
COMP 322: Fundamentals of Parallel Programming

Lecture 18: Midterm Review

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

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



Worksheet #17:

Critical Path Length for Computation with Signal Statement

Name: _____ Netid: _____

Compute the WORK and CPL values for the program shown below. How would they be different if the `signal()` statement was removed? (WORK = 204, CPL = 102)

```
1. finish(() -> {
2.   final HjPhaser ph = newPhaser(SIG_WAIT);
3.   asyncPhased(ph.inMode(SIG_WAIT), () -> { // Task T1
4.     A(0); doWork(1);    // Shared work in phase 0
5.     signal();
6.     B(0); doWork(100); // Local work in phase 0
7.     next(); // Wait for T2 to complete shared work in phase 0
8.     C(0); doWork(1);
9.   });
10.  asyncPhased(ph.inMode(SIG_WAIT), () -> { // Task T2
11.    A(1); doWork(1);    // Shared work in phase 0
12.    next(); // Wait for T1 to complete shared work in phase 0
13.    C(1); doWork(1);
14.    D(1); doWork(100); // Local work in phase 0
15.  });
16.}); // finish
```



Async and Finish Constructs for Task Creation and Termination (Lecture 1)


async S

- Creates a new child task that executes statement S

```
// T0 (Parent task)
STMT0;
finish { //Begin finish
  async {
    STMT1; //T1 (Child task)
  }
  STMT2; //Continue in T0
           //Wait for T1
} //End finish
STMT3; //Continue in T0
```

finish S

- Execute S, but wait until *all* asyncs in S's scope have terminated.



Acknowledgments: X10 and Habanero projects



Two-way Parallel Array Sum using HJ-Lib's finish & async API's

```
1. // Start of Task T0 (main program)
2. sum1 = 0; sum2 = 0; // sum1 & sum2 are static fields
3. finish() -> {
4.     async() -> {
5.         // Child task computes sum of lower half of array
6.         for(int i=0; i < X.length/2; i++) sum1 += X[i];
7.     });
8.     // Parent task computes sum of upper half of array
9.     for(int i=X.length/2; i < X.length; i++) sum2 += X[i];
10. });
11. // Parent task waits for child task to complete (join)
12. return sum1 + sum2;
```



Computation Graphs (Lecture 2)

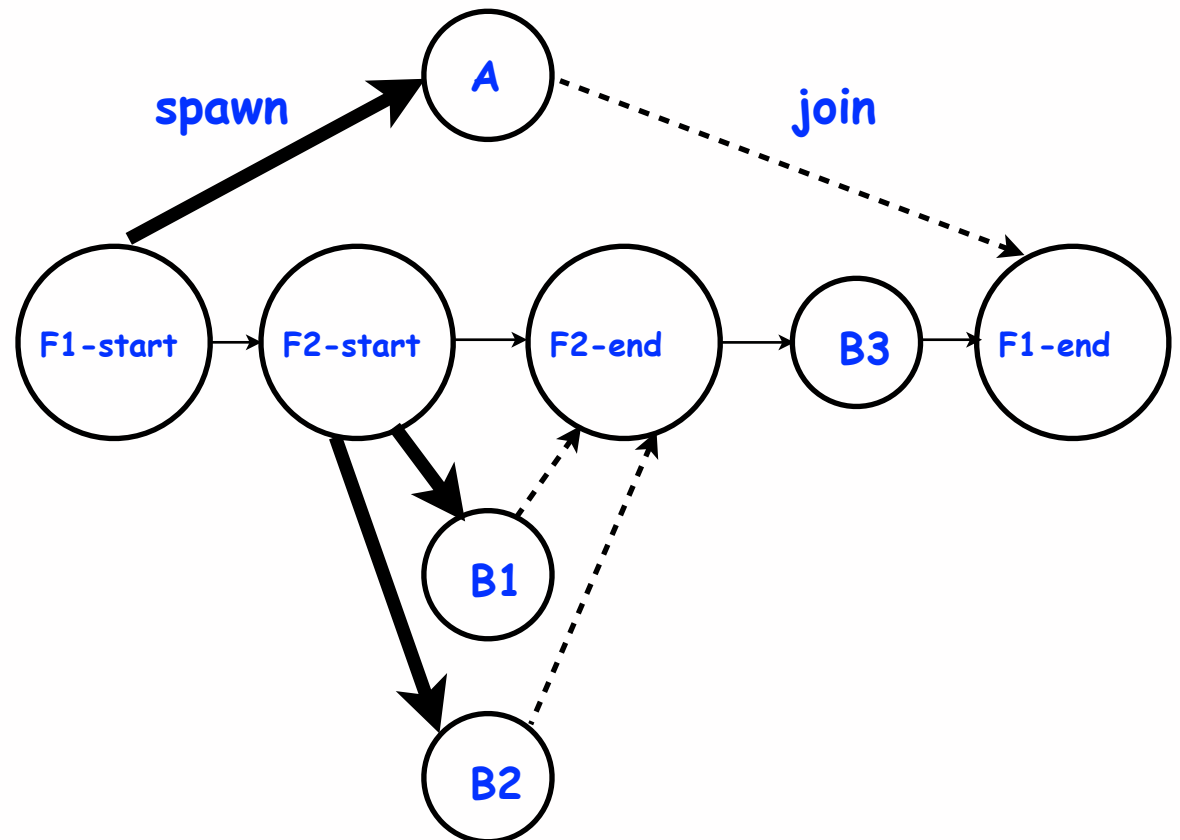
- A Computation Graph (CG) captures the dynamic execution of a parallel program, for a specific input
- CG nodes are “steps” in the program’s execution
 - A step is a sequential subcomputation without any async, begin-finish and end-finish operations
- CG edges represent ordering constraints
 - “Continue” edges define sequencing of steps within a task
 - “Spawn” edges connect parent tasks to child async tasks
 - “Join” edges connect the end of each async task to its IEF’s end-finish operations
- All computation graphs must be acyclic
 - It is not possible for a node to depend on itself
- Computation graphs are examples of “directed acyclic graphs” (dags)



Which statements can potentially be executed in parallel with each other?

```
1.  finish { // F1
2.    async A;
3.  finish { // F2
4.    async B1;
5.    async B2;
6.  } // F2
7.  B3;
8. } // F1
```

Computation Graph



Complexity Measures for Computation Graphs


Define

- $\text{TIME}(N)$ = execution time of node N
- $\text{WORK}(G)$ = sum of $\text{TIME}(N)$, for all nodes N in CG G
 - $\text{WORK}(G)$ is the total work to be performed in G
- $\text{CPL}(G)$ = length of a longest path in CG G , when adding up execution times of all nodes in the path
 - Such paths are called *critical paths*
 - $\text{CPL}(G)$ is the length of these paths (critical path length)
 - $\text{CPL}(G)$ is also the smallest possible execution time for the computation graph



Ideal Parallelism

- Define **ideal parallelism** of Computation G Graph as the ratio, $WORK(G)/CPL(G)$
- Ideal Parallelism is independent of the number of processors that the program executes on, and only depends on the computation graph



Example:

$$WORK(G) = 26$$

$$CPL(G) = 11$$

$$\text{Ideal Parallelism} = WORK(G)/CPL(G) = 26/11 \sim 2.36$$



Bounding the performance of Greedy Schedulers (Lecture 3)

Combine lower and upper bounds to get

$$\max(\text{WORK}(G)/P, \text{CPL}(G)) \leq T_p \leq \text{WORK}(G)/P + \text{CPL}(G)$$

Corollary 1: Any greedy scheduler achieves execution time T_p that is within a factor of 2 of the optimal time (since $\max(a,b)$ and $(a+b)$ are within a factor of 2 of each other, for any $a \geq 0, b \geq 0$).

Corollary 2: Lower and upper bounds approach the same value whenever

- There's lots of parallelism, $\text{WORK}(G)/\text{CPL}(G) \gg P$
- Or there's little parallelism, $\text{WORK}(G)/\text{CPL}(G) \ll P$

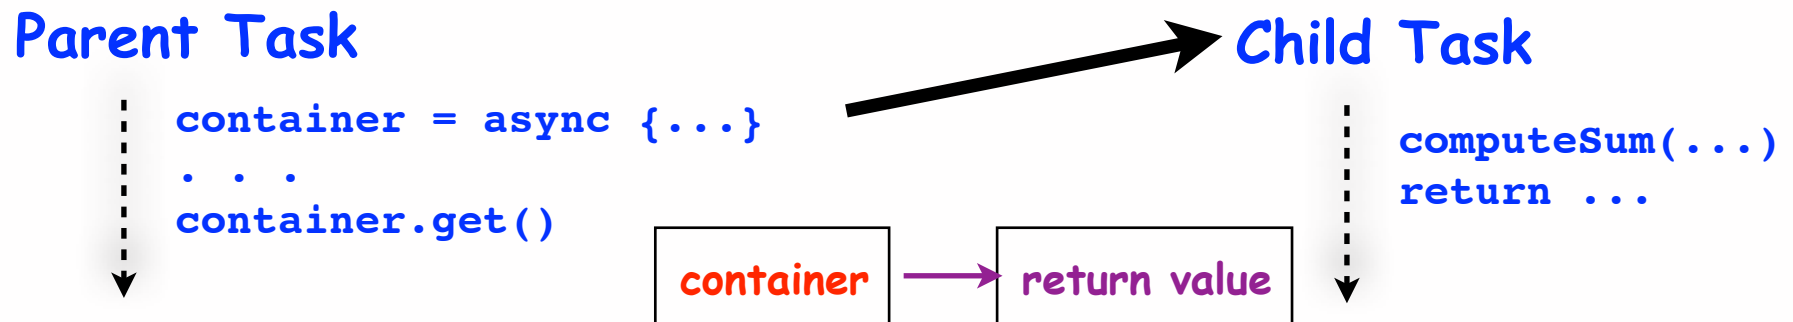


Amdahl's Law [1967] (Lecture 4)

- If $q \leq 1$ is the fraction of WORK in a parallel program that must be executed sequentially for a given input size S , then the best speedup that can be obtained for that program is $\text{Speedup}(S,P) \leq 1/q$.
- Observation follows directly from critical path length lower bound on parallel execution time
 - $\text{CPL} \geq q * T(S,1)$
 - $T(S,P) \geq q * T(S,1)$
 - $\text{Speedup}(S,P) = T(S,1)/T(S,P) \leq 1/q$
- This upper bound on speedup simplistically assumes that work in program can be divided into sequential and parallel portions
 - Sequential portion of WORK = q
 - also denoted as f_s (fraction of sequential work)
 - Parallel portion of WORK = $1-q$
 - also denoted as f_p (fraction of parallel work)
- Computation graph is more general and takes dependences into account

Extending Async Tasks with Return Values (Lecture 5)

- **Example Scenario in PseudoCode**
 1. `// Parent task creates child async task`
 2. `final future container =`
 3. `async { return computeSum(X, low, mid); };`
 4. `. . .`
 5. `// Later, parent examines the return value`
 6. `int sum = container.get();`
- **Two issues to be addressed:**
 - 1) Distinction between **container** and **value** in container (box)
 - 2) Synchronization to avoid race condition in container accesses



HJ Futures: Tasks with Return Values

async { Stmt-Block }

- Creates a new child task that executes **Stmt-Block**, which must terminate with a **return** statement and return value
- Async expression returns a reference to a container of type **future**

Expr.get()

- Evaluates **Expr**, and blocks if Expr's value is unavailable
- Unlike finish which waits for *all* tasks in the finish scope, a get() operation only waits for the specified async expression



Example: Two-way Parallel Array Sum using Future Tasks in HJ-Lib

```
1. // Parent Task T1 (main program)
2. // Compute sum1 (lower half) and sum2 (upper half) in parallel
3. final HjFuture sum1 = future (() -> { // Future Task T2
4.     int sum = 0;
5.     for(int i=0 ; i < X.length/2 ; i++) sum += X[i];
6.     return sum;
7. });
8. final HjFuture sum2 = future (() ->{ // Future Task T3
9.     int sum = 0;
10.    for(int i=X.length/2 ; i < X.length ; i++) sum += X[i];
11.    return sum;
12. });
13. //Task T1 waits for Tasks T2 and T3 to complete
14. int total = sum1.get() + sum2.get();
```



Use of Finish Accumulators to count solutions in Parallel NQueens (Lecture 6)

```
1. final FinishAccumulator ac =
2.     newFinishAccumulator(Operator.SUM, int.class);
3. finish\(ac\) nqueens_kernel(new int[0], 0);
4. System.out.println("No. of solutions = " + ac.get\(\).intValue\(\))
5. . . .
6. void nqueens_kernel(int [] a, int depth) {
7.     if (size == depth) ac.put\(1\);
8.     else
9.         /* try each possible position for queen at depth */
10.        for (int i = 0; i < size; i++) async {
11.            /* allocate a temporary array and copy array a into it */
12.            int [] b = new int [depth+1];
13.            System.arraycopy(a, 0, b, 0, depth);
14.            b[depth] = i;
15.            if (ok(depth+1,b)) nqueens_kernel(b, depth+1);
16.        } // for-async
17. } // nqueens_kernel()
```



Functional vs. Structural Determinism (Lecture 7)

- A parallel program is said to be ***functionally deterministic*** if it always computes the same answer when given the same input
- A parallel program is said to be ***structurally deterministic*** if it always produces the same computation graph when given the same input
- Race-Free Determinism
 - If a parallel program is written using the constructs learned so far (finish, async, futures) and is known to be race-free, *then it must be both functionally deterministic and structurally deterministic*



A Classification of Parallel Programs

Data Race Free?	Functionally Deterministic?	Structurally Deterministic?	Example: String Search variation
Yes	Yes	Yes	Count of all occurrences
No	Yes	Yes	Existence of an occurrence
No	No	Yes	Index of any occurrence
No	Yes	No	“Eureka” extension for existence of an occurrence: do not create more async tasks after occurrence is found
No	No	No	“Eureka” extension for index of an occurrence: do not create more async tasks after occurrence is found

Data-Race-Free Determinism Property implies that it is not possible to write an HJ program with Yes in column 1, and No in column 2 or column 3 (when only using Module 1 constructs)



Map Reduce: Summary (Lecture 8)


- Input set is of the form $\{(k_1, v_1), \dots, (k_n, v_n)\}$, where (k_i, v_i) consists of a key, k_i , and a value, v_i .
 - Assume that the key and value objects are immutable, and that equality comparison is well defined on all key objects.
- Map function f generates sets of intermediate key-value pairs, $f(k_i, v_i) = \{(k_1', v_1'), \dots, (k_m', v_m')\}$. The k_j' keys can be different from k_i key in the input of the map function.
 - Assume that a flatten operation is performed as a post-pass after the map operations, so as to avoid dealing with a set of sets.
- Reduce operation groups together intermediate key-value pairs, $\{(k', v_j')\}$ with the same k' , and generates a reduced key-value pair, (k', v'') , for each such k' , using reduce function g



MapReduce: The Map Step

Input set of
key-value pairs


Flattened intermediate
set of key-value pairs



Source: <http://infolab.stanford.edu/~ullman/mining/2009/mapreduce.ppt>



MapReduce: The Reduce Step



Source: <http://infolab.stanford.edu/~ullman/mining/2009/mapreduce.ppt>



seq clause for async statements (Lecture 10)

```
async seq(cond) <stmt> ≡ if (cond) <stmt> else async <stmt>
```

```
1. // Async task
2. async seq(size < thresholdSize) computeSum(x, lo, mid);
3.
4. // Future example
5. final future<int> sum1 = future seq(size < thresholdSize)
6.                               { return computeSum(x, lo,
   mid); };
```

- **“seq” clause specifies condition under which async should be executed sequentially**
 - **False ⇒ an async is created**
 - **True ⇒ the parent executes async body sequentially**
- **Avoids the need to duplicate code for both cases**



Use of asyncSeq API in HJlib (Quicksort example)

```
1. protected static void quicksort(  
2.     final Comparable[] A, final int M, final int N) {  
3.     if (M < N) {  
4.         // A point in HJ is an integer tuple  
5.         HJPoint p = partition(A, M, N);  
6.         int I = p.get(0);  
7.         int J = p.get(1);  
8.         asyncSeq(I - M <= 5, () -> quicksort(A, M, I));  
9.         asyncSeq(N - J <= 5, () -> quicksort(A, J, N));  
10.    }  
11. }
```

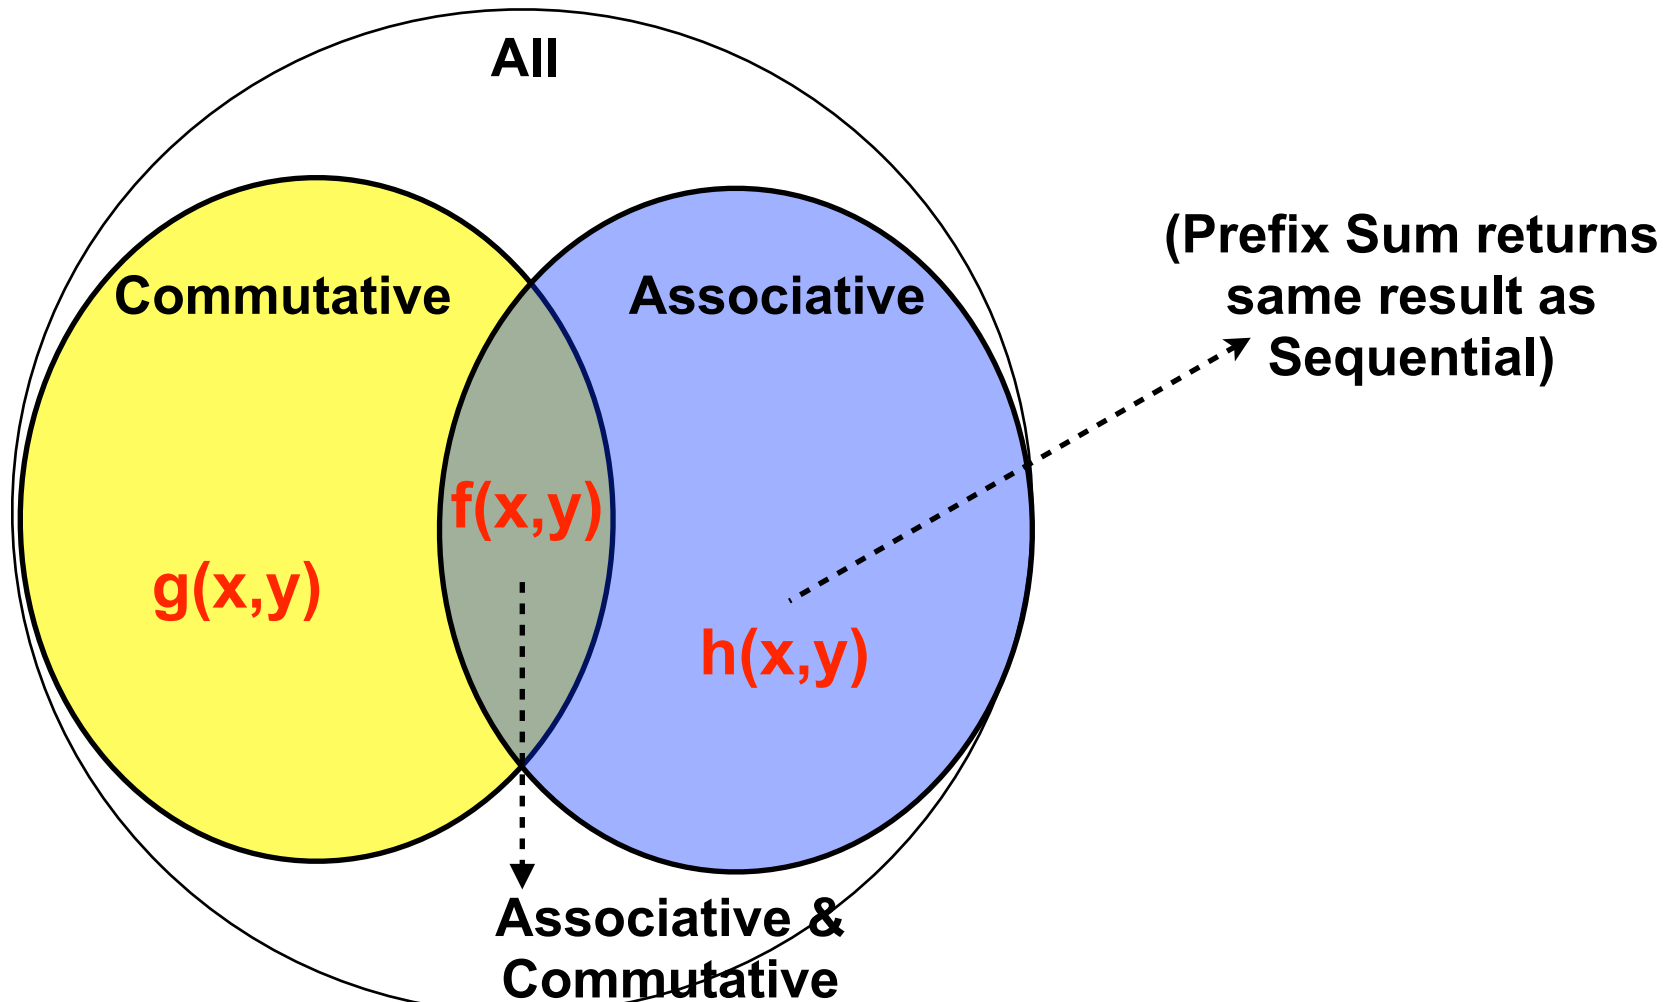


Nqueens example with seq clause

```
1. void nqueensKernel(final int[] a, final int depth,
2.                    final FinishAccumulator ac) {
3.     if (size == depth) {
4.         ac.put(1);
5.         return;
6.     }
7.     /* try each possible position for queen <depth> */
8.     for (int i = 0; i < size; i++) {
9.         final int ii = i;
10.        asyncSeq(depth >= cutoff_value, () -> {
11.            /* allocate a temporary array and copy <a> into it */
12.            final int[] b = new int[depth + 1];
13.            System.arraycopy(a, 0, b, 0, depth);
14.            b[depth] = ii;
15.            if (boardValid((depth + 1), b)) {
16.                nqueensKernel(b, depth + 1, ac);
17.            }
18.        });
19.     }
```



Venn diagram of binary functions



(Prefix Sum & Finish Accumulator return same result as Sequential)



Observations on finish-for-async version (Lecture 11)

- **finish** and **async** are general constructs, and are not specific to loops
 - Not easy to discern from a quick glance which loops are sequential vs. parallel
- Loops in sequential version of matrix multiplication are “perfectly nested”
 - e.g., no intervening statement between “for(i = ...)” and “for(j = ...)”
- The ordering of loops nested between **finish** and **async** is arbitrary
 - They are parallel loops and their iterations can be executed in any order



forall API's in HJlib

(<http://www.cs.rice.edu/~vs3/hjlib/doc/edu/rice/hj/Module1.html>)

- static void
forall(edu.rice.hj.api.HjRegion.HjRegion1D hjRegion,
edu.rice.hj.api.HjProcedureInt1D body)
- static void
forall(edu.rice.hj.api.HjRegion.HjRegion2D hjRegion,
edu.rice.hj.api.HjProcedureInt2D body)
- static void
forall(edu.rice.hj.api.HjRegion.HjRegion3D hjRegion,
edu.rice.hj.api.HjProcedureInt3D body)
- static void forall(int s0, int e0,
edu.rice.hj.api.HjProcedure<java.lang.Integer> body)
- static void forall(int s0, int e0, int s1, int e1,
edu.rice.hj.api.HjProcedureInt2D body)
- static <T> void forall(java.lang.Iterable<T> iterable,
edu.rice.hj.api.HjProcedure<T> body)
- **NOTE: all forall API's include an implicit finish. forasync is like forall, but without the finish.**



forall examples: updates to a two-dimensional Java array

```
// Case 1: loops i,j can run in parallel
forall(0, m-1, 0, n-1, (i, j) -> { A[i][j] = F(A[i][j]);});

// Case 2: only loop i can run in parallel
forall(1, m-1, (i) -> {
    forseq(1, n-1, (j) -> { // Equivalent to "for (j=1;j<n;j++)"
        A[i][j] = F(A[i][j-1]) ;
    }); });

// Case 3: only loop j can run in parallel
forseq(1, m-1, (i) -> { // Equivalent to "for (i=1;i<m;j++)"
    forall(1, n-1, (j) -> {
        A[i][j] = F(A[i-1][j]) ;
    }); });
```



One-Dimensional Iterative Averaging Example

- Initialize a one-dimensional array of $(n+2)$ double's with boundary conditions, $\text{myVal}[0] = 0$ and $\text{myVal}[n+1] = 1$.
- In each iteration, each interior element $\text{myVal}[i]$ in $1..n$ is replaced by the average of its left and right neighbors.
 - Two separate arrays are used in each iteration, one for old values and the other for the new values
- After a sufficient number of iterations, we expect each element of the array to converge to $\text{myVal}[i] = i/(n+1)$
 - In this case, $\text{myVal}[i] = (\text{myVal}[i-1] + \text{myVal}[i+1])/2$, for all i in $1..n$




Illustration of an intermediate step for $n = 8$ (source: Figure 6.19 in Lin-Snyder book)



Example: HJ code for One-Dimensional Iterative Averaging with forseq-forall structure w/ chunking

```
1. int nc = numWorkerThreads();
2. forseq(0, m-1, (iter) -> {
3.     // Compute myNew as function of input array myVal
4.     forallChunked(1, n, n/nc, (j) -> { // Create n tasks
5.         myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
6.     }); // forall
7.     temp=myVal; myVal=myNew; myNew=temp;// Swap myVal & myNew;
8.     // myNew becomes input array for next iteration
9. }); // for
```



Barriers (Lecture 12)

- Question: how can we transform this code so as to ensure that all tasks say hello before any tasks say goodbye?

- Approach 2: insert a “barrier” between the hello's and goodbye's

- “next” statement in HJ's forall loops

```
1. // APPROACH 2
2. forallPhased (0, m - 1, (i) -> {
3.   int sq = i*i;
4.   System.out.println("Hello from task with square = " + sq);
5.   next(); // Barrier
6.   System.out.println("Goodbye from task with square = " + sq);
7. });
```

} Phase 0
} Phase 1

- **next** → each forall iteration suspends at next until all iterations arrive (complete previous phase), after which the phase can be advanced
 - If a forall iteration terminates before executing “next”, then the other iterations do not wait for it
 - Scope of next is the closest enclosing forall statement
 - Special case of “phaser” construct (will be covered later in class)



Observation 1: Scope of synchronization for “next” is closest enclosing forall statement

```
1. forallPhased (0, m - 1, (i) -> {
2.   println("Starting forall iteration " + i);
3.   next(); // Acts as barrier for forall-i
4.   forallPhased (0, n - 1, (j) -> {
5.     println("Hello from task (" + i + "," + j + ")");
6.     next(); // Acts as barrier for forall-j
7.     println("Goodbye from task (" + i + "," + j + ")");
8.   } // forall-j
9.   next(); // Acts as barrier for forall-i
10.  println("Ending forall iteration " + i);
11.}); // forall-i
```



Observation 2: If a forall iteration terminates before “next”, then other iterations do not wait for it

1. `forallPhased (0, m - 1, (i) -> {`
 2. `forseq (0, i, (j) -> {`
 3. `// forall iteration i is executing phase j`
 4. `System.out.println("(" + i + "," + j + ")");`
 5. `next();`
 6. `});`
 7. `});`
- Outer forall-i loop has m iterations, 0...m-1
 - Inner sequential j loop has i+1 iterations, 0...i
 - Line 4 prints (task,phase) = (i, j) before performing a next operation.
 - Iteration i = 0 of the forall-i loop prints (0, 0), performs a next, and then terminates. Iteration i = 1 of the forall-i loop prints (1,0), performs a next, prints (1,1), performs a next, and then terminates. And so on.



HJ code for One-Dimensional Iterative Averaging with grouped forall-foreach structure and barriers (Lecture 13)

```
1. double[] gVal=new double[n+2]; gVal[n+1] = 1;
2. double[] gNew=new double[n+2];
3. HJRegion1D iterSpace = newRectangularRegion1D(1,m);
4. int nc = numWorkerThreads();
5. forallPhased(1, nc, (jj) -> { // Create nc tasks
6.     // Initialize myVal and myNew as local pointers
7.     double[] myVal = gVal; double[] myNew = gNew;
8.     foreach(0, m-1, (iter) -> {
9.         foreach(myGroup(jj,iterSpace,nc), (j) -> {
10.            // Compute MyNew as function of input array MyVal
11.            myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
12.        }); // foreach
13.    next(); // Barrier before executing next iteration of iter loop
14.    // swap local pointers, myVal and myNew
15.    double[] temp=myVal; myVal=myNew; myNew=temp;
16.    // myNew becomes input array for next iter
17. }); // foreach
18. }); // forall
```

This program creates nc async tasks, and performs m barrier operations per task




Single Program Multiple Data (SPMD) Parallel Programming Model

Basic idea

- Run the same code (program) on P workers
- Use the “rank” --- an ID ranging from 0 to $(P-1)$ --- to determine what data is processed by which worker
 - Hence, “single-program” and “multiple-data”
 - Rank is equivalent to index in a top-level “forall (point[i] : [0:P-1])” loop
- Lower-level programming model than dynamic async/finish parallelism
 - Programmer’s code is essentially at the worker level (each forall iteration is like a worker), and work distribution is managed by programmer by using barriers and other synchronization constructs
 - Harder to program but can be more efficient for restricted classes of applications (e.g. for OneDimAveraging, but not for nqueens)
- Convenient for hardware platforms that are not amenable to efficient dynamic task parallelism
 - General-Purpose Graphics Processing Unit (GPGPU) accelerators
 - Distributed-memory parallel machines



Macro-Dataflow Programming (Lecture 14)



- “Macro-dataflow” = extension of dataflow model from instruction-level to task-level operations
- General idea: build an arbitrary task graph, but restrict all inter-task communications to single-assignment variables
 - Static dataflow ==> graph fixed when program execution starts
 - Dynamic dataflow ==> graph can grow dynamically
- Semantic guarantees: race-freedom, determinism
 - Deadlocks are possible due to unavailable inputs (but they are deterministic)



Extending HJ Futures for Macro-Dataflow: Data-Driven Futures (DDFs) and Data-Driven Tasks (DDTs)

```
HjDataDrivenFuture<T1> ddfA = newDataDrivenFuture();
```

- Allocate an instance of a data-driven-future object (container)
- Object in container must be of type T1

```
asyncAwait(ddfA, ddfB, ..., () -> Stmt);
```

- Create a new data-driven-task to start executing **Stmt** after all of **ddfA**, **ddfB**, ... become available (i.e., after task becomes “enabled”)

```
ddfA.put(V) ;
```

- Store object V (of type T1) in **ddfA**, thereby making **ddfA** available
- Single-assignment rule: at most one put is permitted on a given DDF

```
ddfA.get()
```

- Return value (of type T1) stored in **ddfA**
- Can only be performed by **async**'s that contain **ddfA** in their **await** clause (hence no blocking is necessary for DDF gets)



Differences between Futures and DDFs/DDTs

- Consumer task blocks on `get()` for each future that it reads, whereas `async-await` does not start execution till all DDFs are available
- Future tasks cannot deadlock, but it is possible for a DDT to block indefinitely (“deadlock”) if one of its input DDFs never becomes available
- DDTs and DDFs are more general than futures
 - Producer task can only write to a single future object, where as a DDT can write to multiple DDF objects
 - The choice of which future object to write to is tied to a future task at creation time, where as the choice of output DDF can be deferred to any point with a DDT
- DDTs and DDFs can be more implemented more efficiently than futures
 - An “`asyncAwait`” statement does not block the worker, unlike a `future.get()`
 - You will never see the following message with “`asyncAwait`”
 - “ERROR: Maximum number of hj threads per place reached”



Summary of Phaser Construct (Lecture 16)

- **Phaser allocation**
 - `HjPhaser ph = newPhaser(mode);`
 - Phaser `ph` is allocated with registration mode
 - Phaser lifetime is limited to scope of Immediately Enclosing Finish (IEF)
- **Registration Modes**
 - `HjPhaserMode.SIG`, `HjPhaserMode.WAIT`,
`HjPhaserMode.SIG_WAIT`, `HjPhaserMode.SIG_WAIT_SINGLE`
 - NOTE: phaser `WAIT` is unrelated to Java `wait/notify` (which we will study later)
- **Phaser registration**
 - `asyncPhased (ph1.inMode(<mode1>), ph2.inMode(<mode2>), ... () -> <stmt>)`
 - Spawned task is registered with `ph1` in `mode1`, `ph2` in `mode2`, ...
 - Child task's capabilities must be subset of parent's
 - `asyncPhased <stmt>` propagates all of parent's phaser registrations to child
- **Synchronization**
 - `next();`
 - Advance each phaser that current task is registered on to its next phase
All signals are performed, followed by all waits
 - Semantics depends on registration mode
 - Barrier is a special case of phaser, which is why `next` is used for both

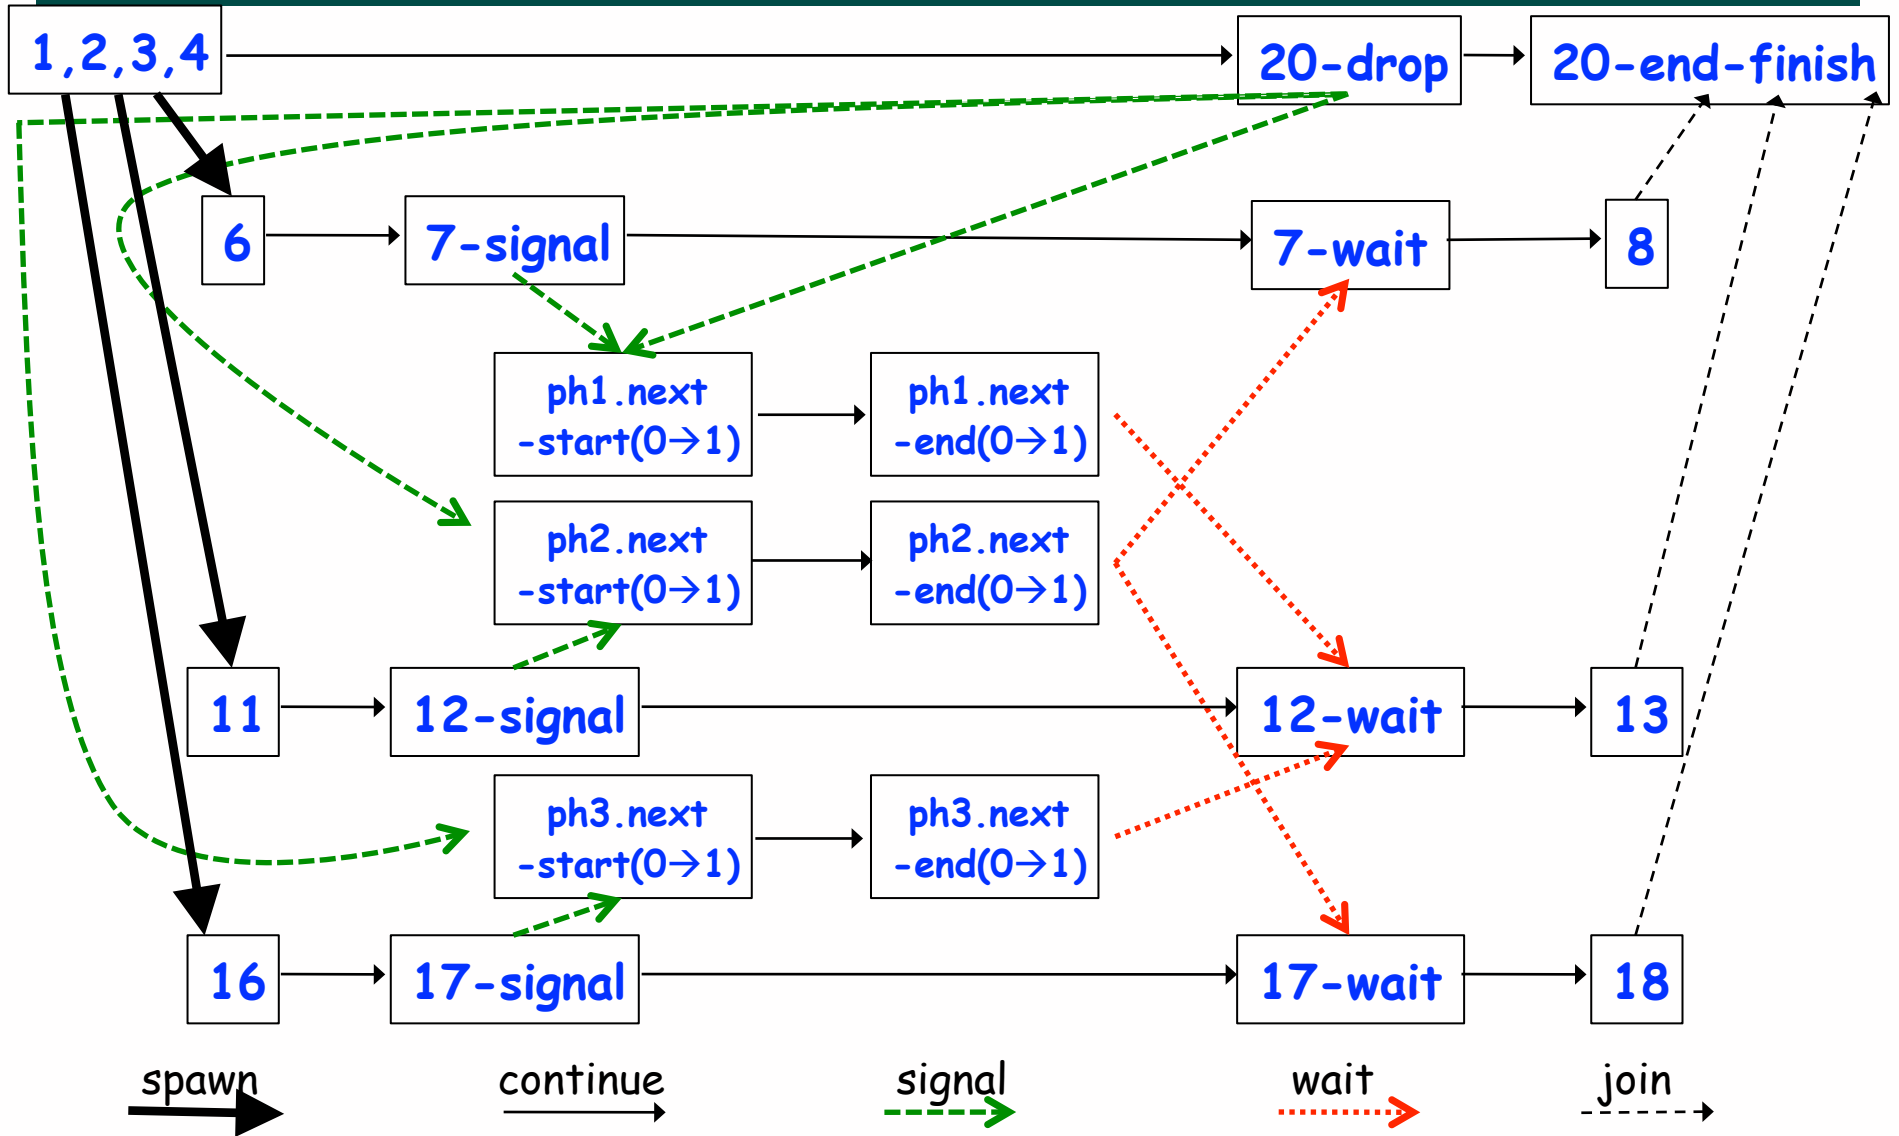


Left-Right Neighbor Synchronization Example for m=3

```
1. finish(() -> { // Task-0
2.     final HjPhaser ph1 = newPhaser(SIG_WAIT);
3.     final HjPhaser ph2 = newPhaser(SIG_WAIT);
4.     final HjPhaser ph3 = newPhaser(SIG_WAIT);
5.     asyncPhased(ph1.inMode(SIG), ph2.inMode(WAIT),
6.         () -> { doPhase1(1);
7.             next(); // signals ph1, waits on ph2
8.             doPhase2(1);
9.         }); // Task T1
10.    asyncPhased(ph2.inMode(SIG), ph1.inMode(WAIT), ph3.inMode(WAIT),
11.        () -> { doPhase1(2);
12.            next(); // signals ph2, waits on ph3
13.            doPhase2(2);
14.        }); // Task T2
15.    asyncPhased(ph3.inMode(SIG), ph2.inMode(WAIT),
16.        () -> { doPhase1(3);
17.            next(); // signals ph3, waits on ph2
18.            doPhase2(3);
19.        }); // Task T3
20.}); // finish
```



Computation Graph for m=3 example



Announcements

- **Take-home midterm exam (Exam 1) will be given after lecture on Wednesday, February 26, 2014**
 - Closed-book, closed computer, written exam that can be taken in any 2-hour duration during that period
 - Will need to be returned to Penny Anderson (Duncan Hall 3080) by 4pm on Friday, February 28, 2014
 - Exam can also be picked up from Penny Anderson starting 2pm on Feb 26th if you're unable to attend lecture.
 - No lecture on Friday, Feb 28th
- **Homework 3 is due by 11:59pm on Wednesday, March 12, 2014**
 - Programming assignment is more challenging than in previous homeworks --- start early!



Scope of Midterm Exam

- **Midterm exam will cover material from Lectures 1 - 17**
 - Lecture 18 (Feb 26th) will be a Midterm review
- **Excerpts from midterm exam instructions**
 - “closed-book, closed-notes, closed-computer”
 - “Record start time when you open the exam, and end time when you finish. The total duration must be at most 2 hours. ”
 - “Since this is a written exam and not a programming assignment, syntactic errors in program text will not be penalized (e.g., missing semicolons, incorrect spelling of keywords, etc) so long as the meaning of your solution is unambiguous.”
 - “If you believe there is any ambiguity or inconsistency in a question, you should state the ambiguity or inconsistency that you see, as well as any assumptions that you make to resolve it.”

