

THE JAVA MEMORY MODEL

J. Manson, W. Pugh, and S. Adve

Brian Gemhardt - Feb 21, 2013
Parallel and Distributed Systems

INTRODUCTION

MEMORY MODEL

Limits on order of memory accesses

Programmer

Compiler
& Hardware

Ease of Use

Optimizations

SEQUENTIAL CONSISTENCY

$x = y = 0$

- Accesses in program order
- Easy to use
- No reordering
- Few optimizations

$r2 = x$ $r1 = y$
 $y = 1$ $x = 2$

r1	0	1	0	1
r2	0	0	2	2

CORRECTLY SYNCHRONIZED

- No data races in *any* sequentially consistent execution

r2 = x r1 = y

- Example always has write next to read

y = 1 x = 2

- Always has a race

DATA RACE FREE

- If you mark synchronization

```
volatile int x, y;
```

- **volatile** and **synchronized**

```
synchronized { synchronized {
```

```
  r2 = x      r1 = y
```

```
  y = 1      x = 2
```

```
}                    }
```

- Then you get Sequentially Consistent behavior

DEFINITIONS

- Lock

```
lock(a)      lock(a)
```

- Unlock

```
  r2 = x      r2 = y
```

- Read volatile

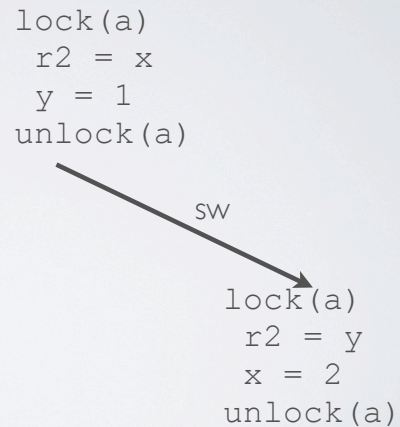
```
  y = 1      x = 2
```

- Write volatile

```
unlock(a)    unlock(a)
```

SYNCHRONIZES WITH

- Order synchronization actions
- Unlocks before locks
- Volatile writes before reads
- Between threads

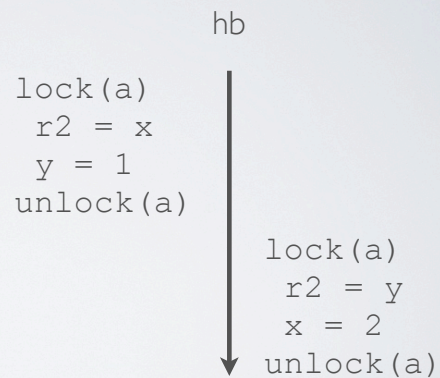


PROGRAM ORDER


- Order of accesses in code
- Only within same thread

HAPPENS-BEFORE

- Synchronizes-with combined with program order
- Transitive Closure
 - If $a \rightarrow b \rightarrow c$
 - Then $a \rightarrow c$



HAPPENS-BEFORE



CONFLICTING ACCESS

- Two or more accesses to same variable

$r1 = x$

$r2 = x$

$x = r3$

- At least one is a write

DATA RACE

- Different threads

- Conflicting access

$r2 = x$ $r1 = y$

$y = 1$ $x = 2$

- No happens-before order

HAPPENS-BEFORE MODEL

HAPPENS-BEFORE MODEL

- Each thread acts as if it was alone
 - Given that it reads the same values
- Non-volatile reads must be happens-before consistent
- Volatile reads must be synchronization order consistent

HAPPENS-BEFORE CONSISTENCY

- Determines when r can read the value from w
 - $r \rightarrow w$ r can't see w
 - $w \rightarrow w' \rightarrow r$ If r sees w , then w' doesn't exist

SYNCHRONIZATION ORDER CONSISTENCY

- Synchronization order is consistent with program order
 - Can't reorder unlocks before locks, etc.
- Volatile read sees the previous write in synchronization order

```
x = 0  
volatile ready = false
```

```
↓ x = 1  
ready = true → if (ready)  
↓ r1 = x
```

```
r1 = x never happens  
or  
r1 == x == 1
```

CAUSALITY

UNDEFINED BEHAVIOR

- Means anything can happen

```
x = y = 0
```

- Hard to understand

```
r2 = x    r1 = y
y = 1    x = 2
```

- Bad for security

- 42 isn't bad

```
r1 = r2 = 42
```

- Password is

CORRECTLY SYNCHRONIZED

```
r1 = x
if (r1 != 0)
    y = 42
```

- Sequentially consistent

- Reads r1 before setting y

- Reads r2 before setting x

- No data race

```
r2 = y
if (r2 != 0)
    x = 42
```

CIRCULAR LOGIC

```
r1 = x
if (r1 != 0)
    y = 42
```

- No happens before order between threads

- Means reads are allowed to see writes

```
r2 = y
if (r2 != 0)
    x = 42
```

- Because reads see writes, writes can occur

- Don't want this

CIRCULAR LOGIC?

```
r1 = a
r2 = a
if (r1 == r2)
    b = 2
r3 = b
a = r3
```

Can `r1 == r2 == r3 == 2`?

OPTIMIZATIONS

Duplicate Read

```
r1 = a
r2 = a
if (r1 == r2)
    b = 2
r3 = b
a = r3
```

OPTIMIZATIONS

```
r1 = a
r2 = r1
if (r1 == r2)
    b = 2
r3 = b
a = r3
```

Always true

OPTIMIZATIONS

```
r1 = a
r2 = r1
if (true);
    b = 2
r3 = b
a = r3
```

OPTIMIZATIONS

```
b = 2
r1 = a
r2 = r1
if (true);
r3 = b
a = r3
```

Now `r1 == r2 == r3 == 2`

WHAT'S THE DIFFERENCE?

<pre>r1 = x if (r1 != 0) y = 42</pre> <hr/>	<pre>r1 = a r2 = a if (r1 == r2) b = 2</pre> <hr/>
<pre>r2 = y if (r2 != 0) x = 42</pre>	<pre>r3 = b a = r3</pre>

With Sequential Consistency:

Writes never happen Write always happens

JAVA MEMORY MODEL

WELL-BEHAVED EXECUTION

- Justifies reordering by iteratively building execution
 - Early execution allowed if not caused by data race
- Maintain consistencies from happens-before model
 - Sequential consistency for proper synchronization
- Well-behaved for races

COMMITTING ACTIONS

- Can commit any reads or writes
- Reads can only see values from committed writes
- Before committing, reads only see happens-before writes
 - Use these values to justify committing more actions
- After committing, read can see any write

JUSTIFYING EARLY WRITE

```
a = b = 0
```

```
r1 = a  
r2 = a  
if (r1 == r2)  
  b = 2
```

- Start with no committed actions

```
r3 = b  
a = r3
```

JUSTIFYING EARLY WRITE

```
a = b = 0
```

```
r1 = a  
r2 = a  
if (r1 == r2)  
  b = 2
```

- Can commit initialization because it happens before everything else
- Can commit all the reads and write of a

```
r3 = b  
a = r3
```

JUSTIFYING EARLY WRITE

```
a = b = 0
```

```
0 r1 = a  
0 r2 = a  
if (r1 == r2)  
  b = 2
```

- Must assume reads see the happens-before writes (initialization)

```
0 r3 = b  
a = r3
```

JUSTIFYING EARLY WRITE

```
a = b = 0
```

```
0 r1 = a  
0 r2 = a  
true if (r1 == r2)  
  b = 2
```

- These reads mean that if condition is true

```
0 r3 = b  
a = r3
```

JUSTIFYING EARLY WRITE

```
a = b = 0
```

```
0 r1 = a
0 r2 = a
true if (r1 == r2)
    b = 2
```

- We can now also commit writing b

```
0 r3 = b
a = r3
```

COMMIT SEES RACE

```
a = b = 0
```

```
r1 = a
r2 = a
if (r1 == r2)
    b = 2
```

- Committed reads can see writes committed at same time.

```
r3 = b
a = r3
```

COMMIT SEES RACE

```
a = b = 0
```

```
r1 = a
r2 = a
if (r1 == r2)
    b = 2
```

- b = 2 committed
- So r3 = b can see 2

```
2 r3 = b
a = r3
```

COMMIT SEES RACE

```
a = b = 0
```

```
r1 = a
r2 = a
if (r1 == r2)
    b = 2
```

- Write of a uses value seen by read into r3

```
2 r3 = b
2 a = r3
```

COMMIT SEES RACE

```
a = b = 0
```

```
2 r1 = a
2 r2 = a
  if (r1 == r2)
    b = 2
```

- Other reads can see committed write of a

```
2 r3 = b
2 a = r3
```

COMMIT SEES RACE

```
a = b = 0
```

```
2 r1 = a
2 r2 = a
  if (r1 == r2)
    b = 2
```

- r1 == r2 == r3 == 2

```
2 r3 = b
2 a = r3
```

IMPLEMENTERS

- Weak control dependence
 - Compiler cannot assume loops will end.
 - Can not move statements ahead of loops
- Most reordering is allowed, within limits
- Can ignore thread-local synchronization

PROGRAMMERS

- Correctly synchronized code acts sequentially consistent
- JVM does not guarantee pre-emption or fairness
 - Due to complications with priorities and real-time
- Java Memory Model created via community discussion
 - Some fine details are mostly matters of taste