
COMP 322: Fundamentals of Parallel Programming

Module 2: Concurrency

©2017 by Vivek Sarkar

February 28, 2017

DRAFT VERSION – PLEASE DO NOT DISTRIBUTE

Contents

5 Mutual Exclusion	2
5.1 Critical Sections and the “isolated” statement	2
5.2 Object-based Isolation	4
5.3 Parallel Spanning Tree Example	7
5.4 Atomic Variables	10
5.5 Read-Write Modes in Object-based Isolation	13
5.6 Serialized Computation Graph for Isolated Statements	13

5 Mutual Exclusion

5.1 Critical Sections and the “isolated” statement

For the programming constructs `async`, `finish`, `future`, `get`, `forall`, the following situation was defined to be a *data race* error — when two accesses on the same shared location can potentially execute in parallel such that at least one access is a write. However, there are many cases in practice when two tasks may legitimately need to perform conflicting accesses to shared locations.

Consider the example shown below in Listing 1. The job of method `deleteTwoNodes()` in line 14 is to delete the first two nodes in doubly-linked list `L`. It does so by calling `L.delete()` and `L.next.delete()` in parallel in lines 16 and 17. Each call to `delete()` needs to perform the book-keeping operations in lines 6 and 7 in *mutual exclusion* from the other, so that there’s no chance of statement instances from the two different calls being interleaved in any way. The term “mutual exclusion” refers to the requirement that lines 6 and 7 are executed as a single indivisible step by each task, thereby *excluding* the other tasks when one task is executing those statements.

```

1  class DoublyLinkedList {
2      DoublyLinkedList 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          . . . // additional work to delete node (mutual exclusion not needed)
10     }
11     . . .
12 }
13 . . .
14 static void deleteTwoNodes(DoublyLinkedList L) {
15     finish {
16         async L.delete();
17         async L.next.delete();
18     }
19 }

```

Listing 1: Example of two tasks performing conflicting accesses

The predominant approach to ensure mutual exclusion proposed many years ago is to enclose a code region such as lines 6 and 7 in a *critical section* [3]. The following definition of critical sections from [12] captures the idea:

“In concurrent programming a critical section is a piece of code that accesses a shared resource (data structure or device) that must not be concurrently accessed by more than one thread of execution. A critical section will usually terminate in fixed time, and a thread, task or process will have to wait a fixed time to enter it (aka bounded waiting). Some synchronization mechanism is required at the entry and exit of the critical section to ensure exclusive use, for example a semaphore.”

The primary mechanisms available to a Java programmer for implementing the synchronization necessary for critical sections is the `synchronized` language construct, and the `java.util.concurrent.locks` library package. You will learn about both mechanisms later in the course. Instead, the Habanero Java (HJ) language offers a simpler construct, `isolated`, that can be used to directly implement critical sections, as shown in lines 5–8 of Listing 1.

Specifically, the HJ construct, `isolated <stmt1>`, guarantees that each instance of `<stmt1>` will be performed in mutual exclusion with all other potentially parallel instances of `isolated` statements `<stmt2>`. Thus, the use of `isolated` in line 5 of Listing 1 is sufficient to ensure that the computations in the two calls to `delete()` (lines 16 and 17) can safely be executed in parallel. After each task executes the `isolated` statement in lines 6 and 7, it can perform the computation in line 9 in parallel with the other task. Unlike `finish`, `isolated` does not dictate the order in which the two `isolated` instances should be executed; it just ensures that they execute in mutual exclusion regardless of the order.

The body of an `isolated` statement may perform any *sequential* Java computation including method calls and exceptions. It is illegal to execute a parallel construct (such as `async`, `finish`, `get`, or `forall`) within an `isolated` statement. If an exception is thrown within an `isolated` statement, `S`, it can be caught by a handler within or outside `S`. If control exits `S` after an exception, then all updates performed by `S` before throwing the exception will be observable after exiting `S`. `isolated` statements may be nested, but an inner `isolated` statement is essentially a no-op since isolation is guaranteed by the outer statement.

There is a trade-off between making `isolated` statements too big or too small in scope. If they are too big, then the parallelism in the program will be limited because interfering `isolated` statements cannot be executed in parallel. (Remember Amdahl's Law?) If they are too small, then they may not provide the desired program semantics. For example, if the `isolated` statement in line 5 of Listing 1 is replaced by two separate `isolated` statements for lines 6 and 7 of Listing 1, we would lose the invariant that lines 6 and 7 execute as an indivisible operation with respect to other tasks. It is also important to note that no combination of `finish`, `async`, and `isolated` constructs can create a *deadlock cycle* among tasks.

5.1.1 Implementations of Isolated Statements

While it is convenient for the programmer to use `isolated` statements, this convenience imposes major challenges on the *implementation* of `isolated` constructs. The discussion in the previous section highlights this issue, especially if we consider the fact that the `n1.delete()` and `n2.delete()` calls may or may not interfere depending on where `n1` and `n2` are located in the linked list.

The Habanero-Java implementation available for COMP 322 takes a simple *single-lock* approach to implementing `isolated` statements. You will learn more about locks later in the class. The idea behind a single-lock approach is to treat each entry of an `isolated` statement as an *acquire()* operation on the lock, and each exit of an `isolated` statement as a *release()* operation on the lock. Though correct, this approach essentially implements `isolated` statements as *critical sections*.

An alternate approach for implementing `isolated` statements being explored by the research community is *Transactional Memory* (TM) [9]. In Software Transactional Memory (STM), a combination of compiler and runtime techniques is used to optimistically execute *transactions* (instances of `isolated` statements) in parallel while checking for possible interference. If an interference is detected, one of the transactions is “rolled back” and restarted. Given the large amount of book-keeping necessary for logging read and write operations to check for interference, it is widely believed that software-only approaches to TM incur too much overhead to be useful in practice. Instead, there have been multiple proposals for Hardware Transactional Memory (HTM) support to assist with this book-keeping. As yet, no computer with HTM support is widely available, but a few hardware prototypes have begun to show promise in this direction *e.g.*, [2].

When comparing implementations of `isolated` statements, the three cases to consider in practice can be qualitatively described as follows:

1. *Low contention*: In this case, `isolated` statements are executed infrequently, and a single-lock approach as in HJ is often the best solution. Other solutions, such as TM, object-based isolation (Section 5.2), and atomic variables (Section 5.4.1), incur additional overhead compared to the single-lock approach because of their book-keeping and checks necessary but there is no visible benefit from that overhead because contention is low.
2. *Moderate contention*: In this case, 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. This is the case that motivates the use of approaches such as TM, object-based isolation, and atomic variables, since the benefit that they offer from reduced serialization outweighs the extra overhead that they incur.

3. *High contention:* In this case, there are phases in the program when interfering `isolated` statements on a single object (often referred to as a “hot spot” object) dominate the program execution time. In such situations, approaches such as TM and atomic variables are of little help since they cannot eliminate the interference inherent in accessing a single object. The best approach in such cases is to find an alternative approach to `isolated` *e.g.*, to use a parallel Array Sum algorithm instead of an `isolated` statement to compute the sum of values generated by different tasks.

5.2 Object-based Isolation

In this section, we introduce the basic functionality of object-based isolation [8]. As mentioned earlier, the focus of object-based isolation is on mutual exclusion rather than strong atomicity *i.e.*, mutual exclusion is only guaranteed between instances of `isolated` statements, unlike strong atomicity where a mutual exclusion guarantees may also exist between `atomic` and non-`atomic` statements. Two given `isolated` statements execute in mutual exclusion if the intersection of their object list is non-empty. Further, object-based `isolated` statements can be combined with global `isolated` statements that enforce mutual exclusion on all objects.

We first review the existing global `isolated` statement in HJ [7]. The HJ construct, `isolated <stmt1>`, guarantees that each instance of `<stmt1>` will be performed in mutual exclusion with all other potentially parallel *interfering* instances of `isolated` statements `<stmt2>`. Two instances of `isolated` statements, `<stmt1>` and `<stmt2>`, are said to interfere with each other if both access the same shared location, such that at least one of the accesses is a write.

The current HJ implementation takes a simple *single-lock* approach to implementing `isolated` statements, by treating each entry of an `isolated` statement as an *acquire()* operation on the lock, and each exit of an `isolated` statement as a *release()* operation on the lock. Though correct, this approach essentially implements `isolated` statements as *critical sections*, thereby serializing interfering and non-interfering `isolated` statement instances. We refer to this approach as *global mutual exclusion*.

The motivation for object-based isolation is that there are many cases when the programmer knows the set of objects that will be accessed in the body of the isolated statement. In that case, they can use a statement of the form, `isolated(obj0, obj1, ...) <stmt1>`, to specify the set of objects involved. (The order of objects is not significant.) We refer to this approach as *partial mutual exclusion*. In this case, two `isolated` statements are only guaranteed to execute in mutual execution if they have a non-empty intersection in their object sets. For convenience, the standard `isolated` statement in HJ is assumed to be equivalent to `isolated(*)` *i.e.*, an object-based `isolated` statement on the universal set of objects.

Figure 1 contains an example that uses global mutual exclusion to implement the `insert` function in the `SortList` benchmark. Lines 14 to 21 contain the critical section that performs node insertion and executes in mutual exclusion with any other critical sections that operate on objects `prev` and `curr`.

Figure 2 shows how the example presented in Figure 1 can be rewritten to use object-based `isolated` statements instead. In this case, the programmer explicitly identifies objects `prev` and `curr` as being involved in the mutual exclusion. As discussed in Section 5.1.1, lock-based implementations of object-based isolation rely on an ability to order the objects. This ordering is in turn used to guarantee an absence of deadlock in the implementation of object-based isolation.

As mentioned earlier, imposing a total order on the isolated objects is the key mechanism to avoid deadlocks. This can be done easily for objects in a single isolated list, but nested isolated constructs can pose a challenge. A sufficient condition for deadlock avoidance with nested isolation is to prohibit an inner `isolated` statement from including an object that was not already acquired by an outer `isolated` statement. (Note that this condition permits a task to re-acquire the same object, as is done with reentrant locks.) Figure 3 contains an example to illustrate this rule.

```
public boolean insert(final int v) {
    while (true) {
        INode curr, prev = null;
        for (curr = first; curr != null;
            curr = curr.getNext()) {
            final int val = curr.getValue();
            // v already exists
            if (val == v) return false;
            else if (val > v) break;
            prev = curr;
        }

        boolean set = false;
        isolated {
            if (validate(prev, curr)) {
                final INode neo = new INode();
                neo.setValue(v);
                link(prev, neo, curr);
                set = true;
            }
        }
        if (set) return true;
    }
}
```

Figure 1: SortList Insert Operation with Global Isolation

```
public boolean insert(final int v) {
    while (true) {
        INode curr, prev = null;
        for (curr = first; curr != null;
            curr = curr.getNext()) {
            final int val = curr.getValue();
            if (val == v) return false; // v already exists
            else if (val > v) break;
            prev = curr;
        }

        boolean set = false;
        if (prev != null && curr != null) {
            isolated(prev, curr) {
                if (validate(prev, curr)) {
                    final INode neo = new INode();
                    neo.setValue(v);
                    link(prev, neo, curr);
                    set = true;
                }
            }
        }
    }
}
```

Figure 2: SortList Insert Operation with Partial Mutual Exclusion

<pre>isolated (obj1, obj2) { isolated (obj3) { S; } }</pre> <p>(a)</p>	<pre>isolated { isolated (obj1, obj2) { S; } }</pre> <p>(b)</p>
--	---

Figure 3: Simple cases for nested isolation.

In Figure 3 (a), if `obj3` is neither an alias of `obj1` nor of `obj2`, then the statement `S` should not be permitted to execute and a runtime exception is thrown at that point. This is because the inner `isolated` region tries to acquire a new object that was not included in the outer `isolated` region, thereby opening the possibility of a deadlock. Figure 3 (b) shows a legal example for nested isolation. The outer `isolated` region is a global mutual exclusion construct (i.e. it acquires all objects), which makes it is legal to acquire any object in the inner `isolated` region. To ensure the correctness of nested isolation, the implementation relies on runtime checking as compile-time checking is undecidable in general.

Finally, a `null` object reference in list for an object-based `isolated` statement is essentially no-op (unlike the Java synchronized statement, which throws a `NullPointerException` in that case). Consider the following example:

```

    
```

`isolated (obj1, obj2) S;`

```

    
```

If `obj1` is `null`, then an implementation of the `isolated` statement only needs to acquire `obj2`, thereby making the previous statement is equivalent to the following:

```

    
```

`isolated (obj2) S;`

```

    
```

If both `obj1` and `obj2` are `null`, then the `isolated` statement degenerates to a no-op and no isolation is imposed on `S`.

5.3 Parallel Spanning Tree Example

In this section, we discuss a more complicated use of `isolated` statements in a program written to find a *spanning tree* of an *undirected graph*. An undirected graph $G = (V, E)$ consists of a finite set of *vertices*, V , and a set of *unordered edges*, E . An edge, $\{u, v\}$ connects vertices $u \neq v$, but the order is not significant. This is unlike directed graphs, such as a computation graph, where the direction of the edge matters. A common representation for undirected graphs is the *adjacency list*. For each vertex u in the graph, an array `neighbors` can be used to store the set of vertices $v \neq u$ such that there is an edge connecting vertices u and v . Figure 4 contains a simple example of an undirected graph.

A *path* from vertex a to vertex b in graph $G = (V, E)$ is a sequence $\langle v_0, v_1, v_2, \dots, v_k \rangle$ of vertices where $v_0 = a$, $v_k = b$, and $\{v_i, v_{i+1}\} \in E$ for $i = 0, 1, \dots, k - 1$. A graph is *connected* if each pair of vertices is connected by at least one path. A *cycle* (or *circuit*) is a path that starts and ends with the same vertex. A *tree* is a special kind of undirected graph that is connected and contains no cycles. A tree always contains $|V| - 1$ edges. A *rooted tree* has a designated vertex that is called the *root*. A rooted tree can be represented compactly by a `parent` field in each non-root vertex such that there is a tree edge $\{u, v\}$ if and only if $u.parent = v$ or $v.parent = u$. A *spanning tree* of a connected undirected graph G is a tree containing all the vertices of G , and a subset of G 's edges. A graph may have multiple spanning trees. Figure 5 shows three possible spanning trees for the graph in Figure 4.

Spanning trees have multiple applications in areas as diverse as telephone networks and transportation, since a spanning tree contains the minimum number of edges ($|V| - 1$) to ensure that the vertices in a graph remain connected. Removing an edge from a spanning tree will cause the resulting graph to become disconnected. If each edge $e = \{u, v\}$ in the input graph is assigned a *weight* (e.g., based on the distance or cost of the connection between the two vertices), then a *minimum spanning tree* is a spanning tree with the smallest total cost, when adding the weights of edges in the tree. The problem of finding the *minimum spanning tree* has received much attention in Computer Science. In this section, we will study the simpler problem of finding any spanning tree of an input graph.




Figure 4: Example undirected graph (source [11])




Figure 5: Three spanning trees for undirected graph in Figure 4 (source [11])

Listing 2 shows the sketch of a Java class `V` that implements a *sequential* program to compute the spanning tree of an undirected graph. Each vertex in the input graph is created as an instance of class `V`, and the its `neighbors` array contains its adjacency list. The algorithm starts by setting `root.parent = root` in line 17. Even though the root vertex does not have a defined parent, this convention simplifies the implementation of the `tryLabeling()` method, as we will see. Line 18 invokes the recursive method, `compute()`, on `root`.

The body of method `compute()` is defined in lines 9 – 13. It iterates through each vertex (named `child`) in the `neighbors` list of the current (`this`) vertex, and attempts to become the parent of `child` by calling `child.tryLabeling()`. Consider an example where the `this` vertex is `a` and the `child` vertex is `b` in line 11. Line 5 in method `tryLabeling()` will then check if `b`'s parent is `null`. If so, it sets `b.parent` to equal `a` and returns `true`. Otherwise, the call to `tryLabeling()` returns `false`. Back in line 11, method `compute()` will recursively call itself on `child a` if `tryLabeling()` returns `true`, thereby ensuring that `compute()` is called exactly once for each vertex. When all recursive calls have completed, the output spanning tree is represented by the `parent` fields.

```

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        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               child.compute(); //escaping async
13        }
14    } // compute
15 } // class V
16 . . .
17 root.parent = root; // Use self-cycle to identify root
18 root.compute();
19 . . .

```

Listing 2: Sequential Algorithm for computing the Spanning Tree of an undirected graph

Listing 3 shows a parallel version of the spanning tree program in Listing 2. The only changes are the addition of `isolated` in line 5, the addition of `async` in line 12, and the addition of `finish` in line 18. The addition of these keywords lead to a parallel program which computes a valid spanning tree of the graph, but without preserving the left-to-right order when traversing neighbors in the original sequential version. This is a nondeterministic algorithm since different executions may yield different spanning trees, all of which are valid solutions.

```

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 . . .

```

Listing 3: Parallel Spanning Tree of an undirected graph

There are some other points worth observing in this example. Method `compute()` in lines 8–14 contains an `async` but no `finish` statement. In such a case, the `async` in line 12 is called an *escaping async*, because its parent task’s method can return before the `async` has completed. Instead, a `finish` is included in line 18 to ensure that all `async`’s terminate before the program proceeds to line 19. The `isolated` statement in line 5 is necessary because multiple neighbors of a vertex `v2` may compete to be its parent. The winning vertex `v1` is the one whose `async` finds `v2.parent == null`, in which case it sets `v2.parent = v1`. All other neighbors will then fail to become the parent of `v2` since `v2.parent == null` will no longer be true.

5.4 Atomic Variables

5.4.1 AtomicInteger and AtomicIntegerArray

The `java.util.concurrent` package [10, 4] (also referred to as `j.u.c`) was introduced over five years ago as part of Java 5.0 to offer a standard set of library utilities for writing parallel programs in Java. You will learn more details of the `j.u.c.` package later in the class, when you graduate from parallel programming in HJ to parallel programming in pure Java. The use of many `j.u.c.` features by the programmer is prohibited in HJ, since they can easily violate HJ’s properties such as deadlock freedom¹.

However, there are two groups of utilities in `j.u.c.` that can be freely used in HJ code for performance and convenience — *atomic variables* and *concurrent collections* [4]. This section focuses on the former, which is embodied in the `java.util.concurrent.atomic` sub-package [1, 10]. Specifically, the `j.u.c.atomic` sub-package provides library calls that can be used by HJ programmers as a more efficient substitute for certain patterns of `isolated` statements. It does so by encapsulating a single variable in an object, and providing certain read, write and read-modify-write operations that are guaranteed to be performed “atomically” on the object *i.e.*, as though they occurred in an `isolated` statement. The source of efficiency arises from the fact that many platforms offer hardware support for executing these read-modify-write operations atomically.

Thus, the basic idea behind atomic variables is to implement common access patterns occurring in `isolated` statements as predefined methods that can be invoked by the programmer, with efficient implementations that avoid the use of locks. Atomic variables provided a restricted solution to scalable implementations of `isolated`. If an `isolated` statement matches an available atomic pattern, then it can be implemented by using an atomic variable; otherwise, the default implementation of `isolated` or object-based isolation has to be used instead.

The operations of interest for two `j.u.c.` atomic classes, `AtomicInteger` and `AtomicIntegerArray`, are summarized in Table 1. Let us start with `AtomicInteger`. It has two constructors, `AtomicInteger()` (with a default initial value value of 0) and `AtomicInteger(init)` (with a specified initial value). A single instance of `AtomicInteger` encapsulates an object with a single integer field, `val`, that can only be read or written using predefined methods, as shown in Table 1. Each such method call is guaranteed to execute in isolation with other methods invoked on the same object. To use these classes in HJ, you will need to include the following statement at the start of your program, `import java.util.concurrent.atomic.*`.

Table 1 shows equivalent HJ `isolated` statements for `AtomicInteger` methods `get()`, `set()`, `getAndSet()`, `addAndGet()`, `getAndAdd()`, and `compareAndSet()`. While the *functionality* of your HJ program will remain unchanged if you use `AtomicInteger` methods as in column 2 or `isolated` statements as in column 3, the *per-*

¹The HJ implementation uses `j.u.c.` features to implement parallel constructs that you have already learned, such as `async`, `finish`, `future` and `phaser`, but this usage is not visible to the HJ programmer.

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

Table 1: Methods in java.util.concurrent atomic classes AtomicInteger and AtomicIntegerArray and their equivalent HJ object-based isolated statements. Variable v refers to a j.u.c.atomic object in column 2 and to an equivalent non-atomic object in column 3. val refers to a field of type int, and arr refers to a field of type int[].

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

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

formance of your program in the *moderate contention* case will usually be superior when using `AtomicInteger` methods for the reasons discussed in Section 5.1.1. The complete list of `AtomicInteger` methods can be found in [1]. The methods omitted from Table 1 include `decrementAndGet()`, `getAndDecrement()`, `incrementAndGet()`, and `getAndIncrement()` (since they are equivalent to the `add` methods in Table 1 with `delta=1` or `delta=-1`), `doubleValue()`, `floatValue()` and `intValue()` (since their functionality can be obtained by simple cast operations), and `weakCompareAndSet()` (since it is identical to `compareAndSet()` in HJ, but differs from `compareAndSet()` in some fine points of the Java Memory Model that are not observed on current hardware). While the *isolated* patterns supported by the `AtomicInteger` methods in Table 1 may appear to be limited, they capture idioms that occur frequently in parallel programs as shown in Table 2. In addition, variants of the `addAndGet()` and `getAndAdd()` methods have been studied extensively for the last 30 years [6, 5] as primitives for building scalable parallel algorithms.

Table 1 also shows constructors and methods for the `AtomicIntegerArray` class. The main advantage of using an instance of `AtomicIntegerArray` instead of creating an array of `AtomicInteger`'s is that an `AtomicIntegerArray` instance occupies less space since its book-keeping overhead is amortized over an entire integer array, whereas a `AtomicInteger[]` array is essentially an array of objects. An element access in `AtomicInteger[]` also incurs an extra indirection relative to an element access in `AtomicIntegerArray`.

The `j.u.c.atomic` sub-package also includes `AtomicLong`, `AtomicLongArray`, and `AtomicBoolean` classes, with methods that are easy to understand once you know the methods available in `AtomicInteger` and `AtomicIntegerArray`.

5.4.2 AtomicReference

It is useful to perform atomic operations on object references, in addition to atomic operations on the primitive data types outlined in Section 5.4.1. Table 3 summarizes the operations available for `AtomicReference` and `AtomicReferenceArray` classes in the `j.u.c.atomic` sub-package. The `compareAndSet()` method can be especially useful in practice. As an example, consider the code in Listing 4 with an object-based *isolated* statement used in the Parallel Spanning Tree example.

```

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          boolean retVal;
6          isolated(this) { if (parent == null) parent=n; retVal = (parent == n);
7              return retVal;
8          } // tryLabeling
9          . . .
10 } // class V

```

Listing 4: Use of `isolated` in Parallel Spanning Tree example

```

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      } // tryLabeling
7      . . .
8  } // class V

```

Listing 5: Use of `compareAndSet()` as a replacement for `isolated` in Listing 4

The `isolated` statement in line 5 of Listing 4 can be replaced by a `compareAndSet()` method if `parent` is stored as an `AtomicReference`, as shown above in Listing 5. There are additional `j.u.c.atomic` classes

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

Table 3: Methods in `java.util.concurrent` atomic classes `AtomicReference` and `AtomicReferenceArray` and their equivalent HJ object-based `isolated` statements. Variable `v` refers to a `j.u.c.atomic` object in column 2 and to an equivalent non-atomic object in column 3. `ref` refers to a field of type `Object`, and `arr` refers to a field of type `Object []`.

available called `AtomicMarkableReference` and `AtomicStampedReference` that support atomic manipulation of a reference+boolean pair or a reference+int pair respectively. Unfortunately, there is no pre-defined method that can support atomic manipulation of multiple objects as in the `isolated` statement in Listing 1.

5.5 Read-Write Modes in Object-based Isolation

TO BE COMPLETED

5.6 Serialized Computation Graph for Isolated Statements

How can the Computation Graph (CG) structure be extended to model `isolated` statements? We start by modeling each instance of an `isolated` statement as a distinct step (node) in the CG. This is permissible since the execution of an `isolated` statement is purely sequential with no internal continuation points. Next, we reason about the order in which interfering `isolated` statements are executed. This is complicated because the order may vary from execution to execution *e.g.*, the `isolated` statement instance invoked by the `async` in line 16 of Listing 1 may execute before the `isolated` statement in line 17 in one execution, and vice versa in another execution.

To solve this dilemma, we introduce a family of *Serialized Computation Graphs* (SCG's) that can be obtained when executing a program for a given input. Each SCG consists of a CG with additional *serialization* edges. Consider a CG being constructed on-the-fly as a parallel program executes. Each time an `isolated` step, S' , is executed, we add a *serialization* edge from S to S' for each `isolated` step, S , that has already executed. For simplicity, we omit *serialization* edges when the source and destination steps belong to the same task, since they will always be redundant.

Each SCG represents a set of executions in which all interfering `isolated` statements execute in the same order. Different SCG's are used to reason about different orders of execution of interfering `isolated` statements. We can use SCG's to reason about properties of parallel programs that you have already studied with respect to a specific CG. For example, the critical path length (CPL) of the execution of a program with `isolated` statements can be obtained by computing the CPL of the corresponding SCG.

Consider the computation graph in Figure 6 (ignoring the red serialization edges) and assume that nodes v_{10} , v_{11} , v_{16} all consist of interfering `isolated` steps as shown in the bottom right of the figure. There are three possible orderings for these nodes, when taking *continue*, *spawn* and *join* edges into account: $v_{16} \rightarrow v_{10} \rightarrow v_{11}$, $v_{10} \rightarrow v_{16} \rightarrow v_{11}$ and $v_{10} \rightarrow v_{11} \rightarrow v_{16}$. Each order leads to a different SCG, when serialization edges are added. The red edges in Figure 6 show serialization edges added when the `isolated` steps execute in the sequence, $v_{10} \rightarrow v_{16} \rightarrow v_{11}$. These edges increase the critical path length of the SCG from 17 nodes (without serialization edges) to 18 nodes (with serialization edges). Alternate SCG's are also possible with serialization edges $v_{10} \rightarrow v_{11} \rightarrow v_{16}$ and $v_{16} \rightarrow v_{10} \rightarrow v_{11}$ which result in a critical path length of 17 nodes.




Figure 6: Serialized Computation Graph with Serialization Edges

5.6.1 Data Races and the Isolated Statement

The following definition of data races studied earlier can also be directly applied to an SCG:

“Formally, a data race occurs on location L in a program execution with computation graph CG if there exist steps S_1 and S_2 in CG such that:

1. S_1 does not depend on S_2 and S_2 does not depend on S_1 i.e., there is no path of dependence edges from S_1 to S_2 or from S_2 to S_1 in CG , and
2. both S_1 and S_2 read or write L , and at least one of the accesses is a write.”

By this definition, there is no data race between (say) node v_{10} and v_{16} in Figure 6, since the insertion of serialization edges ensures that there cannot be a data race between any two interfering `isolated` statement instances. The absence of a **data race** appears reasonable for variable x in Figure 6 (since all executions will result in $x++$ being performed exactly three times in isolation), but what about variable y ? The final value of y depends on the order in which the `isolated` statements are executed. While this appears to be a race of some sort, it is not considered to be a *data race* according to the above definition. In general, data races are considered more harmful than other forms of races in parallel programs, because the semantics of data races depends on intricate details of the underlying *memory consistency model*.

As another example, consider the `forall` loop in Listing 6. The read and write accesses to `rank.count` in line 3 result in data races between all pairs of `forall` iterations.

```
1 rank.count = 0; // rank object contains an int field , count
2 forall (point [i] : [0:m-1]) {
3     int r = rank.count++;
4     StringArray [i] = Hello , World from task with rank =      + r;
5 }
```

Listing 6: Example of a `forall` loop with a data race error

Earlier, you learned how `finish` statements can be inserted to fix data race errors. For this example, an `isolated` statement needs to be used instead, as shown in line 4 of Listing 7.

```
1 rank.count = 0; // rank object contains an int field , count
2 forall (point [i] : [0:m-1]) {
3     int r;
4     isolated { r = rank.count++; }
5     StringArray [i] = Hello , World from task with rank =      + r;
6 }
```

Listing 7: Using an `isolated` statement to fix the data race error in Listing 7

Note that the program in Listing 7 is *nondeterministic* because different executions with the same input can result in different outputs (related to which iteration `i` gets assigned which `rank r`). It is informative to review the following property studied earlier, in light of `isolated` constructs:

Determinism property: if a parallel program with `async`, `finish`, `forall`, `future` and `get` operations can never have a data race, then it must be deterministic with respect to its inputs.

This Determinism property does not hold in general for data-race-free programs that include `isolated` statements, as seen in the examples in Figure 6 and in Listing 7. Thus, the distinction between *deterministic* and *data-race-free* programs needs to be understood with more care, when `isolated` statements are also under consideration.

Finally, we note that the `isolated` construct supports *weak isolation* [9]. With weak isolation, any memory accesses performed outside `isolated` constructs are not checked for interference with accesses within `isolated` constructs. Thus, there is the possibility of the “indivisibility” associated with an `isolated` statement being broken due to data races with non-isolated accesses, such as (say) a non-isolated `x++` operation in node `v17` in Figure 6.

References

- [1] Package `java.util.concurrent.atomic`. URL <http://download.oracle.com/javase/1.5.0/docs/api/java/util/concurrent/atomic/package-summary.html>.
- [2] Shailender Chaudhry, Robert Cypher, Magnus Ekman, Martin Karlsson, Anders Landin, Sherman Yip, Håkan Ziffer, and Marc Tremblay. Rock: A high-performance sparc cmt processor. *IEEE Micro*, 29:6–16, March 2009. ISSN 0272-1732. URL <http://portal.acm.org/citation.cfm?id=1550399.1550516>.
- [3] Edsger W. Dijkstra. The structure of the “the”-multiprogramming system. In *Proceedings of the first ACM symposium on Operating System Principles*, SOSP '67, pages 10.1–10.6, New York, NY, USA, 1967. ACM. URL <http://doi.acm.org/10.1145/800001.811672>.
- [4] B. Goetz. *Java Concurrency In Practice*. Addison-Wesley, 2007.

- [5] A. Gottlieb, R. Grishman, C.P. Kruskal, K.P. McAuliffe, L. Rudolph, and M. Snir. The nyu ultra-computer — designing an mimd shared memory parallel computer. *IEEE Transactions on Computers*, 32:175–189, 1983. ISSN 0018-9340. URL <http://doi.ieeecomputersociety.org/10.1109/TC.1983.1676201>.
- [6] Allan Gottlieb, Boris D. Lubachevsky, and Larry Rudolph. Basic techniques for the efficient coordination of very large numbers of cooperating sequential processors. *ACM Trans. Program. Lang. Syst.*, 5:164–189, April 1983. ISSN 0164-0925. URL <http://doi.acm.org/10.1145/69624.357206>.
- [7] Habanero. Habanero Java programming language. <http://habanero.rice.edu/hj>, Dec 2009.
- [8] Shams Imam, Jisheng Zhao, and Vivek Sarkar. A composable deadlock-free approach to object-based isolation. In Jesper Larsson Trff, Sascha Hunold, and Francesco Versaci, editors, *Euro-Par 2015: Parallel Processing*, volume 9233 of *Lecture Notes in Computer Science*, pages 426–437. Springer Berlin Heidelberg, 2015. ISBN 978-3-662-48095-3. doi: 10.1007/978-3-662-48096-0_33. URL http://dx.doi.org/10.1007/978-3-662-48096-0_33.
- [9] James R. Larus and Ravi Rajwar. *Transactional Memory*. Morgan and Claypool, 2006. URL <http://www.morganclaypool.com/doi/abs/10.2200/S00070ED1V01Y200611CAC002>.
- [10] Qusay H. Mahmoud. Concurrent Programming with J2SE 5.0, March 2005. URL <http://java.sun.com/developer/technicalArticles/J2SE/concurrency/>.
- [11] Rashid Bin Muhammad. Trees. URL <http://www.personal.kent.edu/~rmuhamma/GraphTheory/MyGraphTheory/trees.htm>.
- [12] Wikipedia. Critical section, 2010. URL http://en.wikipedia.org/wiki/Critical_section.