

# WHAT IS A MEMORY MODEL?

JAEHEON YI  
OCTOBER 22, 2008  
TWIGS

# THAT CODE DOESN'T WORK?

---

```
red == blue == 0
```

```
<p1>
```

```
red = 1;
```

```
if (blue == 0)
```

```
    //critical
```

```
    //section
```

```
<p2>
```

```
blue = 1;
```

```
if (red == 0)
```

```
    //critical
```

```
    //section
```

# ARCHITECTURAL TRENDS

---

- Relative cost of memory access expensive
  - 100's of cycles to access
- Processor speed difficult to increase
  - Ridiculous energy consumption

# IMPLICATIONS FOR SOFTWARE

---

- Everybody has multicore, soon manycore
- Dominance of shared-memory paradigm
- Multiple threads of execution for performance
- H/W and compilers try to hide memory latencies


# HARDWARE FUNNY BUSINESS

---

- Some examples:
  - Write buffers
  - Caches at multiple levels
  - Out-of-order execution
  - Speculative execution


# WRITE BUFFERS

---



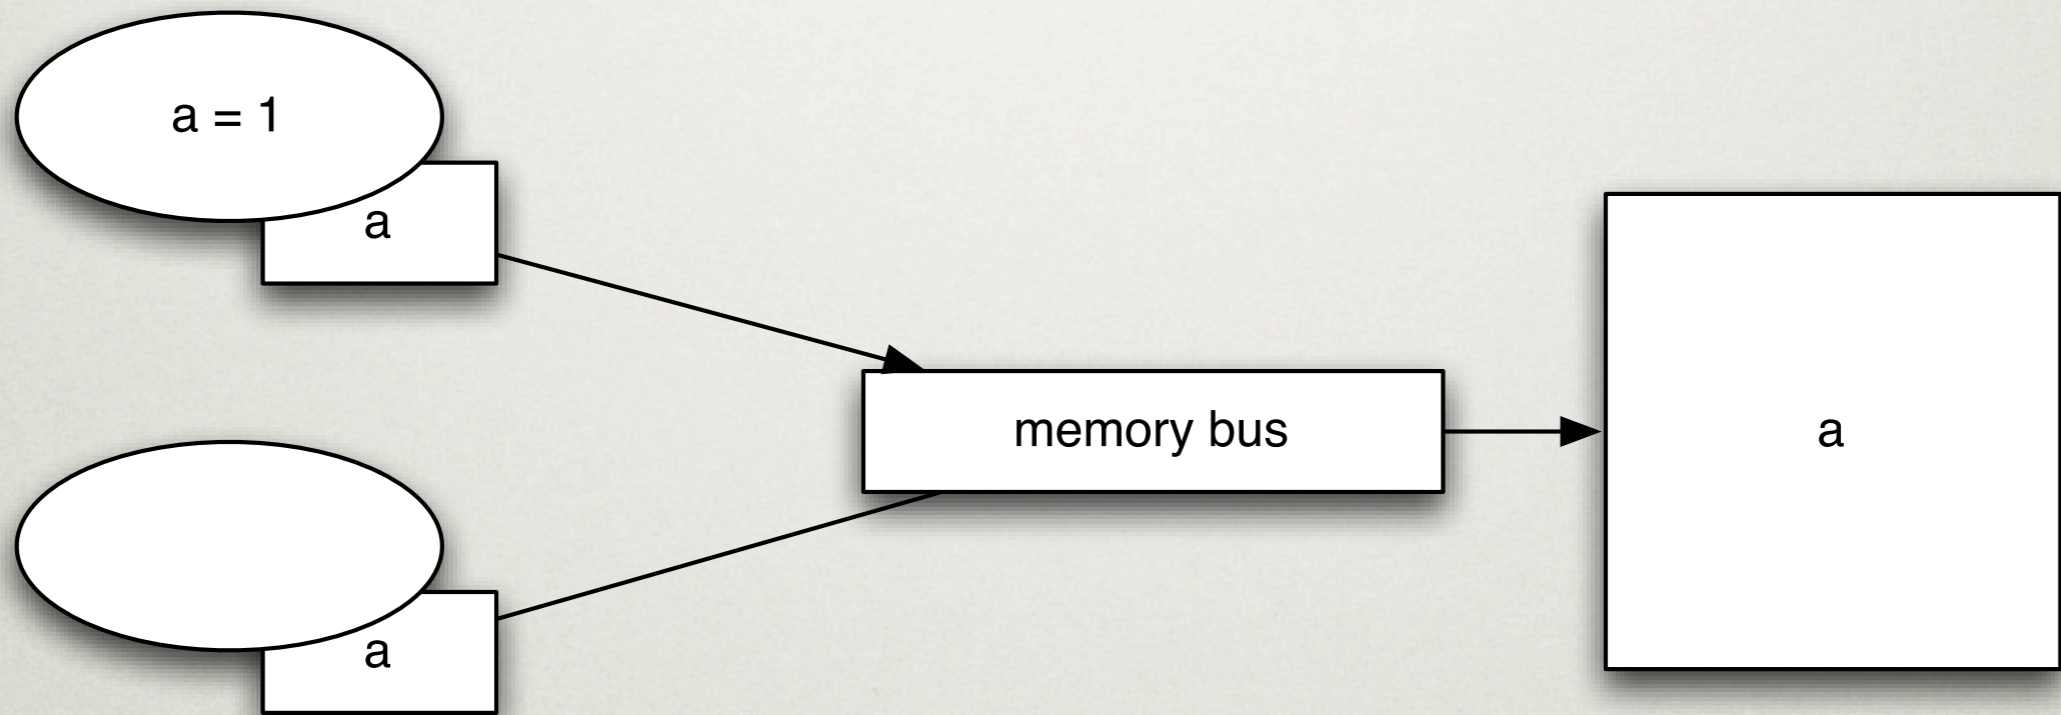
# WRITE BUFFERS

---



# CACHES

---





# OUT-OF-ORDER EXECUTION

---

`r0 = x`      `->`      `r0 = x`  
`x = 3`                      `r1 = 2`  
`r1 = 2`                      `x = 3`

# SPECULATIVE EXECUTION

---

|                        |                    |                                 |
|------------------------|--------------------|---------------------------------|
| <code>if (x==0)</code> | <code>-&gt;</code> | <code>Assume x is 0</code>      |
| <code>    y = 0</code> |                    | <code>y = 0</code>              |
| <code>r1 = 1</code>    |                    | <code>Is x really 0? (y)</code> |
|                        |                    | <code>r1 = 1</code>             |

# COMPILER FUNNY BUSINESS

---

- Some examples:
  - Loop-invariant code motion
  - Register allocation
  - Common sub-expression elimination

# LOOP-INV. CODE MOTION

---

```
while (i < 100)
    a[i] = (b*3+3) + i
```

*becomes*

```
int r0 = b*3+3
while (i < 100)
    a[i] = r0 + i
```

# REGISTER ALLOCATION

---

a = 3 \* a

c = b + 4

->

r0 = a

r0 = r0 \* 3

a = r0

r1 = b

r1 = r1 + 4

c = r1

# COMMON SUBEXPRESSION ELIMINATION

---

```
x = 3*4 + a[i]*a[j]
y = a[i]*a[j] - a[k]
```

*becomes*

```
tmp = a[i]*a[j]
x = 3*4 + tmp
y = tmp - a[k]
```

# THAT CODE DOESN'T WORK.

---

```
red == blue == 0
```

```
<p1>
```

```
red = 1;
```

```
if (blue == 0)
```

```
    //critical
```

```
    //section
```

```
<p2>
```

```
blue = 1;
```

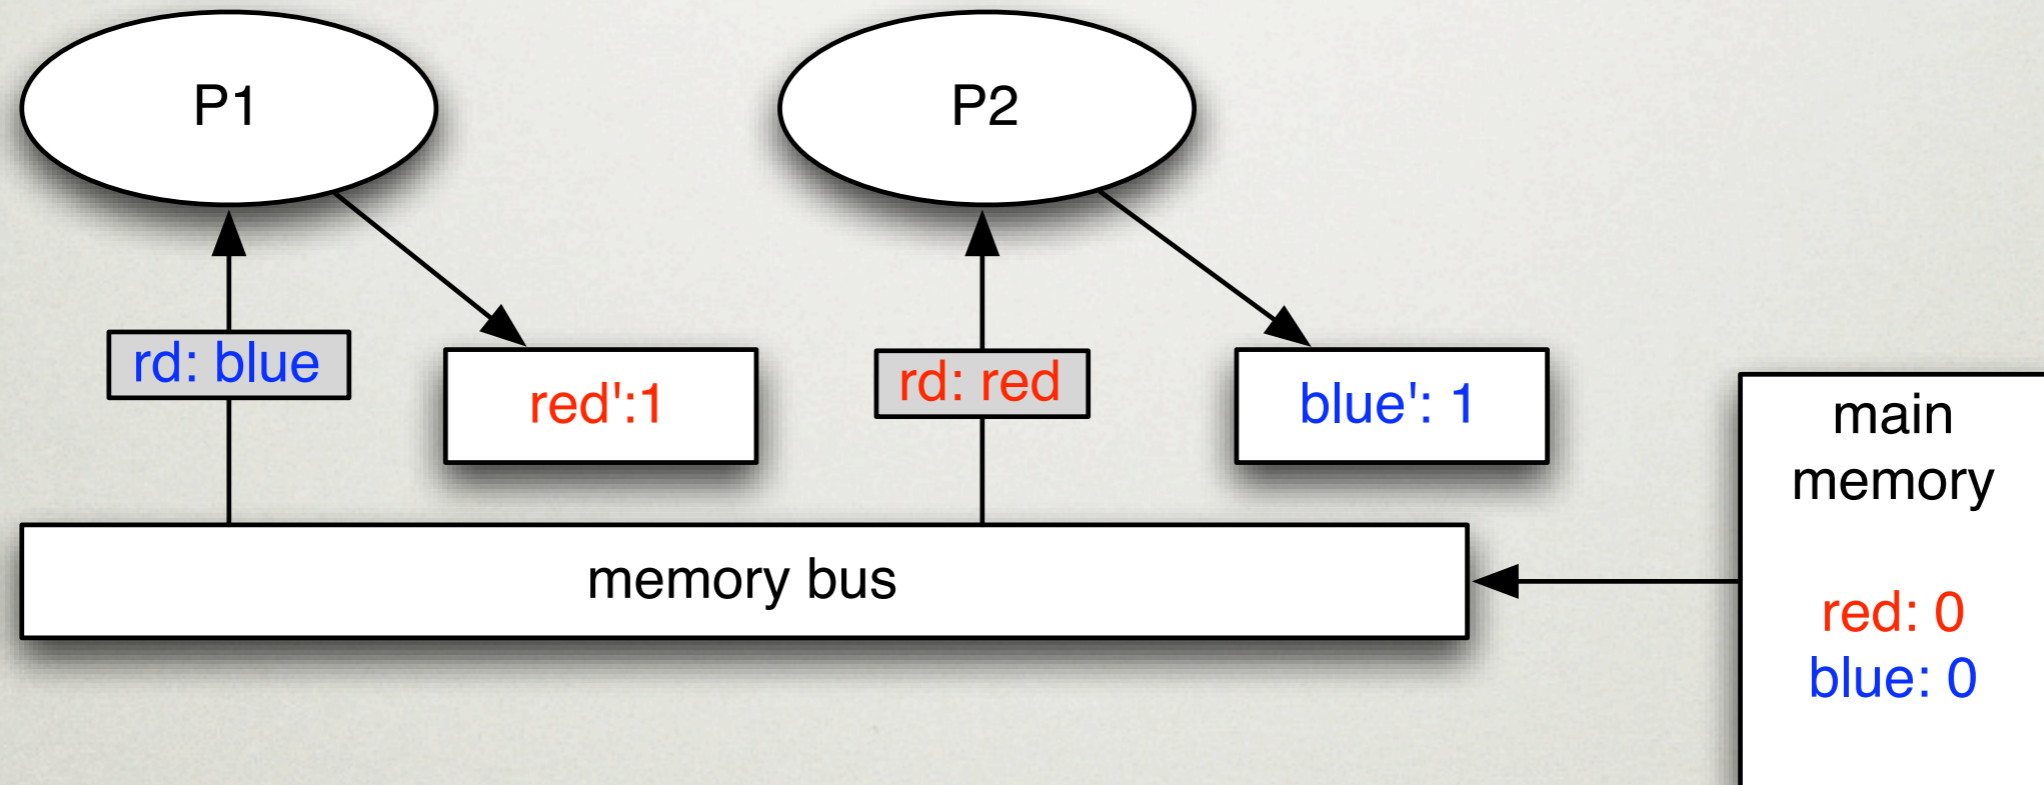
```
if (red == 0)
```

```
    //critical
```

```
    //section
```

# THAT CODE DOESN'T WORK.

---





# THAT CODE DOESN'T WORK.

---

```
red == blue == 0
```

```
<p1>
```

```
r0 = blue;
```

```
red = 1;
```

```
if(r0==0)
```

```
    //critical
```

```
    //section
```

```
<p2>
```

```
r0 = red;
```

```
blue = 1;
```

```
if(r0==0)
```

```
    //critical
```

```
    //section
```

# MEMORY CONSISTENCY MODEL

---

- Defines observable values
- Observation means “reads”
- What can the value be that you read out of memory?

# MEMORY CONSISTENCY MODEL

---

- Architecture level
  - x86 family has a memory model
- Language level
  - Java has one, C doesn't
  - C++ memory model in the works


# SEQUENTIAL CONSISTENCY

---

- In each processor, instructions execute in program order
- Memory operations appear to execute one-at-a-time across entire system

# SEQUENTIAL CONSISTENCY

---




# SEQUENTIAL CONSISTENCY

---

$r0 = a$


$b = 2$

$a: 1$   
 $b: 2$



# SEQUENTIAL CONSISTENCY

---



# SEQUENTIAL CONSISTENCY

---

- Canonical, *intuitive*, most restrictive
- Prevents optimizations
  - Kills performance



# RELAXED MEMORY MODELS

---

- Speeds up processors
- Harder to understand and program
- Many flavors

# WAYS TO RELAX

---

- Program order may be relaxed
- Write atomicity may be relaxed

# WAYS TO RELAX

---

- Program ordering requirement
  - Write(x); Read(y)
  - Write(x); Write(y)
  - Read(x); Read(y) / Write(y)

# WAYS TO RELAX

---

- Write atomicity requirement
  - Write<sub>1</sub>(x); early Read<sub>2</sub>(x)
  - Write<sub>1</sub>(x); early Read<sub>1</sub>(x)

# RELAXING TOO MUCH? RESTORING ORDER

---

- Relaxed memory models often come with special instructions in ISA
- Hardware synchronization primitives
- Expensive to use
- Fence, CAS, LL/SC

# RELAXING TOO MUCH?

## RESTORING ORDER

---

- Memory fence instruction
  - Also called a memory barrier
  - Makes all pending writes visible to issuer
  - May cost 100's of cycles to execute
  - Load fences, store fences, full fences

# RELAXING TOO MUCH? RESTORING ORDER

---

- Compare-And-Swap instruction
  - Universal concurrency primitive
  - Atomically executed write
  - `bool CAS(addr, oldVal, newVal)`

# RELAXING TOO MUCH? RESTORING ORDER

---

- Load-Linked / Store-Conditional instruction pair
  - Universal concurrency primitive
  - LL reads from address x
  - SC attempts write to address x



# RELAXING TOO MUCH? RESTORING ORDER

---

- Synchronization primitives to build
  - Mutex
  - Semaphore
  - Condition Variables
  - Monitors

# EASIER WAY OUT

---

- Reasoning about relaxations is not relaxing!
- Virtually all explicitly defined memory models have a safety hatch built in

# EASIER WAY OUT

---

- Fundamental Property of Memory Models:
  - If your program is free of data races, then it will have sequentially consistent semantics

# EASIER WAY OUT

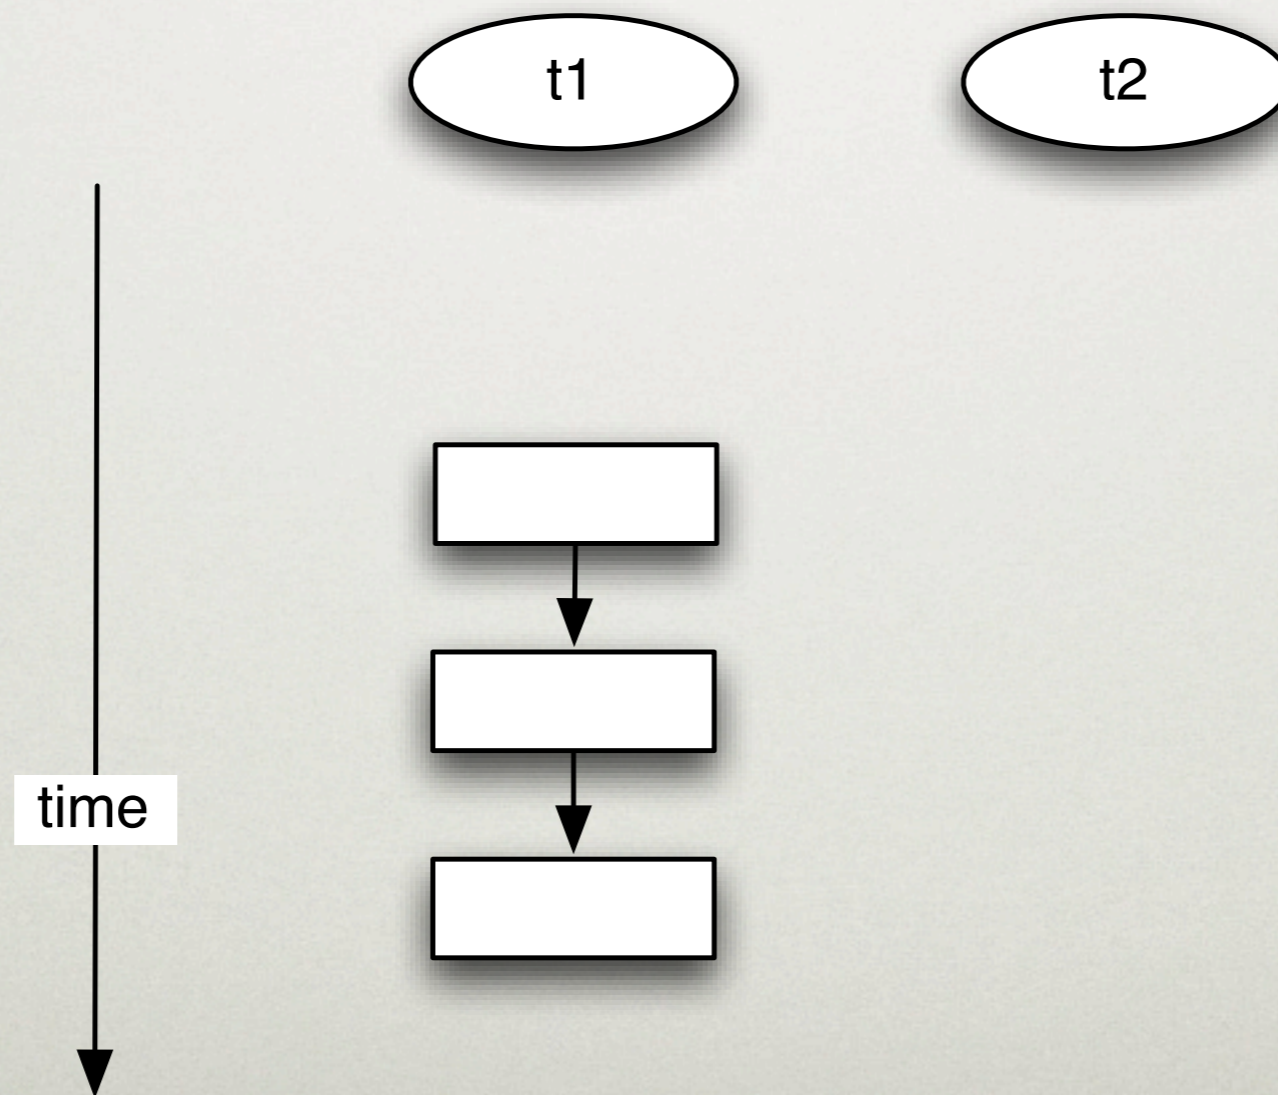
---

- Happens-before relation
  - A binary relation on executed instructions, transitively-closed
  - If  $x$  occurs before  $y$  in time, and  $x$  *conflicts with*  $y$ , then  $x$  *happens-before*  $y$

# EASIER WAY OUT

---

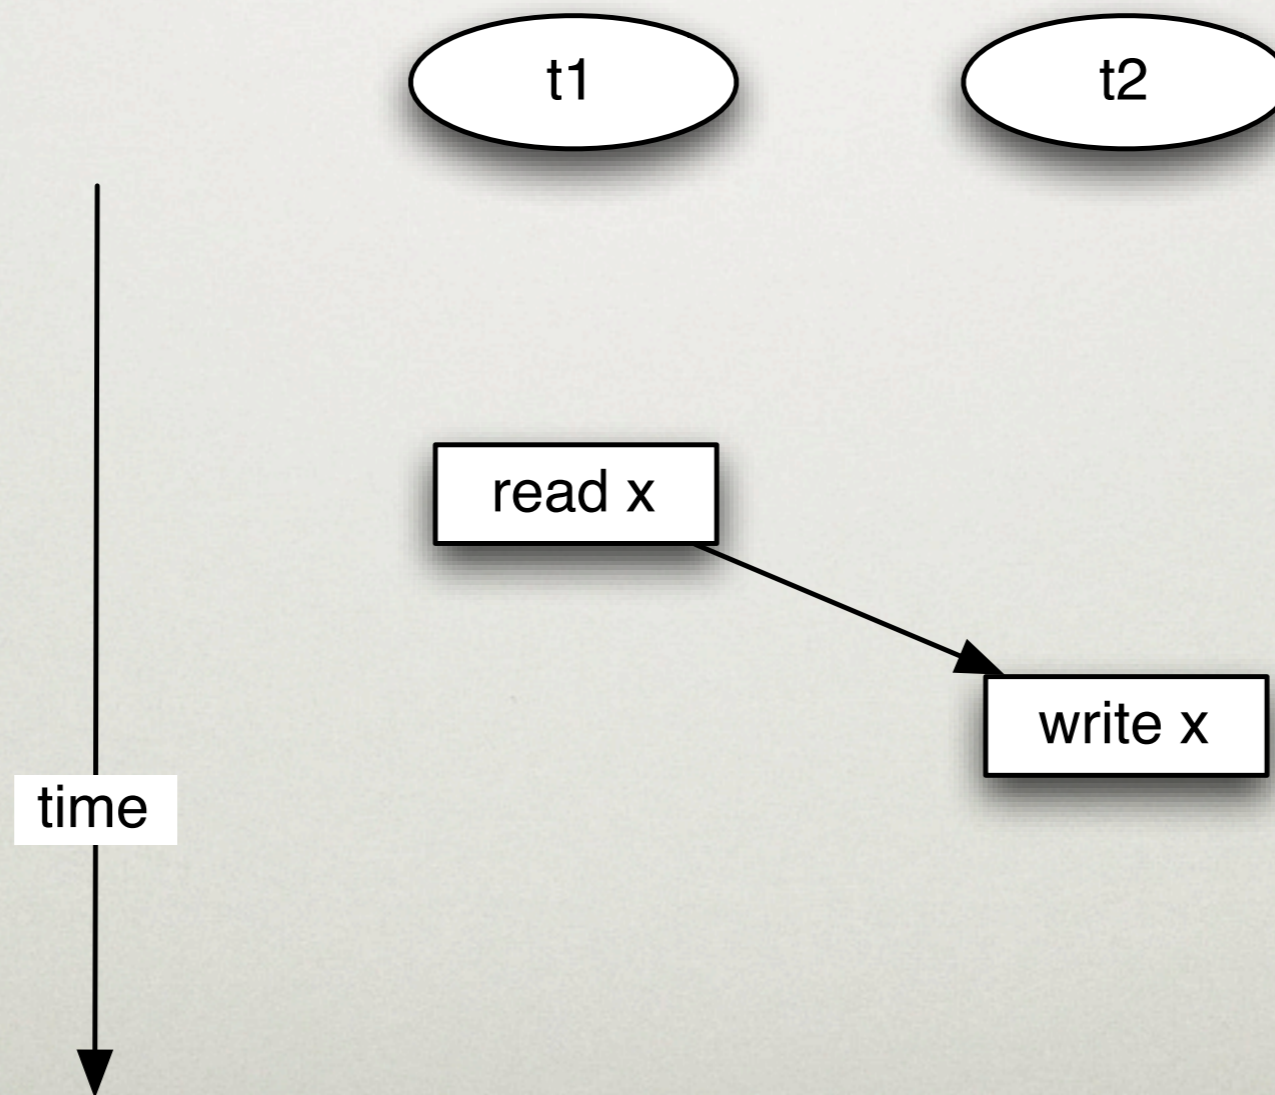
- *x conflicts with y*, program order



# EASIER WAY OUT

---

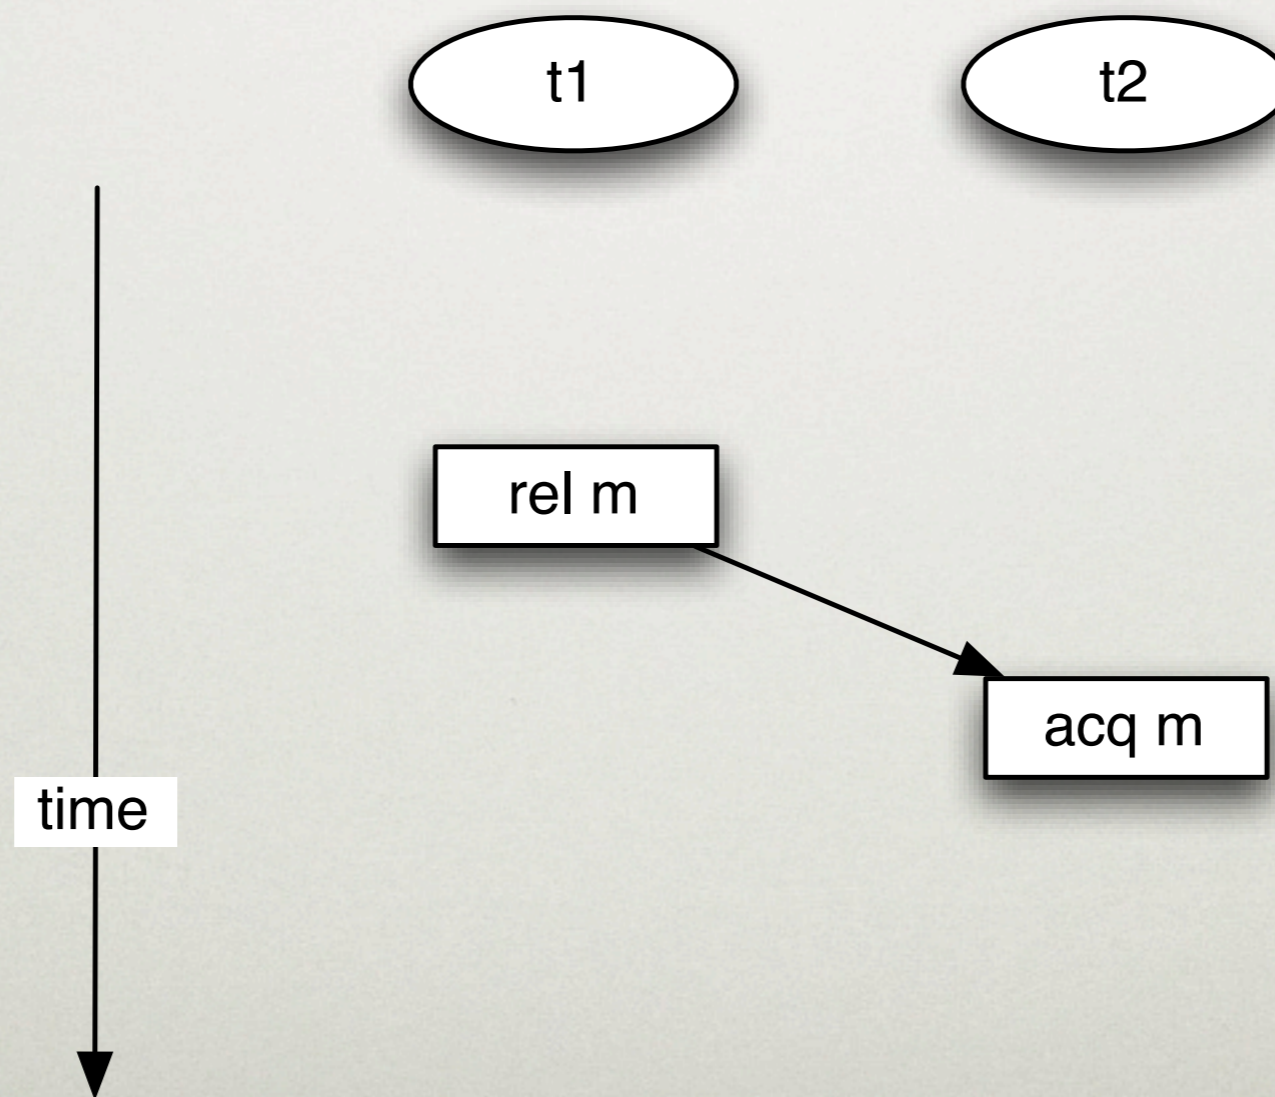
- *x conflicts with y, communication order*



# EASIER WAY OUT

---

- *x conflicts with y, synchronization order*



# EASIER WAY OUT

---

- Data race occurs when:
  - A variable is read by  $>1$  threads
  - It is written by at least one thread
  - Read / write *not* ordered by synchronization edge



# EASIER WAY OUT

---

- Data-Race-Freedom (DRF) property:  
The absence of any data races in a program
- Defines synchronization property at the program level, not the execution trace level

# X86 MEMORY MODEL

---

- Formally defined as of last year
- 8 principles define possible instruction orderings

# X86 MEMORY MODEL

---

1. Loads are *not* reordered with other loads.
2. Stores are *not* reordered with other stores.
3. Stores are *not* reordered with older loads.
4. Loads may be reordered with older stores to different locations but *not* with older stores to the same location.
5. In a multiprocessor system, memory ordering obeys causality (memory ordering respects transitive visibility).
6. In a multiprocessor system, stores to the same location have a total order.
7. In a multiprocessor system, locked instructions have a total order.
8. Loads and stores are *not* reordered with locked instructions.

# X86 MEMORY MODEL

---

- DRF programs will have sequentially consistent semantics

# JAVA MEMORY MODEL

---

- Part of the *language* definition
- First version got it wrong
- Second version worth a PhD thesis
- Tension between allowing efficient optimizations by JIT and processors, and programmer understanding

# JAVA MEMORY MODEL

---

- Three nice properties
  - Consistency guarantees for programmers with DRF programs
  - Guarantee for compiler writers and VM architects
  - Guarantee that programmers can use nonblocking synchronization

# JAVA MEMORY MODEL

---

- DRF guarantee: Correctly synchronized programs will have sequentially consistent semantics
- Correctly synchronized IFF all sequentially consistent executions are DRF

# JAVA MEMORY MODEL

---

- Different *happens-before* relation
  - Program order
  - Synchronization order
  - Volatile order
  - Thread fork / join order
  - Interrupt order
  - Finalizer order
  - Transitivity



# JAVA MEMORY MODEL

---

- If non-DRF program, then what can it do?
- Accounts for substantial part of JMM definition

# WHAT DRF GETS YOU

---

- Can ignore memory models if DRF
- Now need to *ensure* that your program is DRF
  - Static and dynamic methods
- Necessary, but often not sufficient

# OTHER CORRECTNESS PROPERTIES

---

- Atomicity
  - Implementable by transactional memory
- Deterministic Parallelism
  - Generalized atomicity

# CONCLUSION

---

- Memory models matter when dealing with shared memory parallel programs
- If you need performance, system details are crucial
- If you program in a data-race-free manner, many of your worries go away

# REFERENCES

---

- The Art of Multiprocessor Programming (Herlihy, Shavit)
- Java Concurrency in Practice (Goetz)
- Shared Memory Consistency Models: A Tutorial (Adve, Gharachorloo)
- Intel 64 Architecture Memory Ordering White Paper (Intel)
- The Semantics of x86 Multiprocessor Machine Code (Sarkar et al.)
- The Java Memory Model (Manson, Pugh, Adve)
- Wikipedia articles