

---

# COMP 322: Fundamentals of Parallel Programming

## Lecture 19: Java Atomic Variables — a special case of isolated

Vivek Sarkar  
Department of Computer Science  
Rice University  
vsarkar@rice.edu



# Acknowledgments for Today's Lecture

---

- Lecture 19 handout



# HJ isolated statement (Recap)

---

## isolated <body>

- Two tasks executing isolated statements with interfering accesses must perform the isolated statement in mutual exclusion
  - Two instances of isolated statements,  $\langle \text{stmt1} \rangle$  and  $\langle \text{stmt2} \rangle$ , are said to interfere with each other if both access a shared location, such that at least one of the accesses is a write.
  - Weak isolation guarantee: no mutual exclusion applies to non-isolated statements i.e., to (isolated, non-isolated) and (non-isolated, non-isolated) pairs of statement instances
- Isolated statements may be nested (redundant)
- Isolated statements must not contain any other parallel statement: *async*, *finish*, *get*, *forall*
- In case of exception, all updates performed by  $\langle \text{body} \rangle$  before throwing the exception will be observable after exiting  $\langle \text{body} \rangle$



# DoublyLinkedListNode example (Recap)

---

```
1. class DoublyLinkedListNode {
2.     DoublyLinkedListNode prev, next;
3.     . . .
4.     void delete() {
5.         isolated { // start of mutual exclusion region (critical section)
6.             if (this.prev != null) this.prev.next = this.next;
7.             if (this.next != null) this.next.prev = this.prev
8.         } // end of mutual exclusion region (critical section)
9.         . . .
10.    }
11.    . . .
12.}
13. . . .
14. static void deleteTwoNodes(DoublyLinkedListNode n1, n2) {
15.     finish {
16.         async n1.delete();
17.         async n2.delete();
18.     }
19. }
```



# Implementations of isolated statement

---

- isolated statements are convenient for the programmer but pose significant challenges for the language implementation
  - Implementation does not know ahead of time if two dynamic instances of isolated statements will interfere or not
- HJ implementation used in *COMP 322* takes a simple single-lock approach to implementing isolated statements
  - Entry to isolated statement is treated as an `acquire()` operation on the lock
  - Exit from isolated statement is treated as a `release()` operation on the lock
  - Though correct, this approach essentially implements isolated statements as critical sections, thereby serializing all interfering and non-interfering isolated statement instances.
- How can we do better?



# Transactional Memory (TM)

---

- Execution of an isolated statement is treated as a transaction
  - In database systems, a transaction refers to a “unit of work” that has “all-or-nothing” semantics. Each unit of work must either complete in its entirety or have no visible effect.
- A TM system logs all read and write operations performed in a transaction and optimistically permits transactions to run in parallel, speculating that there won't be interference
- At the end of a transaction, a TM system checks if interference occurred with another transaction
  - If not, the transaction can be committed
  - If so, the transaction fails and has to be “retried”
- Both software and hardware implementations of TM have been explored extensively by the research community, but no implementation is suitable for mainstream use as yet



# Three cases of contention among isolated statements

---

1. **Low contention:** when isolated statements are executed infrequently
  - A single-lock approach as in HJ is often the best solution. No visible benefit from other techniques because they incur overhead that is not needed since contention is low.
2. **Moderate contention:** when the serialization of all isolated statements in a single-lock approach limits the performance of the parallel program due to Amdahl's Law, but a finer-grained approach that only serializes interfering isolated statements results in good scalability
  - Atomic variables usually do well in this scenario since the benefit obtained from reduced serialization far outweighs any extra overhead incurred.
3. **High contention:** when interfering isolated statements dominate the program execution time in certain phases
  - Best approach in such cases is to find an alternative algorithm to using isolated



# java.util.concurrent

---

Sub-packages include

- Atomic variables
  - Efficient implementations of special-case patterns of isolated statements
- Concurrent Collections:
  - Queues, blocking queues, concurrent hash map, ...
  - Data structures designed for concurrent environments
- Executors, Thread pools and Futures
  - Execution frameworks for asynchronous tasking
- Locks and Conditions
  - More flexible synchronization control
  - Read/write locks
- Synchronizers: Semaphore, Latch, Barrier, Exchanger, Phaser
  - Tools for thread coordination





# Table 1: Methods in java.util.concurrent atomic classes AtomicInteger and AtomicIntegerArray

| j.u.c.atomic Class and Constructors  | j.u.c.atomic Methods  | Equivalent HJ isolated statements                    |
|--|---|--|
| <b>AtomicInteger</b><br><br><b>AtomicInteger()</b><br>// init = 0<br><br><b>AtomicInteger(init)</b>                          | int j = v.get();  | int j; isolated j = v.val;                           |
|  | v.set(newVal);  | isolated v.val = newVal;                             |
|  | int j = v.getAndSet(newVal);  | int j; isolated { j = v.val; v.val = newVal; }       |
|  | int j = v.addAndGet(delta);   | isolated { v.val += delta; j = v.val; }              |
|  | int j = v.getAndAdd(delta);   | isolated { j = v.val; v.val += delta; }              |
| boolean b =<br>v.compareAndSet<br>(expect,update);   | boolean b;<br>isolated<br>if (v.val==expect) {v.val=update; b=true;}<br>else b = false;       |  |
| <b>AtomicIntegerArray</b><br><br><b>AtomicIntegerArray</b><br>(length) // init = 0<br><br><b>AtomicIntegerArray</b><br>(arr) | int j = v.get(i);   | int j; isolated j = v.arr[i];                        |
|  | v.set(i,newVal);  | isolated v.arr[i] = newVal;                          |
|  | int j = v.getAndSet(i,newVal);  | int j; isolated { j = v.arr[i]; v.arr[i] = newVal; } |
|  | int j = v.addAndGet(i,delta);   | isolated { v.arr[i] += delta; j = v.arr[i]; }        |
|  | int j = v.getAndAdd(i,delta);   | isolated { j = v.arr[i]; v.arr[i] += delta; }        |
| boolean b =<br>v.compareAndSet<br>(i,expect,update);   | boolean b;<br>isolated<br>if (v.arr[i]==expect) {v.arr[i]=update; b=true;}<br>else b = false; |  |



## Table 2: Examples of common isolated statement idioms and their equivalent AtomicInteger implementations

|   |   |
|---|---|
| <p><b>1) Rank computation:</b><br/> <code>rank = new ...; rank.count = 0;</code><br/> <code>. . .</code><br/> <code>isolated r = ++rank.count;</code></p>   | <pre>AtomicInteger rank = new AtomicInteger(); . . . r = rank.incrementAndGet();</pre>  |
| <p><b>2) Work assignment:</b><br/> <code>rem = new ...; rem.count = n;</code><br/> <code>. . .</code><br/> <code>isolated r = rem.count--;</code><br/> <code>if ( r &gt; 0 ) . . .</code></p>   | <pre>AtomicInteger rem = new AtomicInteger(n); . . . r = rem.getAndDecrement(); if ( r &gt; 0 ) . . .</pre>   |
| <p><b>3) Counting semaphore:</b><br/> <code>sem = new ...; sem.count = 0;</code><br/> <code>. . .</code><br/> <code>isolated r = ++sem.count;</code><br/> <code>. . .</code><br/> <code>isolated r = --sem.count;</code><br/> <code>. . .</code><br/> <code>isolated s = sem.count; isZero = (s==0);</code></p> | <pre>AtomicInteger sem = new AtomicInteger(); . . . r = sem.incrementAndGet(); . . . r = sem.decrementAndGet(); . . . s = sem.get(); isZero = (s==0);</pre> |
| <p><b>4) Sum reduction:</b><br/> <code>sum = new ...; sum.val = 0;</code><br/> <code>. . .</code><br/> <code>isolated sum.val += x;</code></p>  | <pre>AtomicInteger sum = new AtomicInteger(); . . . sum.addAndGet(x);</pre>   |



## Table 3: Methods in java.util.concurrent atomic classes AtomicReference and AtomicReferenceArray

| j.u.c.atomic Class and Constructors   | j.u.c.atomic Methods   | Equivalent HJ isolated statements  |
|---|--|--|
| <b>AtomicReference</b><br><br><b>AtomicReference()</b><br>// init = null<br><br><b>AtomicReference(init)</b>                        | Object o = v. <b>get</b> ();                                 | Object o; <b>isolated</b> o = v.ref;   |
|   | v. <b>set</b> (newRef);                                      | <b>isolated</b> v.ref = newRef;  |
|   | Object o =<br>v. <b>getAndSet</b> (newRef);                  | Object o;<br><b>isolated</b> { o = v.ref; v.ref = newRef; }  |
|   | boolean b =<br>v. <b>compareAndSet</b><br>(expect,update);   | boolean b;<br><b>isolated</b><br>if (v.ref==expect) {v.ref=update; b=true;}<br>else b = false;       |
| <b>AtomicReferenceArray</b><br><br><b>AtomicReferenceArray</b><br>(length) // init = null<br><br><b>AtomicIntegerArray</b><br>(arr) | Object o = v. <b>get</b> (i);                                | Object o; <b>isolated</b> o = v.arr[i];  |
|   | v. <b>set</b> (i,newRef);                                    | <b>isolated</b> v.arr[i] = newRef;   |
|   | Object o =<br>v. <b>getAndSet</b> (i,newRef);                | Object o;<br><b>isolated</b> { o = v.arr[i]; v.arr[i] = newRef; }                                    |
|   | boolean b =<br>v. <b>compareAndSet</b><br>(i,expect,update); | boolean b;<br><b>isolated</b><br>if (v.arr[i]==expect) {v.arr[i]=update; b=true;}<br>else b = false; |




# Parallel Depth-First Search Spanning Tree Example revisited

```

1. class V {
2.   V [] neighbors; // adjacency list for input
   graph
3.   V parent;      // output value of parent in
   spanning tree
4.   boolean tryLabeling(V n) {
5.     isolated if (parent == null) parent=n;
6.     return parent == n;
7.   } // tryLabeling
8.   void compute() {
9.     for (int i=0; i<neighbors.length; i++) {
10.      V child = neighbors[i];
11.      if (child.tryLabeling(this))
12.        async child.compute(); //escaping async
13.    }
14.  } // compute
15.} // class V
16.. . .
17.root.parent = root; // Use self-cycle to identify
   root
18.finish root.compute();
19.. . .

```



# Parallel Depth-First Search Spanning Tree Example revisited

---

```
1. class V {
2.     V [] neighbors; // adjacency list for input graph
3.     AtomicReference parent; // output value of parent in
   spanning tree
4.     boolean tryLabeling(V n) {
5.         return parent.compareAndSet(null ,n);
6.     }
7. } // tryLabeling
8. void compute() {
9.     for (int i=0; i<neighbors.length; i++) {
10.        V child = neighbors[i];
11.        if (child.tryLabeling(this))
12.            async child.compute(); //escaping async
13.    }
14. } // compute
15.} // class V
16.. . .
17.root.parent = root; // Use self-cycle to identify root
18.finish root.compute();
19.. . .
```

