COMP 322: Fundamentals of Parallel Programming
Module 2: Concurrency

©2017 by Vivek Sarkar

March 8, 2017

DRAFT VERSION – PLEASE DO NOT DISTRIBUTE

Contents

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 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* [16]. The following definition of critical sections from [33] 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 throw 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) [27]. 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.,* [14].

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 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 [26]. 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 [20]. 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`(obj_0, obj_1, ...) ⟨*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.

```java
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

```java
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

```java
isolated (obj1, obj2) {          isolated {
  isolated (obj3) {                isolated (obj1, obj2) {
    S;                               S;
  }                                }
}                                }
        (a)                              (b)
```

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 $\langle v_0, v_1, v_2, \ldots, v_k \rangle$ of vertices where $v_0 = a$, $v_k = b$, and $\{v_i, v_{i+1}\} \in E$ for $i = 0, 1, \ldots, k-1$. 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 with 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 [29])

Figure 5: Three spanning trees for undirected graph in Figure 4 (source [29])

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 [28, 17] (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 freedom[1].

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* [17]. This section focuses on the former, which is embodied in the `java.util.concurrent.atomic` sub-package [5, 28]. 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 teh 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-*

---

[1]The 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()	int j = v.**getAndSet**(newVal);	int j; isolated(v) { j = v.val; v.val = newVal; }
// init = 0	int j = v.**addAndGet**(delta);	isolated(v) { v.val += delta; j = v.val; }
	int j = v.**getAndAdd**(delta);	isolated(v) { j = v.val; v.val += delta; }
AtomicInteger(init)	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	int j = v.**getAndSet**(i,newVal);	int j; isolated(v) { j = v.arr[i]; v.arr[i] = newVal; }
(length) // init = 0	int j = v.**addAndGet**(i,delta);	isolated(v) { v.arr[i] += delta; j = v.arr[i]; }
	int j = v.**getAndAdd**(i,delta);	isolated(v) { j = v.arr[i]; v.arr[i] += delta; }
AtomicIntegerArray (arr)	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;`	`AtomicInteger rank = new AtomicInteger();`
`.  .  .`	`.  .  .`
`isolated(rank) r = ++rank.count;`	`r = rank.incrementAndGet();`
2) Work assignment:	
`rem = new ...; rem.count = n;`	`AtomicInteger rem = new AtomicInteger(n);`
`.  .  .`	`.  .  .`
`isolated(rem) r = rem.count--;`	`r = rem.getAndDecrement();`
`if ( r > 0 ) .  .  .`	`if ( r > 0 ) .  .  .`
3) Counting semaphore:	
`sem = new ...; sem.count = 0;`	`AtomicInteger sem = new AtomicInteger();`
`.  .  .`	`.  .  .`
`isolated(sem) r = ++sem.count;`	`r = sem.incrementAndGet();`
`.  .  .`	`.  .  .`
`isolated(sem) r = --sem.count;`	`r = sem.decrementAndGet();`
`.  .  .`	`.  .  .`
`isolated(sem) s = sem.count; isZero = (s==0);`	`s = sem.get(); isZero = (s==0);`
4) Sum reduction:	
`sum = new ...; sum.val = 0;`	`AtomicInteger sum = new AtomicInteger();`
`.  .  .`	`.  .  .`
`isolated(sum) sum.val += x;`	`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 [5]. 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 [19, 18] 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.

```java
class V  {
  V [] neighbors; // adjacency list for input graph
  V parent;        // output value of parent in spanning tree
  boolean tryLabeling(V n) {
    boolean retVal;
    isolated(this) { if (parent == null) parent=n; retVal = (parent == n);
    return retVal;
  } // tryLabeling
  . . .
} // class V
```

Listing 4: Use of isolated in Parallel Spanning Tree example

```java
class V  {
  V [] neighbors; // adjacency list for input graph
  AtomicReference parent;        // output value of parent in spanning tree
  boolean tryLabeling(V n) {
    return parent.compareAndSet(null, n);
  } // tryLabeling
  . . .
} // 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 `isolated` statements
AtomicReference	Object o = v.**get**();	Object o; `isolated` o = v.ref;
	v.**set**(newRef);	`isolated` v.ref = newRef;
AtomicReference() // init = null	Object o = v.**getAndSet**(newRef);	Object o; `isolated` { o = v.ref; v.ref = newRef; }
AtomicReference(init)	boolean b = v.**compareAndSet** (expect,update);	boolean b; `isolated` if (v.ref==expect) {v.ref=update; b=true;} else b = false;
AtomicReferenceArray	Object o = v.**get**(i);	Object o; `isolated` o = v.arr[i];
	v.**set**(i,newRef);	`isolated` v.arr[i] = newRef;
AtomicReferenceArray (length) // init = null	Object o = v.**getAndSet**(i,newRef);	Object o; `isolated` { o = v.arr[i]; v.arr[i] = newRef; }
AtomicIntegerArray (arr)	boolean b = v.**compareAndSet** (i,expect,update);	boolean b; `isolated` 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 $v10$, $v11$, $v16$ 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: $v16 \rightarrow v10 \rightarrow v11$, $v10 \rightarrow v16 \rightarrow v11$ and $v10 \rightarrow v11 \rightarrow v16$. 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, $v10 \rightarrow v16 \rightarrow v11$. 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 $v10 \rightarrow v11 \rightarrow v16$ and $v16 \rightarrow v10 \rightarrow v11$ 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 $v10$ and $v16$ 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* [27]. 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.

6   The Actor Model

Acknowledgment: the content in this section was developed in collaboration with Shams Imam.

The Actor Model (AM) promotes a *no-shared mutable state* and an event-driven philosophy. It was first defined in 1973 by Carl Hewitt et al. during their research on Artificial Intelligent (AI) agents [23]. It was designed to address the problems that arise while writing distributed applications. Further work by Henry Baker [24], Gul Agha [12], and others added to the theoretical development of the AM. The AM is different from task parallelism in that it is primarily an asynchronous message-based concurrency model. An actor is the central entity in the AM that defines how computation proceeds. The key idea is to encapsulate mutable state and use asynchronous messaging to coordinate activities among actors.

6.1 Introduction to Actors

An actor is defined as an object that has the capability to process incoming messages. Usually the actor has a mailbox, as shown in Figure 7, to store its incoming messages. Other actors act as producers for messages that go into the mailbox. An actor also maintains local state which is initialized during creation. Henceforth, the actor is only allowed to update its local state using data (usually immutable) from the messages it receives and intermediate results it computes while processing the message. The actor is restricted to process at most one message at a time. This allows actors to avoid data races and to avoid the need for synchronization as there is no other actor contending for access to its local data. There is no restriction on the order in which the actor decides to process incoming messages. As an actor processes a message, it is allowed to change its behavior that affects how it processes the subsequent messages.

Figure 7: Actors store incoming messages in a mailbox, maintain a local state which is not directly exposed to other actors, and process at most one message at a time.

An actor has a well-defined life cycle and restrictions on the actions it performs in the different states. During its life cycle an actor is in one of the following three states:

- *new*: An instance of the actor has been created; however, the actor is not yet ready to receive or process messages.

- *started*: An actor moves to this state from the *new* state when it has been started using the `start` operation. It can now receive asynchronous messages and process them one at a time. While processing a message, the actor should continually receive any messages sent to it without blocking the sender.

- *terminated*: The actor moves to this state from the *started* state when it has been terminated and will not process any messages in its mailbox or new messages sent to it. An actor signals termination by using the `exit` operation on itself while processing some message.

An actor interacts with other actors in two ways as shown in Figure 9. Firstly, it can send and receive messages to and from other actors. The sending and receiving of messages is done asynchronously, i.e. the sending actor can deliver a message without waiting for the receiving actor to be ready to process the message. An actor learns about the existence of other actors by either receiving their addresses in incoming messages or during creation. This brings us to the second manner of actor interaction: an actor can create new actors. This new actor can have its local state initialized with information from the *parent* actor. It is important to note that the network of actors an actor knows about can grow dynamically thus allowing formation of arbitrary connection graphs among actors and a wide range of communication and coordination

Figure 8: Actors have a simple life cycle. The most interesting state is *started* which is where the actor is receiving and processing messages.

patterns between them. In summary, while processing a message an actor may perform a finite combination of the following steps:

1. Asynchronously send a message to another actor whose address is known;

2. Create a new actor providing all the parameters required for initialization;

3. Become another actor, which specifies the replacement behavior to use while processing the subsequent messages [31].


Figure 9: During the processing of a message, actor interactions include exchanging messages with other actors and creating new actors.

6.1.1 Desirable Properties

The only way an actor conveys its internal state to other actors is explicitly via messages and responses to messages. This property obtains benefits similar to encapsulation in object-oriented programming and encourages modularity. The encapsulation of the local state also helps prevent data races because only the actor can modify its local state. Due to the asynchronous mode of communication, the lack of restriction on the order of processing messages sent from different actors, and the absence of synchronization via encapsulation of local data, actors expose inherent concurrency and can work in parallel with other actors.

6.1.2 Habanero-Java Actors

Habanero-Java (HJ) actors are defined by extending the `hj.lang.Actor` base class. Concrete sub-classes are required to implement the method used to process messages. This method is named `process()` in HJ *light* actors. Actors are like other objects and can be created by a `new` operation on concrete classes. An actor is activated by the `start()` method, after which the runtime ensures that the actor's message processing method is called for each message sent to the actor's mailbox. The actor can terminate itself by calling the `exit()` method while processing a message. Messages can be sent to actors from actor code or non-actor code by invoking the actor's `send()` method using a call as follows, `someActor.send(aMessage)`. A `send()` operation is non-blocking and the recipient actor processes the message asynchronously.

All async tasks created internally within an actor are registered on the `finish` scope that contained the actor's `start()` operation. The `finish` scope will block until all actors started within it terminate. This is similar to the `finish` semantics while dealing with `async`s.

```
1
2   import hj.lang.Actor;
3
4   public class HelloWorld {
5       public static void main(final String[] args) {
6           PrintActor actor = new PrintActor();
7           actor.start();
8           actor.send("Hello");
9           actor.send("World");
10          actor.send(PrintActor.STOP_MSG);
11      }
12  }
13
14  class PrintActor extends Actor<Object> {
15      static final Object STOP_MSG = new Object();
16      protected void process(final Object msg) {
17          if (STOP_MSG.equals(msg)) {
18              exit();
19          } else {
20              System.out.println(msg);
21          }
22      }
23  }
```

Listing 8: HelloWorld using HJ actors

Listing 8 shows a `HelloWorld` example using HJ actors. We are guaranteed ordered sends, i.e. though `Hello` and `World` will be processed asynchronously, they will be processed in that order.

```
1
2   import hj.lang.Actor;
3
4   public class SimplePipeline {
5       public static void main(final String[] args) {
6           // rely on the implicit IEF for HJ programs
7           Actor<Object> firstStage =
8               new EvenLengthFilter(
9                   new LowerCaseFilter(
10                      new PrintStage()));
11
12          firstStage.start();
13
```

```
14          firstStage.send("pipeline");
15          firstStage.send("Filter");
16          firstStage.send("example");
17
18          firstStage.send(new StopMessage());
19      }
20  }
21
22  class StopMessage {}
23
24  class EvenLengthFilter extends Actor<Object> {
25      private final Actor<Object> nextStage;
26      EvenLengthFilter(final Actor<Object> nextStage) {
27          this.nextStage = nextStage;
28      }
29      public void start() {
30          super.start();
31          nextStage.start();
32      }
33      protected void process(final Object msg) {
34          if (msg instanceof StopMessage) {
35              nextStage.send(msg);
36              exit();
37          } else if (msg instanceof String) {
38              String msgStr = (String) msg;
39              if (msgStr.length() % 2 == 0) {
40                  nextStage.send(msgStr);
41              }
42          }
43      }
44  }
45  class LowerCaseFilter extends Actor<Object> {
46      private final Actor<Object> nextStage;
47      LowerCaseFilter(final Actor<Object> nextStage) {
48          this.nextStage = nextStage;
49      }
50      public void start() {
51          super.start();
52          nextStage.start();
53      }
54      protected void process(final Object msg) {
55          if (msg instanceof StopMessage) {
56              exit();
57              nextStage.send(msg);
58          } else if (msg instanceof String) {
59              String msgStr = (String) msg;
60              if (msgStr.toLowerCase().equals(msgStr)) {
61                  nextStage.send(msgStr);
62              }
63          }
64      }
65  }
66  class PrintStage extends Actor<Object> {
67      protected void process(final Object msg) {
68          if (msg instanceof StopMessage) {
69              exit();
70          } else if (msg instanceof String) {
```

```
71                 System.out.println(msg);
72             }
73         }
74  }
```

Listing 9: A simple pipeline using HJ actors

Listing 9 shows an example of implementing a static pipeline using HJ actors. The pipeline is built up of three stages:

- *EvenLengthFilter*: The first stage of the pipeline filters out strings which are of odd length. It forwards strings of even length to the next stage in the pipeline.

- *LowerCaseFilter*: The second stage of the pipeline converts all the input strings to lowercase before forwarding to the third (and final) stage of the pipeline.

- *PrintStage*: The final stage of the pipeline prints the input strings. These strings are of even length and in lowercase due to the processing in the earlier stages of the pipeline.

6.1.3   Tips and Pitfalls

- Use an actor-first approach when designing programs that use actors *i.e.,* think about which actors need to be created and how they will communicate with each other. This step will also require you to think about the communication objects used as messages.

- If possible, use immutable objects for messages, since doing so avoids data races and simplifies debugging of parallel programs.

- When overriding the `start()` or `exit()` methods in actor classes, remember to make the appropriate calls to the parent's implementation with `super.start()` or `super.exit()`, respectively,

- The HJ actor `start()` method is not idempotent. Take care to ensure you do not invoke `start()` on the same actor instance more than once. The `exit()` method on the other hand is idempotent, invoking `exit()` multiple times is safe within the same call to `process()`.

- *Always remember to terminate a started actor* using the `exit()` method. If an actor that has been started is not terminated, the enclosing `finish` will wait forever (deadlock).

6.2   Ordering of Messages

Habanero actors preserve the order of messages with the same sender task and receiver actor, but messages from different senders may be interleaved in an arbitrary order. This is similar to the message ordering provided by ABCL [35].

6.3   Sieve of Eratosthenes

Pipelining is used for repetitive tasks where each task can be broken down into independent sub-tasks (also called stages) which must be performed sequentially, one after the other. Each stage partially processes data and then forwards the partially processed result to the next stage in the pipeline for further processing. This pattern works best if the operations performed by the various stages of the pipeline are balanced, i.e. take comparable time. If the stages in the pipeline vary widely in computational effort, the slowest stage creates a bottleneck for the aggregate throughput.

The pipeline pattern is a natural fit with the actor model since each stage can be represented as an actor. The single message processing rule ensures that each stage/actor processes one message at a time before handing off to the next actor in the pipeline. The stages however need to ensure ordering of messages while processing them. In our actor model, this ordering support is provided by default for messages from the

same actor to another actor. However, the amount of parallelism in a full pipeline is limited by the number of stages.

One algorithm that can be solved elegantly using a dynamic pipeline is the Sieve of Eratosthenes [34]. The algorithm incrementally builds knowledge of primes. The basic idea is to create multiple stages of the pipeline that forward a candidate prime number to the next stage only if the stage determines the candidate is locally prime. When the candidate reaches the end of the pipeline, the pipeline may need to be extended. Thus, this is also an example of a dynamic pipeline where the number of stages is not necessarily known in advance. A simple diagrammatic explanation of how the pipeline would work is shown in Figure 10. Note that to reduce the relative overhead in a real implementation, you will need to increase the amount of work done in each stage by having it store and process multiple prime numbers as a batch.


Figure 10: Illustration of Sieve of Eratosthenes algorithm (source: `http://golang.org/doc/sieve.gif`)

6.4 Producer-Consumer with Unbounded Buffer

TO BE COMPLETED

6.5 Producer-Consumer with Bounded Buffer

TO BE COMPLETED

6.6 Integration with Task Parallelism

There is internal concurrency in an actor in that it can be processing a message, receiving messages from other actors and sending messages to other actors at the same time. However, the requirement that the actor must process at most one message at a time is often misunderstood to mean that the processing must be done via sequential execution. In fact, there can be parallelism exposed even while message processing as long as the semantics is equivalent to that of processing at most one message at a time. A unique feature of the Habanero Actor model lies in its integration with task parallelism, so that (for example) async and finish constructs can be used when processing a message.

6.6.1 Using `finish` constructs during message processing

The traditional actor model already ensures that the actor processes one message at a time. Since no additional restrictions are placed on the message processing body, we can achieve parallelism by creating new `async-finish` constructs inside the message processing body. This will mean that we will spawn off new tasks to achieve the parallelism at the cost of blocking the original message processing task at the new `finish`. Since the main message processing task only returns after all spawned tasks have completed,

the invariant that only one message is processed at a time is maintained. Figure 11 shows an example code snippet that achieves this. Note that there is no restriction on the constructs used inside the newly constructed `finish`.

```scala
class ParallelizedProcessingActor() extends HybridActor {
  override def behavior() = {
    case msg: SomeMessage =>
      // preprocess the message
      finish { // finish to ensure all spawned tasks complete
        async {
          // do some processing in parallel
        }
        async {
          // do some more processing in parallel
        }
      }
      // postprocess the message after spawned tasks inside finish complete
      ...
  }
}
```

Figure 11: An actor exploiting the `async-finish` parallelism inside actors message processing body. The nested `finish` ensures no spawned tasks escape away causing the actor to process multiple messages at a time.

6.6.2 Pause and Resume operations

Requiring all spawned asyncs inside the message processing body are captured is too strict. This restriction can be relaxed based on the observation that the *at most one message processing rule* is required to ensure there are no internal state changes of an actor being effected by two or more message processing tasks of the same actor. We can achieve the same invariant by introducing a *paused* state in the actor life cycle and adding two new operations: `pause` and `resume`. In the *paused* state the actor is not processing any messages from its mailbox. The actor is simply idle as in the *new* state, however the actor can continue receiving messages from other actors. The actor will resume processing its messages, at most one at a time, when it returns to the *started* state. The `pause` operation takes the actor from a *started* state to a *paused* state while the `resume` operation achieves the reverse. The actor is also allowed to terminate from the *paused* state using the `exit` operation. Neither of these operations are blocking, they only affect the internal state of the actor which controls when messages are processed from the mailbox.


Figure 12: Actor life cycle from Figure 8 extended with a *paused* state. The actor can now continually switch between a *started* and *paused* state.

With the two new operations, we can now allow spawned tasks to escape the main message processing task. These spawned tasks are safe to run in parallel with the next message processing task of the same actor as long as they are not concurrently affecting the internal state of the actor. The actor can be suspended in a *paused* state while these spawned tasks are executing and can be signaled to resume processing messages once the spawned tasks determine they will no longer be modifying the internal state of the actor and hence not violating the one message processing rule. Figure 13 shows an example where the `pause` and `resume` operations are used to achieve parallelism inside the message processing body.

```scala
class ParallelizedWithEscapingAsyncsActor() extends HybridActor {
  override def behavior() = {
    case msg: SomeMessage =>
      // preprocess the message
      async {
        // do some processing in parallel
      }
      pause // to prevent the actor from processing the next message
      // note that pause/resume is not blocking
      async {
        // do some more processing in parallel
        // it is now safe for the actor to resume processing other messages
        resume
        // some more processing
      }
    ...
  }
}
```

Figure 13: An actor exploiting parallelism via `asyncs` while avoiding an enclosing `finish`. The asyncs escape the message processing body, but the `pause` and `resume` operations are used to control processing of subsequent messages by the actor.

6.6.3 Stateless Actors

The http://www.gpars.org/guide/GPars project, implemented in Groovy, has a notion of *stateless actors* that are allowed to process multiple messages simultaneously when they do not update internal state. As shown in Figure 14 (with Scala syntax!), it is easy to create stateless actors in Habanero's actor model because of its integration with task parallelism. There is no need to use the `pause` operation for stateless actors; the escaping `async` tasks can process multiple messages to the same actor in parallel.

7 Java Concurrency

7.1 Java Threads

Now that you have learned the fundamentals of parallel programming in the context of Habanero constructs, it is time to remove the "training wheels" and introduce you to parallel programming in the real world starting with *Java threads*. A unique aspect of Java compared to most of its predecessor mainstream programming languages (*e.g.,* C, C++, Pascal, FORTRAN, PL/1, COBOL) is that Java included the notion of threads and synchronization as part of its language definition right from the start, whereas support for parallelism in the predecessor languages is typically achieved through extra-language mechanisms such as *libraries*, *e.g.,* POSIX threads [30], Message-Passing Interface (MPI) [32], and Map-Reduce, or *pseudo-comment "pragmas" e.g.,* OpenMP [15].

```scala
class StatelessActor() extends HybridActor {
  override def behavior() = {
    case msg: SomeMessage =>
      async {
        // process the current message
      }
      if (enoughMessagesProcessed) {
        exit()
      }
      // return immediately to be ready to process the next message
  }
}
```

Figure 14: A simple stateless actor created using the hybrid model. The message processing body spawns a new task to process the current message and returns immediately to process the next message. Because the `async` tasks are allowed to escape, the actor may be processing multiple messages simultaneously.

As we will see, Java essentially takes a library-based approach to parallelism (which in turn made it possible to implement the Habanero-Java library on top of standard Java). The key mechanisms used to communicate task definitions to Java Concurrency libraries were originally `Runnable` and `Callable` objects that provide support for a limited form of closures. Runnable object are reviewed next in Section 7.1.1, and Callable objects in Section 7.1.2. However, with the introduction of lambdas in Java 8, lambdas can capture closures more simply in Java than runnable or callable objects.

Recall the HJ pseudocode below in Listing 10 that you learned in Lecture 1 for computing the sum of an array using two parallel `async` tasks:

```
1  // Start of Task T1 (main program)
2  sum1 = 0; sum2 = 0; // Assume that sum1 & sum2 are fields (not local vars)
3  finish {
4    // Compute sum1 (lower half) and sum2 (upper half) in parallel
5    int len = X.length
6    async for(int i=0 ; i < len/2 ; i++) sum1 += X[i]; // Task T2
7    async for(int i=len/2 ; i < len ; i++) sum2 += X[i]; // Task T3
8  }
9  //Task T1 waits for Tasks T2 and T3 to complete
10 int sum = sum1 + sum2; // Continuation of Task T1
```

Listing 10: Two-way Parallel ArraySum in HJ

In this pseudocode, the two `async` statements in lines 6 and 7 capture the value of local variable, `len`, from their outer scope, and the actual implementation of these statements using Java lambdas effectively creates closures for teh bodies of the async statements.

7.1.1 Using Runnable objects with Java threads

As a historical alternative to lambdas, the *Runnable* interface [9] in Java embodies a convention used by programmers to create a limited form of *closures* in Java. Any instance of a class that implements the *Runnable* interface must contain a `run()` method with a `void` return type. As an example, the following Java code in Listing 11 shows how the body of the first `async` in line 6 of Listing 10 can be created as a `Runnable` instance:

```
1  . . .
```

```
2   final int len = X.length;
3   Runnable r = new Runnable() {
4     public void run() {
5       for(int i=0 ; i < len/2 ; i++) sum1 += X[i];
6     }
7   };
8   . . .
```

Listing 11: Example of creating a Java Runnable instance as a closure

The `new` constructor in line 3 creates an instance of an *anonymous inner class* defined in lines 3–7 and assigns a reference to that instance to variable `r`. It is an *inner* class because it is declared within an existing class, and it is *anonymous* because it has been declared without a name. This class will only be accessible only at the point where it is defined. It can access any *local final* variables in the scope that it's defined, such as variable `len` in line 2. Unlike the HJ case, it is an error to access `len` in a Java anonymous inner class if `len` was not declared to be `final`.

Anonymous inner classes were originally intended for advanced Java users, and little attention was paid to simplifying their usage. However, a number of mainstream libraries for event handling (*e.g.,* `java.awt.EventQueue`) and for parallel programming (`java.util.concurrent`) expect client code to create instances of `Runnable`, which is most conveniently achieved via anonymous inner classes. Thus, even though the syntax for anonymous inner classes is quite cryptic in Java, many Java programmers have to learn this syntax so that they can use these libraries.

A summary of the Java Threads functionality can be found in the definition of the standard *java.lang.Thread* class [4], which is defined as follows:

```
1   public class Thread extends Object implements Runnable {
2     Thread() { ... } // Creates a new Thread
3     Thread(Runnable r) { ... } // Creates a new Thread with Runnable object r
4     void  run() { ... } // Code to be executed by thread
5       // Case 1: If this thread was created using a Runnable object,
6       //         then that object's run method is called
7       // Case 2: If this class is subclassed, then the run() method
8       //         in the subclass is called
9     void start() { ... } // Causes this thread to start execution
10    void join() { ... } // Wait for this thread to die
11    void  join(long m) // Wait at most m milliseconds for thread to die
12    static Thread currentThread() // Returns currently executing thread
13    . . .
14  }
```

Listing 12: java.lang.Thread class

Execution of a Java program begins with an instance of `Thread` created by the Java Virtual Machine (JVM) that executes the program's `main()` method. Parallelism can be introduced by creating additional instances of class `Thread` that execute as parallel threads. The static method, `Thread.currentThread()`, returns a reference to the currently executing thread.

Threads have richer functionality than `async` tasks, but are also more expensive. Earlier, we saw that the overhead of Java threads can be more than $10\times$ greater than that of `async` tasks implemented with *work-sharing* schedulers, and more than $1000\times$ greater than that of `async` tasks implemented with *work-stealing* schedulers. As a result, it is usually not practical to create large numbers of Java threads in a single program execution.

There are two ways to specify the code to be executed by a Java thread:

1. Define a class that implements the Runnable interface and pass an instance of that class to the Thread constructor in line 3 of Listing 12. It is common to create an instance of an anonymous inner class that implements `Runnable` for this purpose, as discussed in Section 7.1.1. In this case, the `Runnable` instance defines the work to be performed, and the `Thread` instance identifies the worker that will perform the work.

2. Subclass `Thread` and override the `run()` method. This is usually inconvenient in practice.

When an instance of `Thread` is created, it does not start executing right away. Instead, it can only start executing when its `start()` method is invoked. (`start()` can be invoked at most once per `Thread` instance.) As with `async`, the parent thread can immediately move to the next statement after invoking `t.start()`.

The operation analogous to `finish` is `join()`. A `t.join()` call forces the invoking thread to wait till thread `t` completes. This is a lower-level primitive than `finish` since it only waits for a single thread rather than a collection of threads. Further, there is no restriction on which thread performs a `join` on which thread, so it is easy for a programmer to erroneously create a deadlock cycle with `join` operations. Likewise, there is no analog of an *Immediately Enclosing Finish* for Java threads

Putting it all together, Listing 13 shows a Java threads version of the HJ code in Listing 10. It is instructive to compare the two versions and see the differences. There are three steps in initiating a thread to execute a computation analogous to `async`. The first step is to create a `Runnable` instance as in lines 5–7 for `r1` and lines 10–12 for `r2`. The bodies of the `run()` methods in lines 6 and 11 specify the work that was done in the two `async` statements in Listing 10. The next step is to create a `Thread` instance as in lines 8 and 13 for `t1` and `t2`. The third step is to start the threads as in line 9 and 14. After the threads have been created, the parent waits for their completion by performing the `join()` calls in line 16.

```
1   // Start of Task T1 (main program)
2   sum1 = 0; sum2 = 0; // Assume that sum1 & sum2 are fields (not local vars)
3   // Compute sum1 (lower half) and sum2 (upper half) in parallel
4   final int len = X.length;
5   Runnable r1 = new Runnable() {
6     public void run(){ for(int i=0 ; i < len/2 ; i++) sum1 += X[i];}
7   };
8   Thread t1 = new Thread(r1);
9   t1.start();
10  Runnable r2 = new Runnable() {
11    public void run(){ for(int i=len/2 ; i < len ; i++) sum2 += X[i];}
12  };
13  Thread t2 = new Thread(r2);
14  t2.start();
15  // Wait for threads t1 and t2 to complete
16  t1.join(); t2.join();
17  int sum = sum1 + sum2;
```

Listing 13: Two-way Parallel ArraySum using Java threads

7.1.2 Using Callable objects with Java threads

In the previous section, you learned how the `start()` and `join()` operations in the Java `Thread` class can be used to support the functionality of `async` and `finish` statements (albeit with greater overhead). With some additional steps, this approach can be extended to also support `future` tasks. It builds on the `Callable` interface introduced in Java 5 [1]. Any instance of a class that implements the `Callable<V>` interface must contain a `call()` method with a `V` return type. As a simple example of using `Callable`, consider the following sequential Java code in Listing 14:

```
1   ImageData image1 = imageInfo.downloadImage(1);
```

```
2   ImageData image2 = imageInfo.downloadImage(2);
3   . . .
4   renderImage(image1);
5   renderImage(image2);
```

<div align="center">Listing 14: HTML renderer in Java before decomposition into Callable tasks</div>

This code can be rewritten as follows (Listing 15) by encapsulating the two calls to `downloadImage()` in `Callable` objects in lines 1–4. As a result, the actual invocation of these calls is postponed to lines 6 and 7 when the `call()` method is invoked on the `Callable` objects.

```
1   Callable<ImageData> c1 = new Callable<ImageData>() {
2     public ImageData call() {return imageInfo.downloadImage(1);}};
3   Callable<ImageData> c2 = new Callable<ImageData>() {
4     public ImageData call() {return imageInfo.downloadImage(2);}};
5   . . .
6   renderImage(c1.call());
7   renderImage(c2.call());
```

<div align="center">Listing 15: HTML renderer in Java after decomposition into Callable tasks</div>

However, this use of `Callable` objects does not directly lead to parallelism because there is no `Thread` constructor that takes a `Callable` object as input. Instead, we need to wrap the `Callable` object in an object that implements the `Runnable` interface. A convenient adapter class for this purpose is `FutureTask<V>` [3]. It has a constructor, `FutureTask(Callable<V> c)`, that takes a `Callable` instance as a parameter. With this adapter, the following sequence of steps can be used to execute a `Callable` object in a separate thread:

1. Create a parameter-less callable closure using a statement like
   "`Callable<Object> c = new Callable<Object>() {public Object call() { return ...; }}; `".

2. Encapsulate the closure as a task using a statement like
   "`FutureTask<Object> ft = new FutureTask<Object>(c);`".

3. Start executing the task in a new thread by issuing the statement, "`new Thread(ft).start();`".

4. Wait for the task to complete, and get its result by issuing the statement, "`Object o = ft.get();`".

These steps can be used to transform the code to Listing 15 to a parallel version shown below in Listing 16:

```
1    Callable<ImageData> c1 = new Callable<ImageData>() {
2      public ImageData call() {return imageInfo.downloadImage(1);}};
3    FutureTask<Object> ft1 = new FutureTask<Object>(c1);
4    new Thread(ft1).start();
5    Callable<ImageData> c2 = new Callable<ImageData>() {
6      public ImageData call() {return imageInfo.downloadImage(2);}};
7    FutureTask<Object> ft2 = new FutureTask<Object>(c2);
8    new Thread(ft2).start();
9    . . .
10   renderImage(ft1.get());
11   renderImage(ft2.get());
```

<div align="center">Listing 16: HTML renderer in Java after parallelization of Callable tasks</div>

Note that the above code in Listing 16 is just equivalent to the following HJ code in Listing 17 using `future` tasks (other than the fact that `Thread` creation in Java is more expensive than task creation in HJ):

```
1  future<ImageData> ft1 = async<ImageData>{return imageInfo.downloadImage(1);};
2  future<ImageData> ft2 = async<ImageData>{return imageInfo.downloadImage(2);};
3  . . .
4  renderImage(ft1.get());
5  renderImage(ft2.get());
```

Listing 17: Equivalent HJ code for the parallel Java code in Listing 16

7.2 Structured Locks: Synchronized Statements and Methods in Java

The previous sections showed how `async` and `future` tasks in HJ can be implemented in standard Java using threads and instances of `Runnable` and `Callable` objects. Now let us consider how to implemented HJ's `isolated` statement. You have already learned how the `java.util.concurrent.atomic` package can be used to implement special cases of `isolated` statements. The key mechanism that can be used to implement more general forms of `isolated` statements (including critical sections) is *locks*. The basic idea behind the use of locks is to implement a statement such as `isolated <stmt>` as follows:

1. Acquire lock L_i

2. Execute `<stmt>`

3. Release lock L_i

The responsibility for ensuring that the choice of locks correctly implements the semantics of `isolated` falls entirely on the shoulders of the programmer. The main guarantee provided by locks is that only one thread can hold a lock at a time, and the thread is blocked if the lock is unavailable.

The Java language attempts to simplify the use of locks by providing the following built-in support:

- Every Java object and Java class has an associated lock which can be accessed by using the `synchronized` statement.

- The statement, `synchronized(p) <stmt>` acquires the lock for object `p`, executes `<stmt>`, and releases the lock for object `p`. (The lock is released even if an exception is thrown in `stmt`.) There is no constraint that the data accesses in `<stmt>` be confined to members of object `p`.

- A method can also be declared to be `synchronized`. If it is a virtual method, then the method is executed as if its body was enclosed in a "`synchronized (this)`" statement *i.e.,* it is guarded by the lock for the underlying object. If it is a static method, then it is executed as if its body was enclosed in a "`synchronized (A.class)`" statement *i.e.,* it is guarded by the lock for the underlying class (`A`).

- `synchronized` statements may be nested.

- Reentrant ("recursive") locking is permitted if two nested `synchronized` statements executed in the same thread attempt to acquire the same lock.

7.3 Unstructured Locks: Java's Lock Library

Coordination and synchronization of parallel tasks is a major source of complexity in parallel programming. These constructs take many forms in practice including directed barrier and point-to-point synchronizations, termination detection of child tasks, and mutual exclusion in accesses to shared resources. A read-write lock is a synchronization primitive that supports mutual exclusion in cases when multiple reader threads are permitted to enter a critical section concurrently (read-lock), but only a single writer thread is permitted in the critical section (write-lock). Although support for reader threads increases ideal parallelism, the read-lock functionality typically requires additional mechanisms, including expensive atomic operations, to handle multiple readers. It is not uncommon to encounter cases in practice where the overhead to support read-lock

operations overshadows the benefits of concurrent read accesses, especially for small critical sections. As a result, the use of read-write locks can often degrade overall performance relative to standard fine-grained locks, due to the overheads involved.

TO BE COMPLETED

7.4   Linearizability of Concurrent Objects

A *concurrent object* is an object that can correctly handle methods invoked in parallel by different tasks or threads. Concurrent objects were originally referred to as *monitors* [25], and are also informally referred to as "thread-safe objects". For simplicity, it is usually assumed that the body of each method in a concurrent object is itself sequential *i.e.,* it does not create child `async` tasks. A classical approach to making a class thread-safe (that dates back to the original concept of monitors) is to enclose each of its methods in a critical section to guarantee mutual exclusion. However, this approach leads to scalability bottlenecks since it effectively serializes all method invocations on the same object.

Examples of concurrent objects include `java.util.concurrent` collection classes, such as `Concurrent-HashMap`, `ConcurrentLinkedQueue` and `CopyOnWriteArraySet`. Instances of these classes are concurrent objects that are designed for high scalability in the presence of method invocations from parallel tasks. In addition, these classes ensure that all their methods are nonblocking. In general, it is very challenging to design and implement concurrent objects with these scalability and nonblocking properties, while retaining semantics equivalent to that of monitors. In this section, we discuss an important correctness property of concurrent objects called *linearizability* [22, 21], which guarantees semantics equivalent to that of methods in critical sections.

A formal definition of linearizability can be found in [22, 21]. An informal definition is as follows:

1. A *linearizable execution* is one in which the semantics of a set of method calls performed in parallel on a concurrent object is equivalent to that of some legal linear sequence of those method calls.

2. A *linearizable concurrent object* is one for which all possible executions are linearizable.

To understand this definition of linearizability, consider a canonical example of a concurrent object — a FIFO (First In, First Out) queue. A FIFO queue, q, typically supports two methods: `q.enq(o)` inserts object `o` at the *tail* of the queue and `q.deq()` removes and returns the item at the head of the queue. The `q.deq()` method call will typically throw an exception (*e.g.,* `EmptyException`) if the queue is empty. For simplicity, let us assume that we have unbounded space to store the queue. If (as in real-world situations) we have a bounded amount of space for the queue, then the `q.enq(o)` method call will throw an exception if the queue is full.

Consider two `async` tasks, *A* and *B*, repeatedly performing `enq()` and `deq()` operations on the FIFO queue. Table 4 shows one possible execution of a monitor-based implementation of a FIFO queue using critical sections. The critical sections ensure that none of the `enq()` and `deq()` methods run in parallel with each other. This is a linearizable execution because its semantics is equivalent to the following sequence:

```
q.enq(x) ; q.enq(y) ; q.deq() returns x
```

Now consider a concurrent (rather than a monitor-based) implementation of a FIFO queue, such as Java's `ConcurrentLinkedQueue`. One possible parallel execution is shown as a schedule in Table 5. This is a linearizable execution because its semantics is equivalent to the following sequence:

```
q.enq(x) ; q.enq(y) ; q.deq() returns x
```

In fact, the execution in Table 5 would also be linearizable if the `q.deq()` operation returned `y` instead of `x` at time $t = 5$ in Table 5. In this case, its semantics would be equivalent to the following sequence:

```
q.enq(y) ; q.enq(x) ; q.deq() returns y
```

Time	Task A	Task B
0	Invoke `q.enq(x)`	
1	Work on `q.enq(x)`	
2	Work on `q.enq(x)`	
3	Return from `q.enq(x)`	
4		Invoke `q.enq(y)`
5		Work on `q.enq(y)`
6		Work on `q.enq(y)`
7		Return from `q.enq(y)`
8		Invoke `q.deq()`
9		Return x from `q.deq()`

Table 4: Example of a linearizable execution of method calls on a monitor-based implementation of FIFO queue `q` by parallel tasks A and B

Time	Task A	Task B
0	Invoke `q.enq(x)`	
1	Work on `q.enq(x)`	Invoke `q.enq(y)`
2	Work on `q.enq(x)`	Return from `q.enq(y)`
3	Return from `q.enq(x)`	
4		Invoke `q.deq()`
5		Return x from `q.deq()`

Table 5: Example of a linearizable execution of method calls on a concurrent FIFO queue `q` by parallel tasks A and B (source: Figure 1 in [22])

We have the flexibility of considering different sequences for the schedule in Table 5 because of the parallel overlap between operations `q.enq(x)` and `q.enq(y)`.

Another way of understanding linearizability is that each method call of a linearizable concurrent object should appear to take effect "instantaneously" at some point in time between its invocation and return. For the schedule in Table 5, we can assume that (say) `q.enq(x)` took effect at time $t = 1$ and `q.enq(y)` took effect at time $t = 2$. For the alternate schedule, we can assume that (say) `q.enq(x)` took effect at time $t = 3$ and `q.enq(y)` took effect at time $t = 2$. We can assume that `q.deq()` took effect at time $t = 4$ in both cases. For the monitor-based implementation in Table 4, we have no choice but to pick times consistent with the `q.enq(x) ; q.enq(y) ; q.deq()` sequence.

Now, consider the parallel execution shown in Table 6. In this case, there is no parallel overlap between `q.enq(x)` and `q.enq(y)` since the return from `q.enq(x)` precedes the invocation of `q.enq(y)`. Thus,

Time	Task A	Task B
0	Invoke `q.enq(x)`	
1	Return from `q.enq(x)`	
2		Invoke `q.enq(y)`
3	Invoke `q.deq()`	Work on `q.enq(y)`
4	Work on `q.deq()`	Return from `q.enq(y)`
5	Return y from `q.deq()`	

Table 6: Example of a non-linearizable execution of method calls on a concurrent FIFO queue `q` by parallel tasks A and B (source: Figure 1 in [22])

Figure 15: A Computation Graph for the execution in Table 4. i-begin and i-end refer to the beginning and end of isolated statements. `q.deq():x` refers to a dequeue operation that returns `x`.

`q.enq(x)` must precede `q.enq(y)` in all linear sequences of method calls invoked on `q`, and it would be illegal for the next `q.deq()` operation to return `y`. Hence, the execution in Table 6 is non-linearizable.

Even with this brief discussion, it should be clear that linearizability is a desirable correctness property for concurrent objects.

7.4.1 Relating Linearizability to the Computation Graph model

How does the definition of linearizability in Section 7.4 relate to the Computation Graph (CG) model that we have studied in the class? Since each method executed on a concurrent object is assumed to be sequential, the method's execution can be represented in the CG as a sequence of nodes (steps) connected by *continue* edges. Some of these nodes will be instances of `isolated` statements, and each such node may be the source and/or destination of *serialization* edges as a result.

Figure 15 shows a CG for the monitor-based FIFO execution shown in Table 4. Since it's a monitor-based implementation, we assume that each `enq()` and `deq()` method is enclosed in an `isolated` statement. The `i-begin` and `i-end` nodes refer to the beginning and end of isolated statements, and `q.deq():x` refers to a dequeue operation that returns x. As before, the CG contains *continue* edges and *serialization* edges to represent orderings within and across tasks.

Recall the following rule for adding serialization edges to a CG:

> 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.*

That is why we have two serialization edges from Task A to Task B in Figure 15, even though the second edge is redundant (transitive) if we take the *continue* edges in Task B into account. (For simplicity, we omit serialization edges when the source and destination steps belong to the same task, so there is no serialization edge from the first `i-end` node to the second `i-begin` node in Task B.) It is clear from this CG that the only legal sequence for the operations on `q` in this execution is `q.enq(x);q.enq(y);q.deq():x`. However, different executions with different serialization edges could lead to sequences `q.enq(y);q.enq(x);q.deq():y` and `q.enq(y);q.deq():y;q.enq(x)`.

The concept of a method call taking effect "instantaneously" can be translated to the CG by creating a reduced method-level version of the CG as shown in Figure 16. This reduction can be defined as follows: take the sequence of CG nodes a_i corresponding to a method call of interest and "collapse" it into a single node b to obtain a new graph, CG′. Each edge with source a_i in CG is redirected to source b in CG′. Likewise, each edge with destination a_i in CG is redirected to destination b in CG′. Finally, if these redirections result in any self-loop edges of the form $b \rightarrow b$ in CG′, those self-loop edges are removed from CG′.

Given a reduced CG, a *sufficient condition* for linearizability is that the *reduced CG is acyclic* as in Figure 16. This means that if the reduced CG does not contain a directed cycle, then the underlying execution must
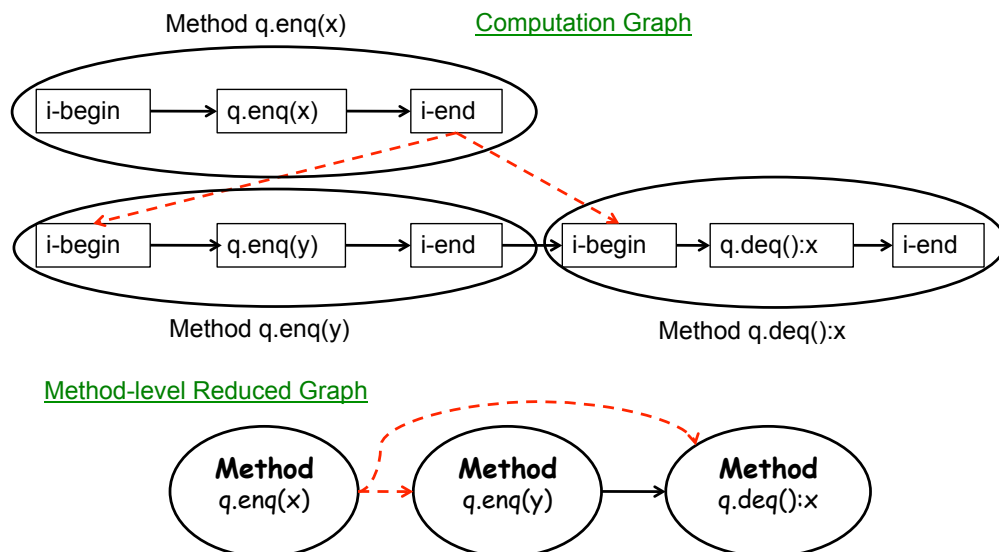
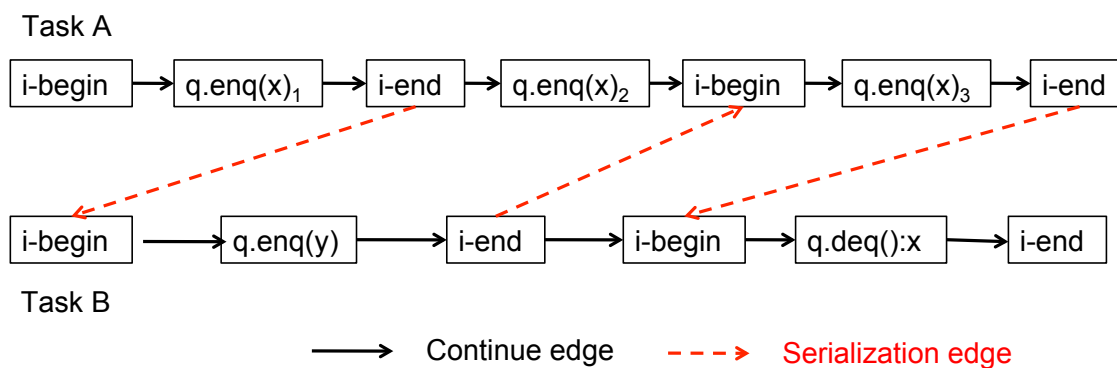Figure 16: Reduced method-level graph for CG in Figure 15

Figure 17: A Computation Graph for the execution in Table 5. $q.enq(x)_1$, $q.enq(x)_2$ and $q.enq(x)_3$ refer to three steps in the `q.enq(x)` method, the first and third of which are enclosed in isolated statements.

be linearizable. However, the converse is not necessarily true, as we will see.

Figure 17 shows a CG for the concurrent FIFO execution shown in Table 5. For this example, we assumed that the concurrent implementation split the the `q.enq(x)` method into three steps — $q.enq(x)_1$, $q.enq(x)_2$ and $q.enq(x)_3$ — the first and third of which are enclosed in isolated statements. However, the `q.enq(y)` and the `q.deq():x` operations were implemented as single steps. Note that Figure 17 shows serialization edges that are consistent with the execution in Table 5.

Figure 18 shows the reduced method-level graph for the CG in Figure 17. This graph has a cycle because it contains edges in both directions between `q.enq(x)` and `q.enq(y)`. However, this execution is still linearizable, as discussed earlier. This confirms that it is not necessary that the reduced CG of a linearizable execution be acyclic.

7.5 Liveness

TO BE COMPLETED

Figure 18: Reduced method-level graph for CG in Figure 17

7.6 The Dining Philosophers Problem

TO BE COMPLETED

A Java Concurrent Collections

One of the biggest successes in object-oriented programming has been the creation of *frameworks* that encapsulate reusable code in different domains. A foundational domain that has received a lot of attention in this regard is *collection data structures*. A collection (container) is a data structure that "collects" multiple element objects into a single group. The Java Collections Framework [13] includes a number of collection classes, some of which you may have already encountered when learning how to write sequential programs in Java (*e.g.,* `Vector`, `ArrayList`, `Stack`, `TreeSet`, and `HashMap`). These classes implement interfaces such as `List`, `Queue`, `Set`, `SortedSet`, and `Map`.

While the original Java Collections Framework (JCF) has been very successful for sequential programs, they pose a number of limitations for parallel programs. Many of the original JCF classes attempted to be "thread-safe" by ensuring that each public method in the class is "synchronized" *i.e.,* acquires and releases a single lock on the collection each time a public method is invoked on the collection. Though simple, the one-lock-per-collection approach proved to be limiting in multiple ways — it did not offer a convenient way to lock multiple collections at the same time, it ended up serializing more computations than needed (thereby exacerbating the effect of Amdahl's Law), and it did not provide a clean deadlock avoidance guarantee as with HJ `isolated` statements.

Instead, it is more desirable to provide collection classes that are designed for concurrency up-front, and can provide the simplicity of `isolated` semantics coupled with the performance and scalability advantages akin to those offered by Java atomic variables. The `java.util.concurrent` collection classes [28, 17] address this requirement. We will discuss three of these classes in this section: `ConcurrentHashMap`, `ConcurrentLinkedQueue` and `CopyOnWriteArraySet`. All operations in these three classes are nonblocking and can be freely used in any HJ program, just like the `j.u.c.atomic.*` classes. Note that there

are many `j.u.c.` classes that include blocking operations, and should not be used in an HJ program *e.g.,* `ArrayBlockingQueue`, `CountDownLatch`, `CyclicBarrier`, `DelayQueue`, `Exchanger`, `FutureTask`, `LinkedBlockingQueue`, `Phaser PriorityBlockingQueue`, `Semaphore`, and `SynchronousQueue`. However, most of these classes provide functionalities that are already available in HJ *e.g.,* `CyclicBarrier`, `Phaser`, `CountDownLatch`, and `Semaphore` are all subsumed by HJ *phasers*.

A.1   java.util.concurrent.ConcurrentHashMap

The `Map` interface in Java [6] is used to define *key-value mappings* for a given domain. It includes basic operations such as `put(key,value)`, `get(key)` and `remove(key)`. There are many applications where it is important to keep track of key-value mappings including *software caches* (where a program stores key-value maps obtained from an external source such as a database), *dictionaries*, and even *sparse arrays*.

There are multiple implementations for the `Map` interface available in a Java runtime environment. The `HashMap` class is completely unsynchronized and assumes that all necessary synchronization is handled by the caller. In the absence of detailed knowledge of the internals of the `HashMap` implementation, any client will likely play safe and ensure that an entire operation such as `put` or `get` is performed in a `isolated` statement or while some lock is held. However, these operations may take more than constant time due to traversals of multiple objects and the need to call an `equals()` method on each object traversed. As a result, this approach could add significantly to the serialization in a parallel program.

The `j.u.c.ConcurrentHashMap` class package implements the `Map` interface with more parallelism (less serialization) in a parallel context than the approaches discussed above. In fact, it implements the `ConcurrentMap` sub-interface of `Map` which includes four additional methods — `putIfAbsent(key,value)`, `remove(key,value)`, `replace(key,value)`, and `replace(key,oldValue,newValue)`.

An instance of `ConcurrentHashMap` permits retrieval operations, such as `get()`, to overlap with update operations, such as `put()` and `remove()`. The result of a `get()` operation will depend on which update operations have completed prior to the start of the `get()`. Thus, single-element operations such as `put()`, `get()`, and `remove()` all appear to be atomic with respect to each other.

However, `ConcurrentHashMap` also supports aggregate operations such as `putAll()` and `clear()`. It is not practical to make these operations appear atomic in a parallel context since they involve multiple elements. As a result, a single-element operation may observe an intermediate state where only some of the entries being inserted by a `putAll()` operation have been inserted, or where only some of the entries being removed by a `clear()` operation have been removed.

Given the widespread use of multithreaded and parallel applications in mainstream software, the `ConcurrentHashMap` is increasing in popularity. Listing 18 shows a sample usage of `ConcurrentHashMap` from an open source project.

```
1   public abstract class BaseDirectory extends BaseItem implements Directory {
2     Map files = new ConcurrentHashMap();
3     . . .
4     public Map getFiles() {
5       return files;
6     }
7     public boolean has(File item) {
8       return getFiles().containsValue(item);
9     }
10    public Directory add(File file) {
11      String key = file.getName();
12      if (key == null) throw new Error(. . .);
13      getFiles().put(key, file);
14      . . .
15      return this;
16    }
17    public Directory remove(File item) throws NotFoundException {
```

```
18        if ( has ( item ) ) {
19          getFiles ().remove( item.getName () );
20          . . .
21        } else throw new NotFoundException (" can ' t _remove _ unrelated _item " );
22      }
23   }
```

Listing 18: Example usage of ConcurrentHashMap in org.mirrorfinder.model.BaseDirectory [2]

A.2 java.util.concurrent.ConcurrentLinkedQueue

The Queue interface in Java [7] defines a collection of elements typically inserted and removed in FIFO (first-in-first-out) order, using operations such as add() and poll(). The j.u.c.ConcurrentLinkedQueue class implements the Queue interface with more parallelism (less serialization) than a synchronized queue. Consistent with FIFO order, new elements are inserted at the tail of the queue, and the queue retrieval operations obtain elements at the head of the queue.

There is no bound on the size of a ConcurrentLinkedQueue, so an add(o) operation for a non-null object always succeeds (assuming that the program does not run out of memory). offer(o) is an alternative name for add(o).

The semantics for removing an object from a queue leads to some interesting possibilities. There is a remove(o) operation that attempts to remove a specified object and returns null if the object is in the queue. The ConcurrentLinkedQueue class also has a poll() operation that returns null if the queue is empty and removes the head of the queue otherwise. This is in contrast to the LinkedBlockingQueue.take() operation which blocks when the queue is empty.

Like ConcurrentHashMap, ConcurrentLinkedQueue is also widely used in practice. Listing 19 shows a sample usage of ConcurrentLinkedQueue from an open source project.

```
 1   class BufferPool15Impl implements BufferPool.BufferPoolAPI {
 2     protected int maxSize;
 3     protected AtomicInteger size = new AtomicInteger (0);
 4     protected ConcurrentLinkedQueue queue = new ConcurrentLinkedQueue ();
 5     . . .
 6     public XByteBuffer getBuffer (int minSize, boolean discard) {
 7       XByteBuffer buffer = (XByteBuffer) queue.poll ();
 8       if ( buffer != null ) size.addAndGet(−buffer.getCapacity ());
 9       if ( buffer == null ) buffer = new XByteBuffer (minSize,discard);
10       else if ( buffer.getCapacity () <= minSize ) buffer.expand (minSize);
11       . . .
12       return buffer;
13     }
14     public void returnBuffer (XByteBuffer buffer) {
15       if ( (size.get () + buffer.getCapacity ()) <= maxSize ) {
16         size.addAndGet( buffer.getCapacity ());
17         queue.offer ( buffer );
18       }
19     }
20   }
```

Listing 19: Example usage of ConcurrentLinkedQueue in org.apache.catalina.tribes.io.BufferPool15Impl [8]

A.3 java.util.concurrent.CopyOnWriteArraySet

Like a *set* in Mathematics, the Set interface in Java [10] defines a collection of elements that is prohibited from containing any duplicates, as defined by the equals() method on the objects in the collection. The

j.u.c.CopyOnWriteArraySet class implements the Set interface with more parallelism (less serialization) than a synchronized set data structure by using an array-based data structure. Unlike ConcurrentHashMap and ConcurrentLinkedQueue, the CopyOnWriteArraySet class achieves thread safety by making copies of the set as needed. This makes it well suited for applications in which sets are not large, and read operations dominate update operations in frequency. This is because update operations such as add() and remove() involve making copies of the array. However, iterators can traverse array "snapshots" efficiently without worrying about changes during the traversal.

Listing 20 shows a sample usage of CopyOnWriteArraySet from an open source project.

```
1  public class DefaultTemplateLoader implements TemplateLoader, Serializable
2  {
3    private Set resolvers = new CopyOnWriteArraySet();
4    public void addResolver(ResourceResolver res)
5    {
6      resolvers.add(res);
7    }
8    public boolean templateExists(String name)
9    {
10       for (Iterator i = resolvers.iterator(); i.hasNext();) {
11         if (((ResourceResolver) i.next()).resourceExists(name)) return true;
12       }
13       return false;
14   }
15   public Object findTemplateSource(String name) throws IOException
16   {
17     for (Iterator i = resolvers.iterator(); i.hasNext();) {
18       CachedResource res = ((ResourceResolver) i.next()).getResource(name);
19       if (res != null) return res;
20     }
21     return null;
22   }
23  }
```

Listing 20: Example usage of CopyOnWriteArraySet in org.norther.tammi.spray.freemarker.DefaultTemplateLoader [11]

References

[1] java.lang.Callable interface. URL `http://download.oracle.com/javase/1.5.0/docs/api/java/util/concurrent/Callable.html`.

[2] ConcurrentHashMap Java Source Code example from org.mirrorfinder.model.BaseDirectory. URL `http://www.javadocexamples.com/java_source/org/mirrorfinder/model/BaseDirectory.java.html`.

[3] java.util.concurrent.FutureTask class. URL `http://download.oracle.com/javase/1.5.0/docs/api/java/util/concurrent/FutureTask.html`.

[4] java.lang.Thread. URL `http://download.oracle.com/javase/1.5.0/docs/api/java/lang/Thread.html`.

[5] Package java.util.concurrent.atomic. URL `http://download.oracle.com/javase/1.5.0/docs/api/java/util/concurrent/atomic/package-summary.html`.

[6] The Map Interface. URL `http://download.oracle.com/javase/tutorial/collections/interfaces/map.html`.

[7] The Queue Interface, . URL `http://download.oracle.com/javase/tutorial/collections/interfaces/queue.html`.

[8] ConcurrentLinkedQueue Java Source Code example from org.apache.catalina.tribes.io.BufferPool15Impl, . URL `http://www.javadocexamples.com/java_source/org/apache/catalina/tribes/io/BufferPool15Impl.java.html`.

[9] java.lang.Runnable interface. URL `http://download.oracle.com/javase/1.5.0/docs/api/java/lang/Runnable.html`.

[10] The Set Interface, . URL `http://download.oracle.com/javase/tutorial/collections/interfaces/set.html`.

[11] CopyOnWriteArraySet Java Source Code example from org.norther.tammi.spray.freemarker.DefaultTemplateLoader, . URL `http://www.javadocexamples.com/java_source/org/norther/tammi/spray/freemarker/DefaultTemplateLoader.java.html`.

[12] Gul Agha. *Actors: a model of concurrent computation in distributed systems*. MIT Press, Cambridge, MA, USA, 1986. ISBN 0-262-01092-5.

[13] Joshua Bloch. Java Collections Framework. URL `http://download.oracle.com/javase/tutorial/collections`.

[14] Shailender Chaudhry, Robert Cypher, Magnus Ekman, Martin Karlsson, Anders Landin, Sherman Yip, Håkan Zeffer, 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`.

[15] L. Dagum and R. Menon. OpenMP: An industry standard API for shared memory programming. *IEEE Computational Science & Engineering*, 1998.

[16] 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`.

[17] B. Goetz. *Java Concurrency In Practice*. Addison-Wesley, 2007.

[18] 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`.

[19] 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`.

[20] Habanero. Habanero Java programming language. http://habanero.rice.edu/hj, Dec 2009.

[21] Maurice Herlihy and Nir Shavit. *The art of multiprocessor programming*. Morgan Kaufmann, 2008. ISBN 978-0-12-370591-4. URL `http://www.elsevierdirect.com/companion.jsp?ISBN=9780123705914`.

[22] Maurice P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12:463–492, July 1990. ISSN 0164-0925. URL `http://doi.acm.org/10.1145/78969.78972`.

[23] Carl Hewitt, Peter Bishop, and Richard Steiger. Artificial Intelligence A Universal Modular ACTOR Formalism for Artificial Intelligence. Proceedings of the 3rd International Joint Conference on Artificial Intelligence, Stanford, CA, August 1973.

[24] Hewitt, Carl and Baker, Henry G. Actors and Continuous Functionals. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA, February 1978.

[25] C. A. R. Hoare. *Monitors: an operating system structuring concept*, pages 272–294. Springer-Verlag New York, Inc., New York, NY, USA, 2002. ISBN 0-387-95401-5. URL `http://portal.acm.org/citation.cfm?id=762971.762982`.

[26] 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`.

[27] James R. Larus and Ravi Rajwar. *Transactional Memory*. Morgan and Claypool, 2006. URL `http://www.morganclaypool.com/doi/abs/10.2200/S00070ED1V01Y200611CAC002`.

[28] Qusay H. Mahmoud. Concurrent Programming with J2SE 5.0, March 2005. URL `http://java.sun.com/developer/technicalArticles/J2SE/concurrency/`.

[29] Rashid Bin Muhammad. Trees. URL `http://www.personal.kent.edu/~rmuhamma/GraphTheory/MyGraphTheory/trees.htm`.

[30] POSIX. *Threads Extension for Portable Operating Systems*, 1994.

[31] Natarajan Raja and Rudrapatna K. Shyamasundar. Actors as a Coordinating Model of Computation. In *Proceedings of the 2nd International Andrei Ershov Memorial Conference on Perspectives of System Informatics*, pages 191–202. Springer-Verlag, 2004. ISBN 3-540-62064-8.

[32] Anthony Skjellum, Ewing Lusk, and William Gropp. *Using MPI: Portable Parallel Programming with the Message Passing Iinterface*. MIT Press, 1999.

[33] Wikipedia. Critical section, 2010. URL `http://en.wikipedia.org/wiki/Critical_section`.

[34] Wikipedia, The Free Encyclopedia. Sieve of Eratosthenes. `http://en.wikipedia.org/wiki/Sieve_of_Eratosthenes`. URL `http://en.wikipedia.org/wiki/Sieve_of_Eratosthenes`.

[35] Akinori Yonezawa, Jean-Pierre Briot, and Etsuya Shibayama. Object-Oriented Concurrent Programming in ABCL/1. In *Conference proceedings on Object-oriented programming systems, languages and applications*, OOPLSA '86, pages 258–268, New York, NY, USA, 1986. ACM. ISBN 0-89791-204-7. doi: 10.1145/28697.28722. URL `http://doi.acm.org/10.1145/28697.28722`.