# Memory Model

COS 597C 10/5/2010

$$a = Flag = 0$$

#### **Thread**

$$a = 26;$$

$$Flag = 1;$$

$$a = Flag = 0$$

#### **Thread**

$$a = 26;$$

$$Flag = 1;$$



Compiler Transformation

$$Flag = 1;$$



$$a = 26;$$

a = Flag = 0
Thread 1
Thread 2

a = 26;
while (Flag != 1)
{};
Flag = 1;
b = a;

#### What is the value of b after execution?

a = Flag = 0
Thread 1
Thread 2

a = 26;
while (Flag != 1)
{};
Flag = 1;
b = a; 26?

#### What is the value of b after execution?

#### What is the value of b after execution?

Compilers can reorder instructions

$$a = Flag = 0$$

#### Thread 1

#### Thread 2

$$b = a$$

Compilers can reorder instructions

Lets disable compiler reordering. How about now?

Lets disable compiler reordering. How about now?

#### Hardware out-of-order execution

$$a = Flag = 0$$

#### Thread 1

#### a = 26;

$$Flag = 1;$$

Reorder buffer of P1

#### Thread 2

$$b = a; 0!$$

#### Hardware out-of-order execution

$$a = Flag = 0$$

#### Thread 1

a = 26;

Flag = 1;

Reorder buffer of P1

Thread 2

while (Flag != 1)
{};

b = a; 0!

Things could go crazy .....

If we don't define what is a valid optimization

# What is Memory (Consistency) Model?

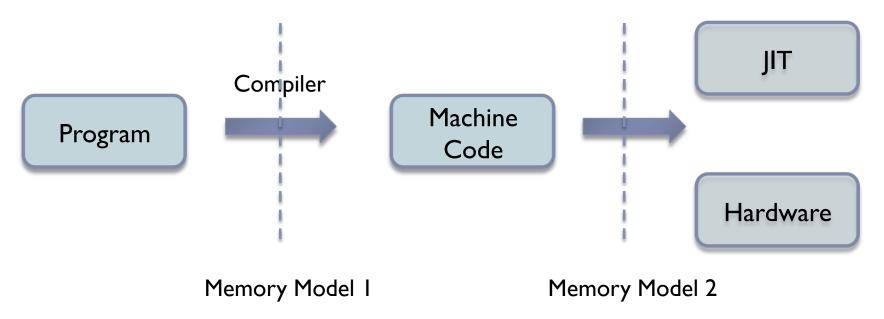
"A formal specification of how the memory system will appear to the programmer, eliminating the gap between the behavior expected by the programmer and the actual behavior supported by a system." [Adve' 1995]

#### Memory model specifies:

- How threads interact through memory
- What value a read can return
- When does a value update become visible to other threads
- What assumptions are allowed to make about memory when writing a program or applying some program optimization

# Why do We Care?

- Memory model affects:
  - Programmability
  - Performance
  - Portability



# The Single Thread Model

- ▶ Memory access executes one-at-a-time in program order
- Read returns value of last write
- For hardware & compiler reordering
  - Optimization must respect data/control dependences
  - Memory operations must follow the order the program is written
- Easy to program and optimize

# Strict Consistency Model

Any read to memory location X returns the value stored by the latest write to X

Thread 1	Thread 2
x = 1;	R1 = X; R2 = X;

R1	1
R2	1
X	1



# Strict Consistency Model

Any read to memory location X returns the value stored by the latest write to X

Thread 1	Thread 2
x = 1;	R1 = X; R2 = X;

R1	0
R2	1
X	1



# Strict Consistency Model

Any read to memory location X returns the value stored by the latest write to X

Thread 1	Thread 2
x = 1;	R1 = X; R2 = X;

R1	0
R2	1
X	1



# Sequential Consistency

- Definition: [Lamport' 1979] the result of any execution is the same as:
  - The operations of each thread appears in program order
  - Operations of all threads were executed in some sequential order atomically

### Atomicity

- Isolation : no one sees partial memory update
- Serialization : memory access appear to occur at the same time for everyone

# Under Sequential Consistency Model

- The operations of each thread appears in program order
- Operations of all threads were executed in some sequential order atomically

Thread 1	Thread 2
x = 1;	R1 = X; R2 = X;

R1	0
R2	1
X	1



# Under Sequential Consistency Model

- The operations of each thread appears in program order
- Operations of all threads were executed in some sequential order atomically

Thread 1	Thread 2
x = 1;	R1 = X; R2 = X;

R1	1
R2	0
X	1



# Dekker's algorithm for critical sections

$$Flag1 = Flag2 = 0;$$

Thread 1	Thread 2
Flag1 = 1;	Flag2 = 1;
if (Flag2 == 0) critical	<pre>if (Flag1 == 0)   critical</pre>

# Dekker's algorithm for critical sections

$$Flag1 = Flag2 = 0;$$

Thread 1	Thread 2
Flag1 = 1;	Flag2 = 1;
<pre>if (Flag2 == 0)   critical</pre>	<pre>if (Flag1 == 0)   critical</pre>

Flags1	1
Flags2	0

# Dekker's algorithm for critical sections

$$Flag1 = Flag2 = 0;$$

Thread 1	Thread 2
<pre>Flag1 = 1; if (Flag2 == 0)   critical</pre>	Flag2 = 1;  if (Flag1 == 0)  critical

Flags1	1
Flags2	1

# Dekker's algorithm for critical sections

$$Flag1 = Flag2 = 0;$$

Thread 1	Thread 2
<pre>Flag1 = 1; if (Flag2 == 0)   critical</pre>	Flag2 = 1;  if (Flag1 == 0)  critical

Flags1	0
Flags2	1

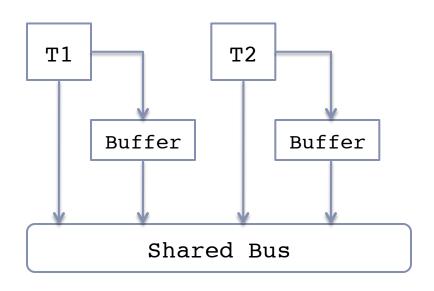
#### Violation !!!

# How do we violate sequential consistency?

# Very EASY!

Lets take a look at several hardware/ compiler optimizations that are commonly used for uniprocessor

### Write buffers with bypassing

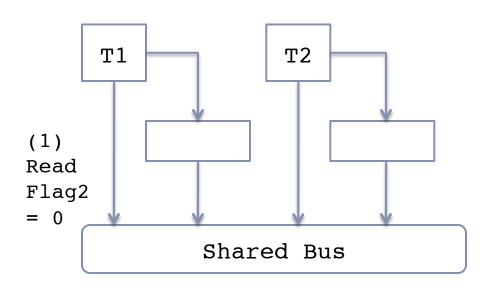


Thread 1 Thread 2

Flag1 = 1; Flag2 = 1; if (Flag2 ==0) critical critical

Flag1	0
Flag2	0

### Write buffers with bypassing

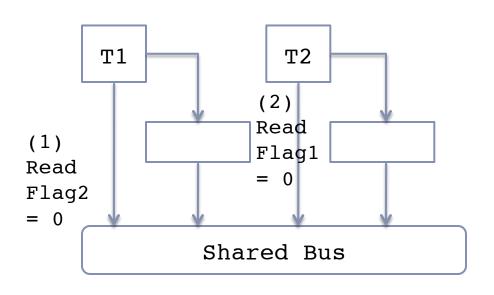


Thread 1 Thread 2

```
Flag1 = 1; Flag2 = 1;
if (Flag2 ==0) if (Flag1 ==0)
  critical critical
```

Flag1	0
Flag2	0

### Write buffers with bypassing

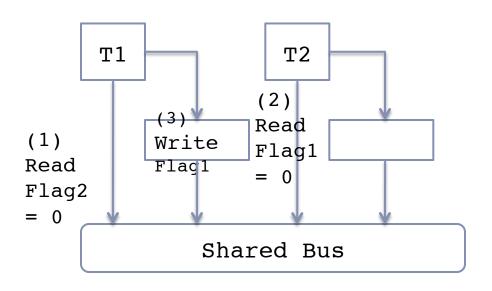


Thread 2

Thread 1

Flag1	0
Flag2	0

# Write buffers with bypassing



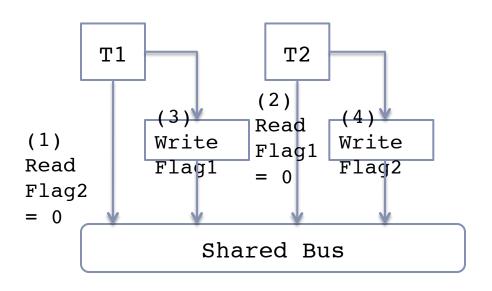
Flag1 = 1; Flag2 = 1;
if (Flag2 ==0) if (Flag1 ==0)
 critical critical

Thread 2

Thread 1

Flag1	0
Flag2	0

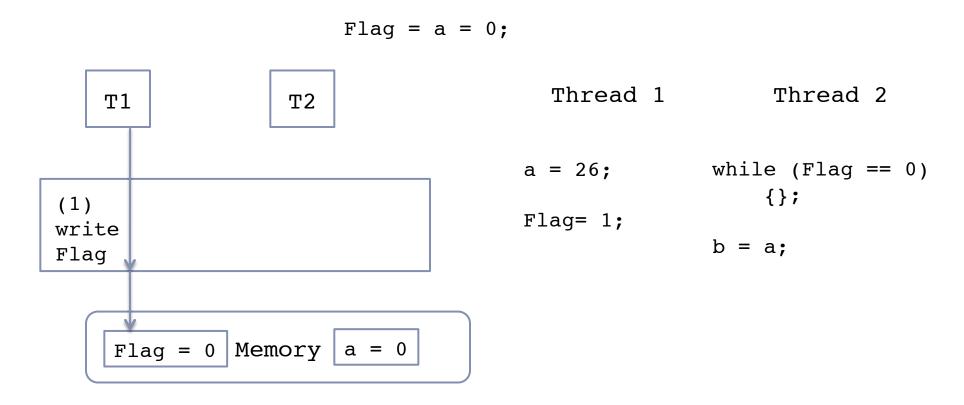
# Write buffers with bypassing

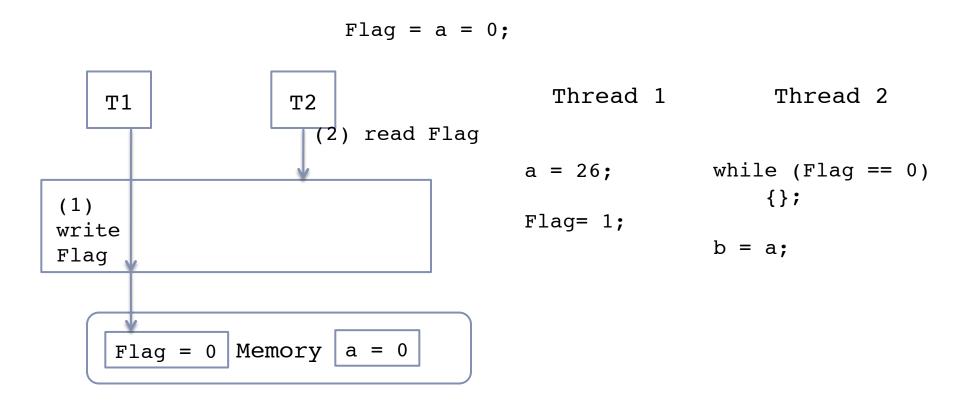


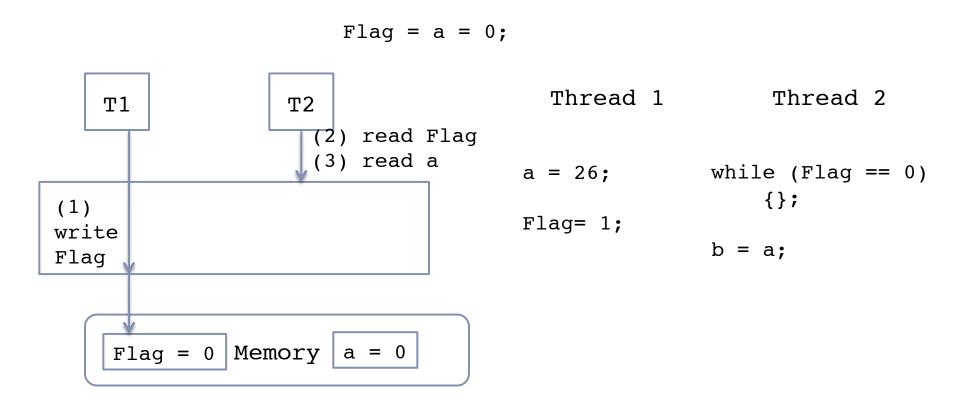
Thread 1 Thread 2

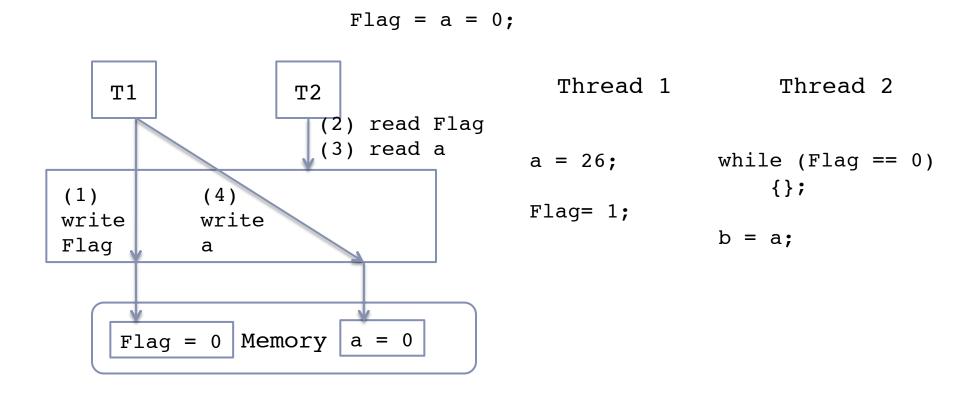
Flag1 = 1; Flag2 = 1; if (Flag2 ==0) critical critical

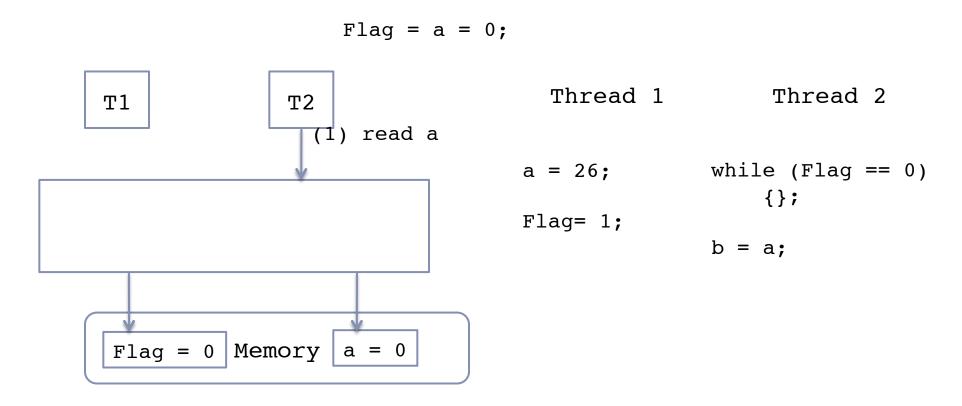
Flag1	0
Flag2	0

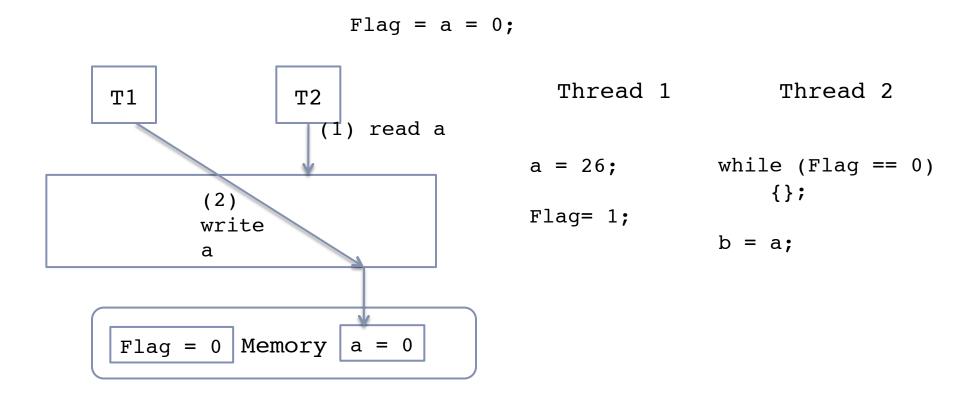


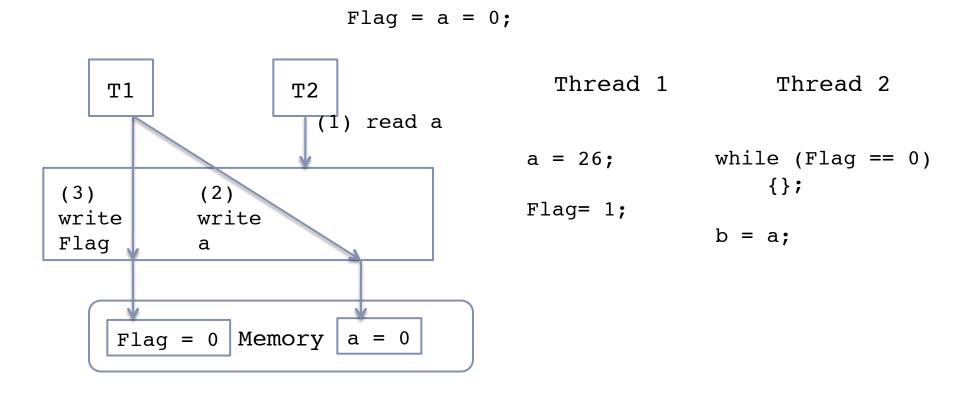


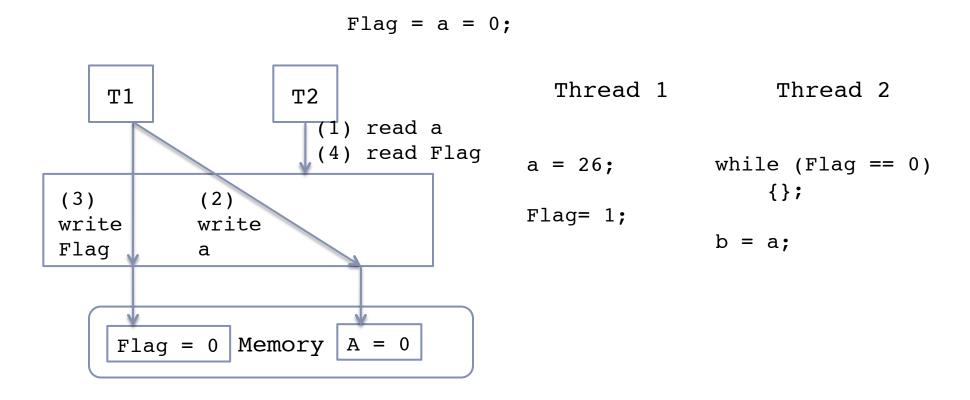












### Architecture with Private Caches

### To comply with Sequential Consistency, we need:

- Cache coherency protocol
  - A write is eventually made visible to all processors
  - Writes to the same location appear to be seen in the same order by all processors (serialization) [Gharachorloo'90]
- Ability to detect the completion of write operations
  - Acknowledgement messages
  - Invalid or update messages
- The illusion of atomic writes

## Write atomicity

$$A = B = C = 0;$$

Thread 1	Thread 2	Thread 3	Thread 4	
A = 1; B = 1;	A = 2; C = 1;	<pre>while (B!=1) {}; while (C!=1) {}; R1 = A;</pre>	<pre>while (B!=1) {}; while (C!=1) {}; R2 = A;</pre>	

What is the value of R1 and R2 after execution?

M M I M COC FOTO F 11 20 10

# Write Atomicity

$$A = B = C = 0;$$

Thread 1	Thread 2	Thread 3	Thread 4
A = 1; B = 1;	A = 2; C = 1;	<pre>while (B!=1) {}; while (C!=1) {}; R1 = A;</pre>	<pre>while (B!=1) {}; while (C!=1) {}; R2 = A;</pre>

$$R1 = 1$$

$$R2 = 1$$



# Write Atomicity

$$A = B = C = 0;$$

Thread 1	Thread 2	Thread 3	Thread 4
A = 1; B = 1;	A = 2; C = 1;	<pre>while (B!=1) {}; while (C!=1) {}; R1 = A;</pre>	<pre>while (B!=1) {}; while (C!=1) {}; R2 = A;</pre>

$$R1 = 2$$

$$R2 = 2$$



## Write Atomicity

$$A = B = C = 0;$$

Thread 1	Thread 2	Thread 3	Thread 4
A = 1; B = 1;	A = 2; C = 1;	<pre>while (B!=1) {}; while (C!=1) {}; R1 = A;</pre>	<pre>while (B!=1) {}; while (C!=1) {}; R2 = A;</pre>

R1 = 1

R2 = 2

#### Violation !!!

Sequential Consistency: operation from all threads must appear In some sequential order atomically



## Compilers Optimization that Violates SC

- Compiler reordering must respect data and control dependencies
- Code motion

Load from a cannot be moved out of the loop

## Compilers Optimization that Violates SC

- Compiler reordering must respect data and control dependencies
- Code motion
- Common sub-expression elimination

```
Thread 1 Thread 2 c = a - 1; \qquad (a-1) \text{ cannot be eliminated for assignment of b} \text{Thread 2} \text{Thread 2} \text{Thread 2} \text{Thread 2} \text{Thread 2} \text{Thread 2} \text{C = a - 1;} \text{while (Flag == 0) {};} \text{assignment of b} \text{Description of a b} \text{Thread 2} \text{Thread 2} \text{Thread 2} \text{C = a - 1;} \text{Thread 2} \text{Thread 2} \text{C = a - 1;} \text{Thread 3} \text{C = a - 1;} \text{Thread 3} \text{Thread 2} \text{C = a - 1;} \text{Thread 3} \text{C = a - 1;} \text{Thread 3} \text{C = a - 1;} \text{Thread 3} \text{Thread 3} \text{Thread 3} \text{Thread 3} \text{C = a - 1;} \text{Thread 3} \text{Thread
```

## Compilers Optimization that Violates SC

- Compiler reordering must respect data and control dependencies
- Code motion
- Common sub-expression elimination
- Register allocation

```
Thread 1 Thread 2 Flag \text{ cannot be allocated to a register} a = 6; while (Flag == 0) \{\}; register b = a;
```

## Sequential Consistency: Summary

Sequential consistency does not guarantee data race free

Thread 1	Thread 2
a _ 1.	n - 0.
A = 1;	A = 2;
B = 1;	C = 1;

#### Data race:

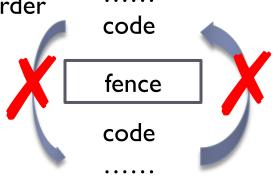
- two memory access to the same location
- one is a write
- they can occur simultaneously
- Possible hardware/compiler optimizations allowed
  - Hardware/software prefetching
  - Speculating read values
- Determining which instructions are allowed to be reordered remain an open question

## Relaxed Memory Models

- Key points
  - Program order for different memory addresses
  - Write atomicity
- Possible relaxations
  - Relaxation on program order (different memory locations)
    - ▶ Relax write to read program order
    - ▶ Relax write to write program order

Relax read to read and read to write program order

- Relaxation on write atomicity
  - Read other's write early
  - Read own write early
- Safety nets, such as fence



# Major Relaxed Hardware Models

Relax	W->R	W->W	R->RW	Read others' write early	Read own write early	Safety Net
SC					<b>✓</b>	
IBM 370	<b>✓</b>					Serial inst
TSO(x86)	<b>✓</b>				<b>✓</b>	RMW, fence
PC	<b>✓</b>			<b>✓</b>	<b>✓</b>	RMW
PSO	<b>✓</b>	<b>✓</b>			<b>✓</b>	RMW
WO	<b>✓</b>	<b>✓</b>	<b>✓</b>		<b>✓</b>	synch
RCsc	<b>✓</b>	<b>✓</b>	<b>✓</b>		<b>✓</b>	lock, nsync,
RCpc	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	RMW
Alpha	<b>✓</b>	<b>✓</b>	<b>✓</b>		<b>✓</b>	MB, WMB
RMO	✓	<b>✓</b>	<b>✓</b>		<b>✓</b>	MEMBAR
PowerPC	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	synch

## **Processor Consistency**

- Writes done by a single processor are received by other processors in the same order as they are issued.
- Writes from different processors may be seen in different order by different processors.

$$A = B = C = 0;$$

Thread 1	Thread 2	Thread 3	Thread 4
A = 1; B = 1;	A = 2; C = 1;	<pre>while (B!=1) {}; while (C!=1) {}; R1 = A;</pre>	<pre>while (B!=1) {}; while (C!=1) {}; R2 = A;</pre>

$$R1 = 1$$

R2 = 2



## Weak Ordering Model [Dubois' 86]

### Classification of memory operations

- Data operations: load, store...
- Synchronization operations: lock unlock etc

### How does it work?

- All pre-issued operations must complete on all processors before executing a synchronization operation
- Execution of synchronization operations must follow program order
- Memory operations between synchronization operations can be reordered

### Data-Race-Free-0 Model

- A program is data-race-free on a particular input if no sequential consistent execution results in a data race
- A new definition of weak ordering [Adve'90 ISCA]
- Advantage:
  - Simple programmability of sequential consistency
  - Implementation flexibility of relaxed models
- Sequential consistency for DRF is widely used
  - C++ memory model

## Relaxed Memory Model: Summary

- Relaxed memory model
  - Relaxes restrains on the order of some memory operations
  - Allows some hardware/compiler optimization
- Why do we use relaxed memory model : performance
- Why do we not use relaxed memory model : complexity

- Adaption of DRF0
  - Sequential consistency for race-free programs
  - Behavior of a program with data race is undefined (no benign data races in C++)
- Data operations:

load, store

- Synchronization operations:
   lock, unlock, atomic load, atomic store, atomic read-modify-write
- Atomic operations must appear sequentially consistent

[Boehm's PLDI 2008: Foundations of the C++ Concurrency Memory Model]

Compiler code reordering allowed when:

For memory operations MI and M2

- MI is a data operation and M2 is a read synchronization operation
- ▶ MI is write synchronization and M2 is data
- MI and M2 are both data with no synchronization sequenceordered between them.
- MI is data and M2 is the write of a lock operation
- ▶ MI is unlock and M2 is either a read or write of a lock.
- Hardware optimization allowed for non-atomic writes

### Semantic of trylock

Can the program assert?

### Semantic of trylock

Can the program assert?

### Semantic of trylock

### Semantic of trylock

We can use a fence, but it is unfair for properly used trylock

### Semantic of trylock

Solution: in C++ memory model, trylock does not guarantee to reveal anything about the state of the lock

## CASE STUDY: The JAVA Memory Model

- JAVA: the first language specification attempts to incorporate memory model
- What JAVA should do?
  - Define semantics of all programs
  - Support execution of untrusted "sandboxed" code
- Sequential consistency for DRF
- Synchronization implemented using monitors
  - Volatile
  - Synchronized primitive
- JAVA memory model does not guarantee deadlock free

## CASE STUDY: The JAVA Memory Model

- JAVE bugs found historically
  - Detached thread
  - Double-checked locking

```
Helper helper;
Helper getHelper() {
 if (helper==null) {
   synchronized(this) {
       if (help==null)
         helper=new Helper();
  return helper;
```

## Lessons Learnt from C++/JAVA

- SC for DRF is the minimal baseline
- Specifying semantics for programs with data races is extremely HARD
- Simple optimization may introduce unintended consequences
- State-of-the-art is still broken
  - Abandon shared memory?
  - Hardware co-designed with high-level memory models?
  - Any volunteer for fixing the whole thing?

### Conclusion

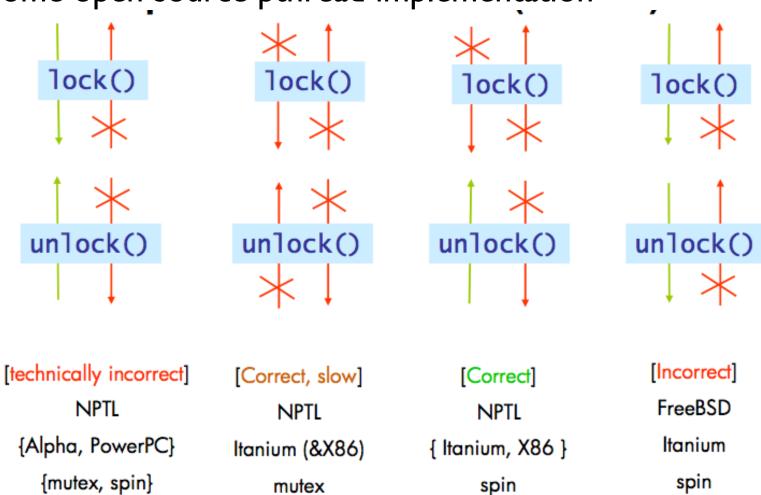
- Memory model is very important and confusing
- Memory model specifies what hardware/compiler can do and cannot do
- Sequential consistency is very intuitive yet prohibits performance
- Relaxed memory models allows some optimization but introduces programming complexity
- Don't try to be Clever, unless you are Clever enough

## **Advanced Topics**

- Why threads cannot be implemented as a library
- Ongoing projects:
  - Deterministic Parallel JAVA (DPJ)
  - Functional languages
  - DeNoVo hardware project

### Pthreads

Some open source pthread implementation



### References

- Boehm's "Foundations of the C++ Concurrency Memory Model"
- Pugh's "Fixing the JAVA Memory Model"
- Adve's "Shared Memory Consistency Models: A Tutorial"
- Dubois' "Memory Access Buffering in Multiprocessors"
- ▶ Bohem's "Threads cannot be implemented as a library"