

---

# COMP 322: Fundamentals of Parallel Programming

## Lecture 22: Read-write Isolation, Atomic Variables

Vivek Sarkar, Eric Allen  
Department of Computer Science, Rice University

Contact email: [vsarkar@rice.edu](mailto:vsarkar@rice.edu)

<https://wiki.rice.edu/confluence/display/PARPROG/COMP322>

---

COMP 322

Lecture 22

13 March 2015



---

## HJ isolated construct (Recap)

---

`isolated (() -> <body> );`

- Isolated construct identifies a critical section
- Two tasks executing isolated constructs are guaranteed to perform them in mutual exclusion
  - Isolation guarantee applies to (isolated, isolated) pairs of constructs, not to (isolated, non-isolated) pairs of constructs
- Isolated constructs may be nested
  - An inner isolated construct is redundant
- Blocking parallel constructs are forbidden inside isolated constructs
  - Isolated constructs must not contain any parallel construct that performs a blocking operation e.g., `finish`, `future get`, `next`
  - Non-blocking async operations are permitted, but isolation guarantee only applies to creation of async, not to its execution
- Isolated constructs can never cause a deadlock
  - Other techniques used to enforce mutual exclusion (e.g., locks) can lead to a deadlock, if used incorrectly



# Object-based isolation (Recap)

---

`isolated(obj1, obj2, ..., () -> <body>)`

- In this case, programmer specifies list of objects for which isolation is required
- Mutual exclusion is only guaranteed for instances of isolated constructs that have a common object in their object lists

—Standard isolated is equivalent to “isolated(\*)” by default i.e., isolation across all objects

- Example:

- `isolated(a,b,()->{..})` and `isolated(c,d,()->{..})` can execute in parallel
- `isolated(a,b,()->{..})` and `isolated(b,c,()->{..})` cannot execute in parallel



## DoublyLinkedListNode Example revisited with Object-Based Isolation

---

```
1. class DoublyLinkedListNode {
2.     DoublyLinkedListNode prev, next;
3.     . . .
4.     void delete() {
5.         isolated(this.prev, this, this.next, () -> { // object-based isolation
6.             this.prev.next = this.next;
7.             this.next.prev = this.prev;
8.         });
9.     . . .
10. }
11. } // DoublyLinkedListNode
12. . . .
13. static void deleteTwoNodes(final DoublyLinkedListNode L) {
14.     finish(() -> {
15.         DoublyLinkedListNode second = L.next;
16.         DoublyLinkedListNode third = second.next;
17.         async(() -> { second.delete(); });
18.         async(() -> { third.delete(); });
19.     });
20. }
```



# Read-Write Object-based isolation in HJ

```
isolated(readMode(obj1),writeMode(obj2), ..., () -> <body> );
```

- Programmer specifies list of objects as well as their read-write modes for which isolation is required
- Not specifying a mode is the same as specifying a write mode (default mode = read + write)
- Mutual exclusion is only guaranteed for instances of isolated statements that have a non-empty intersection in their object lists such that one of the accesses is in writeMode
- Sorted List example

```
1. public boolean contains(Object object) {
2.     return isolatedWithReturn( readMode(this), () -> {
3.         Entry pred, curr;
4.         ...
5.         return (key == curr.key);
6.     });
7. }
8.
9. public int add(Object object) {
10.    return isolatedWithReturn( writeMode(this), () -> {
11.        Entry pred, curr;
12.        ...
13.        if (...) return 1; else return 0;
14.    });
15. }
```



## java.util.concurrent library

- 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
- WARNING: only a small subset of the full java.util.concurrent library can safely be used with HJlib
  - Atomic variables are part of the safe subset
  - We will study the full library later this semester as part of Java Concurrency



# java.util.concurrent.atomic.AtomicInteger

---

- Constructors
  - new `AtomicInteger()`
    - Creates a new `AtomicInteger` with initial value 0
  - new `AtomicInteger(int initialValue)`
    - Creates a new `AtomicInteger` with the given initial value
- Selected methods
  - int `addAndGet(int delta)`
    - Atomically adds `delta` to the current value of the atomic variable, and returns the new value
  - int `getAndAdd(int delta)`
    - Atomically returns the current value of the atomic variable, and adds `delta` to the current value
- Similar interfaces available for `LongInteger`

---

7

COMP 322, Spring 2015 (V.Sarkar, E.Allen)



## Work-Sharing Pattern using AtomicInteger

---

```
1. import java.util.concurrent.atomic.AtomicInteger;
2. . . .
3. String[] X = ... ; int numTasks = ...;
4. int[] taskId = new int[X.length];
5. AtomicInteger a = new AtomicInteger();
6. . . .
7. finish(() -> {
8.     for (int i=0; i<numTasks; i++ )
9.         async(() -> {
10.            do {
11.                int j = a.getAndAdd(1);
12.                // can also use a.getAndIncrement()
13.                if (j >= X.length) break;
14.                taskId[j] = i; // Task i processes string X[j]
15.                . . .
16.            } while (true);
17.        });
18.}); // finish-for-async
```

---

8

COMP 322, Spring 2015 (V.Sarkar, E.Allen)



## java.util.concurrent.AtomicInteger methods and their equivalent isolated constructs (pseudocode)

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

Methods in `java.util.concurrent.AtomicInteger` class and their equivalent HJ isolated statements. Variable `v` refers to an `AtomicInteger` object in column 2 and to a standard non-atomic Java object in column 3. `val` refers to a field of type `int`.



## java.util.concurrent.atomic.AtomicReference

- Constructors
  - `new AtomicReference()`
    - Creates a new `AtomicReference` with initial value 0
  - `new AtomicReference(Object init)`
    - Creates a new `AtomicReference` with the given initial value
- Selected methods
  - `int getAndSet(Object newRef)`
    - Atomically get current value of the atomic variable, and set value to `newRef`
  - `int compareAndSet(Object expect, Object update)`
    - Atomically check if current value = `expect`. If so, replace the value of the atomic variable by `update` and return true. Otherwise, return false.



## java.util.concurrent.AtomicReference methods and their equivalent isolated statements

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

Methods in java.util.concurrent.AtomicReference class and their equivalent HJ isolated statements. Variable v refers to an AtomicReference object in column 2 and to a standard non-atomic Java object in column 3. ref refers to a field of type Object.

AtomicReference<T> can be used to specify a type parameter.



## Parallel Spanning Tree Algorithm using AtomicReference

```
1. class V {
2.     V [] neighbors; // adjacency list for input graph
3.     AtomicReference<V> parent; // output value of parent in spanning tree
4.     boolean tryLabeling(final V n) {
5.         return parent.compareAndSet(null, n);
6.     };
7.     } // tryLabeling
8.     void compute() {
9.         for (int i=0; i<neighbors.length; i++) {
10.            final 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. . . .
```



## COMP 322 Worksheet 21 solution: Eureka-style Speculative Parallelism.

The code snippet below performs a eureka-style search on a 2-D array with a fixed number of tasks. Each task uses the `next()` operation to ensure the computation progresses in a lock step manner where all tasks execute one call to `check()` before performing the equality (`==`) comparison. What does the program print when it completes execution? Also, print the number of `==` comparisons performed by the program. Remember that once the search eureka has been resolved (via a call to `offer()`), subsequent calls to `check()` will cause the task to terminate.

```
final int numRows = 10;
final int numCols = 100;
final int[][] dataArray = new int[numRows][numCols];
for (int i = 0; i < numRows; i++) {
    for (int j = 0; j < numCols; j++) {
        dataArray[i][j] = 100 * i + j;
    }
}

final int searchElement = 625;
final HJSearchEureka<int[]> eureka = new SearchEureka();
finish(eureka, ()-> {
    forasynCPhased(0, numRows - 1, (i) -> {
        for (int j = 0; j < numCols; j++) {
            final int[] elemIndex = {i, j};
            eureka.check(elemIndex);
            next(); // barrier
            if (dataArray[i][j] == searchElement) {
                eureka.offer(elemIndex);
            }
            next(); // barrier
        } // for
    }); // forasynCPhased
}); // finish
final int[] index = eureka.get();
System.out.println("Result = " + Arrays.toString(index));
```

### Answer:

Result: [6, 25] (found by task i=6 in iteration j=25)

Number of `==` comparisons: 260

Due to the presence of the barriers, each of the 10 tasks performs a comparison operation before the next comparison operation is performed by any of the other tasks. Since the call to `offer()` occurs at the 26<sup>th</sup> iteration (j=25), the total number of comparisons is  $10 \times 26 = 260$ .