

Data Races and Synchronization

Atomic Operations, Non-Blocking Algorithms

Effective Software — Lecture 8

David Šišlák

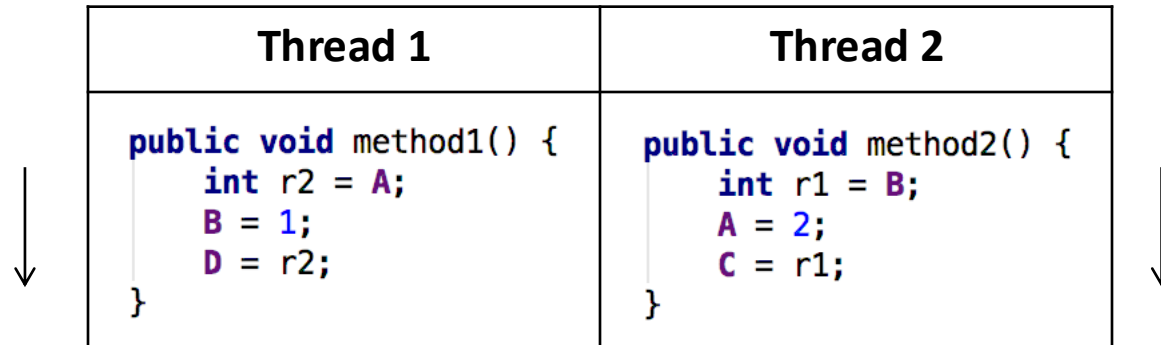
david.sislak@fel.cvut.cz

- [1] Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming. Elsevier, 2008.
- [2] Fog, A.: The microarchitecture of Intel, AMD and VIA CPU, 2016.
- [3] Russell, K., Detlefs, D.: Eliminating Synchronization-Related Atomic Operations with Biased Locking and Bulk Rebiasing in OOPSLA'06. ACM, USA 2006.
- [4] Oaks, S.: Java Performance: 2nd Edition. O'Reilly, USA 2020.

- » Data races
 - Superscalar execution in CPUs
 - Memory barriers – volatile variables
- » Synchronization
 - Reentrant locks
- » Atomic operations
 - Java support
 - Atomic operations on arrays
 - Atomic complex types
- » Non-blocking algorithms
 - LIFO
 - ConcurrentHashMap

Data Races – Multi-threaded Environments

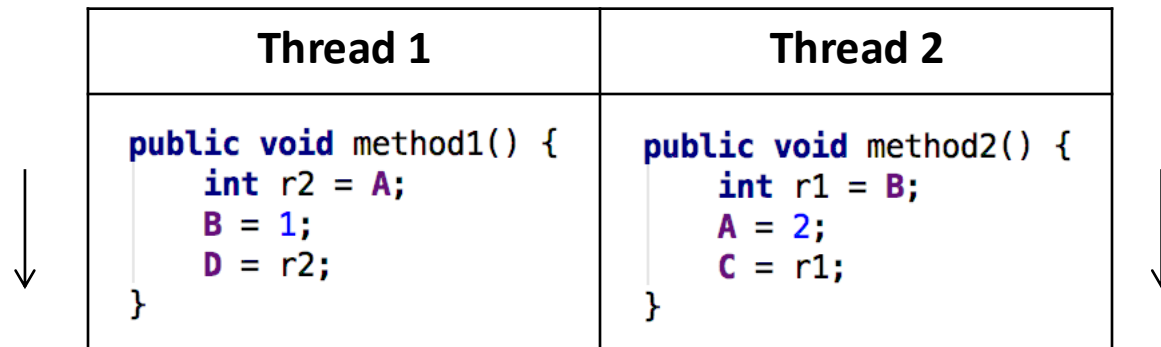
```
public int A = 0;
public int B = 0;
public int C = 0;
public int D = 0;
```



» What are the possible values of C and D?

Data Races – Multi-threaded Environments

```
public int A = 0;
public int B = 0;
public int C = 0;
public int D = 0;
```



» What are the possible values of C and D?


- C=0, D=0
- C=1, D=0
- C=0, D=2
- Any other possibilities?

Data Races – Disassembled Method and Assembly Code

```
public void method1() {
    int r2 = A;
    B = 1;
    D = r2;
}

0: aload_0
1: getfield      #2 // Field A:I
4: istore_1
5: aload_0
6: iconst_1
7: putfield     #3 // Field B:I
10: aload_0
11: iload_1
12: putfield     #5 // Field D:I
15: return
```

Heap object structure:



Klass – internal JVM representation of class Metadata
4B – 32bit, or 64bit <32GB heap
8B – 64bit no compressed OOP

Instructions reordered by the C2 compiler:

RSI points to the object instance

```
0x000000010639924c: movl    $0x1,0x10(%rsi)    ;*putfield B
                                ; - datarace.DataRace::method1@7 (line 11)

0x0000000106399253: mov     0xc(%rsi),%r11d
0x0000000106399257: mov     %r11d,0x18(%rsi)  ;*putfield D
                                ; - datarace.DataRace::method1@12 (line 12)
```

Note: All machine code examples are from JVM 8 (64-bit, <32GB heap), Intel Haswell CPU using AT&T syntax


- » **The same reordering happens in method2, resulting in the fourth outcome**
 - **C=1, D=2**

Data Races – CPU Execution Pipelining

- » Simplified instruction pipelining in a **single core**

IF: Instruction fetch
EX : Execution
WB : Write back

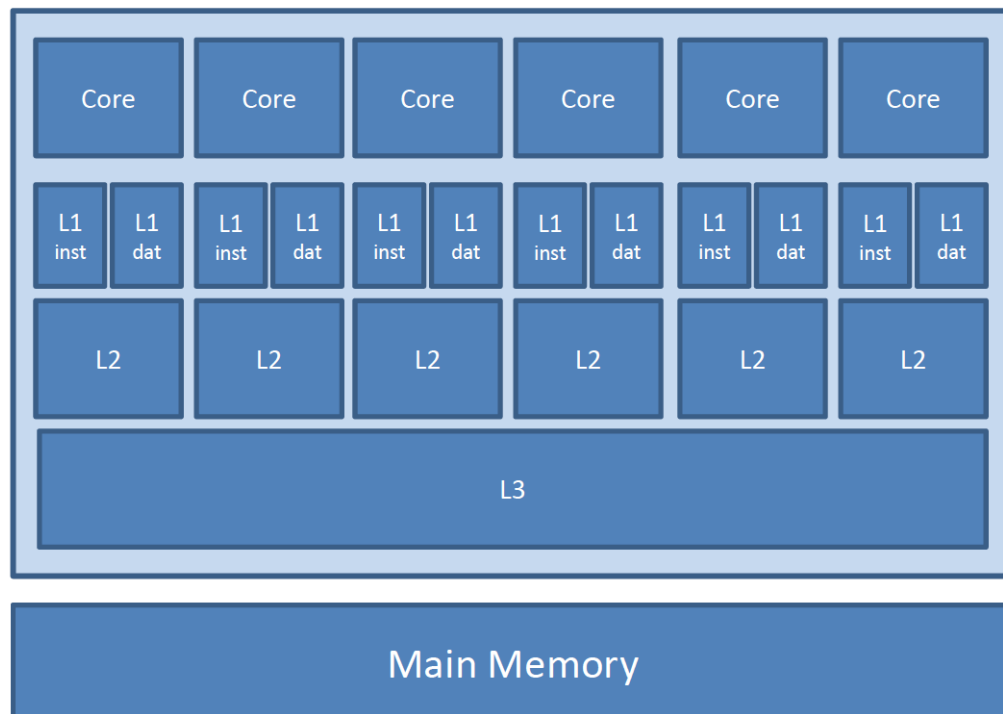
ID : Instruction decode
MEM: Memory access



- » Each stage is also parallelized (e.g., Haswell can execute up to 4 instructions per cycle) - execution depends on type and independence of instructions

Data Races – CPU Memory Model

» CPU vs. core vs. thread

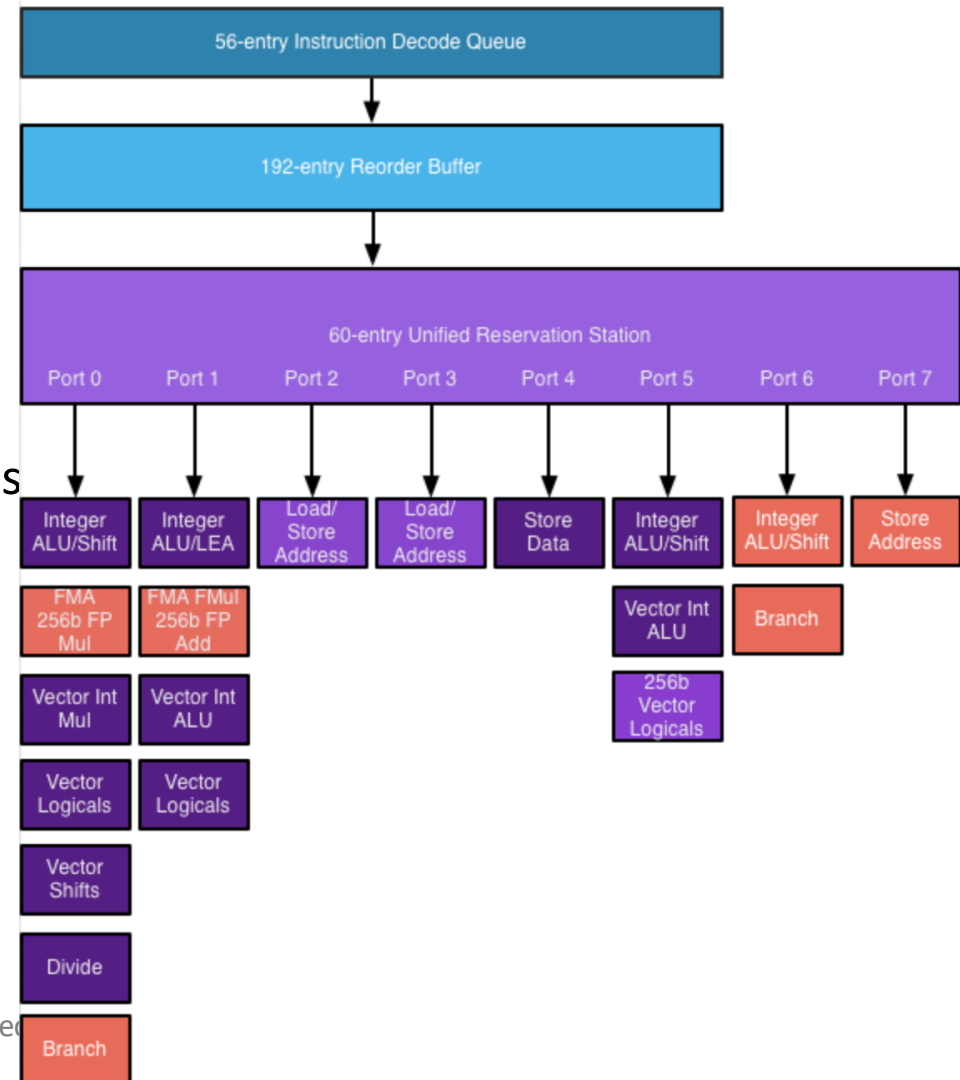


| L1 Data Cache | | | |
|----------------------|-----------|---------|---------------|
| Size | Line Size | Latency | Associativity |
| 32 KB | 64 bytes | 4 ns | 8-way |
| L1 Instruction Cache | | | |
| Size | Line Size | Latency | Associativity |
| 32 KB | 64 bytes | 4 ns | 4-way |
| L2 Cache | | | |
| Size | Line Size | Latency | Associativity |
| 256 KB | 64 bytes | 10 ns | 8-way |
| L3 Cache | | | |
| Size | Line Size | Latency | Associativity |
| 12 MB | 64 bytes | 50 ns | 16-way |
| Main Memory | | | |
| Size | Line Size | Latency | Associativity |
| | 64 bytes | 75 ns | |

- » All writes to main memory are done in **write-back** cache mode
- Standard writes require data to be cached (expensive cache miss)
 - Non-temporal writes (especially useful for large block writes)
 - Data is written directly to memory without being cached
 - Prefetch instructions are available

Data Races – CPU Execution Pipelining – Superscalar Execution

- » Modern CPUs have multiple execution units **in each core** (8 in Intel Haswell)
 - Units have various capabilities (4x integer ALU, 2x FPU mul, 2x mem read, ...)
 - Multiple **μops** with different latencies are executed **in parallel** in each cycle
- » Independent instructions can be **executed out-of-order** or in parallel
 - As long as they do not use the same register or memory address
- » **Memory reads are not reordered**
 - Parallel independent reads
- » **Later independent reads may be reordered and executed before writes**
 - Serialized writes only



Volatile Variable – Memory Barrier

Making A and B volatile:

```
public volatile int A = 0;
public volatile int B = 0;
public int C = 0;
public int D = 0;

public void method1() {
    int r2 = A;
    B = 1;
    D = r2;
}
```

Results in the following assembly code:

```
0x000000010710e08c: mov    0xc(%rsi),%r11d
0x000000010710e090: movl   $0x1,0x10(%rsi)
0x000000010710e097: lock  addl $0x0,(%rsp)
0x000000010710e09c: mov    %r11d,0x18(%rsi)
```

| |
|----------------------|
| 8B - mark word |
| 4B / 8B – Klass ref. |
| ... object data |

- » Memory operations around write to volatile var are **not reordered** in C1/C2
- » The **lock prefix** prevents instruction reordering and ensures all previous writes are **visible** to other CPUs
- » *lock addl \$0x0,(%rsp)* is the fastest **write memory barrier** (no actual operation)
- » No need for **read barriers** – not reordered during CPU execution

Volatile Variable

- » **Never cached locally in a thread** – all accesses go directly to main memory
- » Guarantees **atomic read and write** operations (defines write memory barrier)
- » Can be used for both primitives and references to objects
- » Does not block thread execution
- » BUT:
 - Volatile writes are significantly slower due to cache flushing (~100×)
 - Volatile reads (**when writes occur**) are also slower (~25×, depending on CPU cores)
 - due to invalidated cache
 - Still faster than synchronization/locks
- » Not necessary for:
 - Immutable objects
 - Variable accessed by only one thread (context switches already flush the cache)
 - Where variable is within complex synchronized operation

Counter Example - Volatile

```
public class VolatileCounter {  
    private volatile int cnt=0;  
  
    public int get() {  
        return cnt;  
    }  
  
    public void increment() {  
        cnt++;  
    }  
}
```

» Will it work as expected in a multi-threaded environment?

Counter Example - Volatile

```
public class VolatileCounter {  
    private volatile int cnt=0;  
  
    public int get() {  
        return cnt;  
    }  
  
    public void increment() {  
        cnt++;  
    }  
}
```

increment assembly code:

RSI is this

```
0x0000000010911544c: mov    0xc(%rsi),%edi  
  
0x0000000010911544f: inc   %edi  
0x00000000109115451: mov   %edi,0xc(%rsi)  
0x00000000109115454: lock addl $0x0,(%rsp)
```

| |
|----------------------|
| 8B - mark word |
| 4B / 8B – Klass ref. |
| ... object data |

» Will it work as expected in a multi-threaded environment?

NO

» **volatile**

- **Not suitable for read-update-write operations**

- **Useful for one-thread write** (e.g. termination flag)

- must be used if the flag is set by a different thread otherwise C2 compiler could create **infinite loop** without testing

Volatile Arrays

```
public class VolatileIntArray {  
    private volatile int[] array;  
  
    public VolatileIntArray(int capacity) {  
        array = new int[capacity];  
    }  
  
    public int get(int index) {  
        return array[index];  
    }  
  
    public void put(int index, int value) {  
        array[index] = value;  
    }  
}
```

» Is a write (**put operation**) to an array element handled as volatile?

Volatile Arrays

```
public class VolatileIntArray {
    private volatile int[] array;

    public VolatileIntArray(int capacity) {
        array = new int[capacity];
    }

    public int get(int index) {
        return array[index];
    }

    public void put(int index, int value) {
        array[index] = value;
    }
}
```

| |
|----------------------|
| 8B - mark word |
| 4B / 8B – Klass ref. |
| 4B – array length |
| sequence of values |

» Is a write (**put operation**) to an array element handled as volatile?

NO – as shown in the assembly code, there is no cache synchronization via a lock

```
# this:    rsi:rsi    = 'datarace/VolatileIntArray'
# parm0:   rdx       = int
# parm1:   rcx       = int
```

```
0x0000000011170bbcc: mov     0xc(%rsi),%esi
0x0000000011170bbcf: shl     $0x3,%rsi          ;*getfield array
                                ; - datarace.VolatileIntArray::put@1 (line 15)
```

```
0x0000000011170bbd3: movslq %edx,%rdi
0x0000000011170bbd6: cmp     0xc(%rsi),%edx     ; implicit exception: dispatches to 0x0000000011170bbef
0x0000000011170bbd9: jae     0x0000000011170bbf9 — ArrayOutOfBoundsException
0x0000000011170bbdf: mov     %ecx,0x10(%rsi,%rdi,4) ;*iastore
                                ; - datarace.VolatileIntArray::put@6 (line 15)
```

Volatile Arrays - Solution

```
private volatile int[] array;  
public void put(int index, int value) {  
    array[index] = value;  
    array = array;  
}
```

| |
|----------------------|
| 8B - mark word |
| 4B / 8B – Klass ref. |
| ... object data |

- » Only the array reference is volatile
- » An additional (unnecessary) **array** reference update introduces extra assembly instructions

```
0x000000010db21a67: mov    %r8d,0xc(%rsi)
```

```
0x000000010db21a80: lock addl $0x0,(%rsp)    ;*putfield array  
                        ; - datarace.VolatileIntArray::put@12 (line 16)
```

- » Instruction **lock prefix** forbids all instruction reordering around it and synchronizes previous writes to be visible by all other CPUs
- » **Not suitable for read-update-write operations**

Counter Example – Synchronized and ReentrantLock

```
public class SynchronizedCounter {
    private int cnt=0;

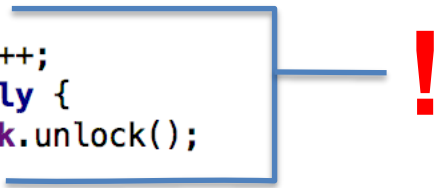
    public int get() {
        return cnt;
    }

    public synchronized void increment() {
        cnt++;
    }
}
```

```
public class ReentrantCounter {
    private int cnt=0;
    private ReentrantLock lock = new ReentrantLock();

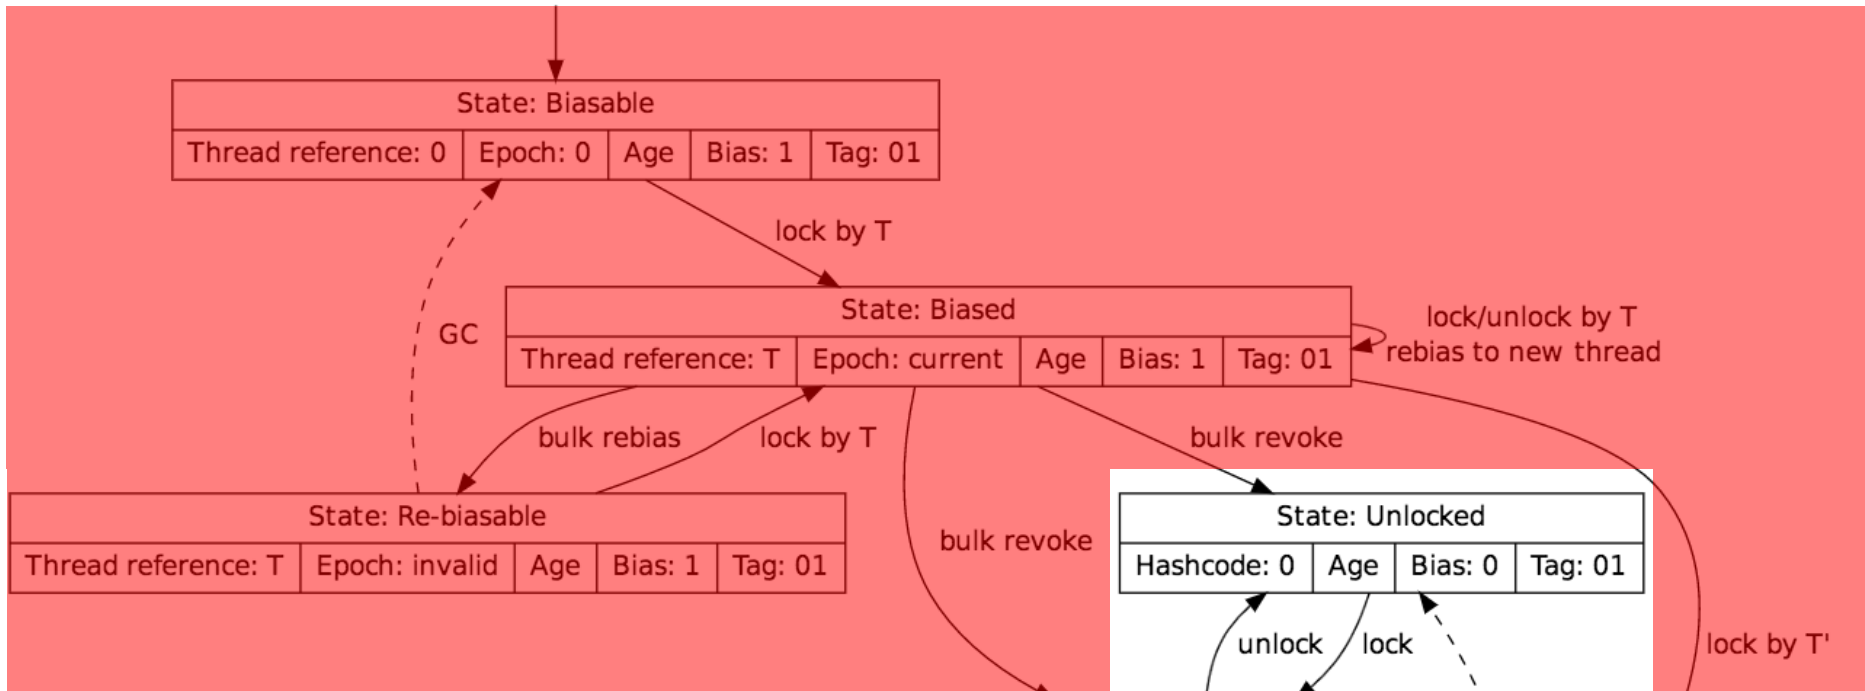
    public int get() {
        return cnt;
    }

    public void increment() {
        lock.lock();
        try {
            cnt++;
        } finally {
            lock.unlock();
        }
    }
}
```



- » No issues with **read-modify-write** operations
- » synchronized
 - method vs. block
 - object instance vs. class instance (static methods)

JVM - Synchronize Implementation



removed since Java 15
» assembly code optimized for thin locking

Reentrant Locks

- » Provides extended locking operations compared to **synchronized**
 - lock(), unlock()
 - lockInterruptibly() throws InterruptedException
 - boolean tryLock()
 - boolean tryLock(long timeout, TimeUnit unit) throws InterruptedException
- » **Fairness**
 - Blocked threads are **ordered** for fair locking
 - **new** ReentrantLock(boolean fair), by default unfair
 - **synchronized** is unfair
 - Unfair ReentrantLocks are slightly faster than synchronized
 - but another instance in HEAP
 - Fair locks are slower (~100x)

Counter Example – AtomicInteger

```
public class AtomicCounter {
    private AtomicInteger cnt = new AtomicInteger( initialValue: 0);

    public int get() {
        return cnt.get();
    }

    public void increment() {
        cnt.incrementAndGet();
    }
}
```

AtomicInteger implementation

```
private static final long valueOffset;
```

```
static {
    try {
        valueOffset = unsafe.objectFieldOffset
            (AtomicInteger.class.getDeclaredField( name: "value"));
    } catch (Exception ex) { throw new Error(ex); }
}
```

```
private volatile int value;
```

```
public final int getAndAddInt(Object var1, long var2, int var4) {
    int var5;
    do {
        var5 = this.getIntVolatile(var1, var2);
    } while(!this.compareAndSwapInt(var1, var2, var5, var5: var5 + var4));

    return var5;
}
```

```
public final int getAndIncrement() {
    return unsafe.getAndAddInt( o: this, valueOffset, i: 1);
}
```




**Non-blocking
pattern
(CAS loop)**

Counter Example – AtomicInteger – Assembly Code

C2 compiler assembly code for AtomicInteger::increment

RSI is this, R12=0

```
0x0000000010b108d4c: mov    0xc(%rsi),%r11d    ;*getfield cnt
                                ; - datarace.AtomicCounter::increment@1 (line 13)
0x0000000010b108d50: test  %r11d,%r11d
0x0000000010b108d53: je    0x0000000010b108d68
0x0000000010b108d55: lock addl $0x1,0xc(%r12,%r11,8) ;*invokevirtual getAndAddInt
                                ; - java.util.concurrent.atomic.AtomicInteger::incrementAndGet@8 (line 186)
                                ; - datarace.AtomicCounter::increment@4 (line 13)
```

 **null pointer check with exception**

- » While loop optimized and replaced with **single instruction**
- » Instruction **lock prefix** forbids all reordering around and synchronizes previous writes to be visible by all other CPUs
- » Instruction **lock prefix** ensures that core has exclusive ownership of the appropriate cache line for the duration of the operation
 - Cache coherency using **MESIF** (Haswell) with fallback to mem bus lock
- » **AtomicInteger-based counter is fastest of all for multi-threaded usage**

Atomic Operations

- » 32-bit CPUs support 64-bit CAS operations
 - **cmpxchg** src_operand, dst_operand – implicit instruction lock prefix
- » 64-bit CPUs support 128-bit CAS operations
 - **cmpxchg16b** – works with RDX:RAX and RCX:RBX register pairs
- » Java uses only 64-bit CAS operations in `java.util.concurrent.atomic`
 - AtomicBoolean
 - AtomicInteger
 - AtomicLong
 - AtomicReference
 - AtomicIntegerArray
 - AtomicLongArray
 - AtomicReferenceArray

Atomic Field Updaters

- » **Suitable for a large number of objects** of the given type – it saves memory
 - Do not require a single instance to have an extra object embedded
- » Refer volatile variable directly without getter and setters

```
public class ObjectWithAtomic {
    private final AtomicInteger value =
        new AtomicInteger(0);
    // ...

    public void method1() {
        // ...
        if (value.compareAndSet(1, 2)) {
            // ...
        }
    }
}
```



```
public class ObjectWithAtomic {
    private static AtomicIntegerFieldUpdater<ObjectWithAtomic>
        valueUpdater = AtomicIntegerFieldUpdater.newUpdater(ObjectWithAtomic.class, "value");
    private volatile int value = 0;
    // ...

    public void method1() {
        // ...
        if (valueUpdater.compareAndSet(this, 1, 2)) {
            // ...
        }
    }
}
```

Atomic Field Updaters

- » But **less efficient** operations for atomic field updaters
- » AtomicIntegerFieldUpdater implementation

```
private void fullCheck(T obj) {
    if (!tclass.isInstance(obj))
        throw new ClassCastException();
    if (cclass != null)
        ensureProtectedAccess(obj);
}

public boolean compareAndSet(T obj, int expect, int update) {
    if (obj == null || obj.getClass() != tclass || cclass != null) fullCheck(obj);
    return unsafe.compareAndSwapInt(obj, offset, expect, update);
}
```

- » Existing field updaters
 - AtomicIntegerFieldUpdater
 - AtomicLongFieldUpdater
 - AtomicReferenceFieldUpdater
- » No array field updaters

Atomic Complex Types

- » **AtomicMarkableReference**
 - **object reference** along with a **mark bit**
- » **AtomicStampedReference**
 - **object reference** along with an **integer “stamp”**
- » Notes:
 - Useful for **ABA problem**
 - Change A -> B and then B -> A
 - How can I know that A has been changed since the last observation?
 - Does not use double-wide CAS (CAS2, CASX) -> much slower than simple atomic types due to **object allocation**

Atomic Complex Types – Larger Than 64-bits

» AtomicMarkableReference

- object reference along with a **mark bit**

» AtomicStampedReference

- object reference along with an **integer “stamp”**

```
public class AtomicStampedReference<V> {  
  
    private static class Pair<T> {  
        final T reference;  
        final int stamp;  
        private Pair(T reference, int stamp) {  
            this.reference = reference;  
            this.stamp = stamp;  
        }  
        static <T> Pair<T> of(T reference, int stamp) {  
            return new Pair<T>(reference, stamp);  
        }  
    }  
  
    private volatile Pair<V> pair;  
  
    public boolean compareAndSet(V expectedReference,  
                                 V newReference,  
                                 int expectedStamp,  
                                 int newStamp) {  
        Pair<V> current = pair;  
        return  
            expectedReference == current.reference &&  
            expectedStamp == current.stamp &&  
            ((newReference == current.reference &&  
              newStamp == current.stamp) ||  
             casPair(current, Pair.of(newReference, newStamp)));  
    }  
}
```

Non-blocking Algorithms

- » **Lock-free**, but typically not **wait-free** (due to unbounded retry loops)
 - based on CAS / CMPXCHG and LOCK prefixed instructions
- » Shared resources secured by locks have drawbacks
 - High-priority thread can be blocked (e.g. interrupt handler)
 - Parallelism reduced by coarse-grained locking (unfair locks)
 - Fine-grained locking and fair locks increases overhead
 - Can lead to **deadlocks**, **priority inversion** (low-priority thread holds a shared resource which is required by high-priority thread)
- » **Non-blocking algorithms properties:**
 - Outperform blocking algorithms because most of CAS / CMPXCHG succeeds on the first try
 - Removes cost for synchronization, thread suspension, context switching
- » Note: **Real-time systems require wait-free algorithms** (finite number of steps) and lock-free is not sufficient

Non-blocking stack (LIFO)

» Treiber's algorithm (1986)

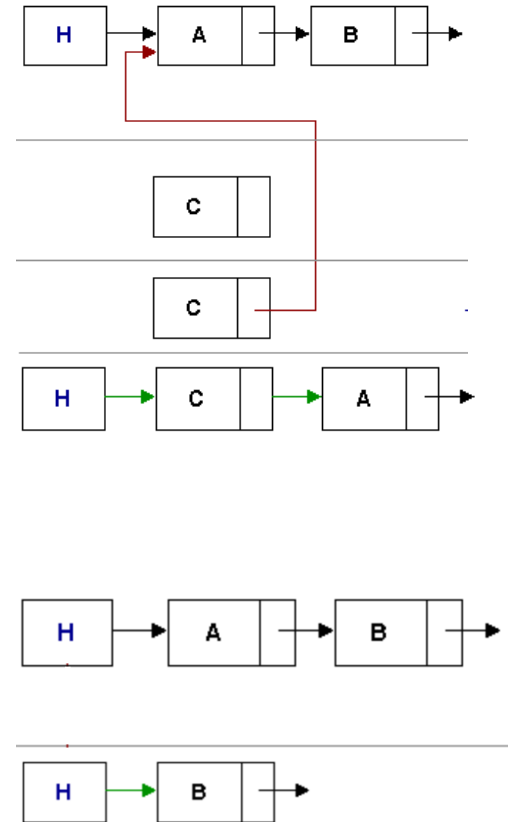
```
static class Node<E> {
    final E item;
    Node<E> next;

    public Node(E item) { this.item = item; }
}

AtomicReference<Node<E>> head = new AtomicReference<Node<E>>();

public void push(E item) {
    Node<E> newHead = new Node<E>(item);
    Node<E> oldHead;
    do {
        oldHead = head.get();
        newHead.next = oldHead;
    } while (!head.compareAndSet(oldHead, newHead));
}

public E pop() {
    Node<E> oldHead;
    Node<E> newHead;
    do {
        oldHead = head.get();
        if (oldHead == null)
            return null;
        newHead = oldHead.next;
    } while (!head.compareAndSet(oldHead, newHead));
    return oldHead.item;
}
```



A push after a pop can cause the ABA problem if the address is reused!

Thread-safe collections and maps

» Blocking collections and maps

- `static<T> Collection<T> Collections.synchronizedCollection(Collection<T> c)`
- `static<T> List<T> Collections.synchronizedList(List<T> list)`
- `static<K,V> Map<K,V> Collections.synchronizedMap(Map<K,V> m)`
- `static<T> Set<T> Collections.synchronizedSet(Set<T> s)`
- also for `SortedSet` and `SortedMap`

» Non-blocking collections and maps

- `ConcurrentLinkedQueue` (interface `Collection`, `Queue`):
 - `E peek()`, `E poll()`, `offer(E)` in FIFO manner
- `ConcurrentLinkedDeque` (interface `Collection`, `Deque`):
 - allow offering, polling and peaking at both ends of the linear collection
- `ConcurrentHashMap` (interface `Map`):
 - `putIfAbsent(K key, V value)`, `remove(Object key, Object value)`
 - `replace(K key, V oldValue, V newValue)`
- `ConcurrentSkipListMap` (interface `SortedMap`), `ConcurrentSkipListSet` (interface `SortedSet`)

» Non-blocking collections and maps are slower for single-threaded access

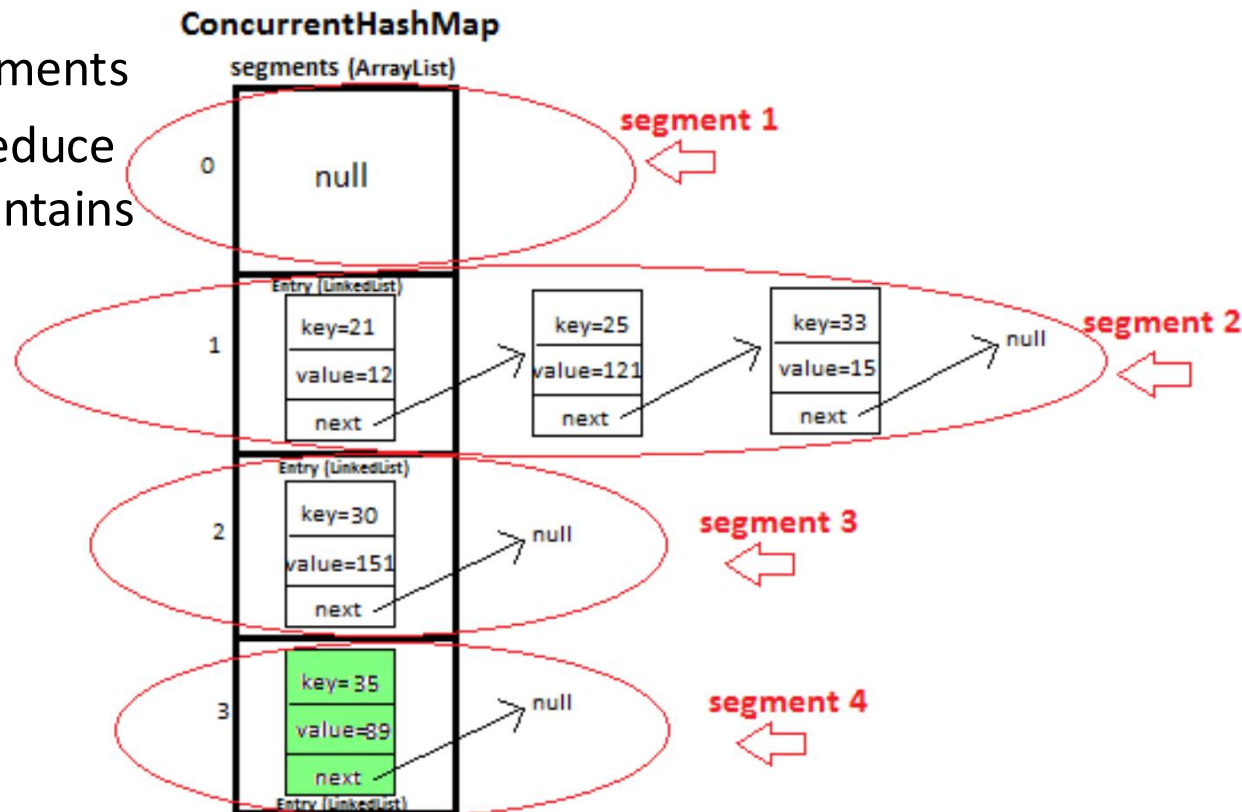
- due to usage of CAS instructions

ConcurrentHashMap

- » **Concurrent reads** – get, iterator
- » **Minimize update contention**
 - Initial concurrency level 16 (can be changed) – number of updating threads
 - Initial insertion into empty segment uses CAS operation
 - Later modifications are based on segment-based locks

- » **Segment contention**

- Use lists for <8 elements
- Balanced tree to reduce search times – maintains next for iteration



ConcurrentHashMap

- » **Table resizing** (occupancy exceeds the load factor 0.75)
 - Power of two expansions
 - Same index or power of two index
 - Reusing internal Node if next is not changed – majority of cases
 - Any thread can help resizing instead of block
 - **Forward nodes** are used to notify users about moved
- » Provide **initialCapacity** if estimate is known