

Semantics, languages and algorithms for multicore programming


Albert Cohen



Luc Maranget



Francesco Zappa Nardelli



Vote: topics for my this lecture



1. The lwarx and stwcx Power instructions [3]

2. Hunting compiler concurrency bugs [12]

3. Operational and axiomatic formalisation of x86-TSO [4]


4. Fence optimisations for x86-TSO [4]

5. The Java memory model [4]

6. The C11/C++11 memory model [17]


7. Static and dynamic techniques for data-race detection [7]

8. The Linux memory model (?!) [18]



1. The Linux memory model (ahem, kinda)

The Linux memory model




Facts:

- abstraction layer over hardware and compilers
- relied upon by kernel developers to write "portable kernel code"
- documented by a text file:

<http://www.kernel.org/doc/Documentation/memory-barriers.txt>

The Linux memory model



Facts:

- abstraction layer over hardware and compilers
- relied upon by kernel developers to write "portable kernel code"
- documented by a text file:

<http://www.kernel.org/doc/Documentation/memory-barriers.txt>



More facts: ...some time ago...

I attempted to understand the doc, and exchanged a few email with Paul Mc Kenney. However I don't understand much...

In the next hour, let's go over the documentation together and see if we can make sense of it...

The Linux memory model

Expected to account for all supported combinations of compiler and hardware memory model...



alpha: Weak ordering. No dependency ordering. "Time does not go backwards" gives guarantees similar to Power/ARM A-cumulativity. Possibly B-cumulativity as well. I am not aware of formalization of this architecture's memory ordering other than Gharachorloo's PhD.

arm: You know at least as much as I do about this one.

avr32: Uniprocessor-only, kernel build failure for SMP.

blackfin: Uniprocessor-only to the best of my knowledge. There are rumored to be some experimental SMP systems that lack cache coherence, and are thus outside of the Linux kernel's remit. See for example: <https://docs.blackfin.uclinux.org/doku.php?id=linux-kernel:smp-like> The system.h file flushes cache when a memory barrier is encountered, which is consistent with an attempt to run the Linux kernel on a non-cache-coherent system...

cris: Uniprocessor-only to the best of my knowledge. Though there appears to be recent addition of some SMP support. Its system.h file is consistent with full sequential consistency. Or extreme optimism on the part of the cris developers.

frv: Uniprocessor-only to the best of my knowledge.

h8300: Uniprocessor-only to the best of my knowledge. There is code in system.h that appears to be intended for SMP, but it looks to me like a (harmless) copy-paste error. Either that or SMP h8300 systems are sequentially consistent.

ia64: Total order of all release operations, which include the "mf" (memory fence) instruction. Memory fences cannot restore sequential consistency.

m32r: Uniprocessor-only to the best of my knowledge. However, there does appear to be some recent multiprocessor support. This is quite strange -- atomic instructions flush cache, but memory barriers are no-ops. Looks quite experimental.

m68k: Uniprocessor-only to the best of my knowledge.

microblaze: Uniprocessor-only to the best of my knowledge. At least one SMP attempt: <http://microblazesmp.blogspot.com/> Its system.h file looks uniprocessor-only.

mips: Multiprocessor. Old SGI MIPS systems were sequentially consistent. Newer systems used for network infrastructure are rumored to have weak memory models similar to Power and ARM. And its system.h file is consistent with a weak memory model.

mn10300: Recent SMP support which I know little about. The system.h file looks uniprocessor only, and contains comments on Intel, so copy-pasted from x86.

parisc: TSO, similar to x86.

powerpc: You know at least as much about this as I do.

s390: TSO, but with self-snooping of store buffer prohibited.

score: Uniprocessor-only to the best of my knowledge.

sh: Recent SMP support which I know little about. Its system.h file is consistent with weak memory ordering.

sparc: TSO, similar to x86. There is documentation about weaker memory models (PSO and RMO), but in practice the hardware is TSO.

tile: Recent SMP CPU which I know little about. Seems to be weakly ordered based on its system.h file.

um: Looks like an x86 knockoff judging by the system.h file.


unicore32: Uniprocessor-only to the best of my knowledge.

x86: You know this one at least as well as do I.

xtensa: Uniprocessor-only -- kernel build failure otherwise.




The Linux memory model



My intuition:

Annoying facts:

The Linux memory model



My intuition:


kinda of lowest common denominator between all hardware memory models of architectures Linux can be compiled to, taking into account also some common gcc optimisations, with some weirdnesses.

Annoying facts:

*semantics of "read barriers" really weak, unclear how to formalise it
compilation of barriers on Itanium looks broken -- hardware might exhibit behaviours prohibited by the MM.*

...let's read the doc...

The Linux memory model: macros



on x86:

```
#define mb()    asm volatile("mfence" ::: "memory")
#define rmb()  asm volatile("lfence" ::: "memory")
#define wmb()  asm volatile("sfence" ::: "memory")
```

in x86TSO lfence is a noop and sfence is like mfence, but things are different in kernel land, eg when performing dma accesses.

on Power:

```
#define mb()    __asm__ __volatile__ ("sync" : : : "memory")
#define rmb()  __asm__ __volatile__ ("sync" : : : "memory")
#define wmb()  __asm__ __volatile__ ("sync" : : : "memory")
#define read_barrier_depends()  do { } while(0)
```

So I still stick with my earlier statements:

- o `smp_mb()` provides transitivity, but is not guaranteed to restore
- o `smp_`
- o `smp_`

Question:

Initially `x=`

Thread 0:

Thread 1:

Thread 2:

to be forbi

And a few emails
exchanged last year...

Not forbidden. If thread 2 did `smp_mb()` instead of `smp_rmb()`, then it would be forbidden.



So I still stick with my earlier statements:

- o `smp_mb()` provides transitivity, but is not guaranteed to restore sequential consistency.
- o `smp_rmb()` simply orders reads. It does not provide transitivity.
- o `smp_wmb()` simply orders writes. It does not provide transitivity.



Question: is `WRC+smp_mb+smp_rmb`, i.e.

Initially `x=0, y=0`

Thread 0: `Wx1`

Thread 1: `Rx1; smp_mb(); Wy1`

Thread 2: `Ry1; smp_rmb(); Rx0`



to be forbidden or not?

Not forbidden. If thread 2 did `smp_mb()` instead of `smp_rmb()`, then it would be forbidden.



On Itanium, both `rmb` and `mb` are compiled to Itanium's `mf`, so there should be no difference in outcome.



However, looking at the Itanium memory model, I do not see how Itanium would forbid the bad outcome in `WRC+mb+mb` because `mf` only imposes thread-local ordering, so thread 2 can see `Wx1` much later (in particular, after the read of `x`).

To make the example work, the `Wx1` would have to be `st.rel`, no?



If this is the case, I suspect that the Linux kernel has a few possible failure modes when running on Itanium hardware. Which it might well have...


...it looks like you might need a significantly different linux mm to the one you've sketched, with weaker barriers and with release/acquire primitives, and to rewrite any WRC-like code using them, no?



Challenging research direction:

Sort out what the REAL Linux memory model is

Yes. Of course, if people come up with lots of situations where the more-complex programming model would help significantly, then it might be worth revisiting this.



Actually: how to design a high-level programming language memory model that does not assign undefined behaviour to racy programs?



2. The C++11 memory model

a good example of an axiomatic memory model



The C++11 memory model

1300 page prose specification defined by the ISO.

The design is a detailed compromise:

hardware/compiler implementability

useful abstractions

broad spectrum of programmers

Welcome to the official home of



JTC1/SC22/WG21 - The C++ Standards Committee

2011-09-15: [standards](#) | [projects](#) | [papers](#) | [mailings](#) | [internals](#) | [meetings](#) | [contacts](#)

News 2011-09-11: The new C++ standard - C++11 - is published!

The syntactic divide

```
// for regular programmers:
```

```
atomic_int x = 0;
```

```
x.store(1);
```

```
y = x.load();
```

```
// for experts:
```

```
x.store(2, memory_order);
```

```
y = x.load(memory_order);
```

```
atomic_thread_fence(memory_order);
```

where *memory_order* is one of the following:

mo_seq_cst mo_release mo_acquire

mo_acq_rel mo_consume mo_relaxed

How may a program execute?

Two layer semantics:

1) a denotational semantics processes programs, identifying memory actions, and constructs candidate executions (E_{opsem});

$$P \longrightarrow E_1, \dots, E_n$$

2) an axiomatic memory model judges E_{opsem} paired with a memory ordering $X_{witness}$

$$E_i \longrightarrow X_{i1}, \dots, X_{im}$$

3) searches the consistent executions for races and unconstrained reads

is there an X_{ij} with a race?

Relations

An E_{opsem} part containing:

sb sequenced before, program order

asw additional synchronizes with, inter-thread ordering

An X_{witness} part containing:

rf relates a write to any reads that take its value

sc a total order over mo_seq_cst and mutex actions

mo modification order, per location total order of writes

From these, compute synchronise-with (sw) and happens-before (hb).

We ignore *consume* atomics, which enables us to live in a simplified model.

Full details in Batty et al., POPL 11.


Formally

```
cpp_memory_model_opsem (p : program) =  
  let pre_executions =  
    { (Eopsem, Xwitness) . opsem p Eopsem ∧  
      consistent execution (Eopsem, Xwitness) }  
  in  
  if ∃X ∈ pre_executions.  
    (indeterminate reads X = {}) ∨  
    (unsequenced races X = {}) ∨  
    (data races X = {})  
  then NONE  
  else SOME pre_executions
```

A single-threaded example

1. sequenced before (sb) - given by opsem


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



A single-threaded example

1. sequenced before (sb) - given by opsem
2. read-from (rf) - part of the witness


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



A single-threaded ex. with undefined behaviour


An unsequenced race.

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



A simple concurrent program

```
int y, x = 2;  
x = 3;          | y = (x==3);
```



We will omit **asw** arrows whenever we are not interested in the initialisation.

Locks and unlocks

```
int x, r;
mutex m;

m.lock();      | m.lock();
x = ...        | r = x;
m.unlock();
```

1. the operational semantics defines the sb arrows

c:L mutex
sb ↓
d:W_{na} x=1
sb ↓
f:U mutex


h:L mutex
sb ↓
i:R_{na} x=1

Locks and unlocks

```
int x, r;
mutex m;

m.lock();      | m.lock();
x = ...        | r = x;
m.unlock();
```

1. the operational semantics defines the sb arrows
2. guess an sc order on Unlock/Lock actions (part of the witness)




Locks and unlocks

```
int x, r;
mutex m;

m.lock();      | m.lock();
x = ...        | r = x;
m.unlock();
```

1. the operational semantics defines the sb arrows
2. guess an sc order on Unlock/Lock actions (part of the witness)
3. the sc order is included in the synchronised-with relation



Locks and unlocks

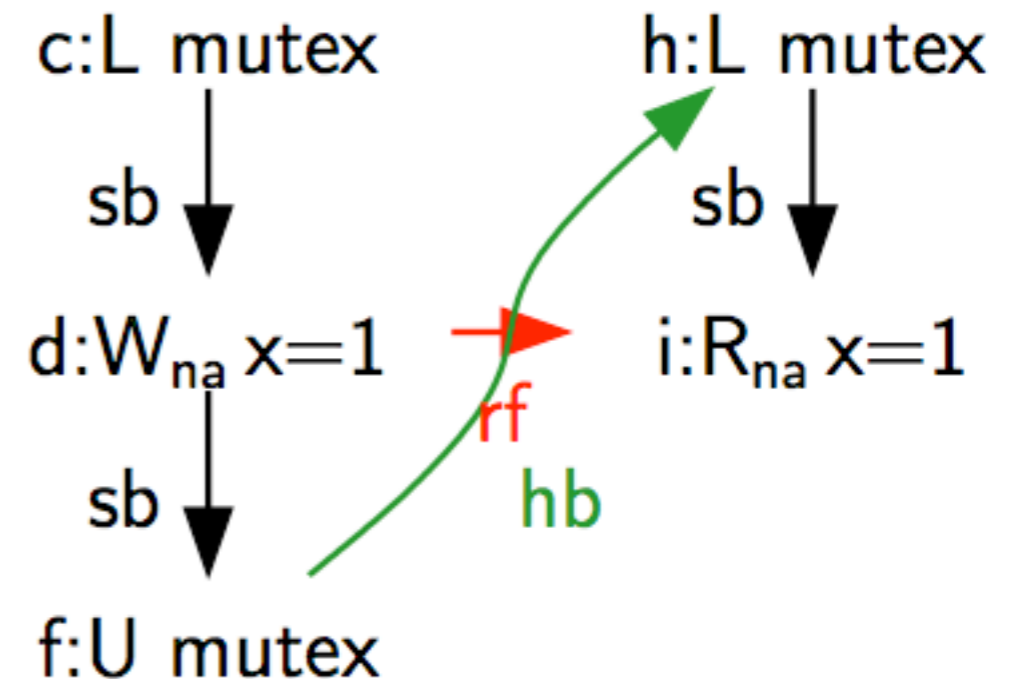
$$\xrightarrow{\text{simple-happens-before}} = \left(\xrightarrow{\text{sequenced-before}} \cup \xrightarrow{\text{synchronizes-with}} \right)^+$$

```
int x, r;
mutex m;

m.lock();
x = ...
m.unlock();
```

```
m.lock();
r = x;
```

1. the operational semantics defines the sb arrows
2. guess an sc order on Unlock/Lock actions (part of the witness)
3. the sc order is included in the synchronised-with relation
4. which in turn defines the happens-before relation...




Happens before

The *happens before* relation is key to the model:


1. non-atomic loads read the most recent write in happens before.
(This is unique in DRF programs)
2. the story is more complex for atomics, as we shall see.
3. data races are defined as an absence of happens before between conflicting actions.

$$\frac{\textit{simple-happens-before}}{\longrightarrow} = \left(\frac{\textit{sequenced-before}}{\longrightarrow} \cup \frac{\textit{synchronizes-with}}{\longrightarrow} \right)^+$$




A data race

```
int y, x = 2;  
x = 3;           | y = (x==3);
```



A data race

```
int y, x = 2;  
x = 3;           | y = (x==3);
```



Here we have two conflicting accesses not related by happens-before.

Data race definition


let $data_races\ actions\ hb =$
 $\{ (a, b) \mid \forall a \in actions\ b \in actions \mid$
 $\neg (a = b) \wedge$
 $same_location\ a\ b \wedge$
 $(is_write\ a \vee is_write\ b) \wedge$
 $\neg (same_thread\ a\ b) \wedge$
 $\neg (is_atomic_action\ a \wedge is_atomic_action\ b) \wedge$
 $\neg ((a, b) \in hb \vee (b, a) \in hb) \}$

Programs with a data race have undefined behaviour (DRF model).

Simple concurrency: Dekker's example and SC

```
atomic_int x = 0;
atomic_int y = 0;

x.store(1, seq_cst); | y.store(1, seq_cst);
y.load(seq_cst);     | x.load(seq_cst);
```




Why is this behaviour forbidden?

Simple concurrency, Dekker's example and SC

```
atomic_int x = 0;
atomic_int y = 0;

x.store(1, seq_cst); | y.store(1, seq_cst);
y.load(seq_cst);     | x.load(seq_cst);
```



The `sc` relation must define a total order over unlocks/locks and `seq_cst` accesses... `sc` is included in `hb`, an `rf` must respect `hb`.

Expert concurrency: the release-acquire idiom

```
// sender
```

```
x = ...
```


```
y.store(1, release);
```

```
// receiver
```

```
while (0 == y.load(acquire));
```

```
r = x;
```

Here we have an `rf` arrow between a pair of release/acquire accesses.



Expert concurrency: the release-acquire idiom

```
// sender
```

```
x = ...
```

```
y.store(1, release);
```


```
// receiver
```

```
while (0 == y.load(acquire));
```

```
r = x;
```

Here we have an **rf** arrow between a pair of release/acquire accesses.

The **rf** arrow between release/acquire accesses induces an **sw** arrow between those accesses.



Expert concurrency: the release-acquire idiom

```
// sender
```

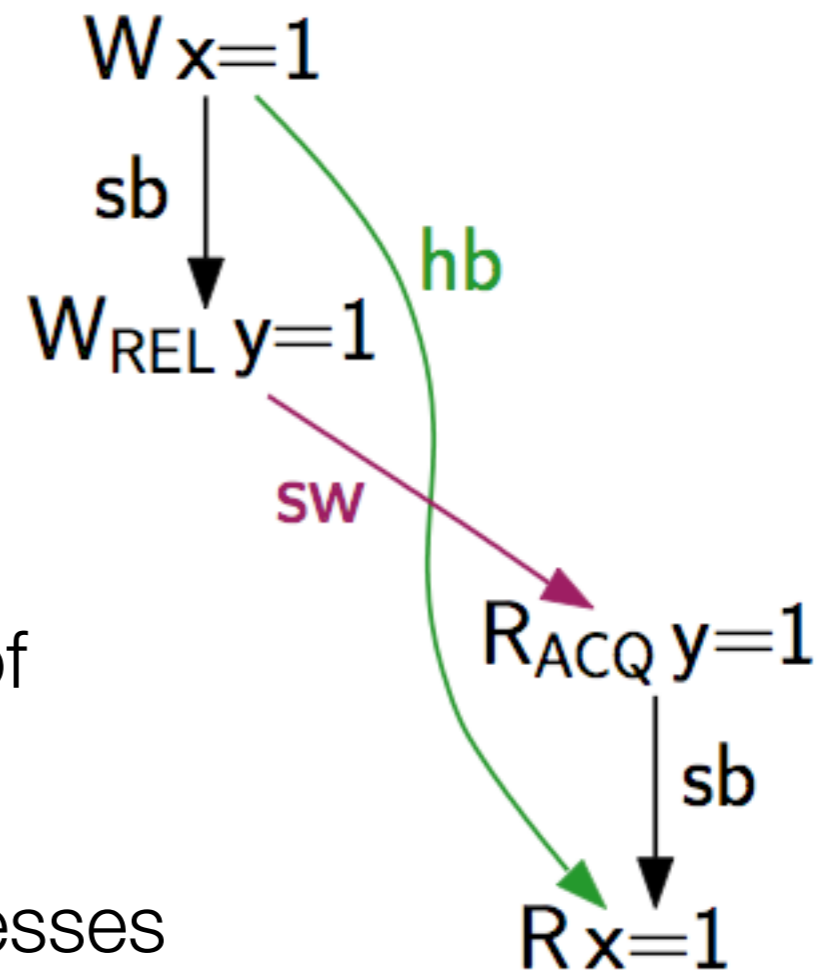
```
x = ...
y.store(1, release);
```

```
// receiver
```

```
while (0 == y.load(acquire));
r = x;
```

Here we have an **rf** arrow between a pair of release/acquire accesses.

The **rf** arrow between release/acquire accesses induces an **sw** arrow between those accesses.




And in turn defines an **hb** constraint.

$$\xrightarrow{\text{simple-happens-before}} = \left(\xrightarrow{\text{sequenced-before}} \cup \xrightarrow{\text{synchronizes-with}} \right)^+$$

Relaxed writes


```
x.load(relaxed);      | y.load(relaxed);  
y.store(1, relaxed); | x.store(1, relaxed);
```



No data-races, no synchronisation cost, but weakly ordered.

Relaxed writes, ctd.

```
atomic_int x = 0;
atomic_int y = 0;
x.store(1, relaxed); | y.store(2, relaxed); | x.load(relaxed); | y.load(relaxed);
                    | y.load(relaxed); | x.load(relaxed);
```



Again, no data-races, no synchronisation cost, but weakly ordered (IRIW).

Expert concurrency: fences avoid excess sync.

```
// sender  
x = ...  
y.store(1, release);
```

```
// receiver  
while (0 == y.load(acquire));  
r = x;
```

```
// sender  
x = ...  
y.store(1, release);
```


```
// receiver  
while (0 == y.load(relaxed));  
fence(acquire);  
r = x;
```

Expert concurrency: fences avoid excess sync.

```
// sender  
x = ...  
y.store(1, release);
```

```
// receiver  
while (0 == y.load(relaxed));  
fence(acquire);  
r = x;
```

Here we have an `rf` arrow between a release write and a relaxed write.




Expert concurrency: fences avoid excess sync.

```
// sender  
x = ...  
y.store(1, release);
```

```
// receiver  
while (0 == y.load(relaxed));  
fence(acquire);  
r = x;
```

Here we have an **rf** arrow between a release write and a relaxed write.



The acquire fence follows the **sb/rf** relations looking for the corresponding release write, adding a **sw** arrow.


Expert concurrency: fences avoid excess sync.

```
// sender  
x = ...  
y.store(1, release);
```

```
// receiver  
while (0 == y.load(relaxed));  
fence(acquire);  
r = x;
```

Here we have an **rf** arrow between a release write and a relaxed write.

The acquire fence follows the **sb/rf** relations looking for the corresponding release write, adding a **sw** arrow.




Happens-before follows as usual...

Modification order

```
atomic_int x = 0;  
x.store(1, relaxed);  
x.store(2, relaxed);
```

 ||





```
x.load(relaxed);  
x.load(relaxed);
```



Modification order is a total order over atomic writes of any memory order.

Coherence and atomic reads

All forbidden:



Idea: atomics cannot read from later writes in happens-before.

Coherence and atomic reads

All forb

a:)

b:)

A pair $E_{\text{opsem}}, X_{\text{witness}}$ (a pre-execution) defines a *consistent execution* when it satisfies the constraints we have sketched on **hb/rf/mo** and is race-free.

b:W x=2

CoWW

d:W x=2

CoRW

Idea: atomics cannot read from later writes in happens-before.

Is C++11 hopelessly complicated?

Programmers cannot be given this model.

However, with a formal definition, we can do proofs!

- Can we compile to x86?

| Operation | x86 Implementation |
|--------------------|--------------------|
| load(non-seq_cst) | mov |
| load(seq_cst) | lock xadd(0) |
| store(non-seq_cst) | mov |
| store(seq_cst) | lock xchg |
| fence(non-seq_cst) | no-op |

- Can we compile to Power?

| C++0x Operation | POWER Implementation |
|------------------|------------------------------|
| Non-atomic Load | ld |
| Load Relaxed | ld |
| Load Consume | ld (and preserve dependency) |
| Load Acquire | ld; cmp; bc; isync |
| Load Seq Cst | sync; ld; cmp; bc; isync |
| Non-atomic Store | st |
| Store Relaxed | st |
| Store Release | lwsync; st |
| Store Seq Cst | sync; st |

Is C++11 hopelessly complicated?

Simplifications:

Full model: *visible sequences of side effects* are unneeded (HOL4)

Derivative models:

- without consume, happens-before is transitive
- DRF programs using only `seq_cst` atomics are SC (false)

```
atomic_int x = 0;
atomic_int y = 0;
if (1 == x.load(seq_cst)) | if (1 == y.load(seq_cst))
    atomic_init(&y, 1);      |    atomic_init(&x, 1);
```

`atomic_init` is a non-atomic write, and in C++11 they race.



The current state of the standard


Fixed:

- in some cases, happens-before was cyclic
- coherence
- `seq_cst` atomics were more broken

Not fixed:

- self satisfying conditional

```
r1 = x.load(mo_relaxed);   |   r2 = y.load(mo_relaxed);  
if (r1 == 42)              |   if (r2 == 42)  
    y.store(r1, mo_relaxed); |       x.store(42, mo_relaxed);
```



- `seq_cst` atomics are still not SC

3. A word on dynamic techniques for data-race detection

Data race detection

Modern high-performance dynamic race detectors are based either on:

happens-before ordering

reconstruct happens-before order
in the current execution
report a race if intersection of two
conflicting accesses are not related
by hb

sound

drawback: misses races
occurring on rare executions

lockset computation

records which locks protect
every memory access
report a race if intersection of all
locksets for a variable is empty

popularised by Eraser (Savage et al.) '97

can detect races not observed in
the execution being monitored

drawback: unsound (false positives)

Examples of lockset computation

| | | |
|-----------|--|-----------|
| lock(b) | | lock(a) |
| lock(a) | | x=2 |
| x=1 | | unlock(a) |
| unlock(a) | | |

1:L(b);1:L(a);1:Wx1;1:U(a);2:L(a);2:Wx2;2:U(a)

locks held: 1:b 1:b,a 2:a

C(x): x:a,b x:a

lockset for x non-empty at the end, no data-race

| | | |
|-----------|--|-----------|
| lock(b) | | lock(c) |
| lock(a) | | x=2 |
| x=1 | | unlock(c) |
| unlock(a) | | |

1:L(b);1:L(a);1:Wx1;1:U(a);2:L(c);2:Wx2;2:U(c)

C(x): x:a,b x:empty

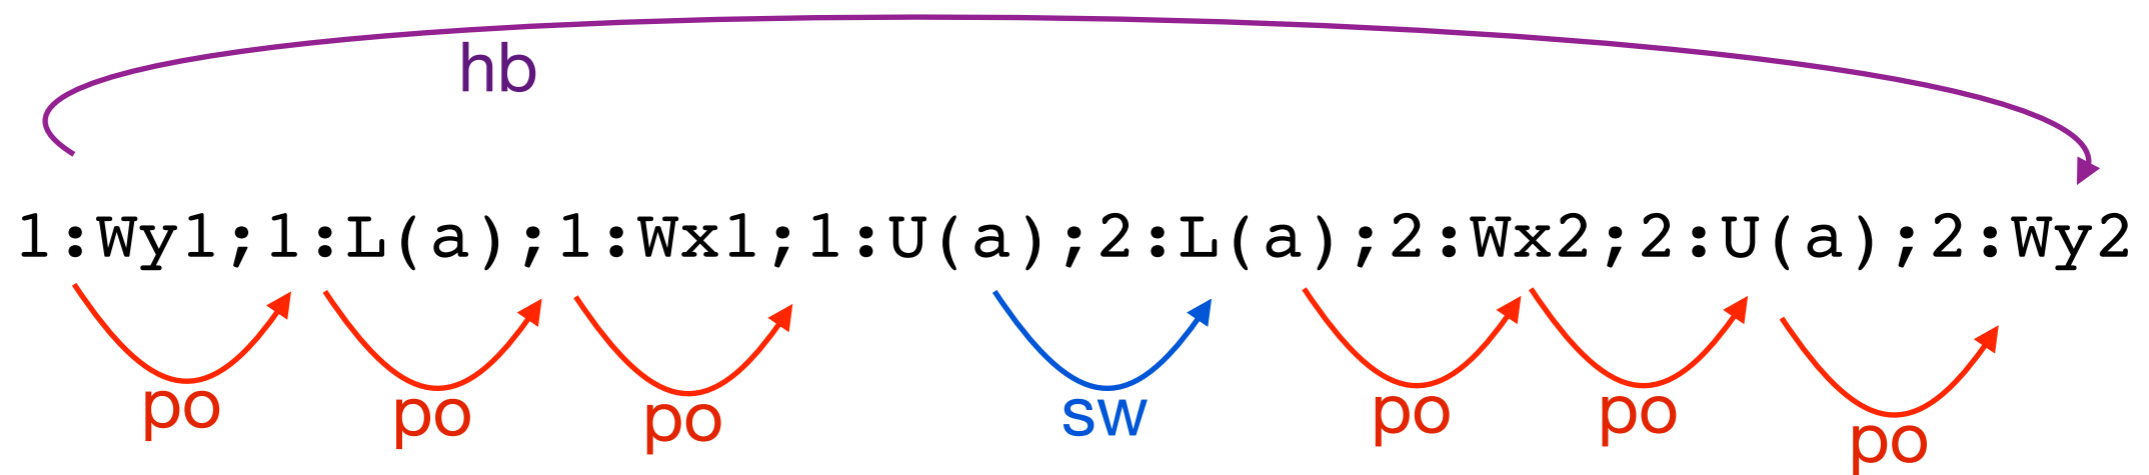
lockset for x empty at the end, possible data-race

lockset vs happens-before

| | | |
|------------------------|--|------------------------|
| <code>y=1</code> | | <code>lock(a)</code> |
| <code>lock(a)</code> | | <code>x=2</code> |
| <code>x=1</code> | | <code>unlock(a)</code> |
| <code>unlock(a)</code> | | <code>y=2</code> |

This program has a race on y

If only the execution below is observed:



happens-before computation does not report a race.

Lockset computation detects instead that accesses to y are unprotected and reports a possible race.

lockset vs happens-before (2)

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

This program instead is DRF.

Happens-before computation will not report a race
(no matter which execution is observed)

Since accesses to y are unprotected, locksets computation reports a false positive.

Data race detection

Modern high-performance dynamic race detectors are based either on:

happens-before ordering

reconstruct happens-before order
in the current execution
report a race if intersection of two
conflicting accesses are not related
by hb

sound

drawback: misses races
occurring on rare executions

lockset computation

records which locks protect
every memory access
report a race if intersection of all
locksets for a variable is empty

popularised by Eraser (Savage et al.) '97

can detect races not observed in
the execution being monitored

drawback: unsound (false positives)

Data race detection

Model

er on:

hap

Current state of the art:

recons

hybrid approaches combining locksets and happens-before ordering + other dynamic annotations

in

report

Helgrind, RaceFuzzer, ThreadSanitizer...

conflict

Impressive:

*<10x slowdown on large applications
found thousands races*

all

y

al.) '97

d in

ed

dr

occu

Still not as reliable as the tool we dream of...

ves)



4. Sketch of an operational formalisation of x86-TSO


...starting with a formalisation of SC

Separate language and memory semantics

```
1  class ArrayWrapper
2  {
3      public:
4          ArrayWrapper (int n)
5              : _p_vals( new int[ n ] )
6                , _size( n )
7          {}
8          // copy constructor
9          ArrayWrapper (const ArrayWrapper& other)
10             : _p_vals( new int[ other._size ] )
11               , _size( other._size )
12          {
13              for ( int i = 0; i < _size; ++i )
14              {
15                  _p_vals[ i ] = other._p_vals[ i ];
16              }
17          }
18          ~ArrayWrapper ()
19          {
20              delete [] _p_vals;
21          }
22      private:
23          int *_p_vals;
24          int _size;
25  };
```

program

semantics defined via an LTS



memory

semantics defined via an LTS

Labels for interaction:

$W_t[a]v$: a write of value v to address a by thread t

$R_t[a]v$: a read of v from a by t by thread t

+ other events for barriers and locked instructions

Separate language and memory semantics

Separate language and state semantics
proved to be a very good choice
in many (unrelated) projects I worked on!

```
1  class Arr  
2  {  
3      public  
4      A  
5  
6  
7      {  
8      /  
9      A  
10  
11  
12      {  
13  
14  
15  
16  
17      }  
18      ~  
19      {  
20  
21      }  
22      priva  
23      int *  
24      int  
25  };
```

semantics defined via an LTS

semantics defined via an LTS

Labels for interaction: $W_t[a]v$: a write of value v to address a by thread t
 $R_t[a]v$: a read of v from a by t by thread t
+ other events for barriers and locked instructions

A tiny language

location, x, m address (or pointer value)

integer, n integer

thread_id, t thread id

k, i, j

| | | |
|----------------------|---------------|------------------------|
| <i>expression, e</i> | ::= | expression |
| | <i>n</i> | integer literal |
| | <i>*x</i> | read from pointer |
| | <i>*x = e</i> | write to pointer |
| | <i>e; e'</i> | sequential composition |
| | <i>e + e'</i> | plus |

| | | |
|-------------------|-------------|----------------------|
| <i>process, p</i> | ::= | process |
| | <i>t:e</i> | thread |
| | <i>p p'</i> | parallel composition |

What can a thread do in isolation?

$e \xrightarrow{l} e'$ e does l to become e'

$$\frac{}{*x \xrightarrow{R\ x=n} n} \quad \text{READ}$$

$$\frac{}{*x = n \xrightarrow{W\ x=n} n} \quad \text{WRITE}$$

$$\frac{e \xrightarrow{l} e'}{*x = e \xrightarrow{l} *x = e'} \quad \text{WRITE_CONTEXT}$$

$$\frac{}{n; e \xrightarrow{\tau} e} \quad \text{SEQ}$$

$$\frac{e_1 \xrightarrow{l} e'_1}{e_1; e_2 \xrightarrow{l} e'_1; e_2} \quad \text{SEQ_CONTEXT}$$

$$\frac{e_1 \xrightarrow{l} e'_1}{e_1 + e_2 \xrightarrow{l} e'_1 + e_2} \quad \text{PLUS_CONTEXT_1}$$

$$\frac{e_2 \xrightarrow{l} e'_2}{n_1 + e_2 \xrightarrow{l} n_1 + e'_2} \quad \text{PLUS_CONTEXT_2}$$

$$\frac{n = n_1 + n_2}{n_1 + n_2 \xrightarrow{\tau} n} \quad \text{PLUS}$$

Observe that we can read an arbitrary value from the memory.

Example

Show that the expression:

$$(*x = *y); *x$$

can perform the following trace:

$$(*x = *y); *x \xrightarrow{R\ y=7} \xrightarrow{W\ x=7} \xrightarrow{\tau} \xrightarrow{R\ x=9} 9$$

Lifting to processes

$p \xrightarrow{l_t} p'$ p does l_t to become p'

$$\frac{e \xrightarrow{l} e'}{t:e \xrightarrow{l_t} t:e'} \quad \text{THREAD}$$
$$\frac{p_1 \xrightarrow{l_t} p'_1}{p_1|p_2 \xrightarrow{l_t} p'_1|p_2} \quad \text{PAR_CONTEXT_LEFT}$$
$$\frac{p_2 \xrightarrow{l_t} p'_2}{p_1|p_2 \xrightarrow{l_t} p_1|p'_2} \quad \text{PAR_CONTEXT_RIGHT}$$

Actions are labelled by the thread that performed the action.

Free interleaving.

A sequentially consistent memory

Take M to be a function from addresses to integers.

$M \xrightarrow{l} M'$ M does l to become M'

$$\frac{M(x) = n}{M \xrightarrow{R x=n} M} \quad \text{MREAD}$$

$$\frac{}{M \xrightarrow{W x=n} M \oplus (x \mapsto n)} \quad \text{MWRITE}$$

SC semantics: whole system transitions

$s \xrightarrow{l_t} s'$ s does l_t to become s'

$$\frac{\begin{array}{l} p \xrightarrow{R_t x=n} p' \\ M \xrightarrow{R x=n} M' \end{array}}{\langle p, M \rangle \xrightarrow{R_t x=n} \langle p', M' \rangle} \quad \text{SREAD}$$


Synchronising between the processes and the memory.

$$\frac{\begin{array}{l} p \xrightarrow{W_t x=n} p' \\ M \xrightarrow{W x=n} M' \end{array}}{\langle p, M \rangle \xrightarrow{W_t x=n} \langle p', M' \rangle} \quad \text{SWRITE}$$

$$\frac{p \xrightarrow{\tau_t} p'}{\langle p, M \rangle \xrightarrow{\tau_t} \langle p', M \rangle} \quad \text{STAU}$$

SC semantics, example

All threads read and write the shared memory. Threads execute asynchronously, the semantics allows any interleaving of the thread transitions.



Each interleaving has a linear order of reads and writes to memory.

...now we just have to define a TSO memory...

A sequentially consistent memory

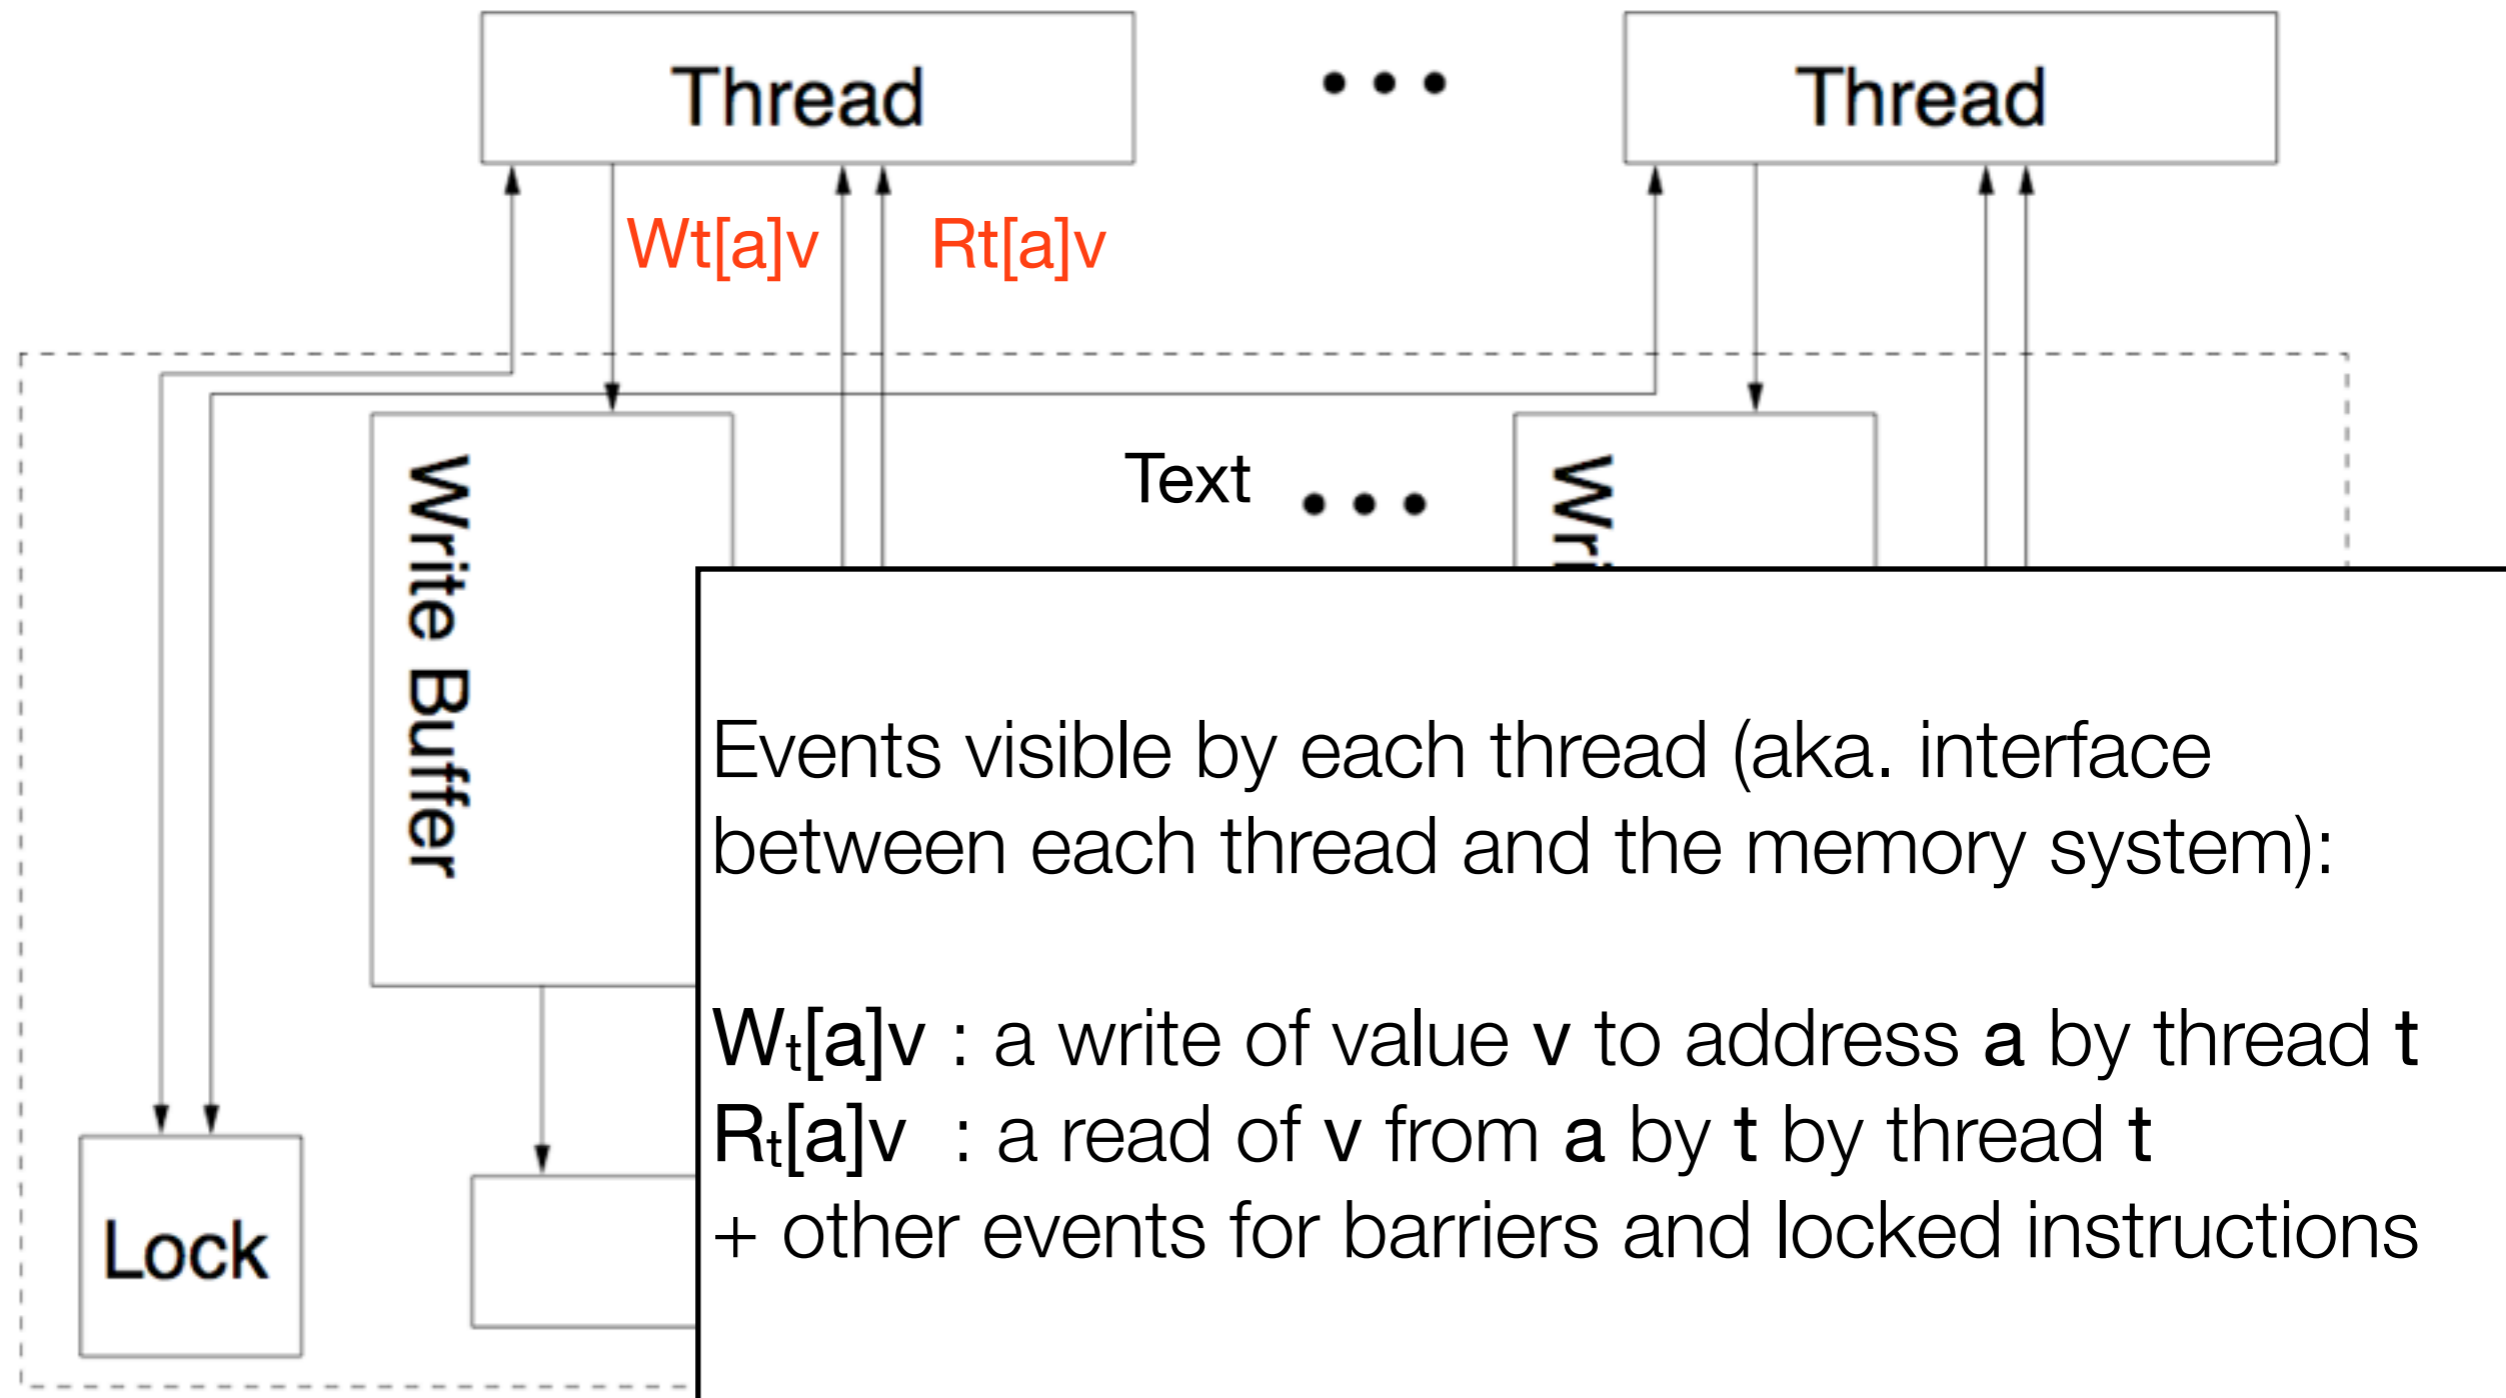
Take M to be a function from addresses to integers.

$$\boxed{M \xrightarrow{l} M'} \quad M \text{ does } l \text{ to become } M'$$

$$\frac{M(x) = n}{M \xrightarrow{R\ x=n} M} \quad \mathbf{MREAD}$$

$$\frac{}{M \xrightarrow{W\ x=n} M \oplus (x \mapsto n)} \quad \mathbf{MWRITE}$$

x86-TSO abstract machine



x86-TSO abstract machine

- The store buffers are FIFO. A reading thread must read its most recent buffered write, if there is one, to that address; otherwise reads are satisfied from shared memory.
- To execute a LOCK'd instruction, a thread must first obtain the global lock. At the end of the instruction, it flushes its store buffer and relinquishes the lock. While the lock is held by one thread, no other thread can read.
- A buffered write from a thread can propagate to the shared memory at any time except when some other thread holds the lock.

ues

x86-tso: a formalisation using an LTS

The machine state s can be represented by a tuple (M, B, L) :

M : address \rightarrow value option

B : tid \rightarrow (address * value) list

L : tid option

where:

M is the shared memory, mapping addresses to values

B gives the store buffer for each thread

L is the global machine lock indicating when a thread has exclusive access to memory (omitted in these slides)

x86-tso abstract machine: selected transition rules

t is *not blocked* in machine state $s = (M, B, L)$ if [...] or] the lock is not held.

In buffer $B(t)$ there are *no pending writes* for address x if there are no (x, v) elements in $B(t)$.

RM: Read from memory

$\text{not_blocked}(s, t)$

$s.M(x) = v$

$\text{no_pending}(s.B(t), x)$

$$s \xrightarrow{R_t x=v} s$$

Thread t can read v from memory at address x if t is not blocked, the memory does contain v at x , and there are no writes to x in t 's store buffer.

x86-tso abstract machine: selected transition rules

RB: Read from write buffer

$\text{not_blocked}(s, t)$

$\exists b_1 b_2. s.B(t) = b_1 ++ [(x, v)] ++ b_2$

$\text{no_pending}(b_1, x)$

$$s \xrightarrow{R_t x=v} s$$

Thread t can read v from its store buffer for address x if t is not blocked and has v as the newest write to x in its buffer;

x86-tso abstract machine: selected transition rules

WB: Write to write buffer

$$s \xrightarrow{W_t x=v} s \oplus \langle B := s.B \oplus (t \mapsto ([x, v] ++ s.B(t))) \rangle$$

Thread t can write v to its store buffer for address x
at any time;

WM: Write from write buffer to memory

not_blocked(s, t)

$s.B(t) = b ++ [x, v]$

$$s \xrightarrow{\tau_t x=v} s \oplus \langle M := s.M \oplus (x \mapsto v) \rangle \oplus \langle B := s.B \oplus (t \mapsto b) \rangle$$

If t is not blocked, it can silently dequeue the oldest write from its store buffer and place the value in memory at the given address, without coordinating with any hardware thread



5. Hunting compiler concurrency bugs



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);
```

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);
```

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);
```


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);
```

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);
```

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);
```

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 1 returns without modifying b.

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 1 returns without modifying b.

Since Thread 1 does not update b, program is *data-race free (DRF)*

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 1 returns without modifying `b`.

Since Thread 1 does not update `b`, program is *data-race free (DRF)*

DRF programs must only exhibit sequentially consistent behaviours

C11/C++11 standard

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 1 returns without modifying b.

Since Thread 1 does not update b, program is *data-race free (DRF)*

DRF programs must only exhibit sequentially consistent behaviours

C11/C++11 standard

This program only prints 42.

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);
```


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);
```

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);
```



...sometimes we get 0 on the screen

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

gcc 4.7 -O2



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

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

gcc 4.7 -O2



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

The outer loop can be (and is) optimised away

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

gcc 4.7 -O2



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

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

gcc 4.7 -O2



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

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

gcc 4.7 -O2



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

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


gcc 4.7 -O2



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

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

.L2:

```
movl    %eax, b(%rip)    # store eax into b  
xorl    %eax, %eax      # store 0 into eax  
ret
```

gcc 4.7 -O2




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

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

The compiled code saves and restores b

Correct in a sequential setting, but...



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

Shared memory

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

Thread 1

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

Thread 2

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

Shared memory

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

Thread 1

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

Thread 2

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

- Read a (1) into edx

Shared memory

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

Thread 1

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

Thread 2

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

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

Shared memory

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

Thread 1

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

Thread 2

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

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

Shared memory

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

Thread 1

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

.L2:

```
movl    %eax, b(%rip)  
xorl    %eax, %eax  
ret
```

Thread 2

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

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

Shared memory

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

Thread 1

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

Thread 2

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

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

The compiled code saves and restores b

Correct in a sequential setting

*Introduces unexpected behaviours
in some concurrent context*

```
ret
.L2:
movl   %eax, b(%rip)
xorl   %eax, %eax
ret
```

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

The compiled code saves and restores `b`

Correct in a sequential setting

*Introduces unexpected behaviours
in some concurrent context*

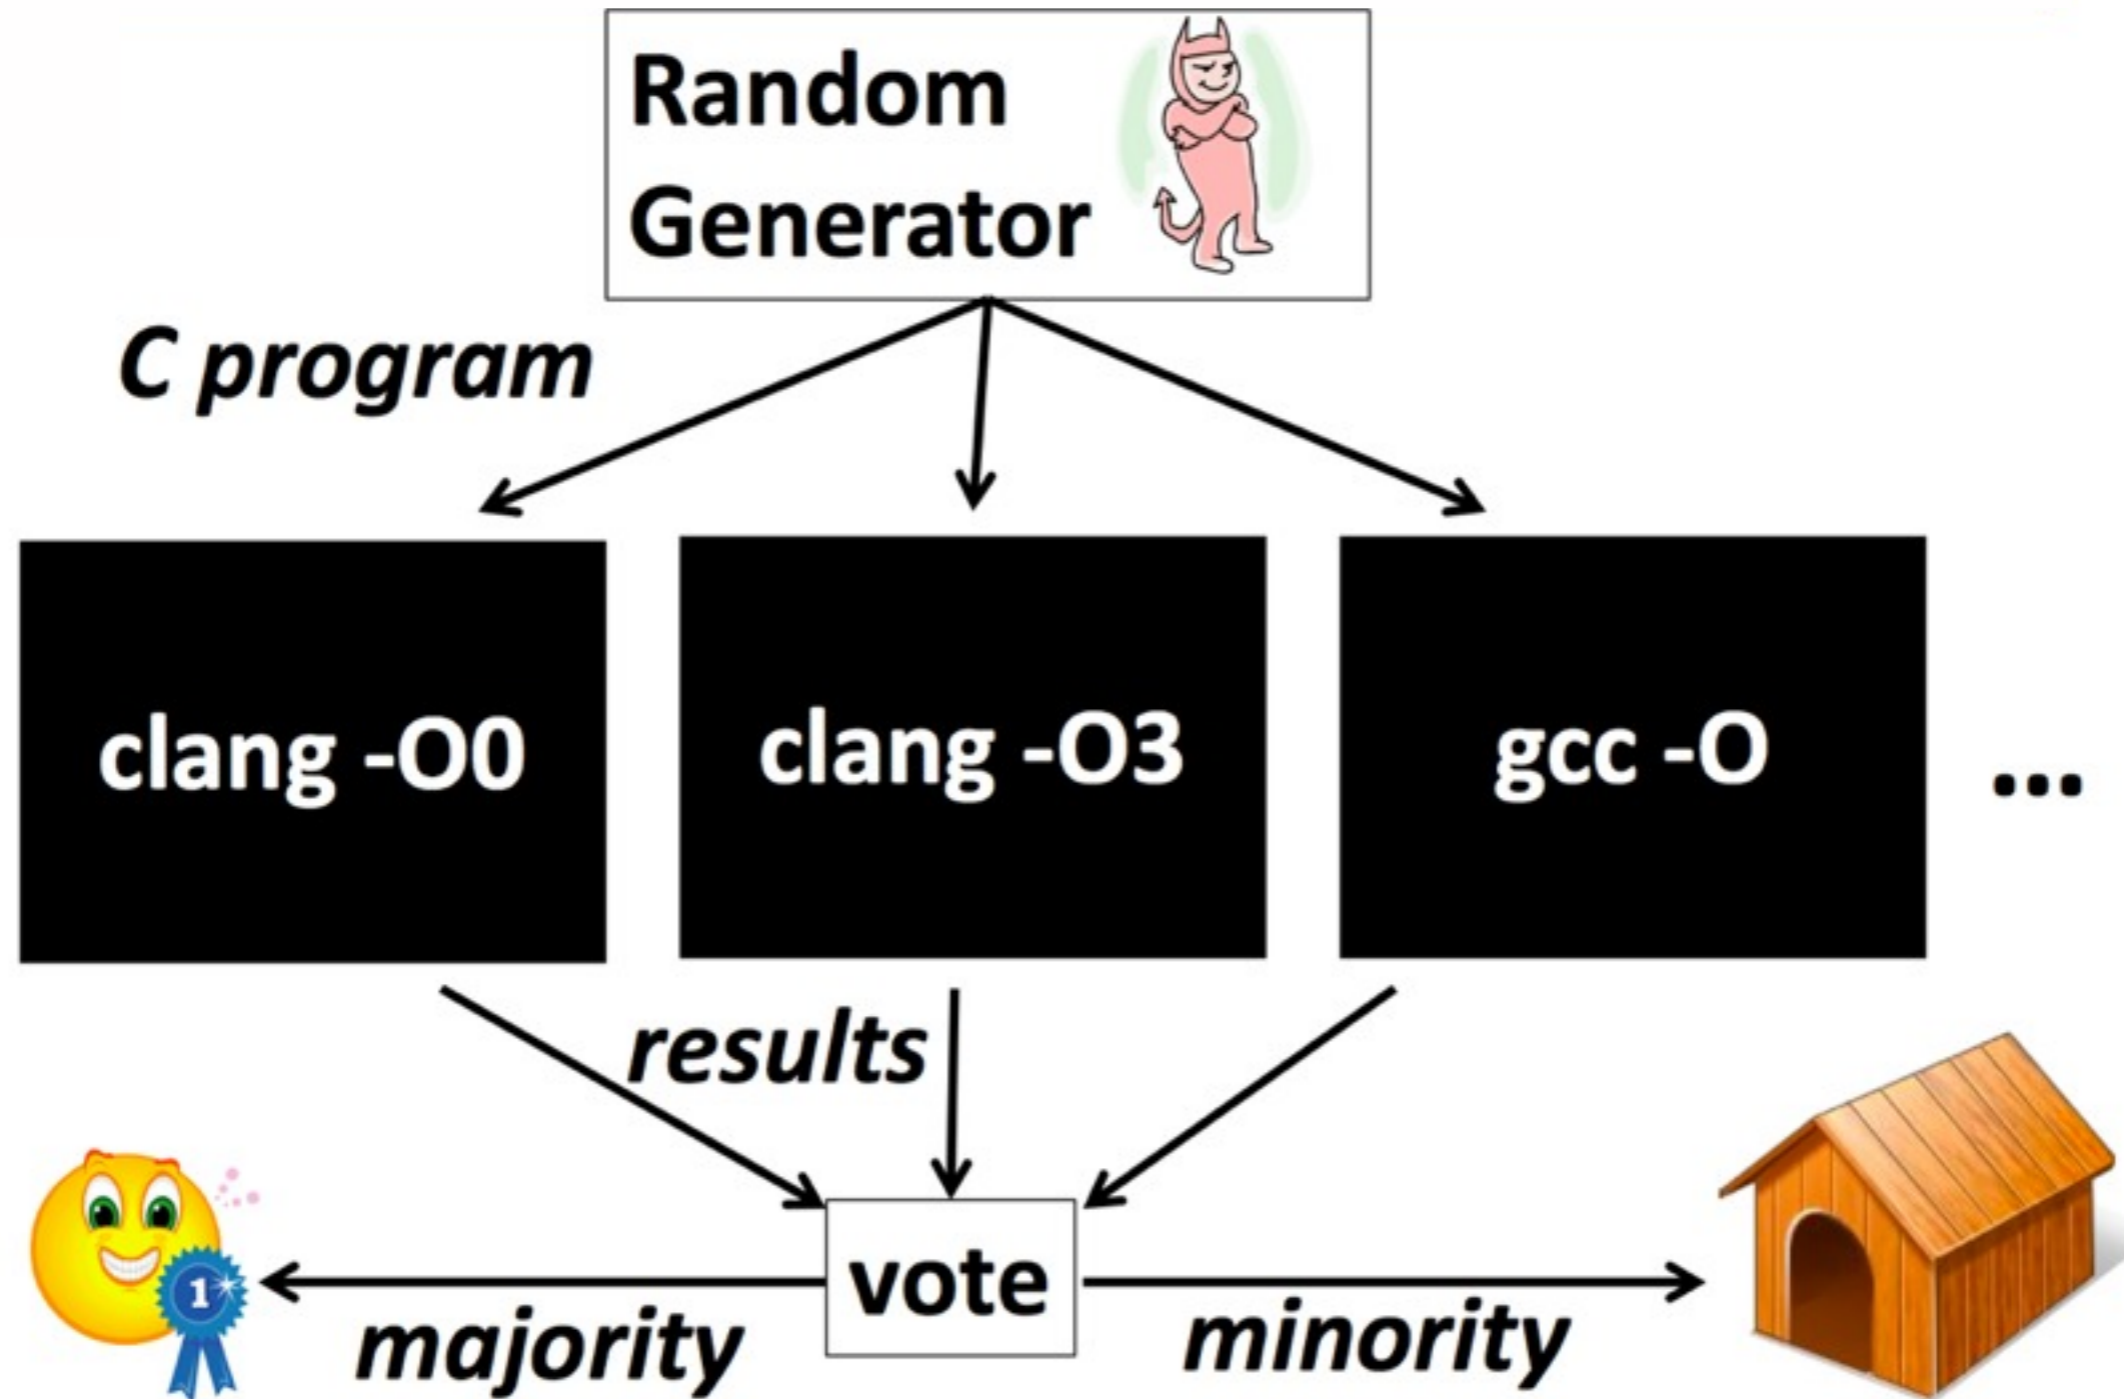
This is a *concurrency compiler bug*

```
movl    %eax, b(%rip)
xorl    %eax, %eax
ret
```

- Read `b (0)` into `eax`
- Store `42` into `b`
- Store `eax (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 -O3

gcc -O

...



majority

vote

minority



Compiler testing: state of the art

Yang, Chen, Eide, Regehr - PLDI 2011

Random



Reported hundreds of bugs

*Cannot catch
concurrency compiler bugs*



majority

vote

minority



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

**Random
Generator**



**SEQUENTIAL
PROGRAM**

*optimising
compiler
under test*

EXECUTABLE

tracing

reference
semantics

**REFERENCE
MEMORY
TRACE**

**MEMORY
TRACE**



**only transformations sound in any
concurrent non-racy context?**

```
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)
        ;
}
```

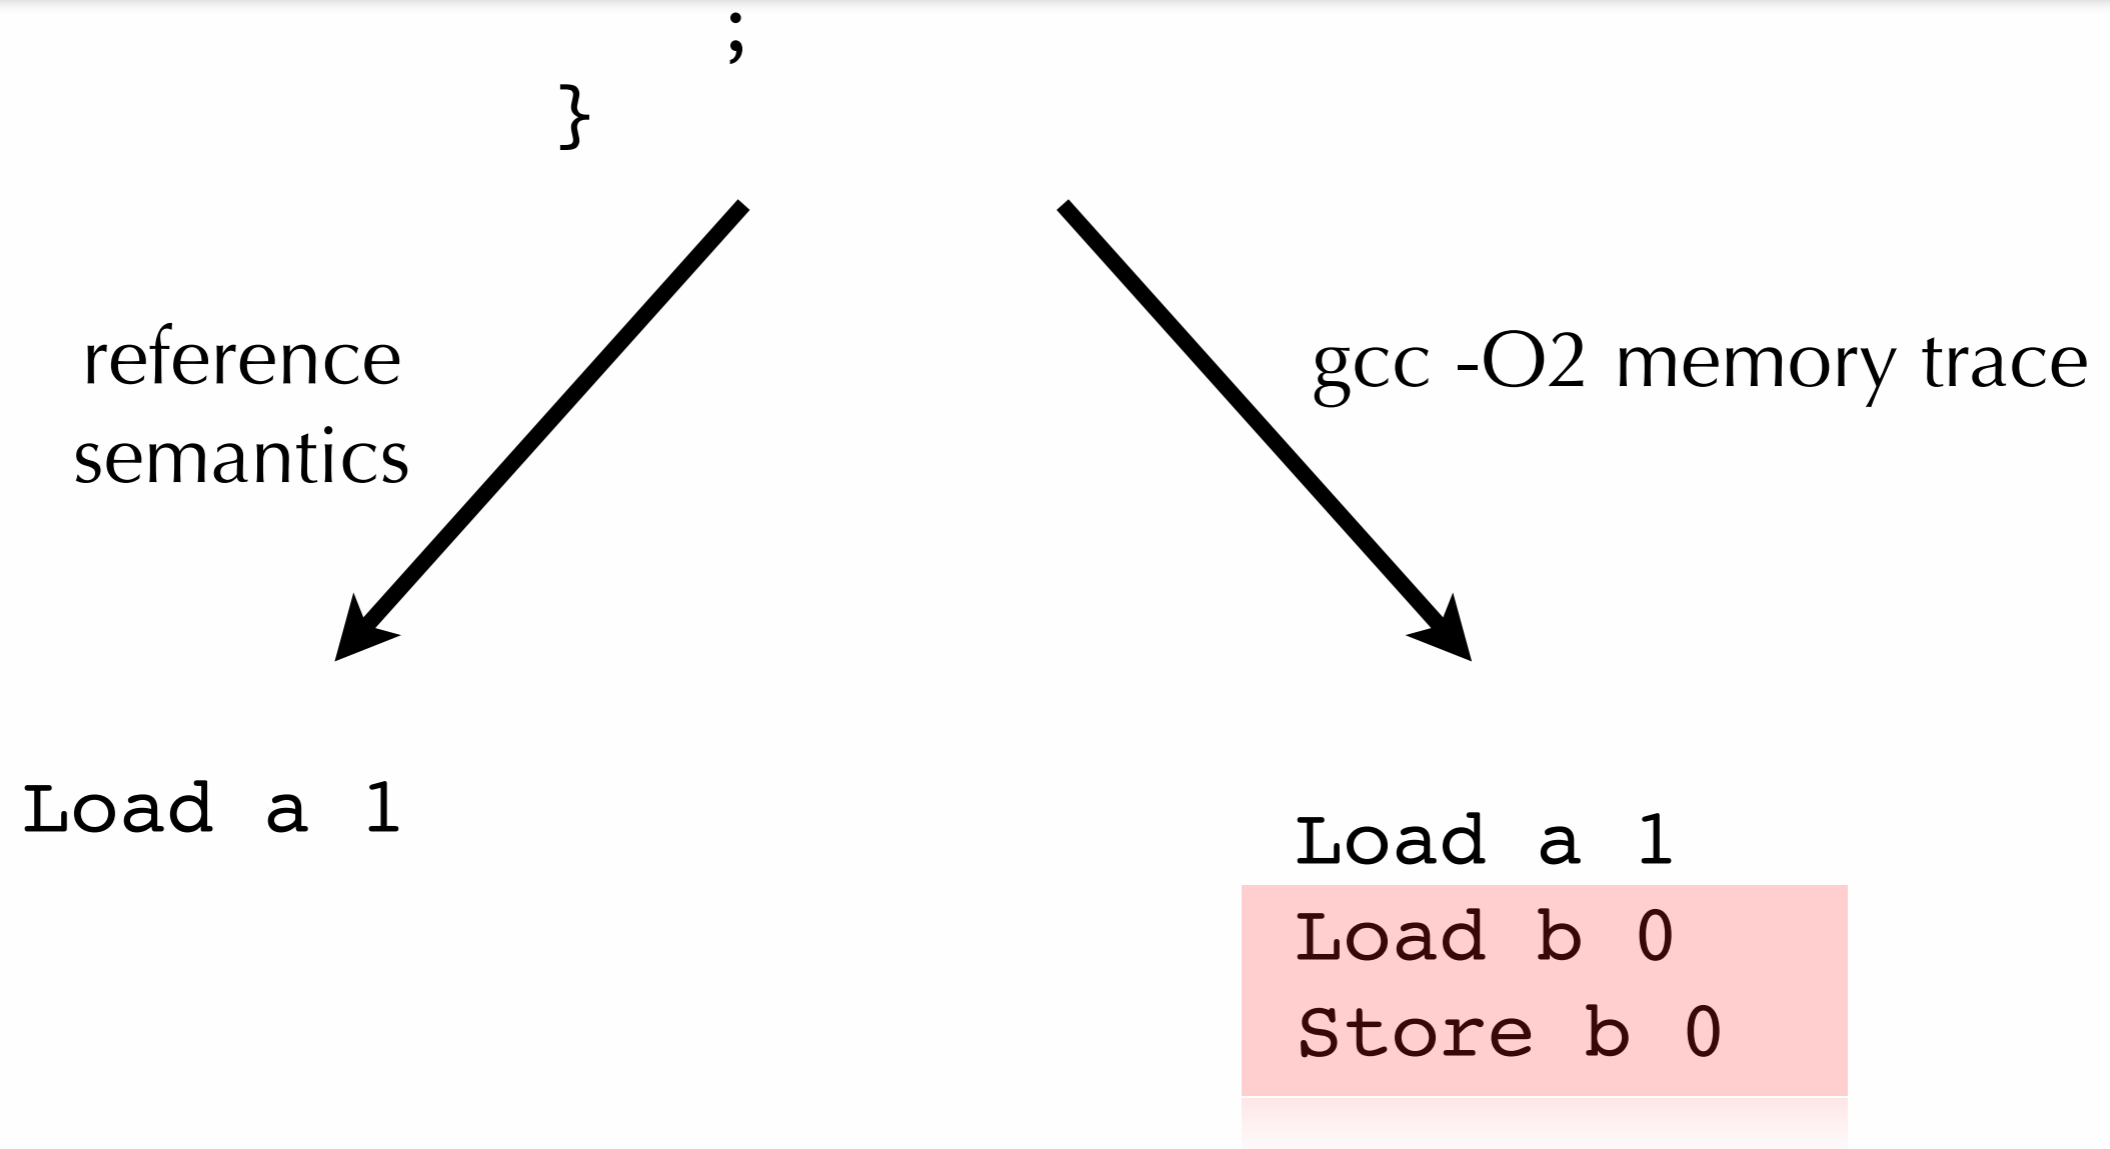
reference
semantics

Load a 1

gcc -O2 memory trace

Load a 1
Load b 0
Store b 0

Cannot match some events → detect compiler bug



Contributions

Sound optimisations in the C11/C++11 memory model
extending Sevcik's work on an idealised DRF model - PLDI 11

A tool to hunt concurrency bugs in C and C++ compilers

Interaction with GCC developers

Sound Optimisations **in the C11/C++11 Memory Model**

Example: loop invariant code motion

Compiler Writer



Semanticist



Example: loop invariant code motion

Compiler Writer



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

Semanticist



Example: loop invariant code motion

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 ;  
}
```


Example: loop invariant code motion

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 ;  
}
```


Example: loop invariant code motion

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 ;  
}
```

Example: loop invariant code motion


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
```

Example: loop invariant code motion


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
```

Elimination of *overwritten writes*



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

Elimination of *overwritten writes*



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

What is the semantics of C11/C++11 concurrent code?

The C11/C++11 memory model

C11/C++11 are based on the DRF approach:

- racy code is undefined
- race-free code must exhibit only sequentially consistent behaviours
- main synchronisation mechanism: lock/unlock

Escape mechanism for experts, *low-level atomics*:

- races allowed
- attributes on accesses specify their semantics:

`MO_SEQ_CST`

`MO_RELEASE/MO_ACQUIRE`

`MO_RELAXED`

MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1,MO_RELEASE);
```

Thread 2

```
while (f.load(MO_ACQUIRE)==0);  
printf ("%d",g)
```


MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1,MO_RELEASE);
```

Thread 2

```
while (f.load(MO_ACQUIRE)==0);  
printf ("%d",g)
```

MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1, MO_RELEASE);
```

Thread 2

```
while (f.load(MO_ACQUIRE) == 0);  
printf ("%d", g)
```

MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1, MO_RELEASE);
```

Thread 2

```
while (f.load(MO_ACQUIRE) == 0);  
printf ("%d", g)
```

MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1, MO_RELEASE);
```

sync →

Thread 2

```
while (f.load(MO_ACQUIRE) == 0);  
printf ("%d", g)
```

MO_RELEASE / MO_ACQUIRE

```
g = 0; atomic f = 0;
```

Thread 1

```
g = 42;  
f.store(1, MO_RELEASE);
```

Thread 2

```
while (f.load(MO_ACQUIRE) == 0);  
printf ("%d", g)
```

sync

The release/acquire synchronisation guarantees that:

- the program is DRF
- 42 is printed at the end of the execution

Remark: unlock \approx release, lock \approx acquire.


Same-thread release/acquire pairs

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 pairs

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



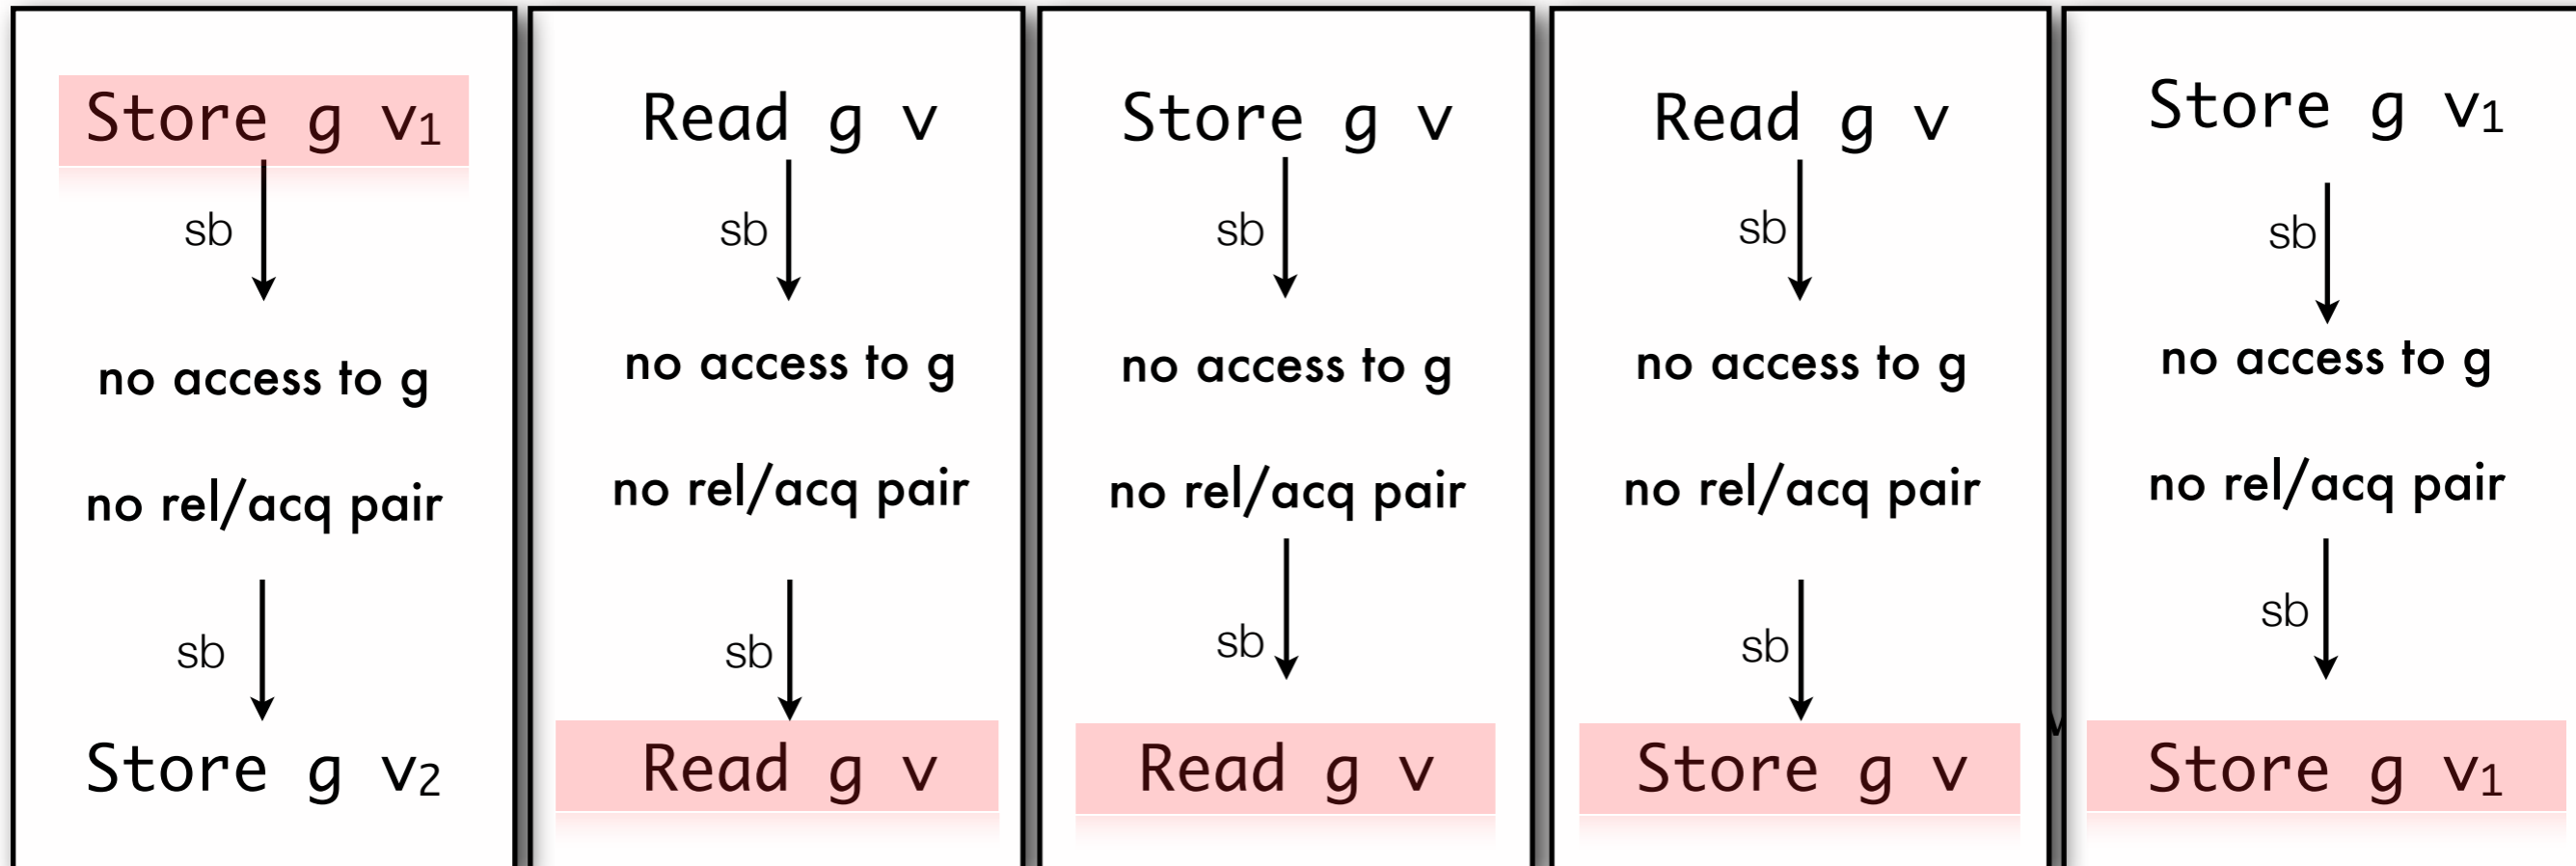
data race



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

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

Elimination: ~~beating~~

Theorem

Soundness proved w.r.t. Batty et al. formalisation of the C11/C++11 memory model (POPL 11)



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

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

Reorderings and introductions

Correctness criterion for reordering events:

- different addresses
- no synchronisations in-between

Roach-motel reordering (reordering across locks) not observed in practice

Read introductions observed in practice (gcc, clang).

Introduction of eliminable reads proved correct.

Introduction of irrelevant reads does not introduce new behaviours, but cannot be proved correct in a DRF model.

The CMMTEST Tool

**Random
Generator**



**SEQUENTIAL
PROGRAM**

*optimising
compiler
under test*

EXECUTABLE

tracing

reference
semantics

**REFERENCE
MEMORY
TRACE**

**MEMORY
TRACE**



**only transformations sound in any
concurrent (non-racy) context?**

CSmith
extended with locks
and atomics

SEQUENTIAL
PROGRAM

*optimising
compiler
under test*

reference
semantics

EXECUTABLE

tracing

REFERENCE
MEMORY
TRACE

MEMORY
TRACE

only transformations sound in any
concurrent (non-racy) context?



CSmith
extended with locks
and atomics

SEQUENTIAL
PROGRAM

*optimising
compiler
under test*

reference
semantics

EXECUTABLE


binary
instrumentation

REFERENCE
MEMORY
TRACE

MEMORY
TRACE

only transformations sound in any
concurrent (non-racy) context?





only transformations sound in any concurrent (non-racy) context?

CSmith
extended with locks
and atomics

SEQUENTIAL
PROGRAM

*optimising
compiler
under test*

gcc/clang -O0

EXECUTABLE

EXECUTABLE

*binary
instrumentation*

REFERENCE
MEMORY
TRACE

binary
instrumentation

MEMORY
TRACE



OCaml tool

1. analyse the traces to detect eliminable actions
2. match reference and optimised traces

CSmith
extended with locks
and atomics



SEQUENTIAL
PROGRAM

*optimising
compiler*

Subtleties:

- *dependencies between eliminable events*
- *some optimisations (e.g. merging of accesses) cannot be expressed in the C11/C++11 formalisation*
- *the tool also ensures that the compilation of atomic accesses is preserved by the optimiser*

TRACE

OCaml tool

1. analyse the traces to detect eliminable actions
2. match reference and optimised traces

Interaction with GCC developers

1. Some GCC bugs

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

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

All promptly fixed.

Remark: these bugs break the Posix thread model too.

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
```



3. Detecting unexpected behaviours

uint16_t g

for (; g==0; 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.

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.



The memory models of modern hardware are better understood.

Still, many open problems...

problems.

attempt

els.

mmers

se



The memory models of modern hardware are better understood.

Still, many research opportunities!

problems.



attempt

els.

mmers

se

All these lectures are based on work done with/by my colleagues. Thank you!



And thank you all for attending these lectures!

Please, fill the course evaluation form, that's important to make a better course next year.

