

JCStress Workshop

or, «One Awful Java Concurrency Test After Another»

Aleksey Shipilëv

shade@redhat.com

@shipilev

Safe Harbor / Тихая Гавань

Anything on this or any subsequent slides may be a lie. Do not base your decisions on this talk. If you do, ask for professional help.

Всё что угодно на этом слайде, как и на всех следующих, может быть враньём. Не принимайте решений на основании этого доклада. Если всё-таки решите принять, то наймите профессионалов.

Workshop Plan (Provisional)

■ Part I: Basic Bits

- JCSstress: Why do it? How does it look?
- JMM: Looking at basic examples

(Coffee Break)

■ Part II: Advanced Bits

- JCSstress: How does it work? Why would you not do it manually?
- JMM: Looking at advanced examples

(Breathing Exercises)

■ Part III: Fun Bits *(optional)*

- JCSstress: Real JVM/JDK bugs discovered
- Breakout: Discussions, Future Work, etc.

Workshop Resources

Most code is available as runnable JCTestress Samples:

```
# Setup -----  
git clone https://github.com/openjdk/jcstress/  
cd jcstress  
export JAVA_HOME=<path-to-jdk-11>  
export PATH=$JAVA_HOME/bin:$PATH  
  
# Build -----  
mvn clean install -DskipTests -T 1C  
  
# Run -----  
java -jar jcstress-samples/target/jcstress.jar ...
```

Workshop Involvement

- JcStress Samples have «How to run this test» section with useful one-liners. `-h` shows some of the runner options that might be useful.
- You **could** run the samples during the workshop. However, you **are not required** to do so. Note interesting cases for yourself, and then run them later.
- If you have questions, ask them in Telegram chat for the talk. I can then modify the samples on the fly, or refer to other samples.

Part I. Basic Bits

JCStress Basics

JCStress Basics: Historical Context

- Circa 2013 (\approx JDK 8), we suddenly realized there are no regular concurrency tests that ask hard questions about JMM conformance
- Attempts to contribute JMM tests to JCK were futile: probabilistic tests
- JVMs are notoriously awkward to test: many platforms, many compilers, many compilation and runtime modes, dependence on runtime profile

JCStress Basics: Scoping

JCStress is `#{value}` testing framework

Where `#{value}` is:







- *Java*: targets low-level JVM work, but touches HW too
- *empirical*: distrust lower layers (JVM, OS, HW) are sane
- *combinatorial*: tests many configurations of the same test
- *experimental*: fluid implementation to fit new techniques

JCStress Basics: Empirical, Not Model Checking

If you want model checking, go for Lincheck workshop:

Day 1. June 15



Time UTC+03:00	Lecture
18:45 Track 2	Workshop. Lincheck: Testing concurrency on the JVM  Maria Sokolova <i>JetBrains</i> <i>#concurrency #jvm #model-checking #testing</i>  
20:30 Track 2	Workshop. Lincheck: Testing concurrency on the JVM (part 2)  Maria Sokolova <i>JetBrains</i> <i>#concurrency #jvm #model-checking #testing</i>  

JCStress Basics: Prior Art

1. Java Compatibility Kit (JCK)

- Developed by Oracle, JCP
- Tests Java Language Specification, Chapter §17
- Limited to normative clauses

2. JSR166 TCK

- Developed by Doug Lea et al.
- Tests `java.util.concurrent.*`

3. Litmus/DIY

- Developed by Peter Sewell et al.¹
- Tests hardware semantics

¹<http://www.cl.cam.ac.uk/~pes20/>

Early Attempts: Concurrency Testing

Concurrency bugs are (data) race (condition) bugs

Need to create a *controllable race condition*:

- large enough, so that threads meet
- small enough, so that we can trust the results
- fast enough, so that timings are not masked

Unfortunately, naive tests do not check all these boxes...

Early Attempts: Try #1

```
volatile int v;  
  
void doTest() {  
    Thread t1 = new Thread(() -> v++);  
    Thread t2 = new Thread(() -> v++);  
    t1.start(); t2.start();  
    t1.join(); t2.join();  
  
    Assert.assertTrue(2, v);  
}
```

«Collision Window» is far too small to capture anything interesting.

Early Attempts: Try #2

```
volatile int v;  
final CountDownLatch l = new CountDownLatch(2);  
  
void doTest() {  
    Thread t1 = new Thread(() -> l.countDown(); l.await(); v++);  
    Thread t2 = new Thread(() -> l.countDown(); l.await(); v++);  
    t1.start(); t2.start();  
    t1.join(); t2.join();  
  
    Assert.assertTrue(2, v);  
}
```

Threads still rarely meet; synchronization costs dominate.

Current Form: Idea At A Glance

```
@JCStressTest
@State
public class MyTest {
    volatile int v;
    @Actor void actor1(I_Result r) { r.r1 = v++; }
    @Actor void actor2(I_Result r) { r.r2 = v++; }
}
```

- Large array of single-use state-bearing objects
- Actors access state objects under race
- Actors save their observations in provided storage
- Test infrastructure counts the observations

Current Form: Idea At A Glance

```
@JCStressTest
@State
public class MyTest {
    volatile int v;
    @Actor void actor1(I_Result r) { r.r1 = v++; }
    @Actor void actor2(I_Result r) { r.r2 = v++; }
}
```

End result: counted (r1, r2) outcomes

RESULT	SAMPLES	FREQ	EXPECT	DESCRIPTION
1, 1	46,946,789	10.1%	Interesting	...
1, 2	110,240,149	23.8%	Acceptable	...
2, 1	306,529,420	66.1%	Acceptable	...


Current Form: JCTest Examples

If you are reading the slides offline,
we are about to look through these examples:

```
https://github.com/openjdk/jcstress/  
tree/master/jcstress-samples/src/main/  
java/org/openjdk/jcstress/samples/api/
```

Current Form: Examples

Switching to JCSstress in 3... 2... 1...



JMM Basics

Spec: ...vs Implementation

Everybody intuitively understands the difference between the *specification* and the *implementation*

```
class Integer {  
    /**  
     * Returns a {@code String} object representing the  
     * specified integer. The argument is converted to signed decimal  
     * representation and returned as a string, exactly as if ...  
     */  
    public static String toString(int i) {  
        // Who cares what is going on here?  
    }  
}
```

Spec: Good Spec Is A Balance

- **Underspecify**, and things become unusable:

```
/**  
 * This method can do whatever it pleases.  
 */  
public void maybeSummonNasalDemons(int count) { ... }
```

- **Overspecify**, and implementation choices are limited:

```
/**  
 * This method checks if Java program halts.  
 */  
public boolean checkHalt(String program) { ... }
```

Spec: Abstract Machines

Language semantics is *specified* by the behavior of the *abstract machine*

<pre>public int m() { int x = 42; int y = 34; int t = x + y; return t; }</pre>	\Rightarrow	<pre>m: ...prolog... mov \$76\$, %rax ...epilog... ret</pre>
--	---------------	--

If the result is not distinguishable from the *abstract machine* behavior, nobody cares how it was achieved!

Spec: JMM Is Part Of Abstract Machine

If the result is not distinguishable from the *abstract machine* behavior, nobody cares how it was achieved!

```
volatile int x;  
public int m() {  
    x = 1;  
    x = 2;  
    return x;  
}
```


⇒

```
m:  
...prolog...  
mov $2$, (mem)  
mov $2$, %rax  
...epilog...  
ret
```

(In practice, not all optimizations are... practical)

JMM: Problem

«Oh, give me 5 minutes to read up on JMM!»



An execution E is described by

- P - a program
- A - a set of actions
- po - program order, which is the order of actions performed by t in A
- so - synchronization order in A

Given a write w , a freeze f , an action a (that is not a read of a `final` field), a read r_1 of the `final` field frozen by f , and a read r_2 such that $hb(w, f)$, $hb(f, a)$, $mc(a, r_1)$, and $dereferences(r_1, r_2)$, then when determining which values can be seen by r_2 , we consider $hb(w, r_2)$. (This *happens-before* ordering does not transitively close with other *happens-before* orderings.)

- Well-formed executions E_1, \dots , where $E_i = \langle P, A_i, po_i, so_i, W_i, V_i, sw_i, hb_i \rangle$.

Given these sets of actions C_0, \dots and executions E_1, \dots , every action in C_i must be one of the actions in E_i . All actions in C_i must share the same relative happens-before order and synchronization order in both E_i and E . Formally:

1. C_i is a subset of A_i

2. $hb_i \subseteq hb$

- There exists a set O of actions such that B consists of a *hang* action plus all the external actions in O and for all $k \geq |O|$, there exists an execution E of P with actions A , and there exists a set of actions O' such that:

- Both O and O' are subsets of A that fulfill the requirements for sets of observable actions.


- $O \subseteq O' \subseteq A$

- $|O'| \geq k$

ne in both E_i and E . Only the
Formally:

JMM: Problem

«Oh, give me 5 minutes to read up on JMM!»



An execution E is described by

- P - a program
- A - a set of actions
- po - program order, which is the order of actions performed by t in A
- so - synchronization order in A

Given a write w , a freeze f , an action a (that is not a read of a `final` field), a read r_1 of the `final` field frozen by f , and a read r_2 such that $hb(w, f)$, $hb(f, a)$, $mc(a, r_1)$, and $dereferences(r_1, r_2)$, then when determining which values can be seen by r_2 , we consider $hb(w, r_2)$. (This *happens-before* ordering does not transitively close with other *happens-before* orderings.)

- Well-formed executions E_1, \dots , where $E_i = \langle P, A_i, po_i, so_i, W_i, V_i, sw_i, hb_i \rangle$.

Given these sets of actions C_0, \dots and executions E_1, \dots , every action in C_i must be one of the actions in E_i . All actions in C_i must share the same relative happens-before order and synchronization order in both E_i and E . Formally:

1. C_i is a subset of A_i

2. $hb_i \subseteq hb$

- There exists a set O of actions such that B consists of a *hang* action plus all the external actions in O and for all $k \geq |O|$, there exists an execution E of P with actions A , and there exists a set of actions O' such that:

- Both O and O' are subsets of A that fulfill the requirements for sets of observable actions.

- $O \subseteq O' \subseteq A$

- $|O'| \geq k$

one in both E_i and E . Only the *happens-before* order. Formally:

JMM: Actions and Executions

Executions \approx Actions \cup Orders \cup Consistency Rules

JMM: Actions and Executions

Executions \approx Actions \cup Orders \cup Consistency Rules

Executions are the behaviors of the **abstract machine**, not the behavior of final implementation. They define all possible ways the Java program can possibly execute.

JMM: Actions and Executions

Executions \approx Actions \cup Orders \cup Consistency Rules

Actions:

- $w(field, V)$ – write value V into $field$
- $r(field) : V$ – read value V from $field$
- $L(monitor)$ – lock the $monitor$
- $UL(monitor)$ – unlock the $monitor$
- ...

JMM: Actions and Executions

Executions \approx Actions \cup Orders \cup Consistency Rules

Orders:

$$w(a, 1) \xrightarrow{\text{hb}} r(a) : 1 \dots w(a, 2)$$

Consistency rules:


- PO consistency
- SO consistency, SO-PO consistency
- HB consistency

JMM: Umm...


When someone explains something to you multiple times but you still have no idea wtf is going on



JMM: Why?



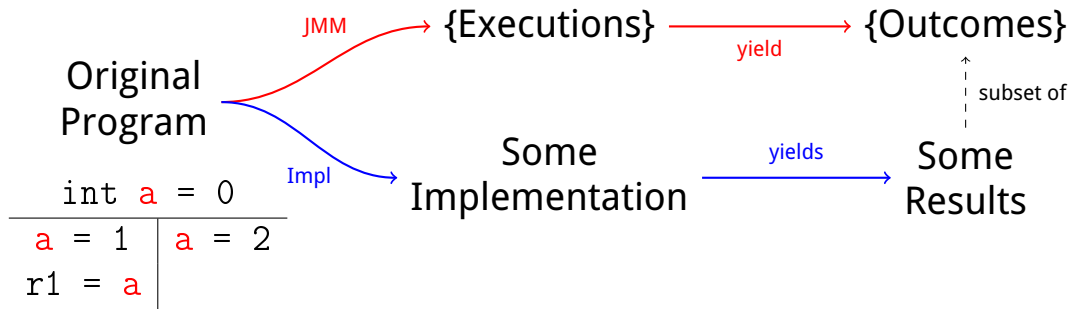
JMM: Why?



JMM: Why?

$w(a, 1) \xrightarrow{\text{hb}} r(a) : 1 \dots w(a, 2)$


$w(a, 1) \xrightarrow{\text{hb}} r(a) : 2 \dots w(a, 2)$



JMM: Why?

$w(a, 1) \xrightarrow{\text{hb}} r(a) : 1 \dots w(a, 2)$
 $w(a, 1) \xrightarrow{\text{hb}} r(a) : 2 \dots w(a, 2)$

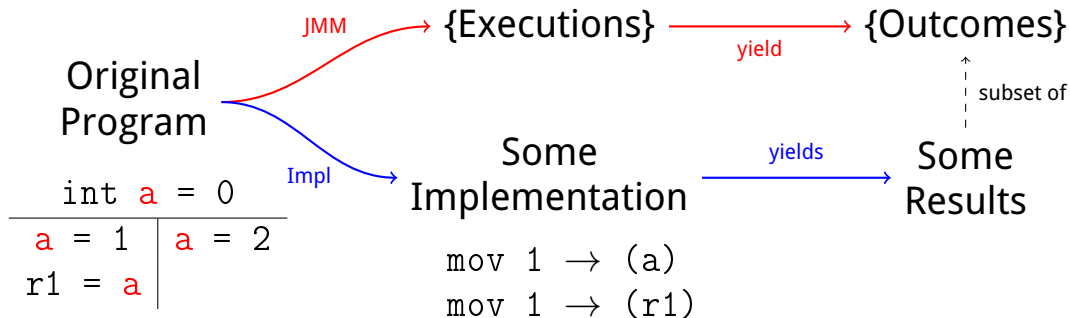
$r1 \in \{1, 2\}$



JMM: Why?

$w(a, 1) \xrightarrow{\text{hb}} r(a) : 1 \dots w(a, 2)$
 $w(a, 1) \xrightarrow{\text{hb}} r(a) : 2 \dots w(a, 2)$

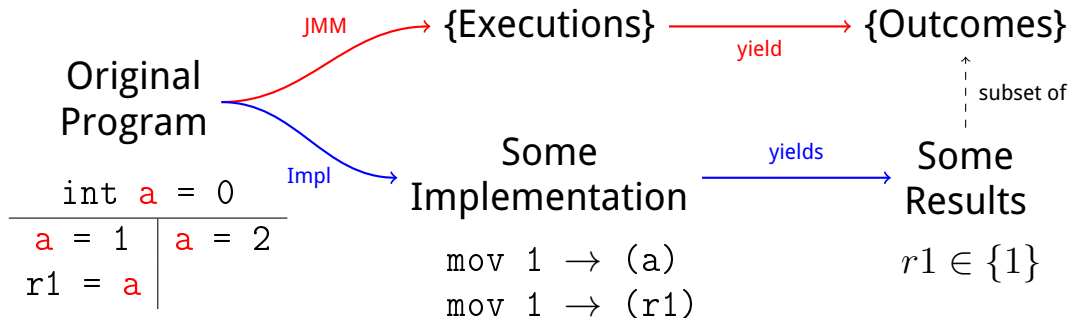
$r1 \in \{1, 2\}$




JMM: Why?

$w(a, 1) \xrightarrow{\text{hb}} r(a) : 1 \dots w(a, 2)$
 $w(a, 1) \xrightarrow{\text{hb}} r(a) : 2 \dots w(a, 2)$

$r1 \in \{1, 2\}$



JMM: Takeaway #1: Studying Implementations




Implementations are allowed to generate the **subset** of allowed outcomes, not all of them

- You can study JSR 133 Cookbook, but take it with a grain of salt



JMM: Takeaway #1: Studying Implementations




Implementations are allowed to generate the **subset** of allowed outcomes, not all of them

- You can study JSR 133 Cookbook, but take it with a grain of salt
- Reductio ad absurdum: Global Interpreter Lock



JMM: Takeaway #2, Interpreting Empirical Tests



The Universe is under no obligation to show you **all** of the outcomes allowed by spec or implementation

- «Not reproducible» does not mean «Not possible»
- Frequency is a *soft* evidence on possibility



JMM: JEP 188

There used to be the JMM Update JEP:

`https://openjdk.java.net/jeps/188`

- Improved formalization
- JVM coverage
- Extended scope
- C11/C++11 compatibility
- Implementation guidance
- ...

JMM: JEP 188



There used to be the JMM Update JEP:

<https://openjdk.java.net/jeps/188>

- Improved formalization
- JVM coverage
- Extended scope
- C11/C++11 compatibility
- Implementation guidance
- ...

JMM: VarHandles

```
java.lang.Object  
  java.lang.invoke.VarHandle
```



```
public abstract class VarHandle  
  extends Object
```

A VarHandle is a dynamically strongly typed reference to a variable, or to a parametrically-defined family of variables, including static fields, non-static fields, array elements, or components of an off-heap data structure. Access to such variables is supported under various *access modes*, including plain read/write access, volatile read/write access, and compare-and-swap.

• • • • •

Access modes control atomicity and consistency properties. Plain read (get) and write (set) accesses are guaranteed to be bitwise atomic only for references and for primitive values of at most 32 bits, and impose no observable ordering constraints with respect to threads other than the executing thread. Opaque operations are bitwise atomic and coherently ordered with respect to accesses to the same variable. In addition to obeying Opaque properties, Acquire mode reads and their subsequent accesses are ordered after matching Release mode writes and their previous accesses. In addition to obeying Acquire and Release properties, all Volatile operations are totally ordered with respect to each other.

JMM: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

JMM: JCStress Examples

If you are reading the slides offline,
we are about to look through these examples:

https:

`//github.com/openjdk/jcstress/tree/
master/jcstress-samples/src/main/java/
org/openjdk/jcstress/samples/jmm/basic`

Data Races: Definitions

- **Conflict:** at least 2 threads accessing the same variable, and at least 1 thread is writer
 - Concurrent only-readers are fine
 - Write-write conflicts are fun

Data Races: Definitions


- **Conflict:** at least 2 threads accessing the same variable, and at least 1 thread is writer
 - Concurrent only-readers are fine
 - Write-write conflicts are fun
- **Data Race:** a **conflict** that is not ordered by synchronization
 - Key problem in low-level concurrency
 - SC-DRF: Sequential Consistency for Data Race Freedom

Data Races: Definitions


- **Conflict:** at least 2 threads accessing the same variable, and at least 1 thread is writer
 - Concurrent only-readers are fine
 - Write-write conflicts are fun
- **Data Race:** a **conflict** that is not ordered by synchronization
 - Key problem in low-level concurrency
 - SC-DRF: Sequential Consistency for Data Race Freedom
- **Race Condition:** system behavior is dependent on the timing of events
 - Not a problem, *unless* some behaviors are undesirable

Data Races: Examples

Switching to JCTestress in 3... 2... 1...



Data Races: Takeaway, #1



In Java, unlike C/C++:

```
int s() {  
    M lm = m;  
    if (lm != null) {  
        return lm.x; // <--- This does not risk NPE  
    }  
    else  
        return 0;  
}
```

This would later become a building block
for so called «benign» data races

Data Races: Takeaway #2

1. Data race behavior is still somewhat deterministic
 - Racy reads are stronger than in other languages
 - Weird stuff still happens, but not completely catastrophic
 - *(This is what allows JCTest to even exist)*
2. Memory-model-wise, there is a difference:

```
int m1() {  
    int x1 = field;  
    int x2 = field;  
    return x1 + x2;  
}
```

```
int m2() {  
    int x1 = field;  
    int x2 = x1;  
    return x1 + x2;  
}
```



Data Races: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Access Atomicity: Definition

For any built-in type T:

```
T t = V1;
```


```
@Actor void actor1() {  
    t = V2;  
}
```

```
@Actor void actor2(T_Result r) {  
    r.r1 = t;  
}
```


$$r1 \in \{V1, V2\}$$


Access Atomicity: Examples

Switching to JCSstress in 3... 2... 1...



Access Atomicity: Takeaway



1. Most built-in types are access atomic
 - Almost all are naturally aligned
 - Unless 32-bit JVMs are present
2. Doing unnatural accesses break atomicity again
 - ByteBuffer-s «compound» operations
 - Unsafe «compound» operations
3. Larger types would break access atomicity again
 - Watch out for value/inline/primitive types

Access Atomicity: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Coherence: Definition

Coherence (*def.*):


The writes to the single memory location
appear to be in a total order
consistent with program order

- Most hardware gives this for free
- Most optimizers give up on this by default (i.e. do not track the order of reads)




Coherence: Examples

Switching to JCSstress in 3... 2... 1...



Coherence: Takeaway



1. Races laugh at our presuppositions about order
 - Most of the time, there is a complete free-for-all
 - Madness usually manifests after code transformations
 - Although hardware can also get us down
2. Coherency, while basic, is not guaranteed, unless...
 - We use `volatile` that is naturally coherent
 - We use weaker forms of `VarHandles` that are coherent
 - We use properly synchronized (non-racy) reads

Coherence: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Causality: Definition

Causality:


If **A** happened, then **B** happened too.

- By far the most **basic** guarantee made by most memory models, extremely hard to accept as the guiding principle
- The cornerstone of most (all?) distributed consistency models

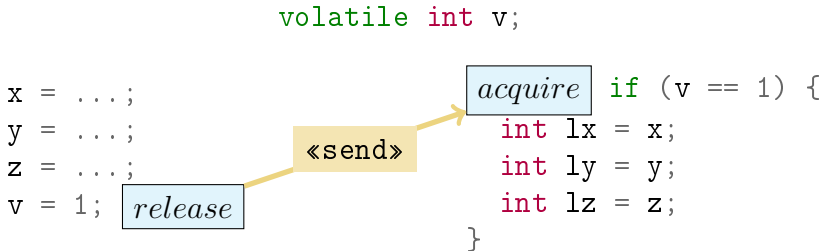


Causality: Examples

Switching to JCSstress in 3... 2... 1...




Causality: Safe Publication



- As if «commits to memory», but only for acq/rel pair
- release «commits», acquire gets the committed
- acquire has to see release witness!

Causality: Takeaway



1. Safe publication is the major (and simple) rule
 - Identify your **acquires** and **releases**
 - Check that **acquires/releases** are on all paths
 - Learn this rule! Then learn it again!
2. The whole thing does not require JMM reasoning
 - Hardly anyone applies «happens-before» correctly
 - Hardly anyone can do it reliably
 - It is very easy to miss the racy access

Causality: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Consensus: Definition

Consensus:


Momentary agreement among threads about program state

- There are different powers of consensus, but even the most basic Consensus-1 is useful.
- Consensus is mostly about *multiple* variables at once. Otherwise, coherence is enough...




Consensus: Examples

Switching to JCSstress in 3... 2... 1...



Consensus: Takeaway #6



1. Consensus is good

- Extremely useful to think about correctness
- Avoid non-SC data races by going `volatile`
- Sprinkle enough `volatiles` around your program, and it eventually becomes data-race-free! /s

2. Consensus is bad

- Extreme costs to get SC in distributed systems
- Most examples so far were fine with just Release/Acquire!
- Relaxing SC is by far the most common optimization technique

Consensus: Overview

		Defined	Atomic	Coherence	Causality	Consensus	Resilient
Java 8	Java 9						
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Finals: Idea

Finals:


«Declared immutable» fields, with additional semantics

- `final`-s are very special: able to hide data races
- The defense-in-depth strategy for concurrent code: work even when external synchronization is broken




Finals: Examples

Switching to JCSstress in 3... 2... 1...



Finals: Takeaway



1. Safe construction rule

- No reason to omit `final` from effectively immutable fields²
- This would be a building block for benign races

2. «Defense in Depth»: extra safety in the face of data races

- Users are known to misread, misinterpret, misuse the docs
- Most API-external objects need to be safely constructed

²Why JVM does not do it itself then, Aleksey?

Finals: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Benign Races: Canonical Form

There are cases when races are very useful:

```
V v; // deliberately non-volatile

public V benignRaceInit() {
    V lv = v;           // RULE 1: Read it once (racily)
    if (lv == null) {    // RULE 2: Check it is fine
        lv = compute(); // RULE 3: Recover by safely constructing
        v = lv;          // Publish unsafely (rely on safe construction)
    }
    return lv;
}
```


Forgo one of the rules, and you get the **non-benign** race.

Benign Races: Real JDK Code


```
public class AbstractMap<K, V> {  
    transient Set<K> keySet; // non-volatile  
  
    public Set<K> keySet() {  
        Set<K> ks = keySet; // RULE 1: Read it once (racily)  
        if (ks == null) { // RULE 2: Check it's fine  
            ks = new KeySet(); // RULE 3: Recover by safely constructing  
            keySet = ks;  
        }  
        return ks;  
    }  
}
```

Benign Races: Examples

Switching to JCTestress in 3... 2... 1...



Benign Races: Takeaway



1. Benign races are useful, albeit dangerous tool
 - Allows avoiding synchronized ops on critical paths!
 - The only sane way to deliberately use races in the program?
2. Works only if three rules are followed:
 - single (racy) read
 - reliability check
 - recovery path that safely constructs

Benign Races: Overview

Java 8	Java 9	Defined	Atomic	Coherence	Causality	Consensus	Resilient
plain	VH Plain	Y	≈	N	N	N	N
-	VH Opaque	Y	Y	Y	N	N	N
-	VH Acq/Rel	Y	Y	Y	Y	N	N
volatile	VH SeqCst	Y	Y	Y	Y	Y	N
final	-	Y	≈	≈	N	N	Y

Summing Up: Rule #1: Safe Publication



Golden Rule:

Thread 1: store everything, then **release**

Thread 2: **acquire**, then read anything

- Automatically happens when publishing via well-designed concurrency primitives
- Has to happen on **all possible** execution paths
- Has to happen in correct order

Summing Up: Rule #2: Safe Construction



Golden Rule:

When in doubt, make all fields `final`.

- Makes the whole thing more resilient to races
- Think «defense in depth»: survive in case some path fails to publish the instance safely

Summing Up: Rule #3: Benign Races



Golden Rule:

Object is safely constructed, and there is single read.

- Exotic optimization technique, rarely needed
- The (only) easy way to avoid synchronization

Summing Up: Rule #4: Exotic Modes



Golden Rule:
Don't.

- Just don't!
- There are cases where performance is so important, you want to have weaker than `volatile`, but stronger than plain – `VarHandles` to rescue!

Summing Up: Workshop Plan (Provisional)

■ Part I: Basic Bits

- JCIstress: Why do it? How does it look?
- JMM: Looking at basic examples

(Coffee Break)

■ Part II: Advanced Bits

- JCIstress: How does it work? Why would you not do it manually?
- JMM: Looking at advanced examples

(Breathing Exercises)


■ Part III: Fun Bits *(optional)*

- JCIstress: Real JVM/JDK bugs discovered
- Breakout: Discussions, Future Work, etc.

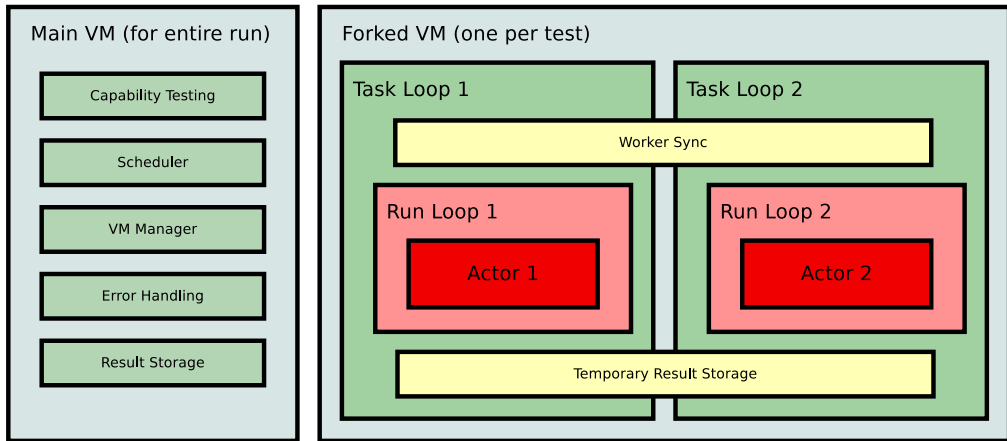
Part II. Advanced Bits

JCStress Advanced Topics

Architecture: Empirical Testing



Architecture: JCTest Architecture



Architecture: Guiding Principles

- Test fidelity:
 - Generate most of the stuff at build time
 - Assume the role of optimizing compiler where possible
- Test speed:
 - Initialize most of the stuff in host VM
 - Force forked VMs to do absolute minimum
- Test coverage:
 - Run forked VMs in all the interesting modes
 - Run forked VMs in all the interesting affinities

Balancing Heap Sizes: Problem


Speed/Fidelity:

JCStress normally runs tens/hundreds of JVMs concurrently

1	[100.0%]	17	[99.4%]	33	[99.4%]	49	[100.0%]				
2	[99.4%]	18	[100.0%]	34	[100.0%]	50	[100.0%]				
3	[100.0%]	19	[100.0%]	35	[99.4%]	51	[99.4%]				
4	[100.0%]	20	[99.4%]	36	[99.4%]	52	[100.0%]				
5	[99.4%]	21	[99.4%]	37	[99.4%]	53	[100.0%]				
6	[99.4%]	22	[100.0%]	38	[99.4%]	54	[99.4%]				
7	[100.0%]	23	[100.0%]	39	[99.4%]	55	[99.4%]				
8	[100.0%]	24	[100.0%]	40	[99.4%]	56	[99.4%]				
9	[100.0%]	25	[100.0%]	41	[99.4%]	57	[99.4%]				
10	[100.0%]	26	[100.0%]	42	[99.4%]	58	[99.4%]				
11	[100.0%]	27	[100.0%]	43	[99.4%]	59	[99.4%]				
12	[100.0%]	28	[100.0%]	44	[99.4%]	60	[100.0%]				
13	[99.4%]	29	[100.0%]	45	[100.0%]	61	[99.4%]				
14	[99.4%]	30	[99.4%]	46	[100.0%]	62	[100.0%]				
15	[100.0%]	31	[100.0%]	47	[100.0%]	63	[99.4%]				
16	[99.4%]	32	[100.0%]	48	[100.0%]	64	[99.4%]				
Mem[11.2G/126G]				Tasks: 157, 1736 thr; 64 running							
Swp[0K/0K]				Load average: 13.02 3.06 1.15							
				Uptime: 9 days, 14:08:35							
PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
107720	shade	20	0	3002M	112M	28104	S	198.	0.1	0:10.58	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107693	shade	20	0	3067M	110M	27760	S	197.	0.1	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107701	shade	20	0	3002M	109M	27980	S	197.	0.1	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107723	shade	20	0	0	0	0	Z	197.	0.0	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107702	shade	20	0	3002M	115M	27800	S	197.	0.1	0:10.57	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107727	shade	20	0	3002M	113M	27856	S	197.	0.1	0:10.54	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107948	shade	20	0	3067M	111M	27908	S	197.	0.1	0:10.51	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107726	shade	20	0	3002M	110M	27916	S	197.	0.1	0:10.59	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107940	shade	20	0	3132M	115M	27936	S	197.	0.1	0:10.53	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107949	shade	20	0	3067M	108M	27556	S	197.	0.1	0:10.52	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107742	shade	20	0	3066M	111M	27868	S	197.	0.1	0:10.46	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic

Balancing Heap Sizes: Examples

Switching to JCTestress in 3... 2... 1...



Balancing Strides: Problem

Fidelity:


When heap size is set, tests may OOME!

```
@JCStressTest
@State
public class MyTest {
    byte[] arr = new byte[1024*1024];

    @Actor void actor1(I_Result r) {
        int s = 0;
        for (byte b : arr) s += b;
        r.r1 = s;
    }
}
```

Balancing Strides: Examples

Switching to JCSstress in 3... 2... 1...



Balancing Thread Counts: Problem


Speed/Fidelity:

JCStress normally runs tens/hundreds of JVMs concurrently

1	[100.0%]	17	[99.4%]	33	[99.4%]	49	[100.0%]				
2	[99.4%]	18	[100.0%]	34	[100.0%]	50	[100.0%]				
3	[100.0%]	19	[100.0%]	35	[99.4%]	51	[99.4%]				
4	[100.0%]	20	[99.4%]	36	[99.4%]	52	[100.0%]				
5	[99.4%]	21	[99.4%]	37	[99.4%]	53	[100.0%]				
6	[99.4%]	22	[100.0%]	38	[99.4%]	54	[99.4%]				
7	[100.0%]	23	[100.0%]	39	[99.4%]	55	[99.4%]				
8	[100.0%]	24	[100.0%]	40	[99.4%]	56	[99.4%]				
9	[100.0%]	25	[100.0%]	41	[99.4%]	57	[99.4%]				
10	[100.0%]	26	[100.0%]	42	[99.4%]	58	[99.4%]				
11	[100.0%]	27	[100.0%]	43	[99.4%]	59	[99.4%]				
12	[100.0%]	28	[100.0%]	44	[99.4%]	60	[100.0%]				
13	[99.4%]	29	[100.0%]	45	[100.0%]	61	[99.4%]				
14	[99.4%]	30	[99.4%]	46	[100.0%]	62	[100.0%]				
15	[100.0%]	31	[100.0%]	47	[100.0%]	63	[99.4%]				
16	[99.4%]	32	[100.0%]	48	[100.0%]	64	[99.4%]				
Mem[11.2G/126G]				Tasks: 157, 1736 thr; 64 running							
Swp[0K/0K]				Load average: 13.02 3.06 1.15							
				Uptime: 9 days, 14:08:35							
PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
107720	shade	20	0	3002M	112M	28104	S	198.	0.1	0:10.58	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107693	shade	20	0	3067M	110M	27760	S	197.	0.1	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107701	shade	20	0	3002M	109M	27980	S	197.	0.1	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107723	shade	20	0	0	0	0	Z	197.	0.0	0:10.60	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107702	shade	20	0	3002M	115M	27800	S	197.	0.1	0:10.57	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107727	shade	20	0	3002M	113M	27856	S	197.	0.1	0:10.54	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107948	shade	20	0	3067M	111M	27908	S	197.	0.1	0:10.51	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107726	shade	20	0	3002M	110M	27916	S	197.	0.1	0:10.59	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107940	shade	20	0	3132M	115M	27936	S	197.	0.1	0:10.53	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107949	shade	20	0	3067M	108M	27556	S	197.	0.1	0:10.52	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic
107742	shade	20	0	3066M	111M	27868	S	197.	0.1	0:10.46	/home/shade/Install/jdk-ea/bin/java -cp tests-custom/target/jcstress.jar -XX:+UnlockDiagnostic

Balancing Thread Counts: Examples

Switching to JCTestress in 3... 2... 1...



VM Modes: Problem

Coverage:


Some tests might fail only with some compilers

Hotspot VM has *at least* three ways to execute Java code:

- Interpreter (-Xint)
- C1 (baseline, client, -XX:-TieredStopAtLevel=1)
- C2 (optimized, server, -XX:-TieredCompilation)

VM Modes: Examples

Switching to JCSstress in 3... 2... 1...



Split Compilation: Problem


Coverage:

Some outcomes only manifest in odd compilation conditions

```
@Actor void actor1(II_Result r) {  
    // Compile with C1, one barrier scheme  
}  
  
@Actor void actor2(II_Result r) {  
    // Compile with C2, another barrier scheme  
    ...  
}
```


Split Compilation: Examples

Switching to JCTest in 3... 2... 1...



Compiler Fuzzing: Problem

Coverage:


Many outcomes are compiler-induced.

Hotspot VM provides a few randomizing flags:

- `-XX:+StressLCM`, `-XX:+StressGCM` – added for JCTestress
- `-XX:+StressIGVN`, `-XX:+StressCCP` – added later
- Host VM probes which ones are available
- Host VM mixes these as the separate configuration

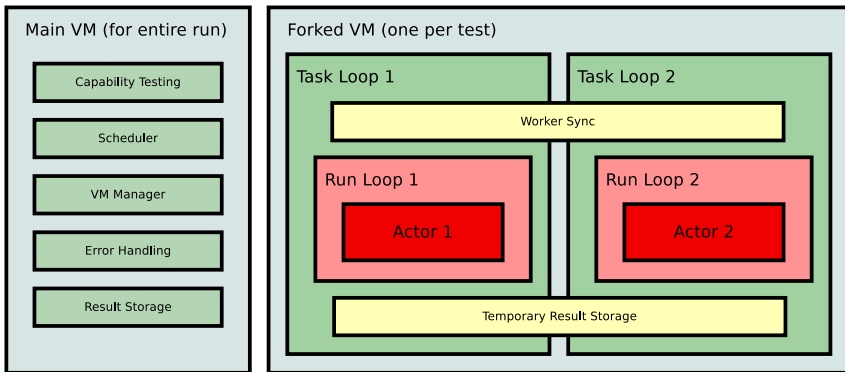
Compiler Fuzzing: Examples

Switching to JCTestress in 3... 2... 1...




Loops: Problem

Fidelity:
JVMs like short and active loops.



Loops: Examples

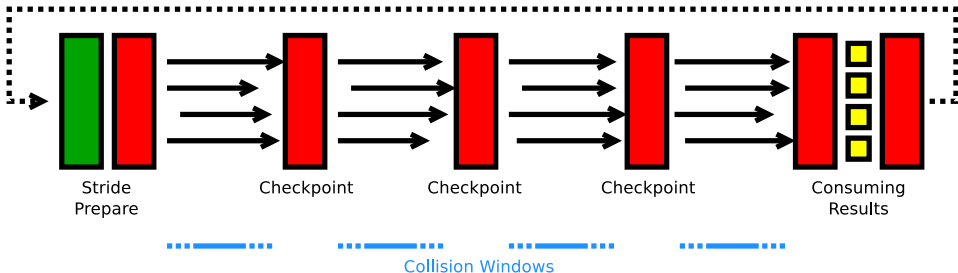
Switching to JCSstress in 3... 2... 1...



Rendezvous: Problem


Fidelity:

Tests are rarely symmetric, so actors outpace each other



Rendezvous: Examples

Switching to JCSstress in 3... 2... 1...

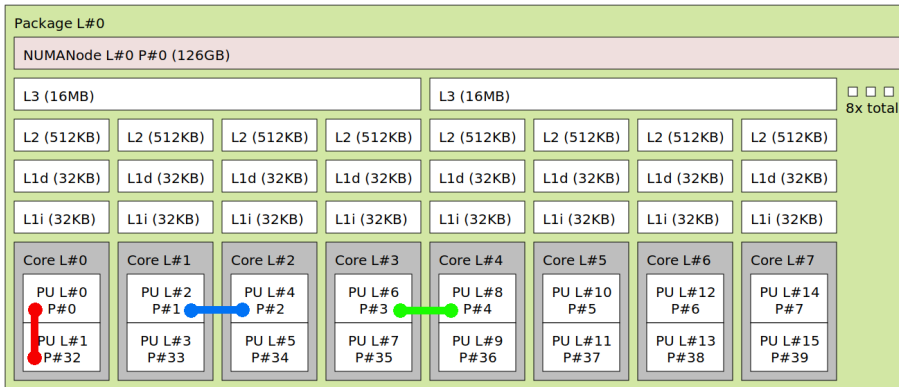


Actor Affinity: Problem

Coverage:


Some outcomes manifest on some test topologies

Machine (126GB total)



Actor Affinity: Examples

Switching to JCSstress in 3... 2... 1...



Busy Waiting: Problem

Speed/Fidelity: Workers often have to wait actively.

HARD

0, 0	1,439,041,804	6.1%
0, 1	10,878,860,414	46.6%
1, 0	10,967,111,106	47.0%
1, 1	23,684,276	0.1%

THREAD_SPIN_WAIT

0, 0	1,368,587,258	5.5%
0, 1	11,458,577,433	46.4%
1, 0	11,834,865,693	47.9%
1, 1	23,107,536	0.1%

THREAD_YIELD


0, 0	728,512,816	2.8%
0, 1	12,766,383,247	49.1%
1, 0	12,448,519,780	47.9%
1, 1	10,079,197	0.1%

LOCKSUPPORT_PARK_NANOS

0, 0	12,051,745	0.1%
0, 1	3,876,507,240	50.1%
1, 0	3,843,372,966	49.7%
1, 1	179,409	<0.1%

Busy Waiting: Examples

Switching to JCTestress in 3... 2... 1...



False Sharing: Problem

Fidelity:

Adjacent fields are false-shared,
decreasing the interesting outcomes frequency


```
volatile int a, b; // false shared?
```

```
@Actor void actor1(II_Result r) {  
    ...  
}
```

```
@Actor void actor2(II_Result r) {  
    ...  
}
```

False Sharing: Examples

Switching to JCTest in 3... 2... 1...



Reusing Objects: Problem

Speed:

API demands @State and @Result objects are one-shot

Without reuse:


RESULT	SAMPLES	FREQ
0, 0	64,113,081	5.9%
0, 1	522,012,200	47.8%
1, 0	504,635,282	46.2%
1, 1	1,608,557	0.1%

With reuse:

RESULT	SAMPLES	FREQ
0, 0	1,368,587,258	5.5%
0, 1	11,458,577,433	46.4%
1, 0	11,834,865,693	47.9%
1, 1	23,107,536	0.1%

Reusing Objects: Examples

Switching to JCTestress in 3... 2... 1...



Null-Pointer Checks: Problem


Fidelity:

NPE semantics induces orderings

```
void actor1(II_Result r) {  
    r.r1 = x; // null-check "r"  
    r.r2 = x;  
}
```


Null-Pointer Checks: Examples

Switching to JCTestress in 3... 2... 1...



JMM Advanced Topics


JMM Advanced Topics: JCTestress Examples

If you are reading the slides offline,
we are about to look through these examples:

```
https://github.com/openjdk/jcstress/  
tree/master/jcstress-samples/src/main/  
java/org/openjdk/jcstress/samples/jmm/  
advanced
```


Synchronized Barriers: Examples

Switching to JCTestress in 3... 2... 1...




Multi-Copy Atomicity: Examples

Switching to JCTestress in 3... 2... 1...




Losing Updates: Examples

Switching to JCTestress in 3... 2... 1...




Misplaced Volatiles: Examples

Switching to JCTest in 3... 2... 1...




Semi-Synchronized: Examples

Switching to JCTestress in 3... 2... 1...




Volatile Arrays: Examples

Switching to JCTestress in 3... 2... 1...




Acquire/Release Orders Wrong: Examples

Switching to JCTest in 3... 2... 1...




Synchronized «Barriers»: Examples

Switching to JCTestress in 3... 2... 1...




Volatile «Barriers»: Examples

Switching to JCTestress in 3... 2... 1...



Volatile != Final: Examples

Switching to JCTestress in 3... 2... 1...



Summing Up: Takeaway

- JMM guarantees are weaker than a single implementation might show
- You have to code to JMM rules, not to implementation behavior!

Part III. Extreme Bits

A Taste Of Real Bugs

Wrong Labels: Problem

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

```
@Actor void actor1() {  
    x = 1;  
    y = 1;  
}
```

```
@Actor void actor2(II_Result r) {  
    // int t = x;  
    r.r1 = y;  
    r.r2 = x;  
}
```

(1, 0) is illegal under JMM rules, but breaks when prior copy
«exists»!

Wrong Labels: C1 Bug³

C1 CSE bug, ignores `volatile` read:

```
t = x;  
r1 = y;  
r2 = x;
```

...so it coalesced the read:

```
t = x;  
r1 = y;  
r2 = t;
```

³<https://bugs.openjdk.java.net/browse/JDK-7170145>

Immortal Referents: Problem

```
final WeakReference<Object> ref = new WeakReference<>(obj);
```

```
@Actor void actor1() {  
    while (ref.get() != null); // wait  
}
```

```
@Actor void actor2(II_Result r) {  
    ref.clear();  
}
```

Expected behavior: test eventually terminates.
Actual behavior: test is stuck.

Immortal Referents: Compiler/GC Bug⁴

```
public abstract class Reference<T> {  
    private T referent;  
    ...  
    public T get() { return referent; }  
}
```

```
// just wait...  
// ...a little  
while (ref.get() != null);  
  
⇒  
if (ref.referent != null) {  
    while (true); // burn !  
}
```

⁴<https://bugs.openjdk.java.net/browse/JDK-7190310>

Stuck Threads: Problem?

```
@Actor void actor1() {  
    while (!Thread.interrupted()); // wait  
}  
  
@Signal void signal(Thread actor1) {  
    actor1.interrupt();  
}
```

Expected behavior: test eventually terminates
Actual behavior: test eventually terminates!

Stuck Threads: Problem!

```
private boolean check() { return Thread.interrupted(); }

@Actor void actor1() {
    while (!check()); // wait
}

@Signal void signal(Thread actor1) {
    actor1.interrupt();
}
```

Actual behavior: test is stuck!

Stuck Threads: C2 Bug⁶

- `Thread.interrupted()` *used to* check a flag in the native

```
@HotSpotIntrinsicCandidate  
private native boolean isInterrupted(boolean ClearInterrupted);
```

- Access was written in C2 IR; effectively a plain read, unless it is specifically written like a «volatile»
- Since then, it was rewritten to plain `volatile` field⁵

⁵<https://bugs.openjdk.java.net/browse/JDK-8229516>

⁶<https://bugs.openjdk.java.net/browse/JDK-8003135>

Eat My Shorts: Problem?

```
short s;  
  
@Actor void actor1() {  
    s = 0xFFFF;  
}  
  
@Actor void actor2(S_Result r) {  
    r.r1 = s;  
}
```

short is supposed to be atomic:

$r1 \in \{0x0000, 0xFFFF\}$, and it is indeed the case.

Eat My Shorts: Problem!

```
short s;
```

```
@Actor void actor1() {  
    s = 0xFFFF;  
}
```

```
@Actor void actor2(BB_Result r) {  
    short t = s;  
    r.r1 = (byte)((t >> 0) & 0xFF);  
    r.r2 = (byte)((t >> 8) & 0xFF);  
}
```

Expected: (0x00, 0x00), (0xFF, 0xFF)

Actual: (0x00, 0x00), (0xFF, 0xFF), (0x00, 0xFF), (0xFF, 0x00)

Eat My Shorts: Gradual Graph Rewrite

// Original code

```
short t = short_load(s.x);  
r.r1 = byte_store(and(shift(t, 0), 0xFF));  
r.r2 = byte_store(and(shift(t, 8), 0xFF));
```

// First round of simplifications

```
short t = short_load(s.x);  
r.r1 = byte_store(t);  
r.r2 = byte_store(shift(t, 8));
```

// Final round of simplifications

```
r.r1 = byte_store(unsigned_short_load(s.x));  
r.r2 = byte_store(shift(signed_short_load(s.x), 8));
```

Eat My Shorts: C2 Bug⁷

```
short t = s.x;  
r.r1 = (byte) ((t >> 0) & 0xFF);  
r.r2 = (byte) ((t >> 8) & 0xFF);
```

Compiles to:

```
movzwl 0xc(%rdx), %r11d    ; read s.x  
mov     %r11b,0xc(%rcx)    ; store r.r1  
movswl 0xc(%rdx), %r10d    ; read s.x again!  
shr     $0x8,%r10d         ; shift  
mov     %r10b,0xd(%rcx)    ; store r.r2
```

⁷<https://bugs.openjdk.java.net/browse/JDK-8000805>

Volatile Clash: Problem

```
volatile double d = 0.0D;  
  
@Actor void actor1() {  
    d = Double.toRawLongBits(-1L);  
}  
  
@Actor void actor2(D_Result r) {  
    r.r1 = d;  
}
```

Expected: $r1 \in \{0x0...0, 0xF...F\}$
Actual: that, plus garbage!

Volatile Clash: Native Unsafe Code

```
#define GET_FIELD_VOLATILE(obj, offset, type_name, v) \  
oop p = JNIHandles::resolve(obj); \  
type_name v = OrderAccess::load_acquire( \  
    (volatile type_name*) index_oop_from_field_offset_long(p, offset));
```

Unsafe_GetDoubleVolatile() compiles to:

```
mov    0x18(%esp),%ebp  
add    %ebp,%eax  
; field offset in %eax  
fldl   (%eax)  
fstpl  0x18(%esp)
```

Volatile Clash: Native Code Bug⁸

```
#define GET_FIELD_VOLATILE(obj, offset, type_name, v) \  
oop p = JNIHandles::resolve(obj); \  
type_name v = OrderAccess::load_acquire( \  
    (volatile type_name*) index_oop_from_field_offset_long(p, offset));
```

Unsafe_GetDoubleVolatile() actually compiled to:

```
mov     0x4(%eax),%edx    ; WHAT  
mov     (%eax),%eax      ; WHAT  
mov     %eax,0x20(%esp)   ; THE  
mov     %edx,0x24(%esp)   ; THE  
fldl    0x20(%esp)        ; HELL  
fstpl   0x18(%ebx)        ; HELL
```

⁸<https://bugs.openjdk.java.net/browse/JDK-8016538>

Power Dekker: Problem

```
volatile int x, y;  
  
@Actor void actor1(II_Result r) {  
    x = 1;  
    r.r1 = y;  
}  
  
@Actor void actor2(II_Result r) {  
    y = 1;  
    r.r2 = x;  
}
```

SC executions: $(r1, r2) \notin \{(0, 0)\}$

Power Dekker: Piece 1: Optimizing For Hardware

- POWER ISA has a lot of registers
 - *(take that, lousy x86)*
- Hotspot PPC port is capitalizing on that
 - Profitable to schedule loads as soon as possible
 - ...unless something prevents it

Power Dekker: Piece 2: Optimizing Barriers Bug⁹

```
x = 1;  
r1 = y;
```

This produces, roughly:

MB \rightarrow store(x, 1) \rightarrow MB \rightarrow load(r1, y) \rightarrow MB

Barrier optimization code mistakenly removes the barrier after volatile store, because it thinks there is a leading membar before volatile load:

MB \rightarrow store(x, 1) \rightarrow load(r1, y) \rightarrow MB

⁹<https://bugs.openjdk.java.net/browse/JDK-8007898>

AArch64 Zero barriers: Problem

```
int x; volatile long y;
```

```
@Actor void actor1() {  
    x = 1;  
    y = 1;  
}
```

```
@Actor void actor2(IJ_Result r) {  
    r.r1 = y;  
    r.r2 = x;  
}
```

Observing (1, 0) on AArch64!

AArch64 Zero barriers: Barrier Selection Bug¹⁰

```
#ifdef ARM
    #define LIGHT_MEM_BARRIER __kernel_dmb()
#else // ARM
    #ifdef PPC
        #define LIGHT_MEM_BARRIER __asm __volatile ("lwsync":::"memory")
    #else // PPC
        #ifdef ALPHA
            #define LIGHT_MEM_BARRIER __sync_synchronize()
        #else // ALPHA
            #define LIGHT_MEM_BARRIER __asm __volatile (""::::"memory")
        #endif // ALPHA
    #endif // PPC
#endif // ARM
```

ARM32 Zero barriers: Problem

Actually, JCTest would not even pass initialization:

```
$ java -jar jcstress-tests-all-20200917.jar  
Java Concurrency Stress Tests
```

```
-----  
Rev: ad66703e2ed0, built by buildbot with 11.0.5-testing at 2020-09-17
```

```
...
```

```
Caused by: java.lang.OutOfMemoryError:
```

```
  Cannot reserve 8192 bytes of direct buffer memory  
  (allocated: 0, limit: -5290888278393214624)
```

```
...
```


ARM32 Zero barriers: Atomicity Bug¹¹

```
static void atomic_copy64(const volatile void *src,
                          volatile void *dst) {
    jlong tmp;
    asm volatile ("ldrexld  %0, [%1]\n"
                  : "=r"(tmp)
                  : "r"(src), "m"(src));
    *(jlong *) dst = tmp;
}
```

¹¹<https://bugs.openjdk.java.net/browse/JDK-8253464>


Conclusions

Conclusions: In One Picture



Backup: JMM Arguments


Data Races: Formal Argument



```
class M { ... }  
M m;
```

<pre>m = new M(); m = null;</pre>	<pre>M lm = m; r1 = (lm != null); r2 = (lm != null);</pre>
---------------------------------------	--

Data Races: Formal Argument



```
class M { ... }  
M m;  
-----  
m = new M();   M lm = m;  
m = null;      r1 = (lm != null);  
               r2 = (lm != null);
```


JMM allows only (F, F) and (T, T)

Data Races: Formal Counter-Argument

Can't compiler «inline» the local variable?

```
class M { ... }  
M m;  
-----  
m = new M();  
m = null;      | r1 = (m != null);  
                | r2 = (m != null);
```

Data Races: Formal Counter-Argument



Can't compiler «inline» the local variable?

```
class M { ... }  
M m;  
-----  
m = new M();  
m = null;    r1 = (m != null);  
              r2 = (m != null);
```

See, there is an obvious execution that yields (T, F) now!

... $r(m) : !null \xrightarrow{\text{po}} r(m) : null$

Data Races: Program Order

Program order (PO) provides the link between the execution and the program in question

- PO – total order for any given thread in isolation
- **PO consistency:** PO is consistent with the source code order in the original program



Data Races: PO And Transformations

Original program:

```
M lm = m;
```

```
r1 = (lm != null);
```

```
r2 = (lm != null);
```

$$\begin{array}{l} w(m, *) \xrightarrow{\text{po}} w(m, null) \\ r(m) : * \end{array}$$

Transformed program:

```
r1 = (m != null);
```

```
r2 = (m != null);
```

$$\begin{array}{l} w(m, *) \xrightarrow{\text{po}} w(m, null) \\ r(m) : * \xrightarrow{\text{po}} r(m) : * \end{array}$$

Data Races: PO And Trans

This execution does not relate to the original program, oops

Original program:

```
M lm = m;  
r1 = (lm != null);  
r2 = (lm != null);
```

$w(m, *) \xrightarrow{\text{po}} w(m, null)$
 $r(m) : *$

Transformed program:

```
r1 = (m != null);  
r2 = (m != null);
```

$w(m, *) \xrightarrow{\text{po}} w(m, null)$
 $r(m) : * \xrightarrow{\text{po}} r(m) : *$

Data Races: PO And Trans

This execution should be used to reason about outcomes for the transformed program

Original program:

```
M lm = m;  
r1 = (lm != null);  
r2 = (lm != null);
```

$w(m, *) \xrightarrow{\text{po}} w(m, \text{null})$
 $r(m) : *$

Transformed program:

```
r1 = (m != null);  
r2 = (m != null);
```

$w(m, *) \xrightarrow{\text{po}} w(m, \text{null})$
 $r(m) : * \xrightarrow{\text{po}} r(m) : *$

Data Races: PO And Transformations

Original program:

```
M lm = m;  
r1 = (lm != null);  
r2 = (lm != null);
```

$w(m, *) \xrightarrow{\text{po}} w(m, null)$
 $r(m); *$




Transformed program:

PO consistency:

Original program has single read?
Relatable executions also have single read!


```
r1 = (m, null)  
r2 = (m); *
```

Coherence: Example 2.1



	int x;
x = 1;	r1 = x; // r1
	r2 = x; // r2

Coherence: Example 2.1




	int x;
x = 1;	r1 = x; // r1
	r2 = x; // r2

JMM allows observing (1, 0), see:

$w(x, 1) \dots r_1(x) : 1 \xrightarrow{\text{po}} r_2(x) : 0$

Coherence: Example 2.1



	int x;
<hr/>	
x = 1;	r1 = x; // r1
	r2 = x; // r2

JMM allows observing (1, 0), see:

$w(x, 1) \dots r_1(x) : 1 \xrightarrow{\text{po}} r_2(x) : 0$

This execution is **PO** consistent, both reads are here!

Coherence: Consistency Rules

PO consistency affects the **structure** of the execution.
What we need: a consistency rule that affects **values**
observed by the actions.

In JMM, there are two of them:

1. Happens-before (HB) consistency
2. Synchronization order (SO) consistency



Coherence: Consistency Rules

PO consistency affects the **structure** of the execution.
What we need: a consistency rule that affects **values**
observed by the actions.

In JMM, there are two of them:

1. Happens-before (**HB**) consistency
2. Synchronization order (**SO**) consistency ← **now!**



Coherence: **SO** – Synchronization Order

SO covers all *synchronization actions*:
volatile read/write, lock/unlock, etc.

- **SO** is a total order («All SA actions relate to each other»)
- **SO-PO consistency**: $\xrightarrow{\text{so}}$ and $\xrightarrow{\text{po}}$ agree
- **SO consistency**: reads see only the latest write in $\xrightarrow{\text{so}}$


Coherence: **SO** – Synchronization Order

SO covers all *synchronization actions*:
volatile read/write, lock/unlock, etc.

- **SO** is a total order («All SA actions relate to each other»)
- **SO-PO consistency**: $\xrightarrow{\text{so}}$ and $\xrightarrow{\text{po}}$ agree
- **SO consistency**: reads see only the latest write in $\xrightarrow{\text{so}}$


Just what coherence wants!

Coherence: Example 2.2



volatile int x;	
x = 1;	r1 = x; // r1
	r2 = x; // r2

Coherence: Example 2.2



volatile int x;	
x = 1;	r1 = x; // r1
	r2 = x; // r2

Valid executions give $(0, 0), (1, 1), (0, 1):^a$

$w(x, 1) \xrightarrow{\text{so}} r_1(x) : 1 \xrightarrow{\text{so}} r_2(x) : 1$

$r_1(x) : 0 \xrightarrow{\text{so}} w(x, 1) \xrightarrow{\text{so}} r_2(x) : 1$

$r_1(x) : 0 \xrightarrow{\text{so}} r_2(x) : 0 \xrightarrow{\text{so}} w(x, 1)$

^aProving no other outcomes exist is left as an exercise for the reader

Causality: SW – Synchronizes-With Order

When one SA «sees» the value of another SA, they are said to be in «synchronizes-with» (SW) relation

- SW is a partial order
- SW connects the operations that «see» each other
- Acts like the «bridge» between the threads




Causality: **HB** – Happens-Before Order

HB is a transitive closure
over the union of **PO** and **SW**

- **HB** is a partial order
(Translation: not everything is connected)
- **HB consistency**: reads observe either:
the last write in $\xrightarrow{\text{hb}}$, or
any other write, not ordered by $\xrightarrow{\text{hb}}$




Causality: Example 3.1



```
int x;  
volatile int y;  
-----  
x = 1; | r1 = y;  
y = 1; | r2 = x;
```

Causality: Example 3.1




```
int x;  
volatile int y;  
-----  
x = 1; | r1 = y;  
y = 1; | r2 = x;
```

We are dealing with this class of executions:

$$w(x, 1) \xrightarrow{\text{po}} w(y, 1) \dots r(y) : * \xrightarrow{\text{po}} r(x) : *$$

Causality: Example 3.1




```
int x;  
volatile int y;  
-----  
x = 1; | r1 = y;  
y = 1; | r2 = x;
```

Racy subclass:

$$\begin{array}{l} w(x, 1) \xrightarrow{\text{hb}} w(y, 1) \dots r(y) : 0 \xrightarrow{\text{hb}} r(x) : 0 \\ w(x, 1) \xrightarrow{\text{hb}} w(y, 1) \dots r(y) : 0 \xrightarrow{\text{hb}} r(x) : 1 \end{array}$$

Causality: Example 3.1




```
int x;  
volatile int y;  
-----  
x = 1; | r1 = y;  
y = 1; | r2 = x;
```

Non-racy subclass:

$$\begin{array}{l} w(x, 1) \xrightarrow{\text{hb}} w(y, 1) \xrightarrow{\text{hb}} r(y) : 1 \xrightarrow{\text{hb}} r(x) : 1 \\ w(x, 1) \xrightarrow{\text{hb}} w(y, 1) \xrightarrow{\text{hb}} r(y) : 1 \xrightarrow{\text{hb}} r(x) : 0 \end{array}$$


Causality: Look Closer, #1

Happens-before is defined over *actions*,
not over statements: notice no **HB** between volatile ops!



Causality: Look Closer, #2

This violates **HB** consistency:




should have seen this!

Causality: Observing the `volatile` store causes observing everything stored before it


Causality: Example 3.2

Notice the order
is different

int x;	
volatile int y;	
<hr/>	
y = 1;	r1 = x;
x = 1;	r2 = y;



Causality: Example 3.2




Notice the order
is different

```
int x;  
volatile int y;  
-----  
y = 1; | r1 = x;  
x = 1; | r2 = y;
```

Hey, look how (1, 0) is allowed:

$w(y, 1) \xrightarrow{\text{hb}} w(x, 1) \dots r(x) : 1 \xrightarrow{\text{hb}} r(y) : 0$

Causality: Example 3.2



Notice the order
is different

```
int x;  
volatile int y;  
-----  
y = 1; | r1 = x;  
x = 1; | r2 = y;
```

Hey, look how (1, 0) is allowed:

$w(y, 1) \xrightarrow{\text{hb}} w(x, 1) \dots r(x) : 1 \xrightarrow{\text{hb}} r(y) : 0$

Look: irrelevant that y is volatile!

Consensus: Example 4.1

volatile int x , y ;			
x = 1;	y = 1;	int r1 = y ;	int r3 = x ;
		int r2 = x ;	int r4 = y ;

HB alone allows seeing (1, 0, 1, 0):

$w(y, 1)$	$\xrightarrow{\text{hb}}$	$r_1(y) : 1$	$\xrightarrow{\text{hb}}$	$r_3(x) : 0$
$w(x, 1)$	$\xrightarrow{\text{hb}}$	$r_2(x) : 1$	$\xrightarrow{\text{hb}}$	$r_4(y) : 0$

Consensus: SC

Sequential Consistency (SC): *(def.)*

«...the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program»



Consensus: **SO** – Synchronization Order

SO covers all *synchronization actions*:
volatile read/write, lock/unlock, etc.

- **SO** is a total order («All SA actions relate to each other»)
- **SO-PO consistency**: $\xrightarrow{\text{so}}$ and $\xrightarrow{\text{po}}$ agree
- **SO consistency**: reads see only the latest write in $\xrightarrow{\text{so}}$

Just what Sequential Consistency wants!

Finals: Example 5.1

```
class M { final int x = 42; }
```

```
M m;
```

```
m = new M()
```

```
M lm = m
```

```
if (lm != null)
```

```
    r1 = lm.x
```

```
else
```

```
    r1 = 1
```

Finals: Example 5.1

```
class M { final int x = 42; }  
M m;
```

<pre>m = new M()</pre>	<pre>M lm = m if (lm != null) r1 = lm.x else r1 = 1</pre>
------------------------	---

JMM guarantees seeing the value of `final` field here:
 $r1 \in \{1, 42\}$

Finals: Example 5.1

```
class M { final int x = 42; }
```

```
M m;
```

```
m = new M()
```

```
M lm = m
```

```
if (lm != null)
```

```
    r1 = lm.x
```

```
else
```

```
    r1 = 1
```

Special rule, if `x` is a `final` field:

$$w(x, 42) \xrightarrow{\text{hb}} r(x) : 42$$

Finals: Example 5.2

```
class M { volatile int x = 42; }  
M m;
```

<pre>m = new M() M lm = m if (lm != null) r1 = lm.x else r1 = 1</pre>

Finals: Example 5.2

```
class M { volatile int x = 42; }
```

```
M m;
```

```
m = new M()
```

```
M lm = m
```

```
if (lm != null)
```

```
    r1 = lm.x
```

```
else
```

```
    r1 = 1
```

JMM allows (0) here:

$$w(cm.x, 42) \xrightarrow{\text{hb}} w(cm, m) \dots r(m) : lm \xrightarrow{\text{hb}} r(lm.x) : 0$$

Finals: Example 5.2

```
class M { volatile int x = 42; }  
M m;
```

<pre>m = new M()</pre>	<pre>M lm = m if (lm != null) r1 = lm.x else r1 = 1</pre>
------------------------	---

`volatile` \notin `final`
`final` \notin `volatile`

Finals: Safe Construction

Special rule for `final` fields:

$$writes_{final} \xrightarrow{hb} reads_{final}$$

The derivation for that rule is complicated.

Two absolutely necessary things:

- Field is `final`
- Constructor does not publish `this`