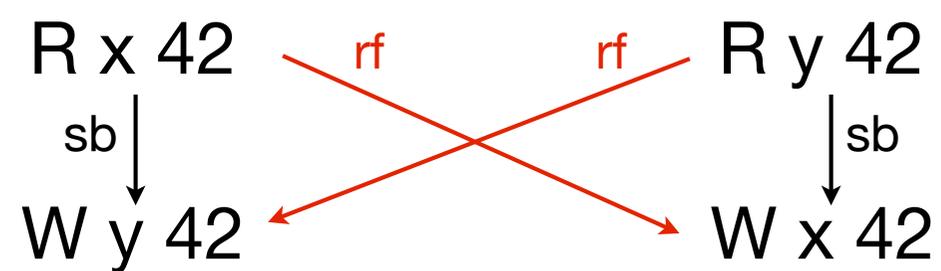


Are there any solutions?

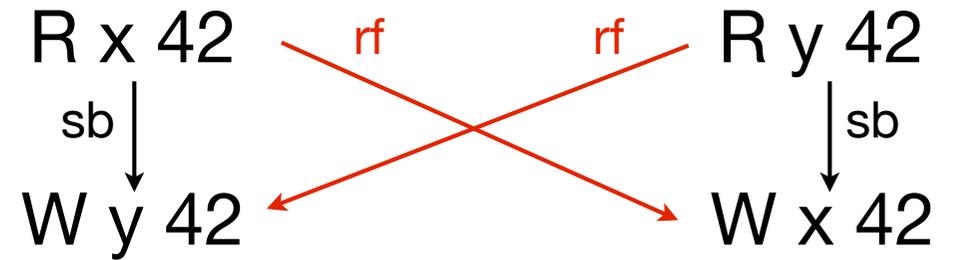
Thread 0	Thread 1
r1 = x y = r1	r2 = y x = 42



Thread 0	Thread 1
r1 = x y = r1	r2 = y x = r2

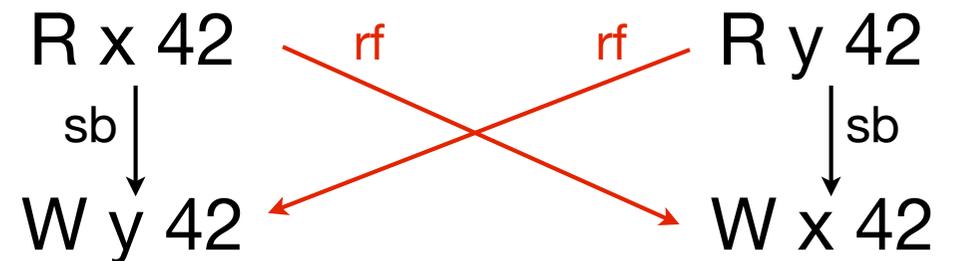


Thread 0	Thread 1
<code>r1 = x</code>	<code>r2 = y</code>
<code>y = r1</code>	<code>x = 42</code>

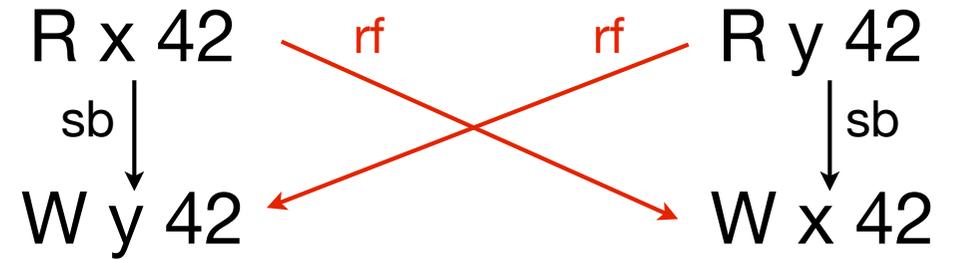


`r1 = r2 = 42.` Can you spot the difference?

Thread 0	Thread 1
<code>r1 = x</code>	<code>r2 = y</code>
<code>y = r1</code>	<code>x = r2</code>

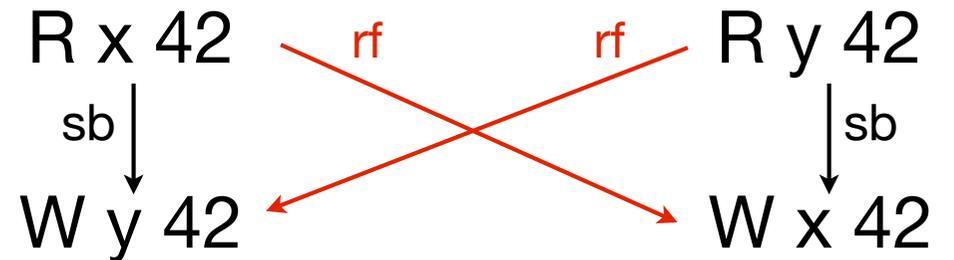


Thread 0	Thread 1
<code>r1 = x</code> <code>y = r1</code>	<code>r2 = y</code> <code>x = 42</code>



The “bad” example has a *cycle of dependencies*.

Thread 0	Thread 1
<code>r1 = x</code> <code>y = r1</code>	<code>r2 = y</code> <code>x = r2</code>

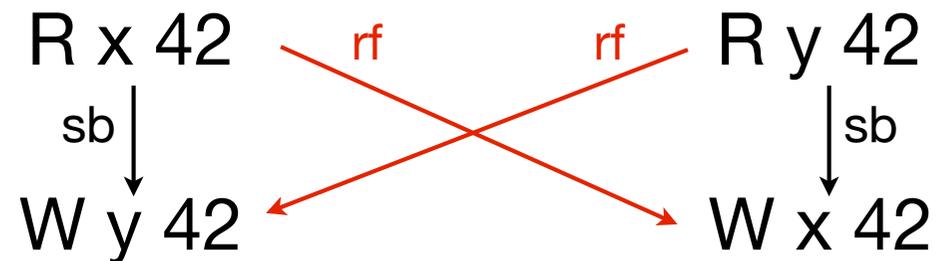


Solution 1.

Prohibit executions with dependency cycles

The “bad” example has a *cycle of dependencies*.

Thread 0	Thread 1
<code>r1 = x</code> <code>y = r1</code>	<code>r2 = y</code> <code>x = r2</code>



**Compiler writers
do not want to track all dependencies**

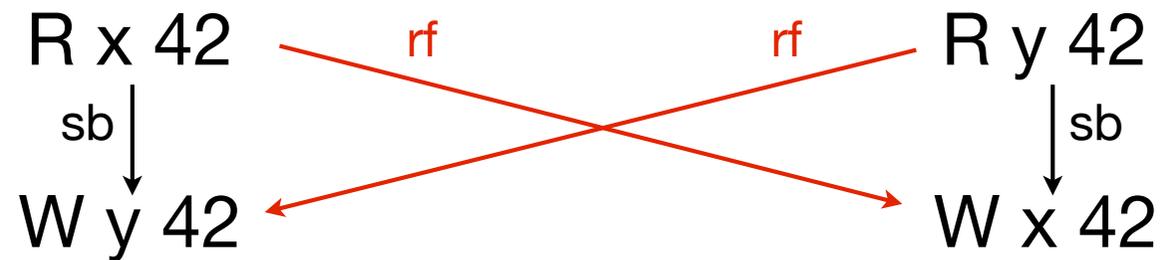
Compiler writers do not want to track all dependencies

```
if (x)
    a[i++] = 1;
else
    a[i++] = 2;
```

Does the store to *i* depend on the load of *x*?

Solution 2. Brute force

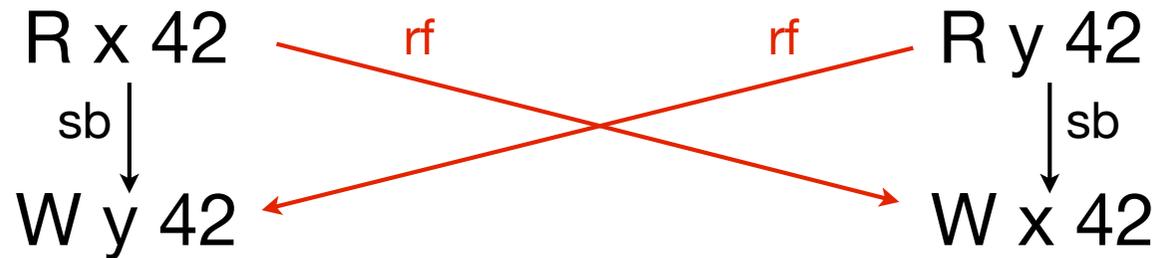
Disallow cycles altogether



$acyclic(\text{hb} \cup \{(a, b) \mid rf(b) = a\})$

Allows all source-to-source optimisations
(except for r/w reordering on atomics)
but expensive on ARM and GPUs

Disallow cycles altogether



$\text{acyclic}(\text{hb} \cup \{(a, b) \mid rf(b) = a\})$

Solution 3. less brute force

Allow cycles but make this racy
by allowing a to read 1

```
if (x.load(rlx)==42) | if (y.load(rlx)==42) | a = 1
  y.write(42,rlx)    |   if (a==1)         |
                     |   x.write(42,rlx)      |
```

Efficient implementation of atomics on ARM/GPUs but all R/W reorderings are unsound

Allow cycles but make this racy
by allowing a to read 1

```
if (x.load(rlx)==42) | if (y.load(rlx)==42) | a = 1
  y.write(42,rlx)    |   if (a==1)      |
                     |   x.write(42,rlx)   |
```

State of the art

“Implementations should ensure that no “out-of-thin-air” values are computed that circularly depend on their own computation.”

Current proposal for C++XX



A word on JSR-133

Goal 1: data-race free programs are sequentially consistent;

Goal 2: all programs satisfy some memory safety requirements;

Goal 3: common compiler optimisations are sound.

Out-of-thin-air

Out-of-thin-air is not so benign for references. Compare:

<u>initially $x = y = 0$</u>			<u>initially $x = y = \text{null}$</u>	
<code>r1 := x</code>	<code>r2 := y</code>	and	<code>r1 := x</code>	<code>r2 := y</code>
<code>y := r1</code>	<code>x := r2</code>		<code>y := r1</code>	<code>x := r2</code>
<code>print r1</code>	<code>print r2</code>			<code>r2.run()</code>

What should `r2.run()` call?

If we allow out-of-thin-air, then it could do anything!

A word on JSR-133



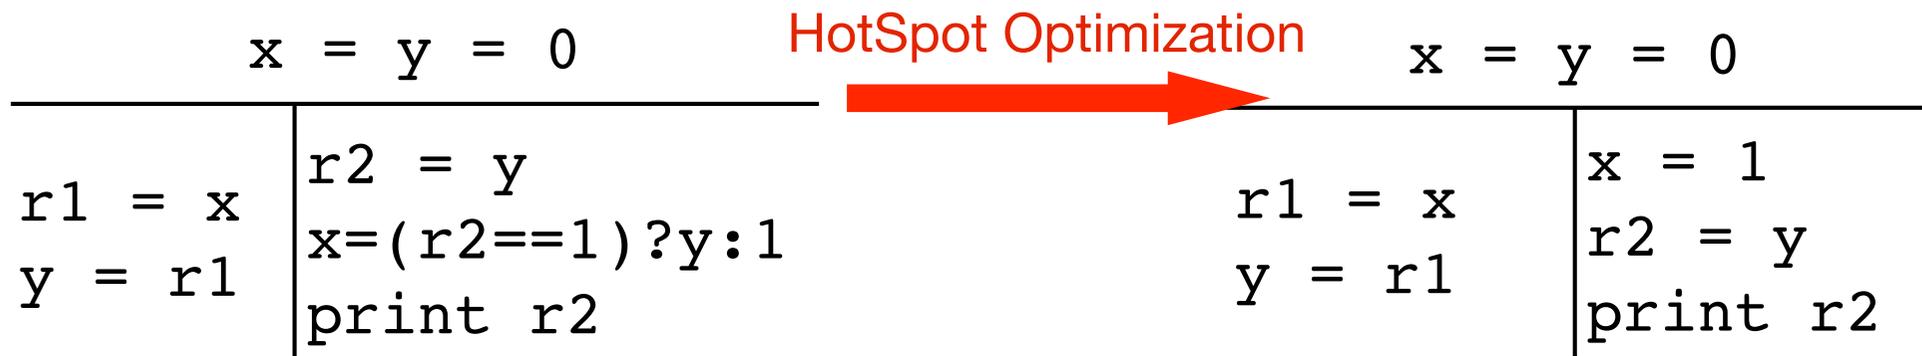
Goal 1: data-race free programs are sequentially consistent;

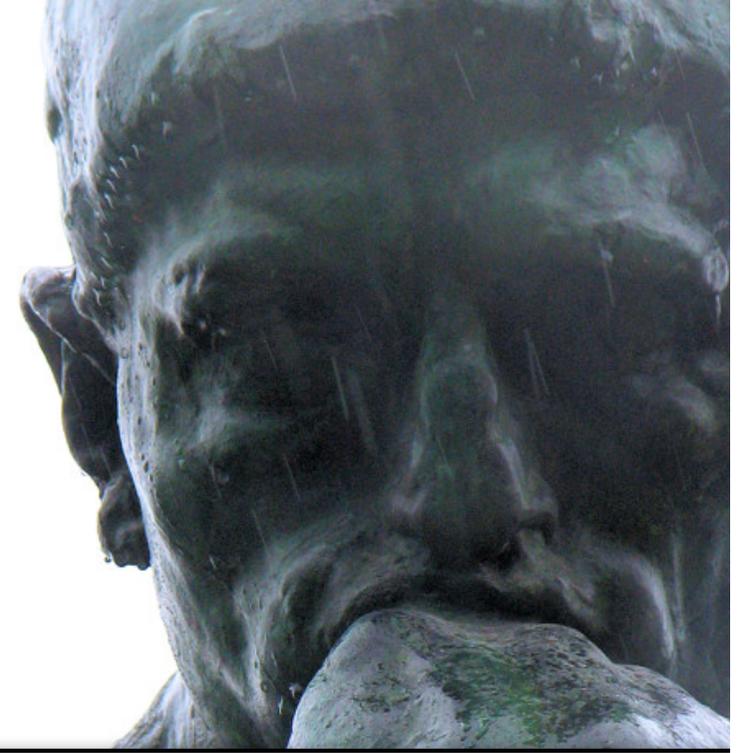
Goal 2: all programs satisfy some memory safety requirements;

Goal 3: common compiler optimisations are sound.

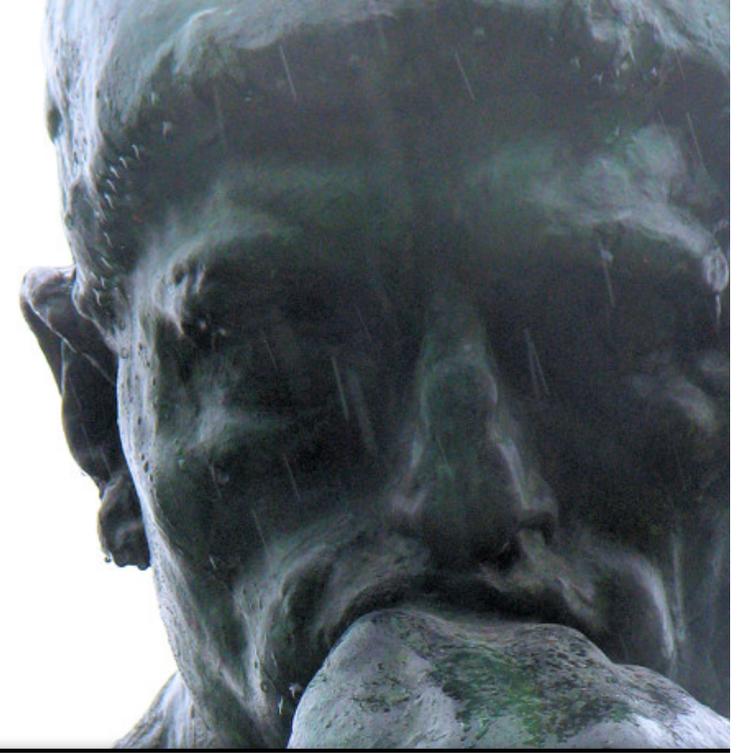
The model is intricate, and fails to meet goal 3.

An example: should the source program print 1? can the optimised program print 1?





*Currently, there is no really satisfactory proposal
for the semantics of a general-purpose
shared-memory concurrent programming language.*



*Currently, there is no really satisfactory proposal
for the semantics of a general-purpose
shared-memory concurrent programming language.*

Remarkable and disturbing.

Resources



<http://www.cl.cam.ac.uk/~pes20/weakmemory/index.html>

Starting point:

J. Sevcik

Safe Optimisations for Shared Memory Concurrent Programs

PLDI 2011

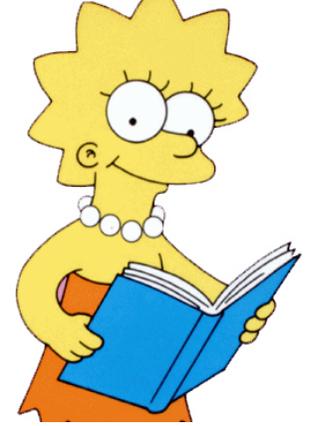
H. Bohem

Threads Cannot Be Implemented as a Library

PLDI 2005

Conclusion

Syllabus



In these lectures we have covered the hardware models of two modern computer architectures (x86 and Power/ARM - at least for a large subset of their instruction set).

We have seen how compiler optimisations can also break concurrent programs and the importance of defining the memory model of high-level programming languages.

We have also introduced some proof methods to reason about concurrency.

After these lectures, you might have the feeling that multicore programming is a mess and things can't just work.



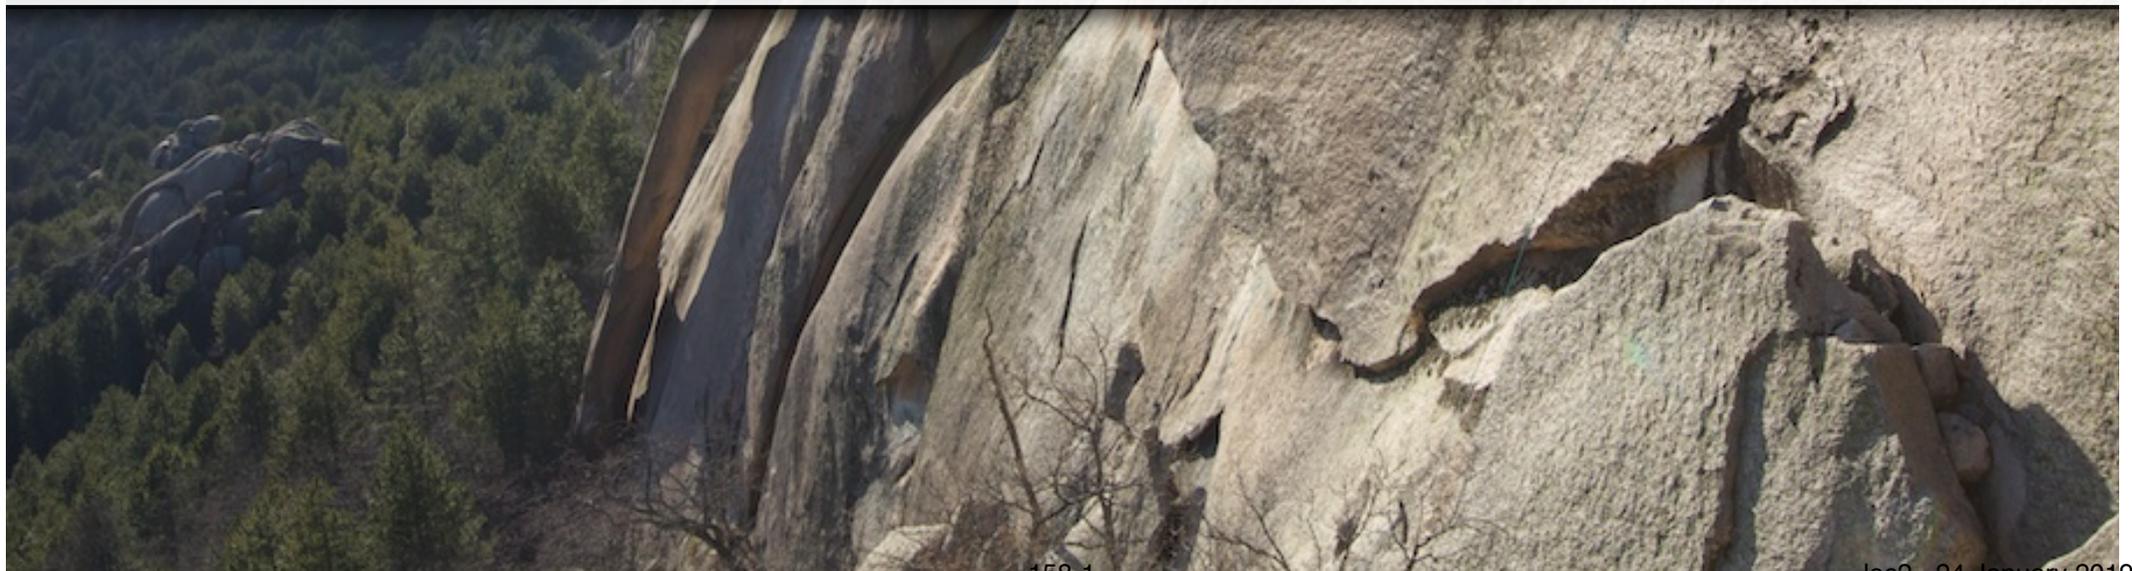
The memory models of modern hardware are better understood.

Programming languages attempt to specify and implement reasonable memory models.

Researchers and programmers are now interested in these problems.



Still, many open problems...





Still, many research opportunities!

