

Stalking the Lost Write: Memory Visibility in Concurrent Java

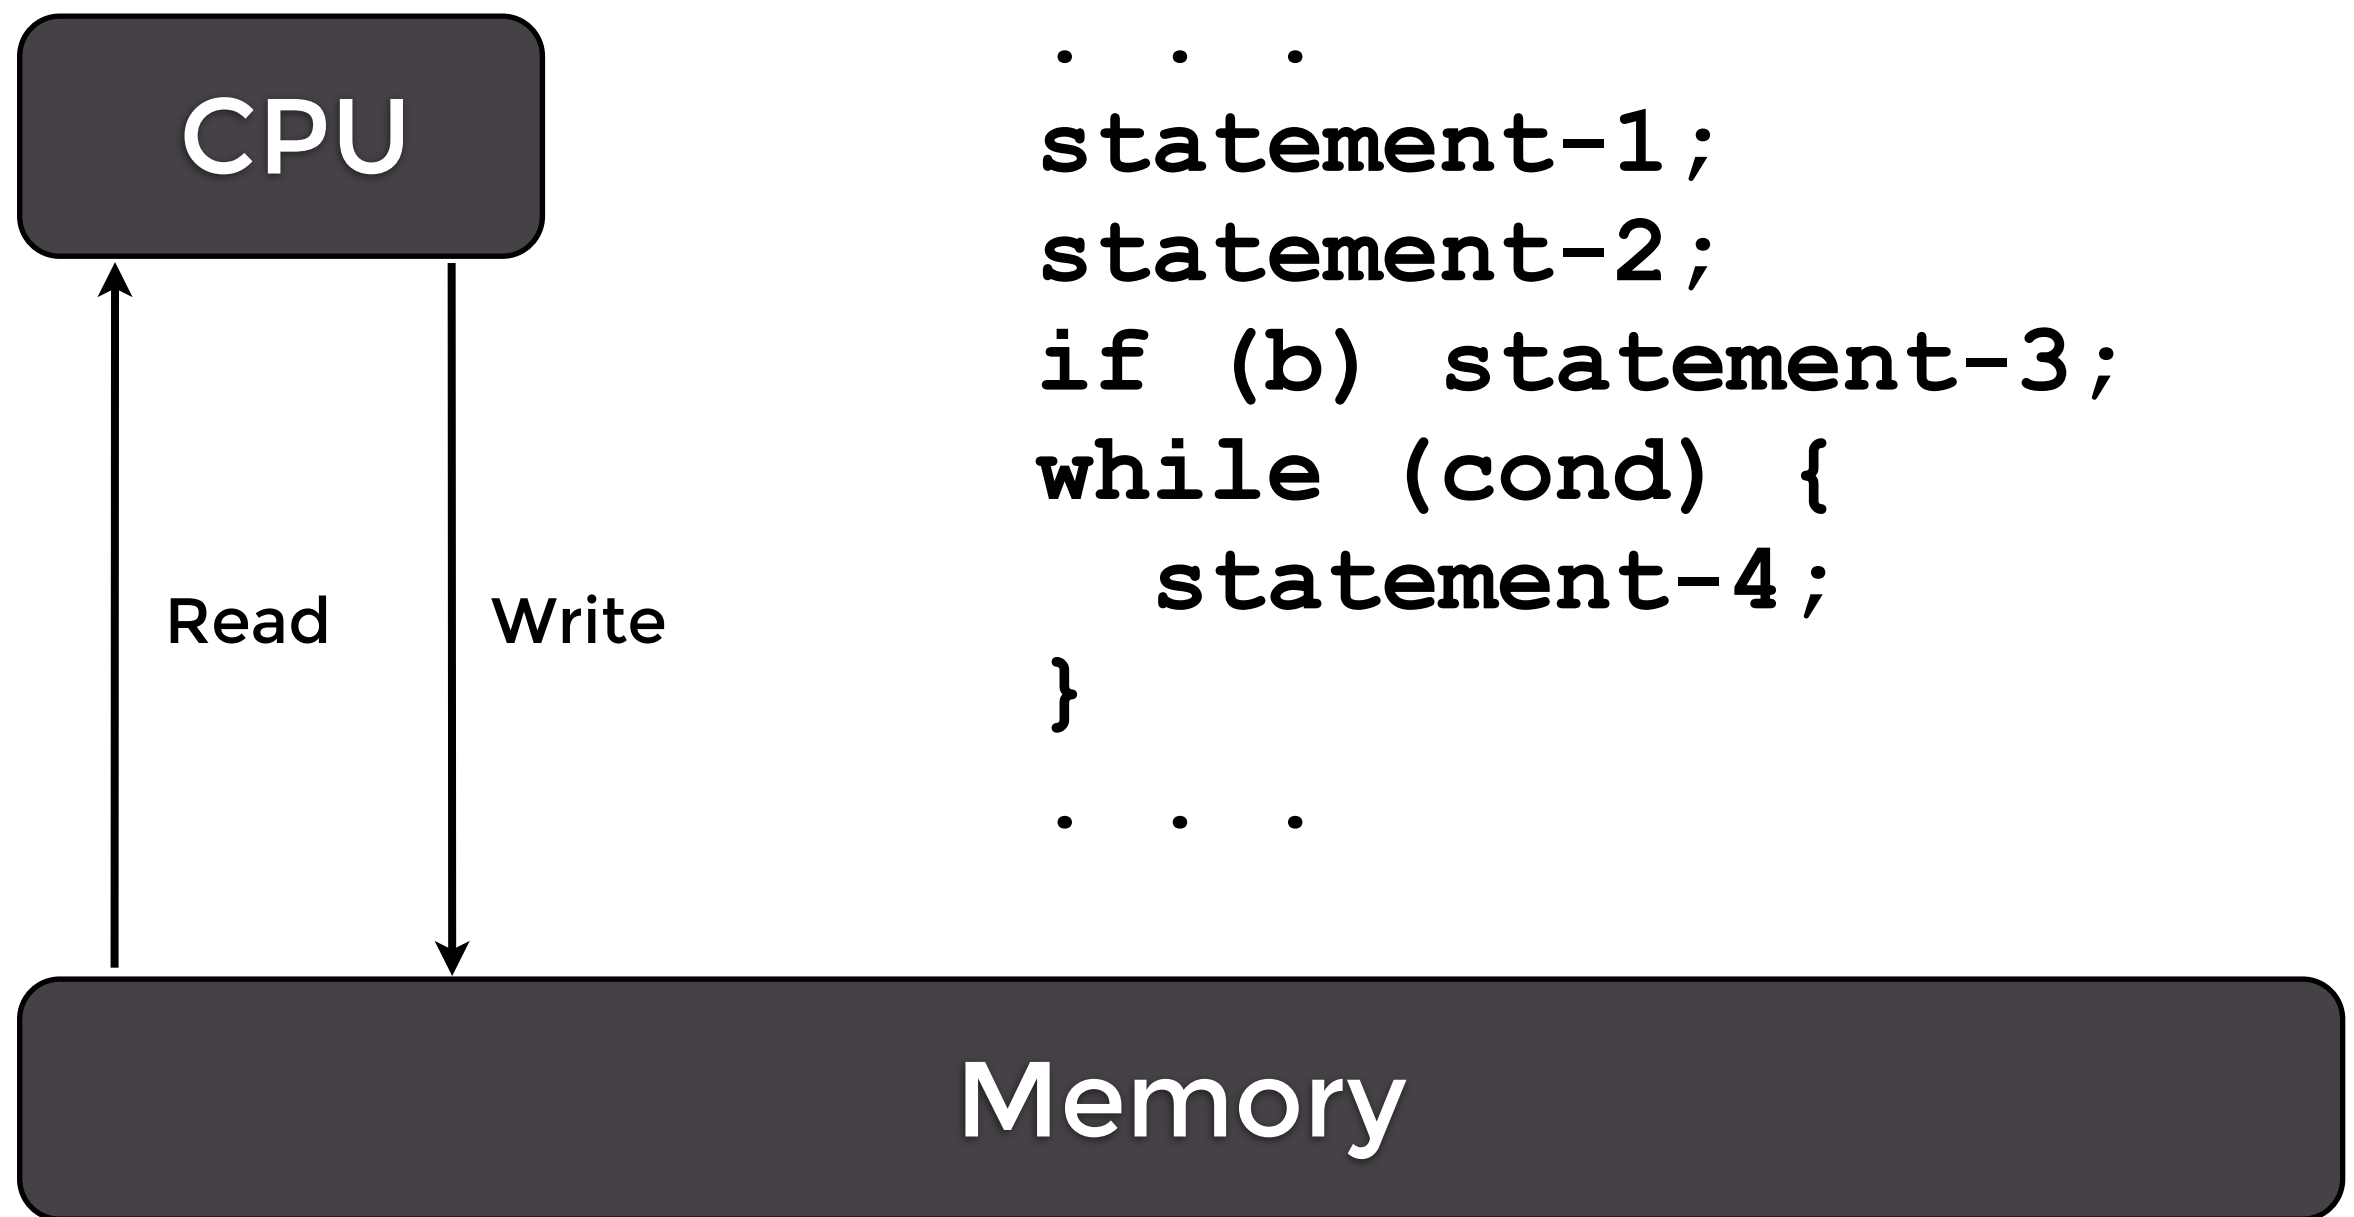
Jeff Berkowitz, New Relic
December 2013

This presentation is is licensed under the Creative Commons
Attribution-NonCommercial-ShareAlike International License 4.0.

It is attributed to Jeff Berkowitz, New Relic Inc.

<http://creativecommons.org/licenses/by-nc-sa/4.0/legalcode>

The Computer We Imagine



The Compiler We Imagine

Java

`x++;`

`y++;`

Assembly Language*

```
mov mem.x, reg1
incr reg1
mov reg1, mem.x
```

```
mov mem.y, reg1
incr reg1
mov reg1, mem.y
```

*** Typical assembly language - no particular CPU**

Monday, December 30, 13

Later we learned that computers didn't really execute what we wrote; rather, a program called a "compiler" translated our program into "machine language". But we could still imagine the intuitive rule of program order for machine language.

The Compiler We Get

Java

x++;

y++;

Assembly Language

**mov mem.x, reg1
mov mem.y, reg2**

**incr reg1
mov reg1, mem.y**

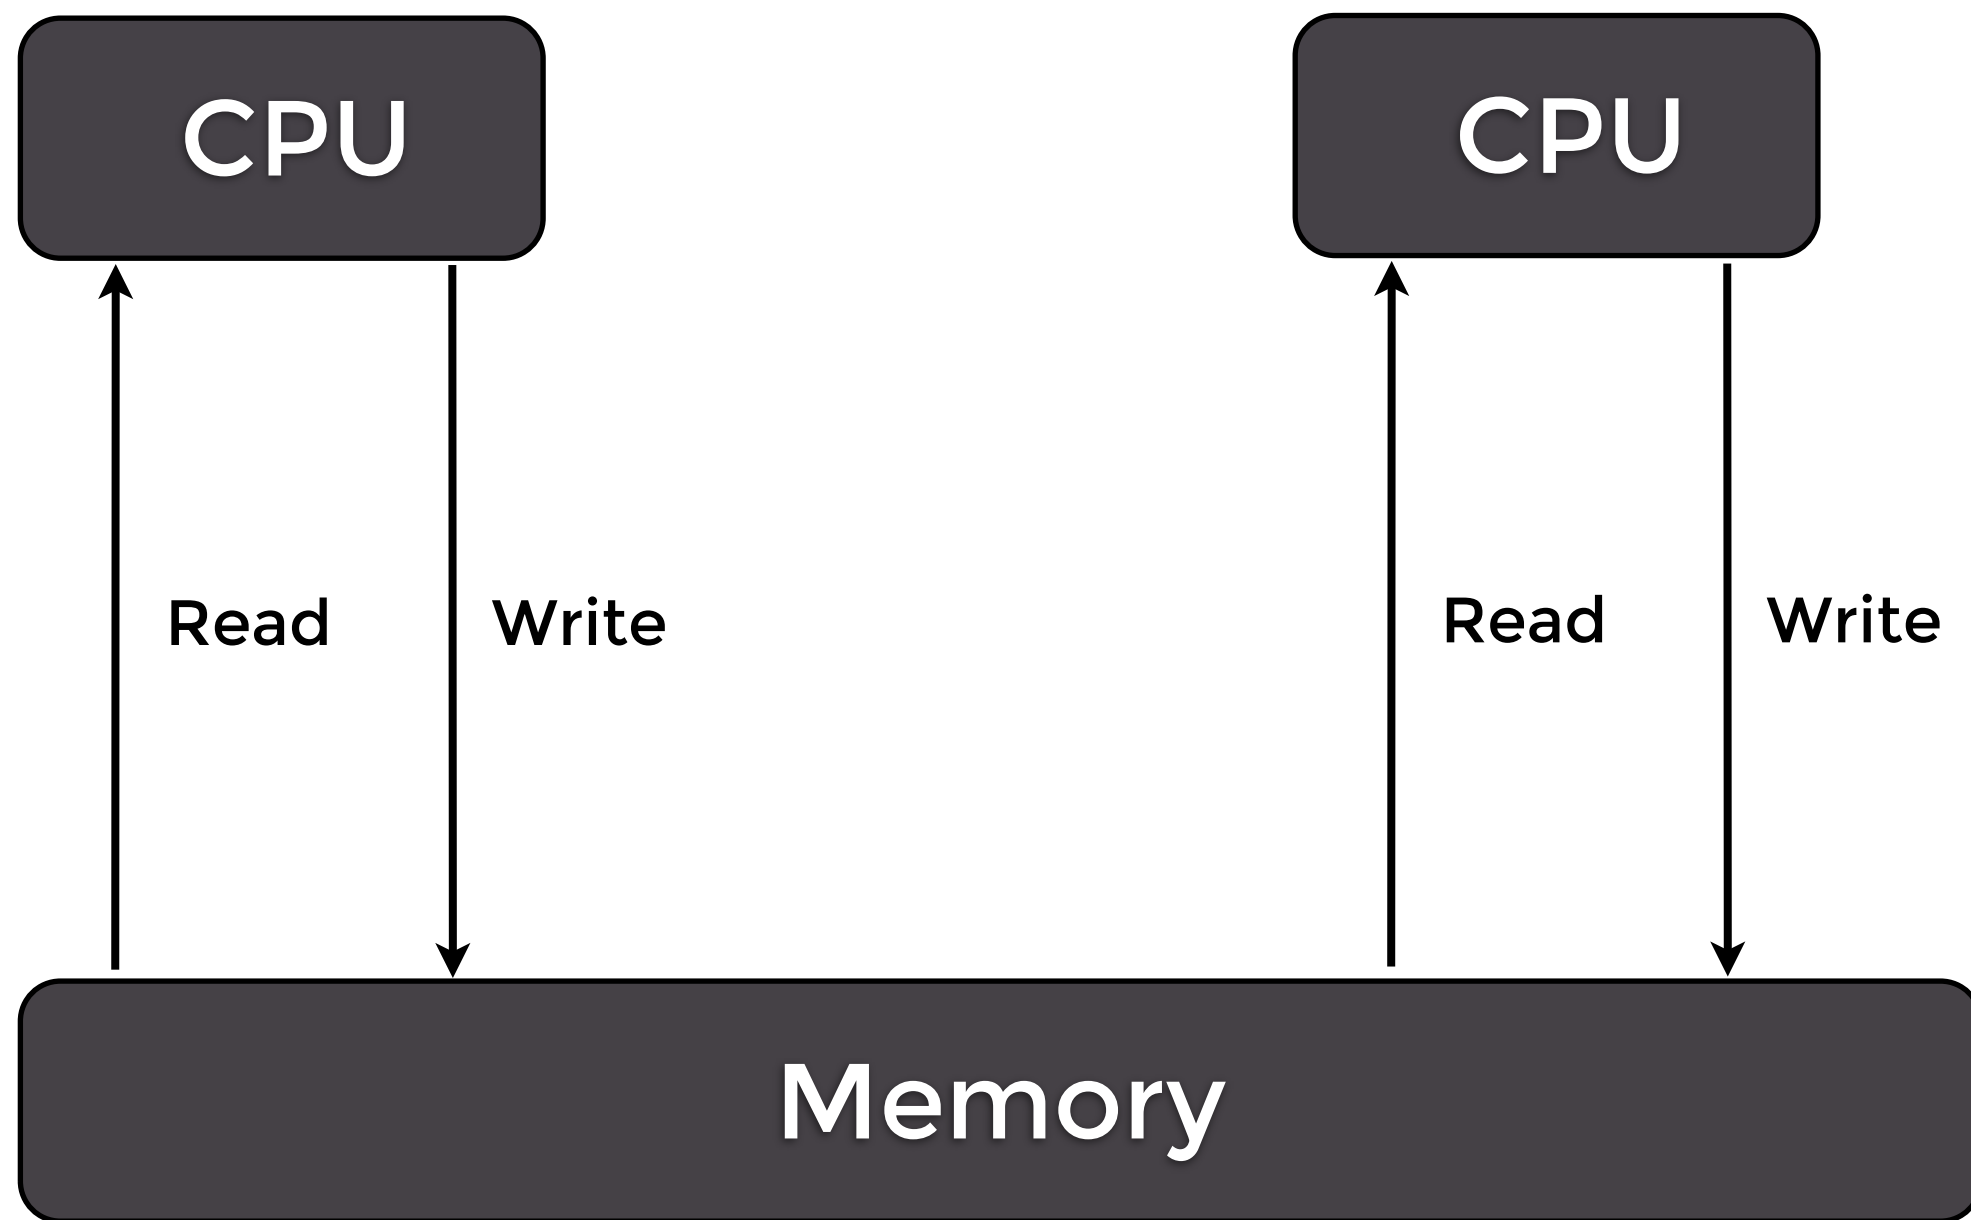
**incr reg2
mov reg2, mem.x**

The End Result

| <u>Java</u> | <u>Assembly Language</u> | <u>Hardware Level*</u> |
|-------------------|--|---|
| <code>x++;</code> | <code>mov mem.x, reg1</code> <code>mov mem.y, reg2</code> | <code>rd.issue(x)</code> <code>rd.issue(y)</code> |
| | <code>incr reg1</code> <code>mov reg1, mem.x</code> | <code>resp.mov(r1)</code> <code>incr r1</code> <code>wr.async(r1, x)</code> |
| <code>y++;</code> | <code>incr reg2</code> <code>mov reg2, mem.y</code> | <code>resp.mov(r2)</code> <code>incr r2</code> <code>wr.async(r2, y)</code> |

*** Typical micro operations - no particular CPU**

The Multiprocessor We Imagine



There are no caches or memory buffering here

Code Example 1

Possible outcomes for x and y?

```
int x, y, a, b; // all zero
```

First CPU ("thread")


```
void m1() {  
    y = a;  
    b = 1;  
}
```

Second CPU

```
void m2() {  
    x = b;  
    a = 2;  
}
```


Do not emulate this example

Possible Execution Trace



Possible Trace #2


Time



Outcome: `x == 0, y == 0`


Possible Trace #3

Time



Outcome: $x == 0, y == 0$

Possible Trace #4



Oh, And #5 and #6

Time

m1()
y = a



b = 1

m2()
x = b



a = 2

m2()
x = b



a = 2

m1()
y = a



b = 1

Outcome: $x == 0, y == 0$

$x == 0, y == 0$

Is That It?

- It looks like x or y must be 0 in the result
- Makes sense: the first statement of $m1()$ grabs a 0, and so does the first statement of $m2()$
- Is our reasoning correct?

Surprisingly, No

```
void m1() {  
    y = a;  
    b = 1;  
}
```


mov #1, mem.b
mov mem.a, mem.y

```
void m2() {  
    x = b;  
    a = 2;  
}
```

mov #2, mem.a
mov mem.b, mem.x

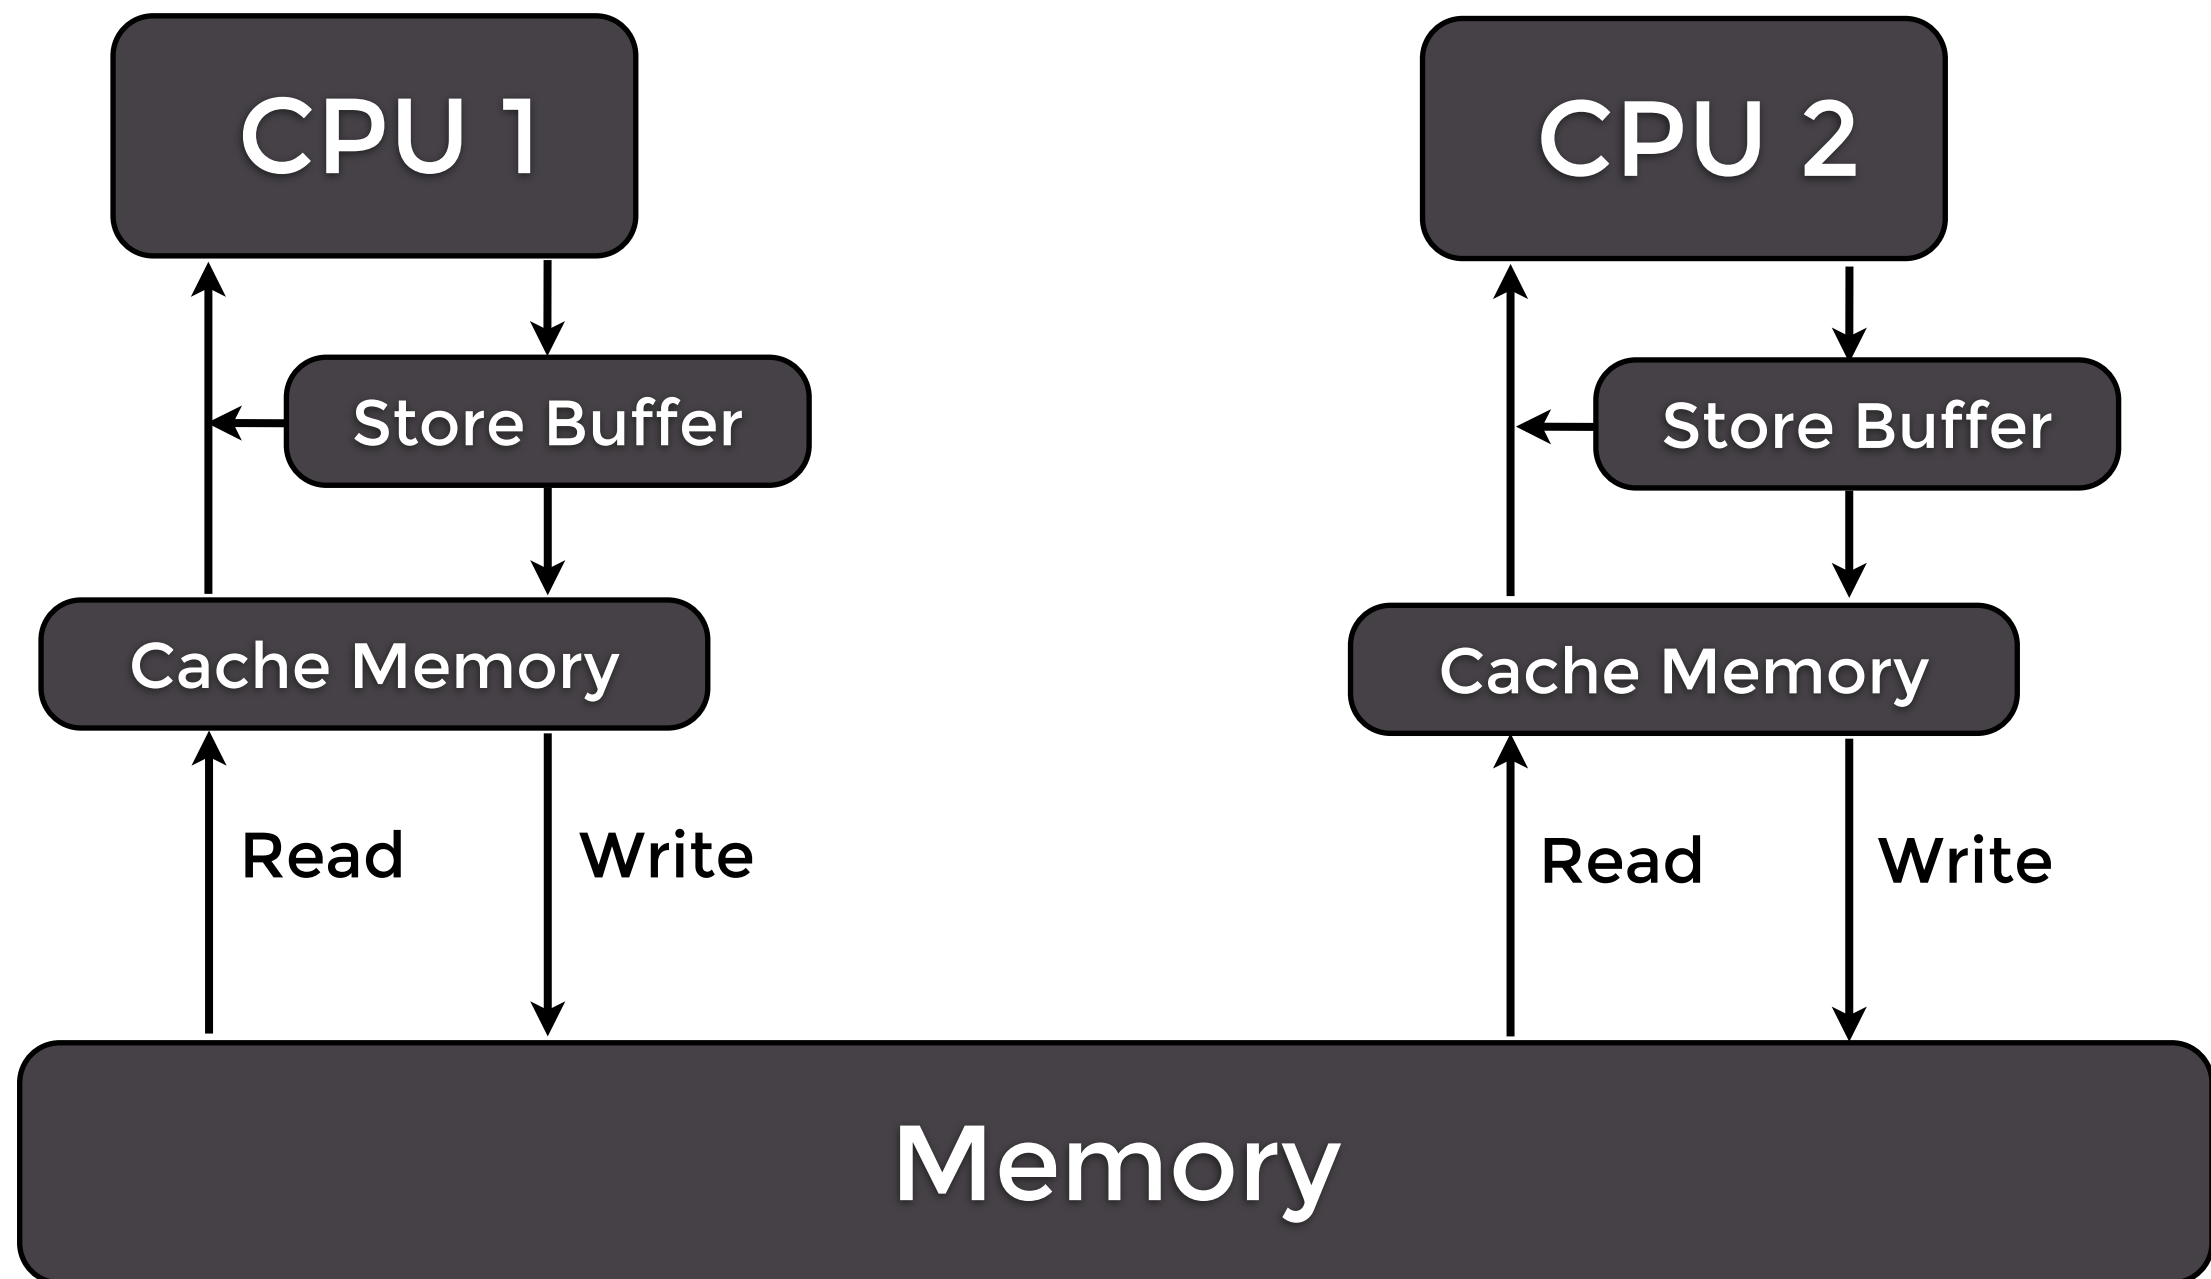
Counterintuitively, the compiler can reverse the order

So This is Also Possible



*

And It Gets Worse...



This is just an example. It's way worse than this.

Monday, December 30, 13

Hardware designers invented caches and store buffers and other performance enhancing hardware tricks. In general, all these optimizations are considered acceptable so long as the intuitive rule of program order isn't violated. We've just seen that the intuitive rule of program order can allow for confusing compiler behavior. Now let's look at the hardware.

Code Example 2

```
int a, b; // both zero
```

First CPU (“thread”)

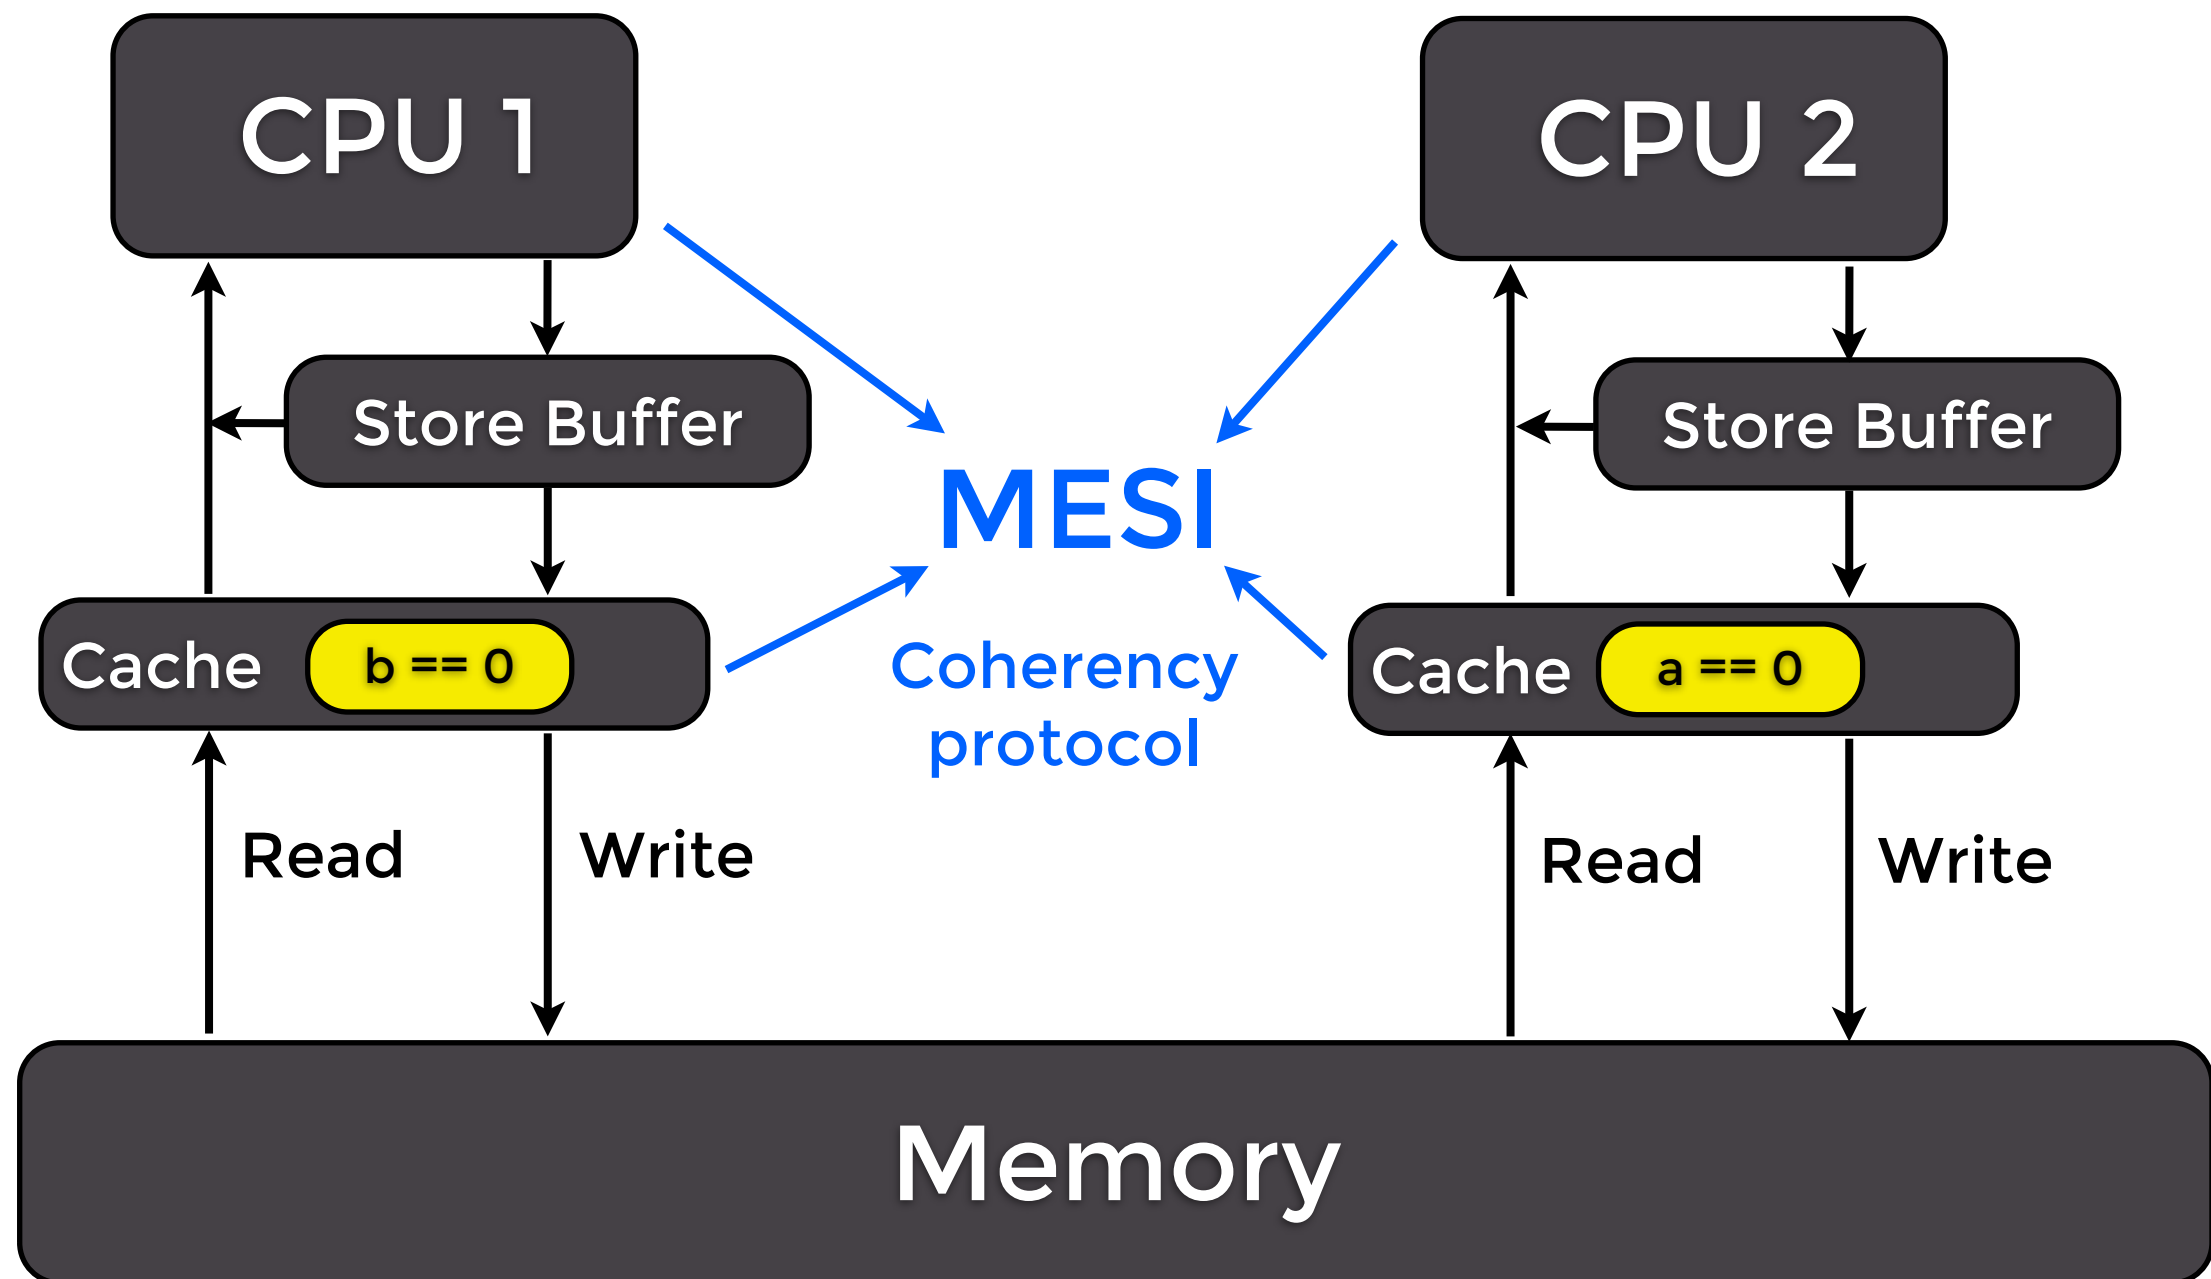
```
void m1() {  
    a = 1;  
    b = 1;  
}
```

Second CPU

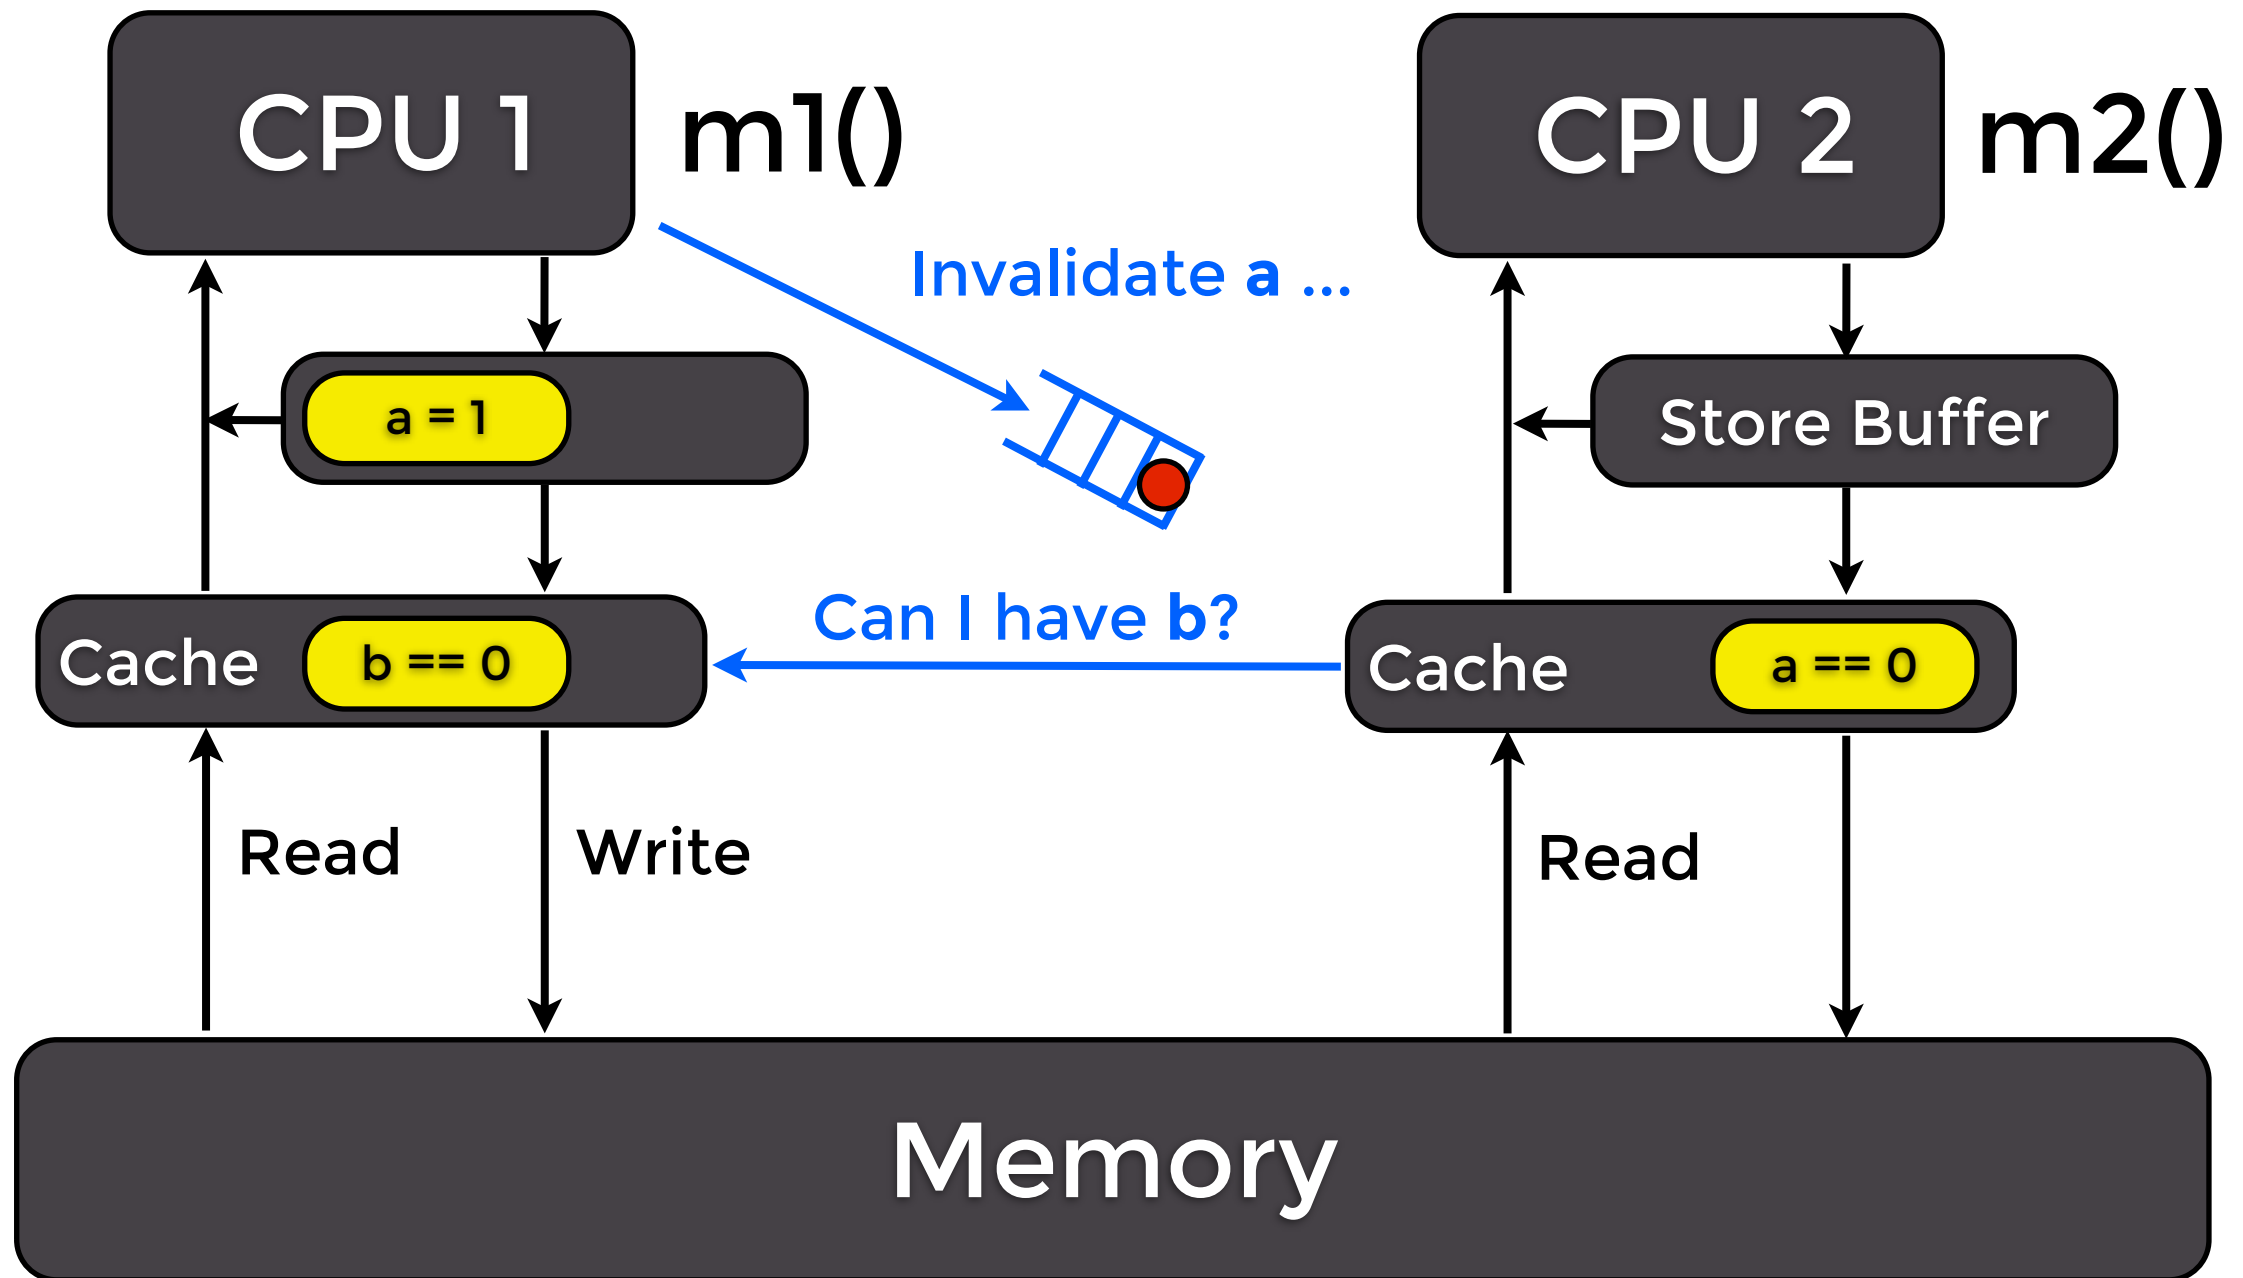
```
void m2() {  
    while (b == 0)  
        ;  
    assert(a == 1);  
}
```

Credit: <http://bit.ly/pjug2013-mckenney-parallel>, Appendix C, p. 231 et seq.

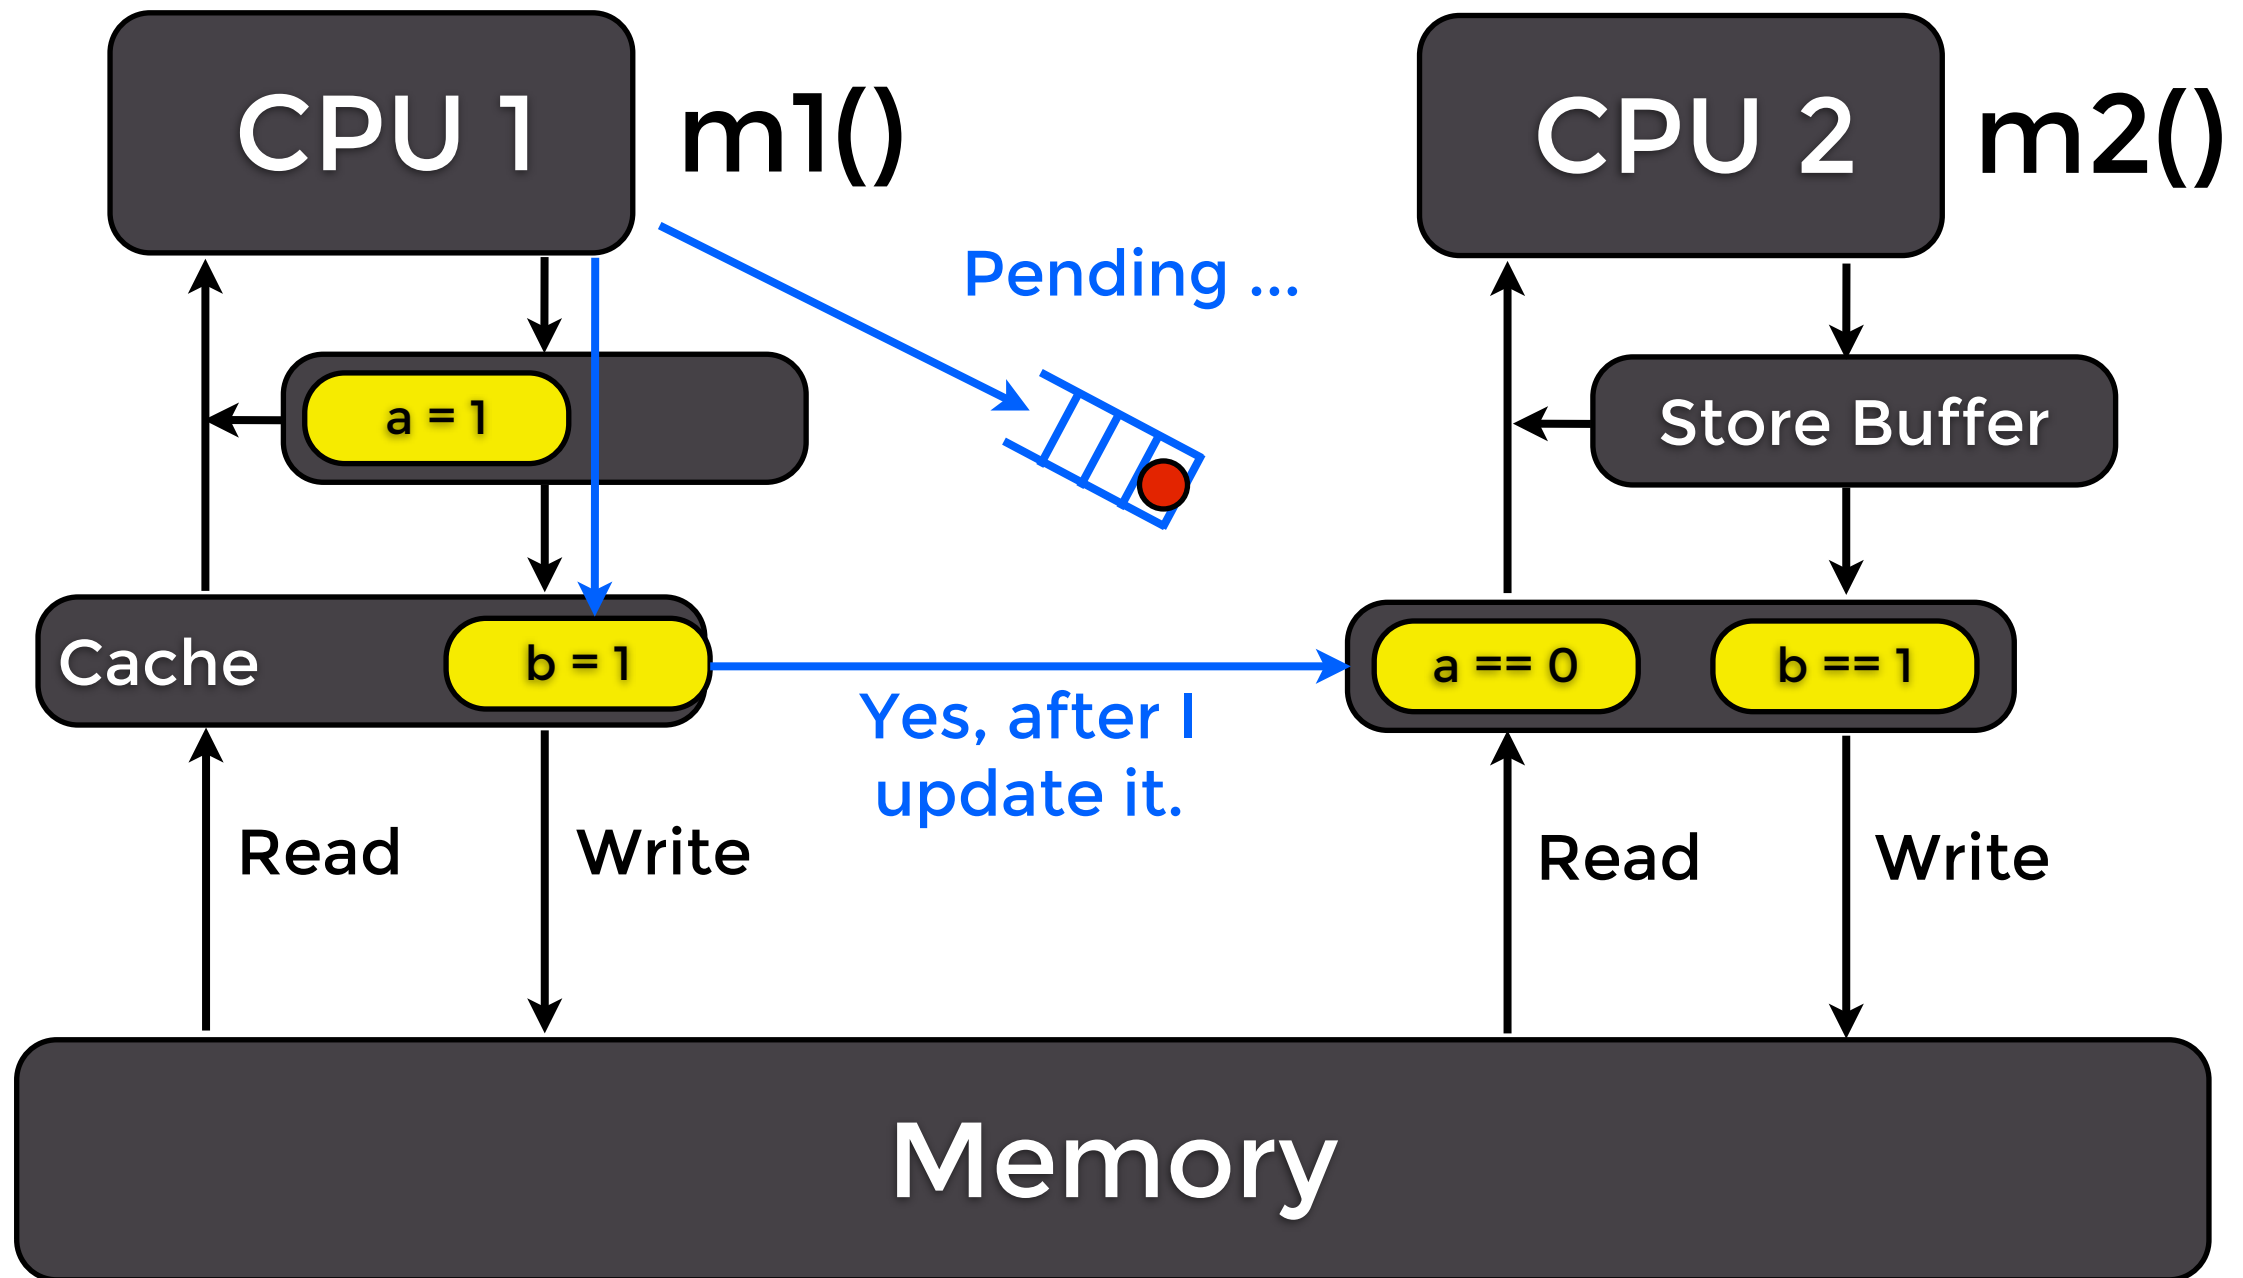
Initial State



Step 1

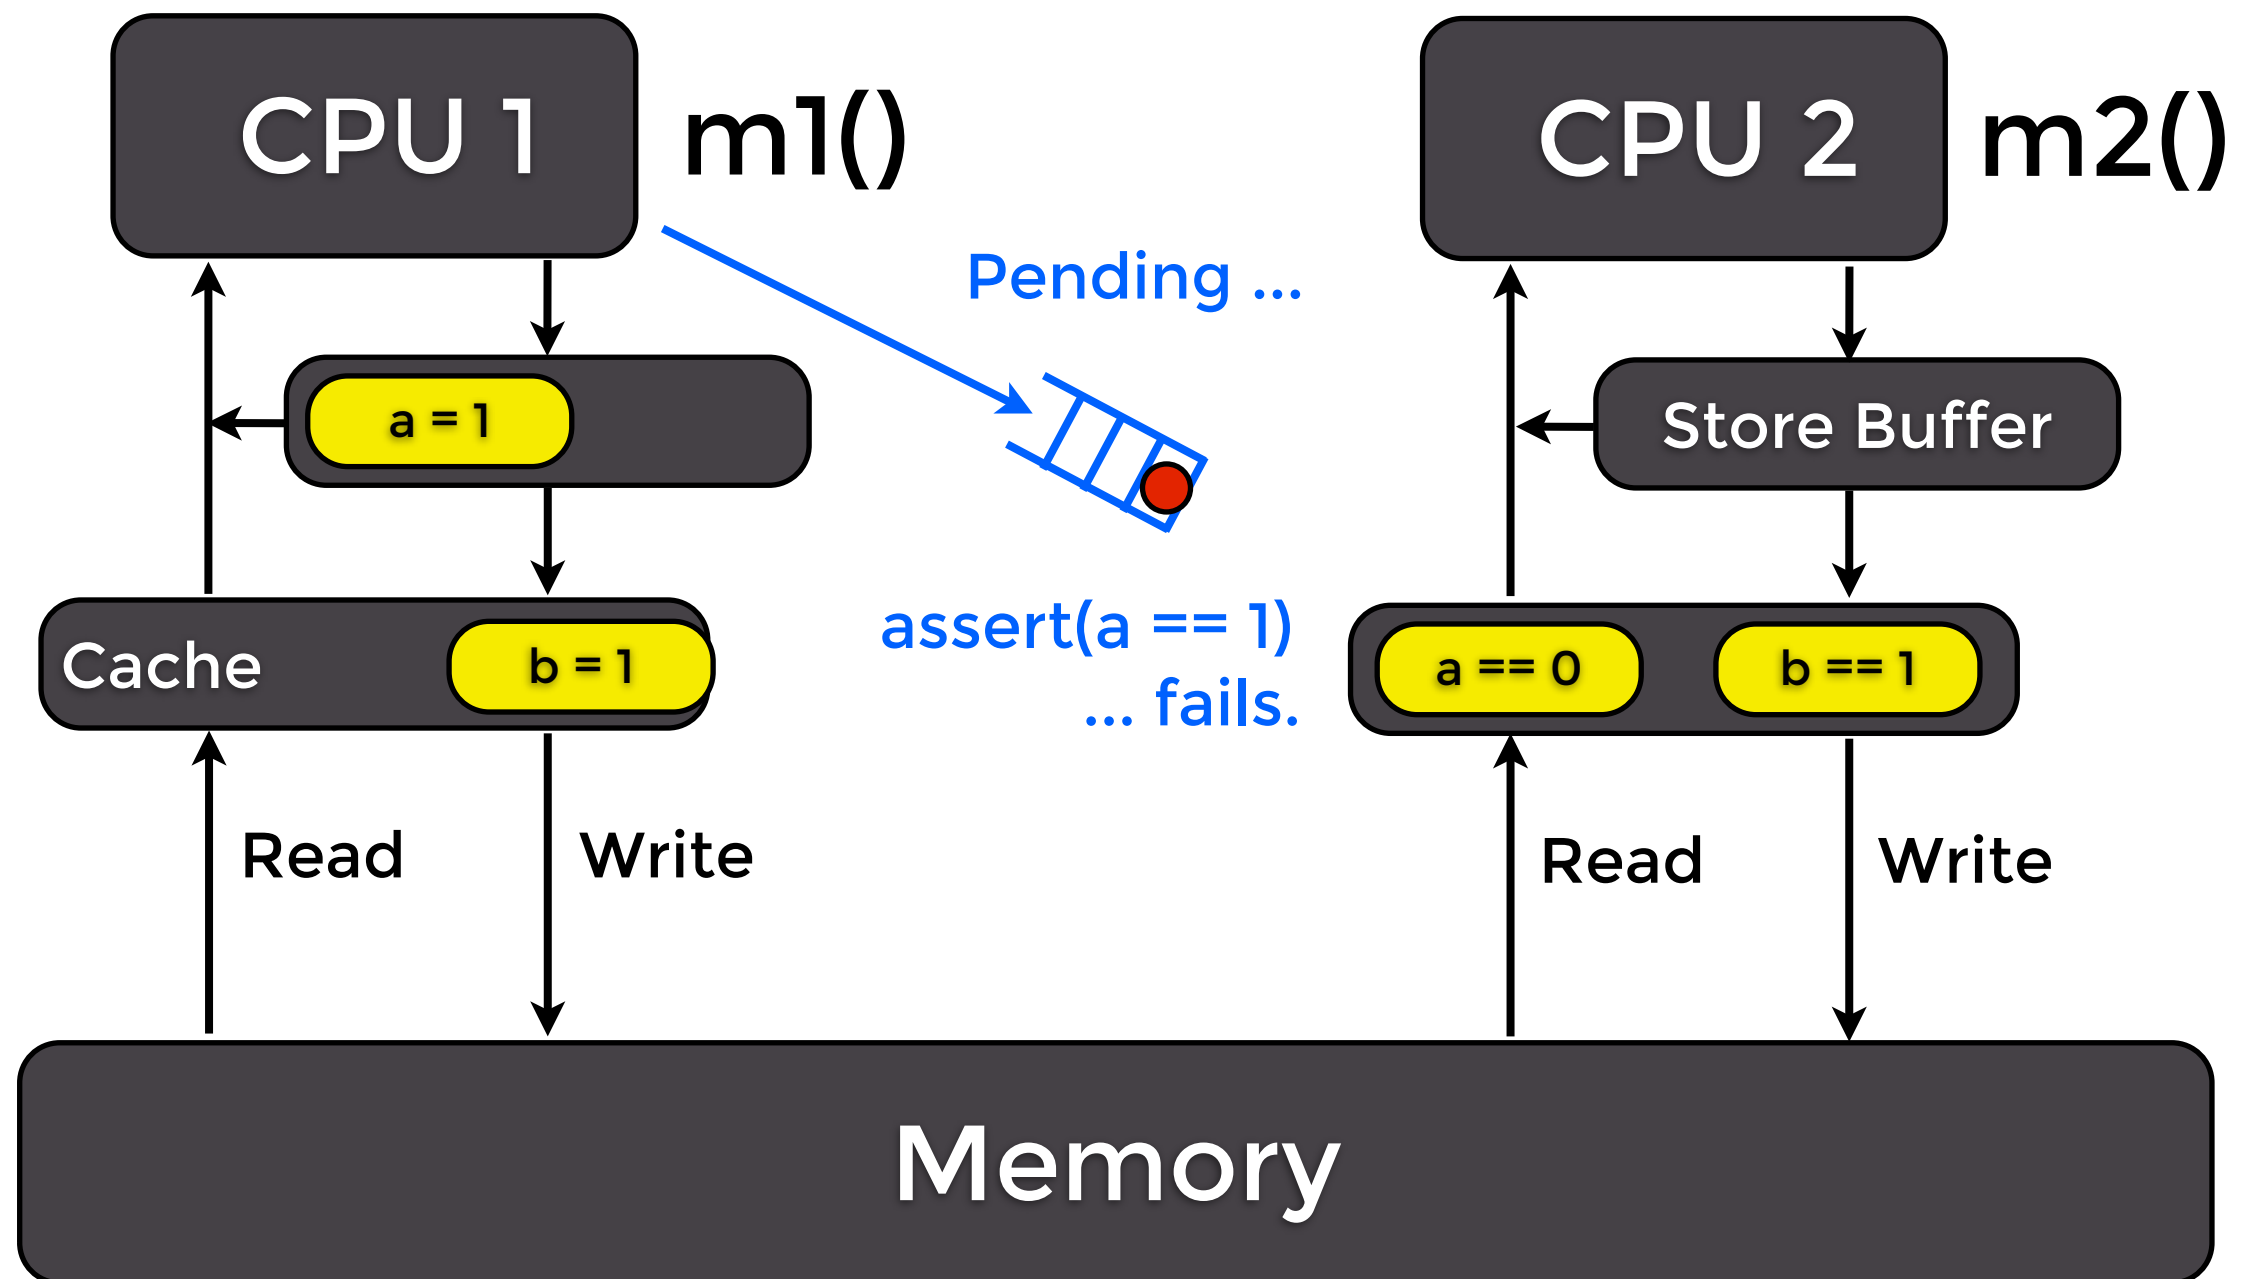


Step 2

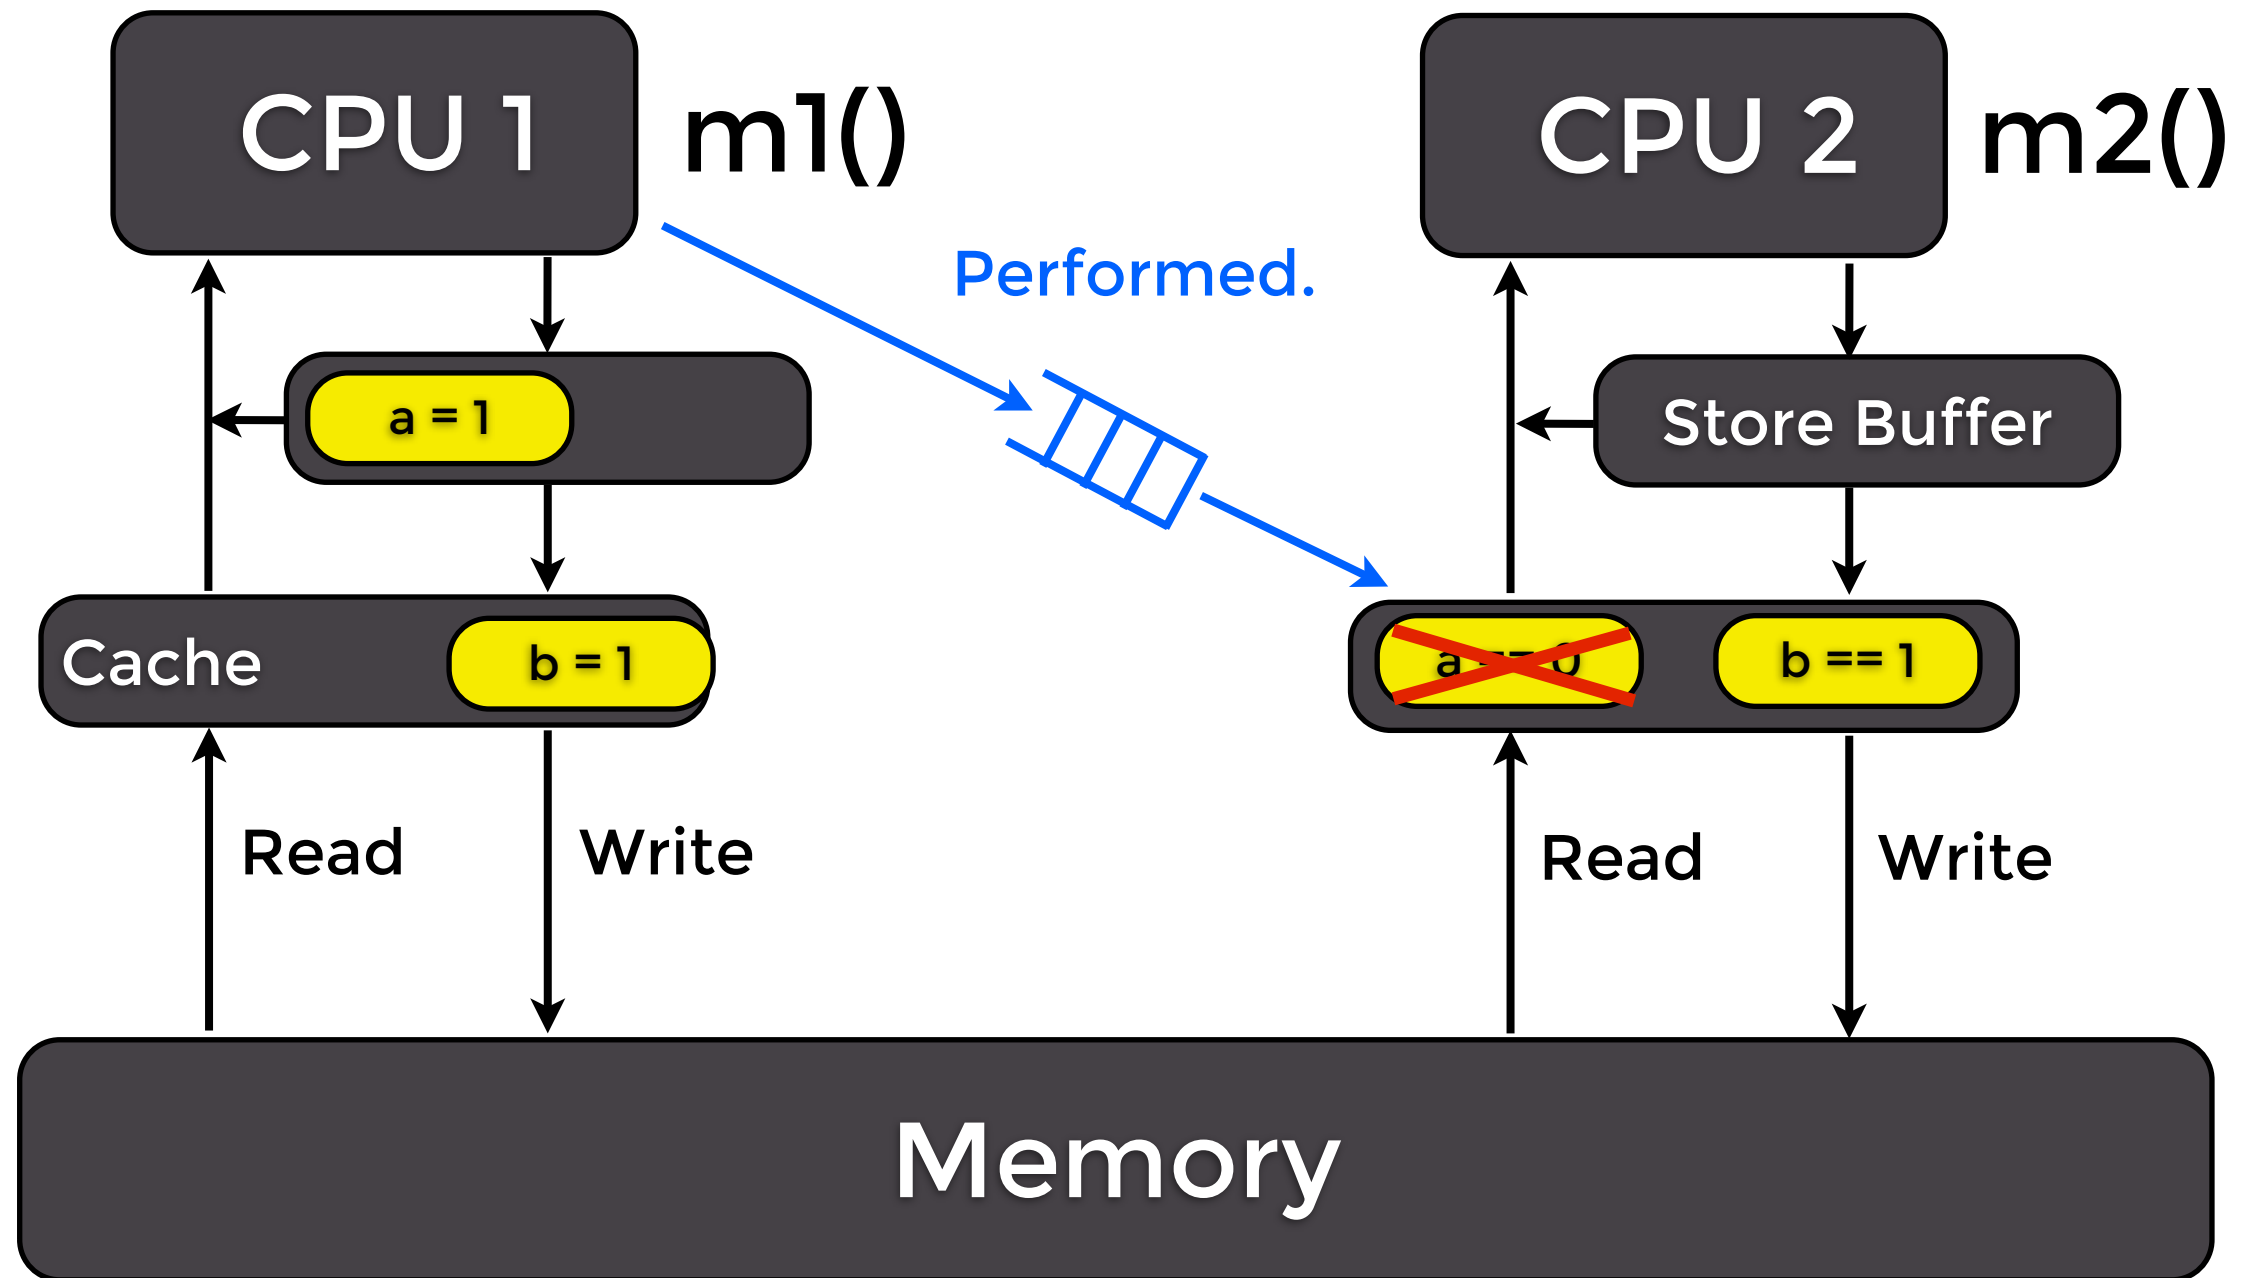


CPU 1 then writes 1 to “b” and CPU 2 tries to test it. The timing happens to work out so that the updated value of “b” is provided to CPU 2. The MESI messages cross like ships in the night – another kind of “reordering”, this time by the hardware.

Sigh



Too Late!



If the compiler doesn't get you, the hardware still can

*

Where Are We?

- The “intuitive rule of program order” is sufficient for single-threaded programs
- Applying the rule to each of multiple threads leads to surprising results
- For multiprocessors, we need better a better rule.

Happily, some of the best and brightest have thought about this.

Monday, December 30, 13

This presentation is about memory “visibility” because writes that occur first on a given CPU do not necessarily become visible in the same order on another CPU. This makes it extremely difficult to reason about our programs.

“Good” and “Bad” Traces

- Why are some traces better than others?
- “Good” traces are equivalent to **some serial execution** of the parallel steps
- A parallel trace that is equivalent to **some serial execution** is said to be **Sequentially Consistent (SC)**.

SC programs are programs we can reason about


Traces (Example 1)

Given:

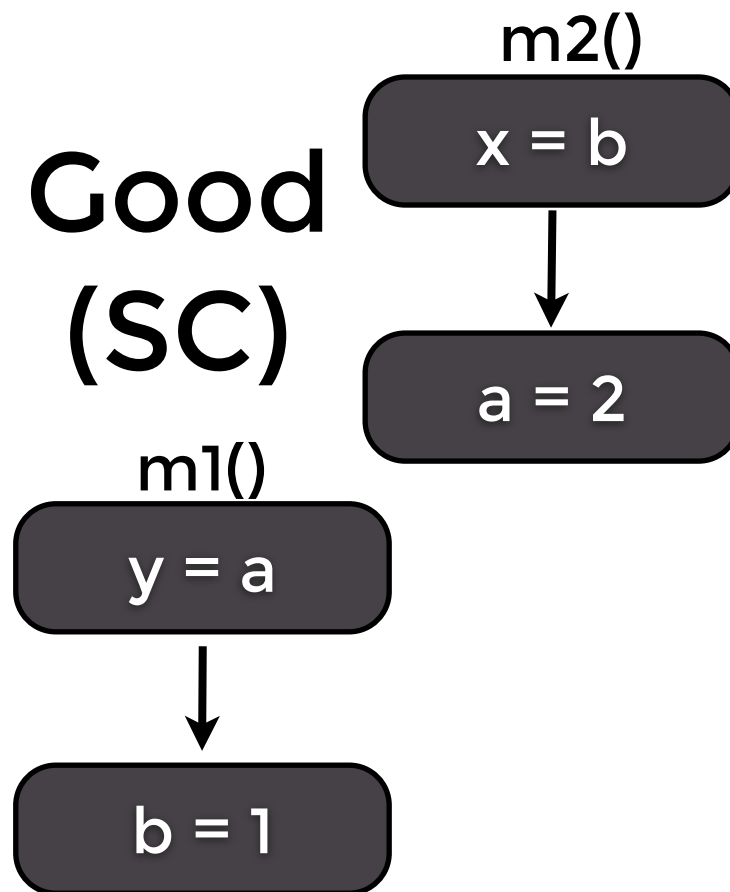
```
void m1() {  
    y = a;  
    b = 1;  
}
```

```
void m2() {  
    x = b;  
    a = 2;  
}
```


Good (SC)



Good (SC)



Bad (not SC)



Language Designer's Choices

- Could force **everything** to be SC ... but
 - Most actions by a thread are irrelevant to other threads
 - Significant performance penalty
 - Resulting language uncompetitive (?)
- Therefore ... as designed, Java requires help from the programmer to ensure SC.

Java Memory Model

- Introduced in Java 1.5 (2004)
- Section 17.4 and 17.5 of JLS
- Based on the concept of a **partial order**
 - Most memory operations are unordered
- **Happens-before** establishes ordering in specific memory operations

<http://bit.ly/PJUG2013-JLS-17>

Two Audiences

- Compiler, JVM, and class library authors
 - **Java Memory Model**
 - “Concurrency Assembly Language”
- All the rest of us
 - Class Library
 - Idioms, patterns, and new languages

*

Example Rules from JMM

“Stmt-1 happens-before stmt-2 if stmt-1 precedes stmt-2 in program order.” [Note: this is “the intuitive rule of program order”]

“All memory operations prior to writing a volatile variable on one thread happen-before a read of the same volatile from another thread.”

It's a Tufte Nightmare! (Too many words.) Sorry!

Modified Example 1

```
volatile int a, b, x, y; // all 0
```

First CPU (“thread”)

```
void m1() {  
    y = a;  
    b = 1;  
}
```

Second CPU

```
void m2() {  
    x = b;  
    a = 2;  
}
```

Another View

```
void m1() {  
    y = a;  
    b = 1;  
}
```

Happens-Before

Happens-Before

```
void m2() {  
    x = b;  
    a = 2;  
}
```


Happens-before is transitive, so if $(y = a) \text{ hb } (b = 1)$ and $(b = 1) \text{ hb } (x = b)$, then $(y = a) \text{ hb } (x = b)$.

Result

The two **happens-before** operations mean that if CPU 2 observes $b = 1$, it must also observe $y = a$.

This rules out the non-SC trace.

Cannot Occur



*

JMM Guarantee

- JMM rule: when one thread writes data another will read, a **happens-before** must separate the write and read.
- Java programs that follow the rule are sequentially consistent.
- Otherwise, the program is said to have a **data race**.
- Programmer's task is to avoid data races.

*Examples 1 and 2 contain **data races**.*

What Does Volatile Do?

- Cause javac to avoid reordering optimizations, so preventing Example 1.
- Generated bytecode does not change
- JIT sees volatile annotation on variables and generates machine-specific barrier instructions to prevent Example 2.
- Also, mutex implementation must cause execution of machine-specific barriers.

<http://bit.ly/PJUG2013-Memory-Barriers>

Monday, December 30, 13

The focus on “volatile” here is, again, just an example. Acquiring and releasing a lock has the same effect on memory visibility, so long as the “other” thread synchronizes on the same lock.

For the Rest of Us

- Idioms
 - E.g. Initialization on demand holder (<http://bit.ly/PJUG2013-Holder>)
- Patterns
 - Proper construction, publication, etc.
- New technology
 - Languages, Frameworks, Java 8, etc.

Key Patterns

- Proper Construction
- Immutability
- Safe Publication
- Concurrent Collections
- Documentation
- **Doing it yourself:** `volatile`,
`synchronized`, `etc.`

Java Concurrency in Practice: <http://jcip.net>

Proper Construction

- The closing curly brace of a constructor is a special point in program execution.
- `this` object must never be published before its constructor is complete
 - Traditional fail: event registration
 - Common remedy: static factory method

Immutability

- Immutable object = all fields final and no way to modify state of contained objects.
- Properly constructed immutable objects are thread safe
- May be passed between threads “willy nilly”

Safe Publication

- Create object in static initializer
- Hold ref in `volatile` **or** `AtomicReference`
- Hold ref in `final` field of some other properly constructed object
- Pass ref through a field guarded by a lock (e.g. a synchronized accessor)

There Are Only Four Ways To Do It.

Use the Class Library

- Learn what's in `java.util.concurrent`!
- Use concurrent collections for safe publication.
- Do not roll your own operations similar to the ones offered by atomics

Documentation

- Some widely-used libraries and frameworks handle this badly
- Describe thread safety of each class and/or method in Javadoc
- Use JSR-305 or similar annotations if they are available to you.

Doing It Yourself

- Use `volatile` (or `atomicType`) for single items of state
 - Same memory visibility guarantees
 - Use `atomic` if methods are helpful
- Intrinsic locking (`synchronized`) when multiple items must be kept consistent
- Threads must refer to the same volatile or lock.

<http://bit.ly/PJUG2013-FAQ>

*

New JVM Languages

- **Scala**
 - **Functional language on JVM**
 - **Potentially huge advantages**
 - **Scala dev team must deal with JMM**
- **Java 8 (not a new language exactly)**
 - **Closures, streams, Spliterators, etc**

<http://bit.ly/PJUG2013-Scala-Issue>

New JVM Frameworks

- Example: Akka
 - Actor (event) framework on JVM
 - Beautiful docs about memory model:
<http://bit.ly/PJUG2013-Akka-Jmm>
- Dalvik (Android “Java” virtual machine)
 - History of issues - maybe better now
 - Details: <http://bit.ly/PJUG2013-Dalvik>

Mention of Akka is an example. There are many others.

Non-JVM Languages

- C
 - You're on your own. Distinct compile and runtime tools. Example follows.
- C++
 - Developing (have?) a memory model
- C#
 - Similar to Java, but docs don't allow for a precise comparison

Explicit Control in C

- Compiler directives/annotations to prevent aggressive compiler reordering
- Linux kernel: macros expand to explicit memory barrier instructions

```
void m1(void) {  
    stmt-1;  
    stmt-2;  
    smp_mb();  
}
```

<http://bit.ly/PJUG2013-C-Linux-Example>

And More Languages

- Go
 - Memory model uses terminology and concepts from JMM
 - <http://bit.ly/PJUG2013-Go-MM>
- Rust? Objective C? GPU code?
 - Left as exercise for the reader. ;-)

Summary

- These issues affect all languages that support programming with threads
- Java community was ahead of the curve in addressing them
- Awareness wins - you may not program against the JMM, but understanding it is powerful.
- Keep learning - avoid “DIY” and use the highest level tools you can.

References

<http://bitly.com/bundles/pdxjib/2>

Contains all the “bit.ly” links
from this presentation

THANK YOU

- Java Agent team and so many others at New Relic for attending my practice talks and providing feedback

Q&A

Followed By

**Implementation of the
Asynchronous Hopped
Products Pattern**