COMP 322: Fundamentals of Parallel Programming

Lecture 19: Midterm Review

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

Contact email: vsarkar@rice.edu

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

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

async S

- Creates a new child task that executes statement S

finish S

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

```
// 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
```

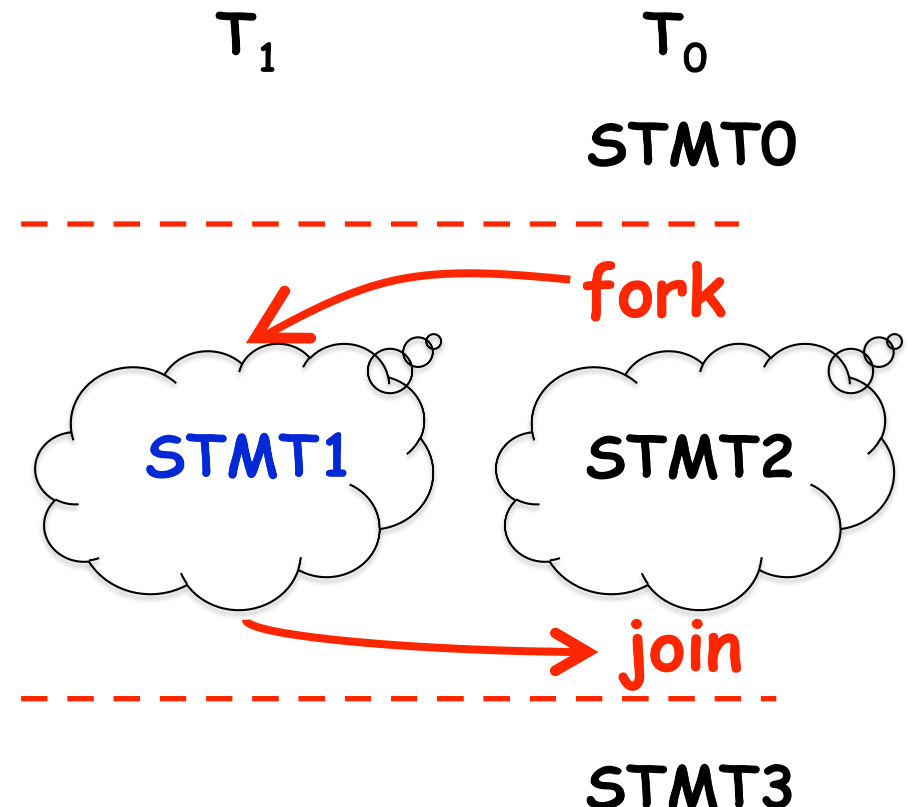

One Possible Solution to Worksheet 1 (Parallel Matrix Multiplication)

```
1.finish {
2.   for (int i = 0 ; i < N ; i++)
3.     for (int j = 0 ; j < N ; j++)
4.       async {
5.         for (int k = 0 ; k < N ; k++)
6.           C[i][j] += A[i][k] * B[k][j];
7.       } // async
8.} // finish
```

This program generates N^2 parallel async tasks, one to compute each C[i][j] element of the output array. Additional parallelism can be exploited within the inner k loop, but that would require more changes than inserting async & finish.

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
```

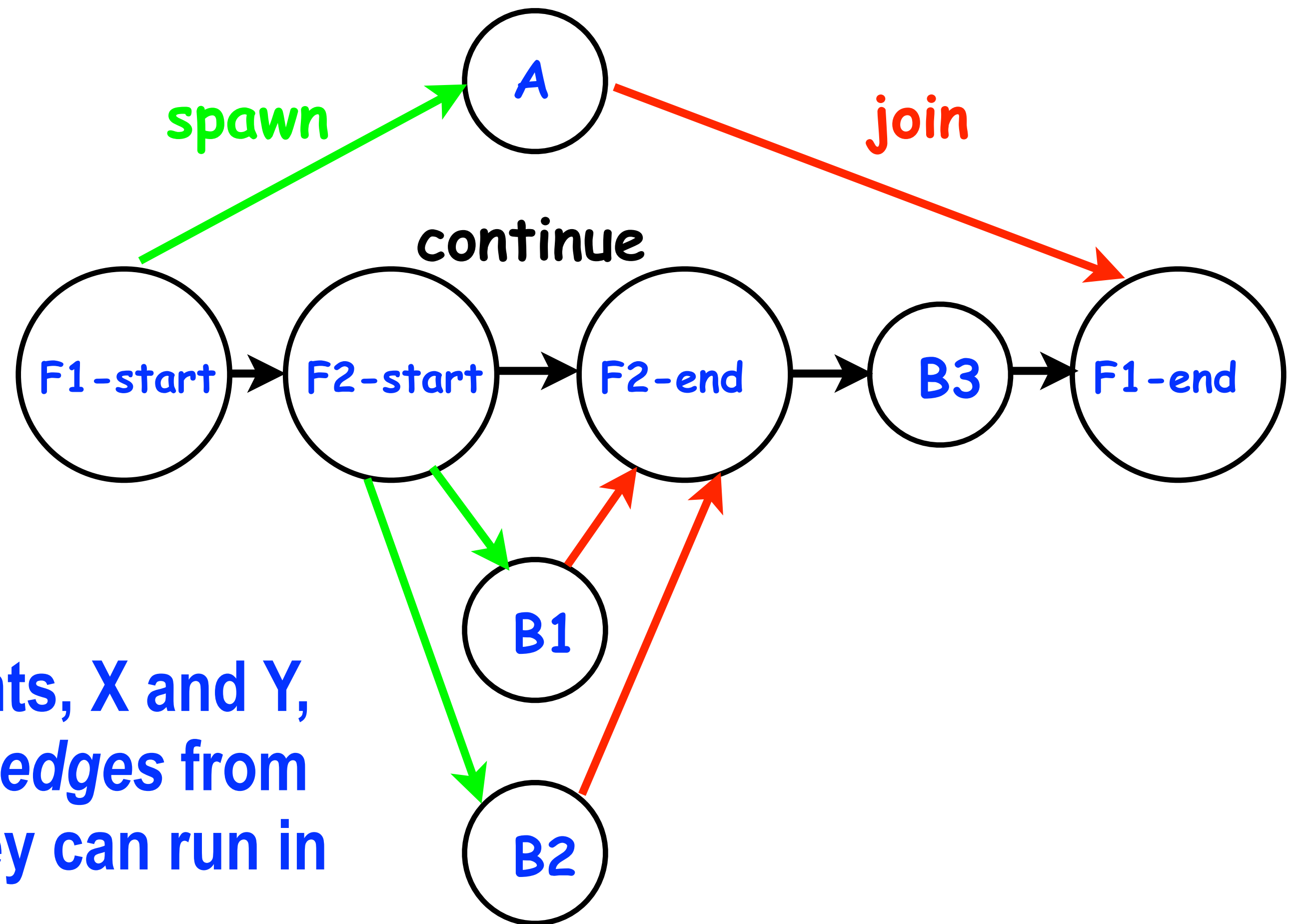

Key idea: If two statements, X and Y, have *no path of directed edges* from one to the other, then they can run in parallel with each other.

Complexity Measures for Computation Graphs

Define

- TIME(N) = execution time of node N
- WORK(G) = sum of TIME(N), for all nodes N in CG G
  - WORK(G) is the total work to be performed in G
- 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*
  - CPL(G) is the length of these paths (critical path length)
  - CPL(G) is also the smallest possible execution time for the computation graph

Which Computation Graph has more ideal parallelism?

Assume that all nodes have TIME = 1, so WORK = 10 for both graphs.

Computation Graph 1

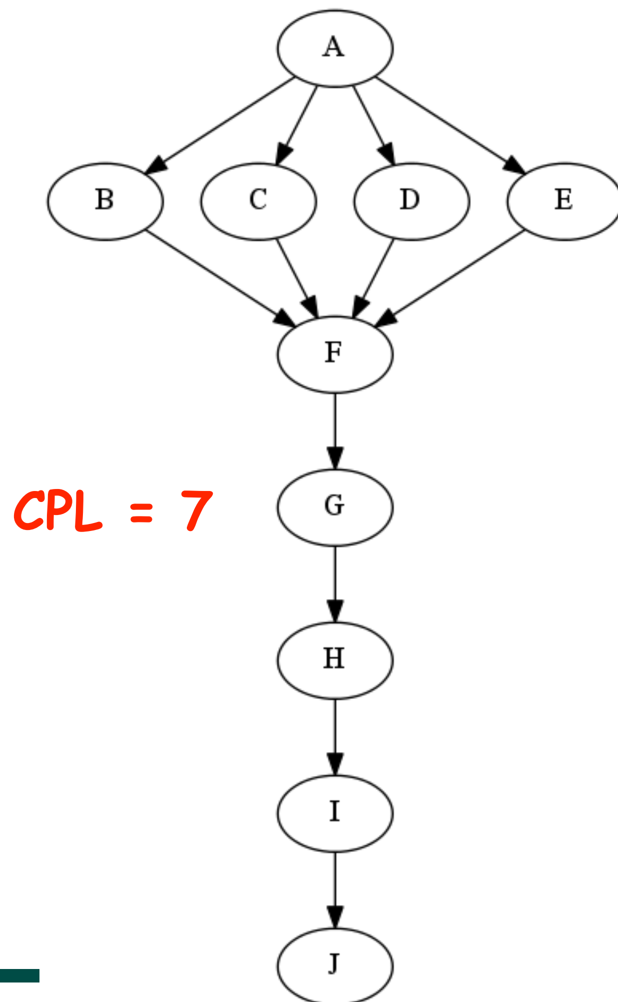

CPL = 7

Computation Graph 2

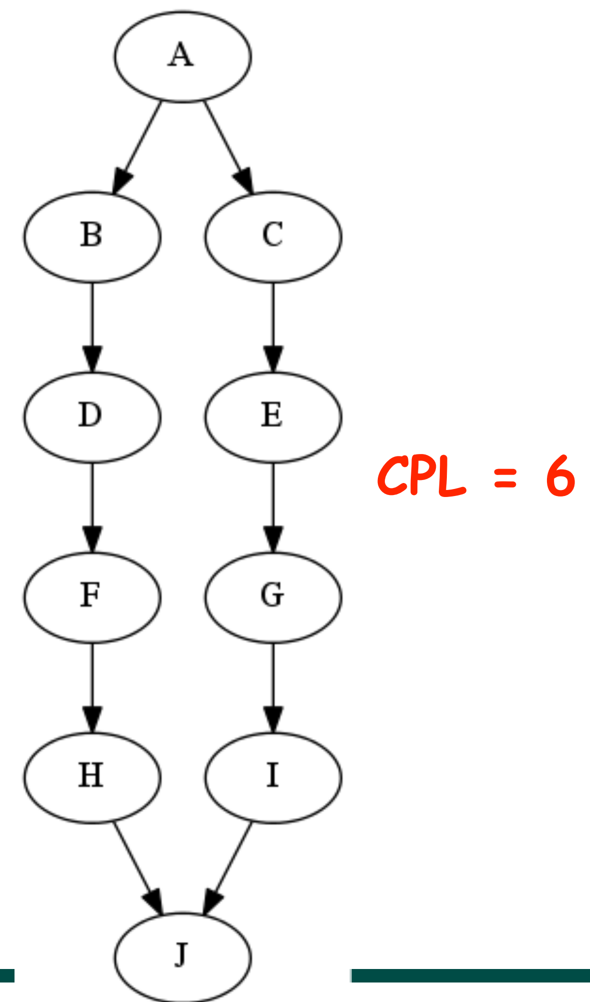

CPL = 6

Data Races

A data race occurs on location L in a program execution with computation graph CG if there exist steps (nodes) S1 and S2 in CG such that:

1. **S1 does not depend on S2 and S2 does not depend on S1, i.e., S1 and S2 can potentially execute in parallel, and**
2. **Both S1 and S2 read or write L, and at least one of the accesses is a write.**

- **A data-race is an error. The result of a read operation in a data race is undefined. The result of a write operation is undefined if there are two or more writes to the same location.**
- **Above definition includes all "potential" data races i.e., we consider it to be a data race even if S1 and S2 execute on the same processor.**

One Possible Solution to Worksheet 2 (Reverse Engineering a Computation Graph)

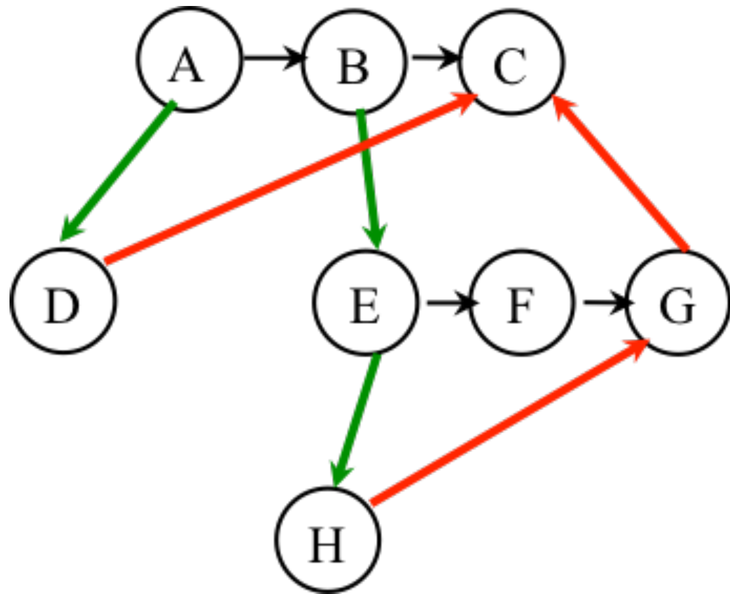


```
1.A();
2.finish { // F1
3.   async D();
4.   B();
5.   async {
6.     E();
7.     finish { // F2
8.       async H();
9.         F();
10.    } // F2
11.    G();
12.  }
13. } // F1
14. C();
```

Observations:

- Any node with out-degree > 1 must be an async (must have an outgoing spawn edge)
- Any node with in-degree > 1 must be an end-finish (must have an incoming join edge
- Adding or removing transitive edges does not impact ordering constraints

Abstract Performance Metrics (Lecture 3)

- Basic Idea
  - Count operations of interest, as in big-O analysis
  - Abstraction ignores many overheads that occur on real systems
- Calls to doWork()
  - Programmer inserts calls of the form, doWork(N), within a step to indicate abstraction execution of N application-specific abstract operation
    - e.g., adds, compares, stencil ops, data structure ops
  - Multiple calls dynamically add to the execution time of current step in computation graph
- Abstract metrics are enabled by calling
  - HjSystemProperty.abstractMetrics.set(true);
- If an HJ program is executed with this option, abstract metrics are printed at end of program execution with WORK(G), CPL(G), Ideal Parallelism = WORK(G) / CPL(G)

One Possible Solution to Worksheet 3 (Multiprocessor Scheduling)

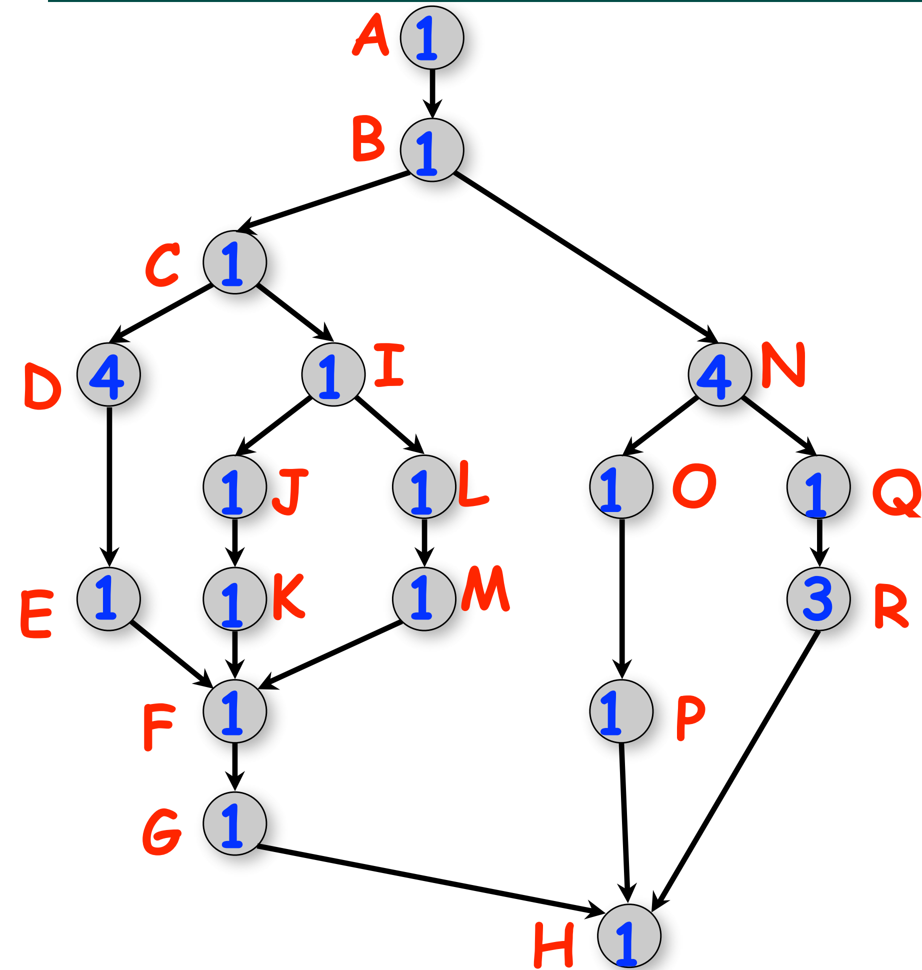

Start time	Proc 1	Proc 2
0	A	
1	B	
2	C	N
3	D	N
4	D	N
5	D	N
6	D	O
7	I	Q
8	J	R
9	L	R
10	K	R
11	M	E
12	F	P
13	G	
14	H	
15		

- As before, WORK = 26 and CPL = 11 for this graph
- T_2 = 15, for the 2-processor schedule on the right
- We can also see that

  $\max(CPL, WORK/2) \le T_2 < CPL + WORK/2$

How many processors should we use? (Lecture 4)

- Define Efficiency(P) = Speedup(P)/ P = $T_1/(P * T_P)$
  - Processor efficiency --- figure of merit that indicates how well a parallel program uses available processors
  - For ideal executions without overhead, 1/P <= Efficiency(P) <= 1
- Half-performance metric
  - $S_{1/2}$ = input size that achieves Efficiency(P) = 0.5 for a given P
  - Figure of merit that indicates how large an input size is needed to obtain efficient parallelism
  - A larger value of $S_{1/2}$ indicates that the problem is harder to parallelize efficiently
- How many processors to use?
  - Common goal: choose number of processors, P for a given input size, S, so that efficiency is at least 0.5

Solution to Worksheet 4

- Estimate T(S,P) ~ WORK(G,S)/P + CPL(G,S) = (S-1)/P + log2(S) for the parallel array sum computation shown in slide 4.
- Assume S = 1024 ==> log2(S) = 10
- Compute for 10, 100, 1000 processors
  - T(P) = 1023/P + 10
  - Speedup(10) = T(1)/T(10) = 1033/112.3 ~ 9.2
  - Speedup(100) = T(1)/T(100) = 1033/20.2 ~ 51.1
  - Speedup(1000) = T(1)/T(1000) = 1033/11.0 ~ 93.7
- Why does the speedup not increase linearly in proportion to the number of processors?
  - Because of the critical path length, log2(S), is a bottleneck

Extending Async Tasks with Return Values (Lecture 5)

Example Scenario (PseudoCode)

```
// Parent task creates child async task
final future container =
                 async { return computeSum(X, low, mid); };
. . .
// Later, parent examines the return value
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

Parent Task

```
container = async {...}
. . .
container.get()
```

Child Task

```
computeSum(...)
return ...
```

container → return value

Worksheet #5 solution: Computation Graphs for Async-Finish and Future Constructs

1) Can you write an HJ program with async-finish constructs that generates a Computation Graph with the same ordering constraints as the graph on the right?
No

2) Can you write an HJ program with future async-get constructs that generates a Computation Graph with the same ordering constraints as the graph on the right? If so, provide a sketch of the program.
Yes, see program sketch in next slide with void futures. A dummy return value can also be used.

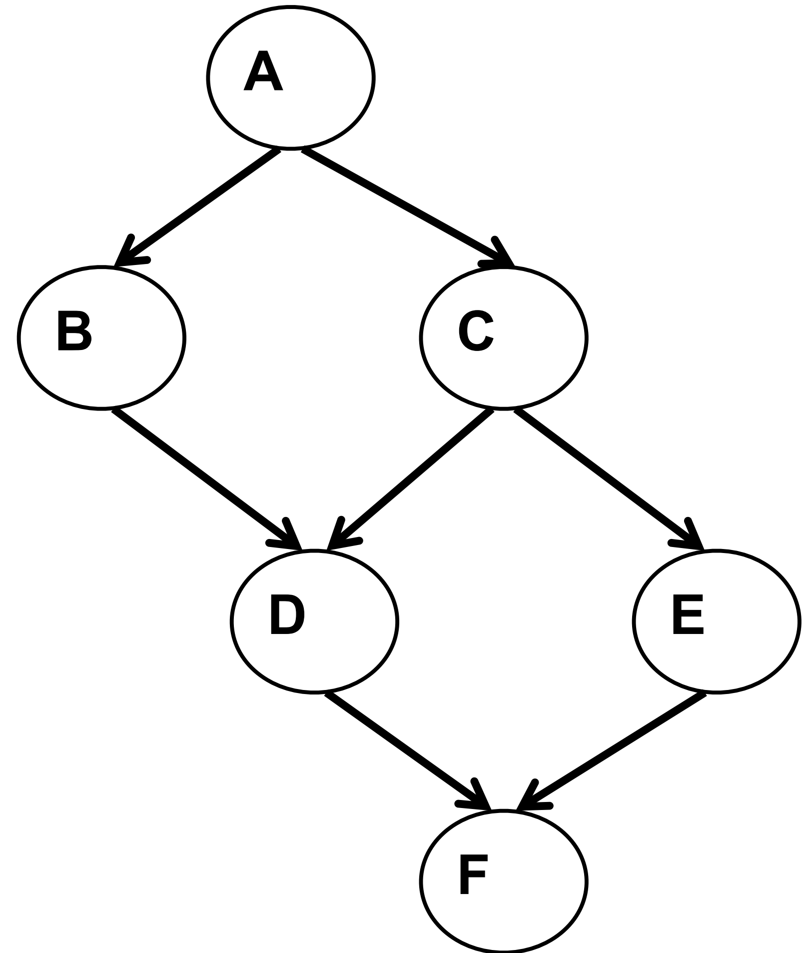

Worksheet #5 solution (contd)

```
1. final HjFuture<Void> A =
2.     future(() -> { return null; });
3. final HjFuture<Void> B =
4.     future(() -> { A.get(); return null; });
5. final HjFuture<Void> C =
6.     future(() -> { A.get(); return null;});
7. final HjFuture<Void> D =
8.     future(() -> { B.get(); C.get(); return null; });
9. final HjFuture<Void> E =
10.    future(() -> {C.get(); return null; });
11.final HjFuture<Void> F =
12.    future(() -> { D.get(); E.get(); return null; });
13.F.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()
```

Worksheet #6 solution: Associativity and Commutativity

Recap:
A binary function f is *associative* if f(f(x,y),z) = f(x,f(y,z)).
A binary function f is *commutative* if f(x,y) = f(y,x).

Worksheet problems:
1) Claim: a Finish Accumulator (FA) can only be used with operators that are *associative and commutative.* Why? What can go wrong with accumulators if the operator is non-associative or non-commutative?
You may get different answers in different executions if the operator is non-associative or non-commutative

2) For each of the following functions, indicate if it is associative and/or commutative.
a) f(x,y) = x+y, for integers x, y, is associative and commutative

b) g(x,y) = (x+y)/2, for integers x, y, is commutative but not associative

c) h(s1,s2) = concat(s1, s2) for strings s1, s2, e.g., h("ab","cd") = "abcd", is associative but not commutative

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
- *Data-Race-Free Determinism Property*
  - If a parallel program is written using the constructs learned so far (finish, async, futures) and is known to be data-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 {(k1, v1), . . . (kn, vn)}, where (ki, vi) consists of a key, ki, and a value, vi.
  - 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(ki,vi) = {(k1' ,v1'),...(km',vm')}. The kj' keys can be different from ki 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', vj')} with the same k', and generates a reduced key-value pair, (k',v''), for each such k', using reduce function g

Worksheet #8 solution: Analysis of Map Reduce Example

Analyze the total WORK and CPL for the Map reduce example in the previous slide, under the following assumptions:

- **Assume that each Map step has WORK = number of input words, and CPL=1**
- **Assume that each Reduce step has WORK = number of input word-count pairs, and CPL = log2(# occcurences for input word with largest # pairs)**

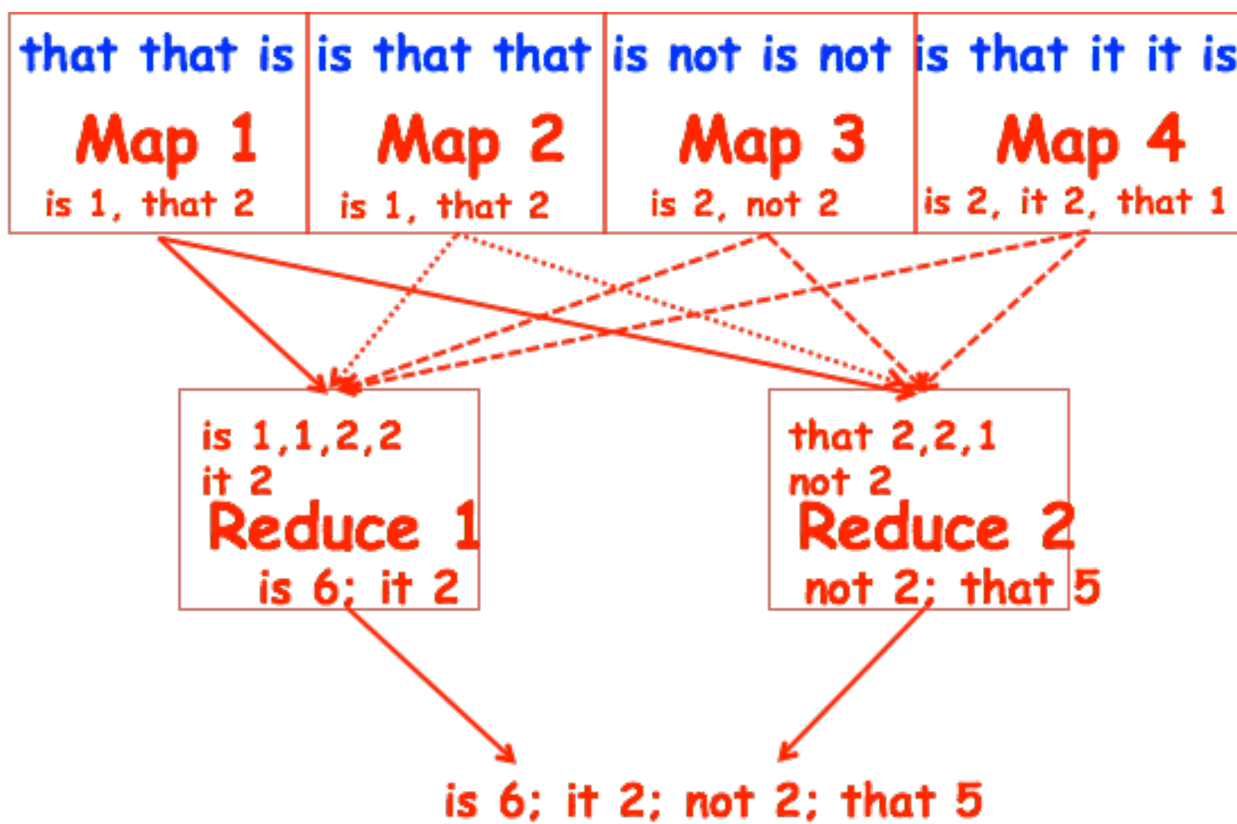

WORK/CPL for all Map steps:

- **WORK = 15**
- **CPL = 1**

WORK/CPL for Reduce 1 step:

- **WORK = 5**
- **CPL = log2(4) = 2**

WORK/CPL for Reduce 2 step:

- **WORK = 4**
- **CPL = log2(3) = 2 (round up)**

Total WORK and CPL

- **WORK = 15+5+4 = 24**
- **CPL = 1 + 2 = 3**

Memoization (Lecture 9)

- Memoization is the idea of saving and reusing previously computed values of a function rather than recomputing them
  - A space-time tradeoff
- A function can only be memoized if it is *referentially transparent*, i.e. functional
- Related to caching
  - memoized function "remembers" the results corresponding to some set of specific inputs
  - memoized function populates its cache of results transparently on the fly, as needed, rather than in advance

Example: Binomial Coefficient (parallel memoized version w/ futures)

```
1.  final Map<Pair<Int, Int>, future<Int>> cache = new ...;
2.  int choose(final int N, final int K) {
3.      final Pair<Int, Int> key = Pair.factory(N, K);
4.      if (cache.contains(key)) {
5.          return cache.get(key).get();
6.      }
7.      future<Int> f = future {
7.        if (N == 0 || K == 0 || N == K) return 1;
8.        int left = future { return choose (N-1, K-1); }
9.        int right = future { return choose (N-1, K); }
12.       return left.get() + right.get();
13.     }
14.     cache.put(key, f);
15.     return f.get();
16. }
```

- Assumes availability of a “thread-safe” cache library, e.g., ConcurrentHashMap

Worksheet #9 solution: Parallelizing Pascal's Triangle with Futures and Memoization

There are four variants of the Binomial Cooefficients program provided in four different HJlib methods in the next page:

a. Sequential Recursive without Memoization (chooseRecursiveSeq())

b. Parallel Recursive without Memoization (chooseRecursivePar())

c. Sequential Recursive with Memoization (chooseMemoizedSeq())

d. Parallel Recursive with Memoization (chooseMemoizedPar())

Your task is to analyze the WORK, CPL, and Ideal Parallelism for these four versions, for the input N = 4, and K = 2. Assume that each call to ComputeSum() has COST = 1, and all other operations are free. Complete all entries in the table:

Variant	Work	CPL	Ideal Parallelism
chooseRecursiveSeq	5	5	1
chooseRecursivePar	5	3	5/3 = 1.67
chooseMemoizedSeq	4	4	1
chooseMemoizedPar	4	3	4/3 = 1.33

One-Dimensional Iterative Averaging Example (Lecture 10)

- Initialize a one-dimensional array of (n+2) double's with boundary conditions, myVal[0] = 0 and myVal[n+1] = 1.
- In each iteration, each interior element 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 myVal[i] = (myVal[i-1]+myVal[i+1])/2, for all i in 1..n

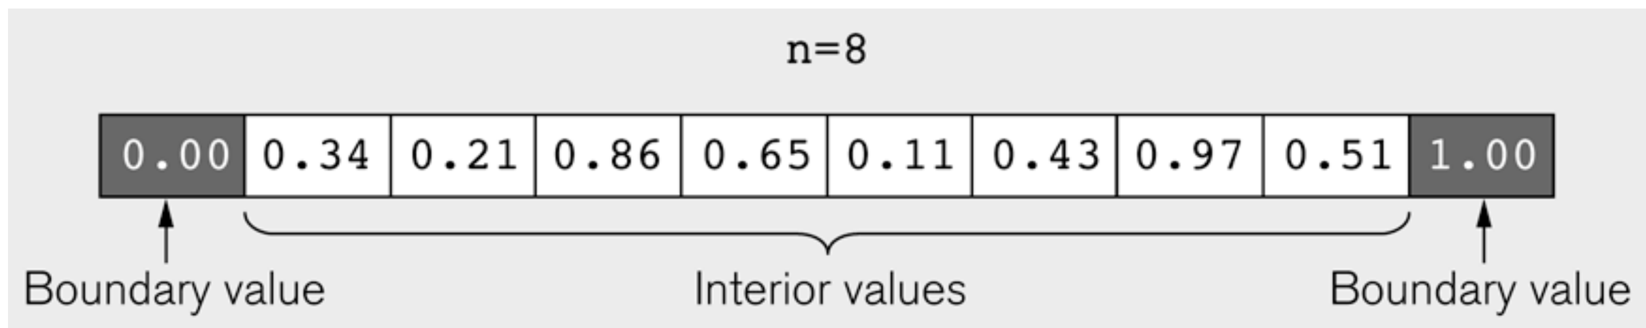

Illustration of an intermediate step for n = 8 (source: Figure 6.19 in Lin-Snyder book)

HJ code for One-Dimensional Iterative Averaging using nested forseq-forall structure

```
1. float[] myVal = new float[n+2];
2. float[] myNew = new float[n+2];
3.  … // Intialize myVal, m, n
4. forseq(0, m-1, (iter) -> {
5.    // Compute MyNew as function of input array MyVal
6.      forall(1, n, (j) -> { // Create n tasks
7.          myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
8.      }); // forall
9.    // What is the purpose of line 10 below?
10.   float[] temp=myVal; myVal=myNew; myNew=temp;
11.   // myNew becomes input array for next iteration
12.}); // for
```

Solution to Worksheet #10: One-dimensional Iterative Averaging Example

1) Assuming n=9 and the input array below, perform a "half-iteration" of the iterative averaging example by only filling in the blanks for odd values of j in the myNew[] array (different from the real algorithm). Recall that the computation is "myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;"

index, j	0	1	2	3	4	5	6	7	8	9	10
myVal	0	0	0.2	0	0.4	0	0.6	0	0.8	0	1
myNew	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1

2) Will the contents of myVal[] and myNew[] change in further iterations?

No, this represents the converged value (equilibrium/fixpoint).

3) If m is large enough, write the formula for the final value of myNew[i] as a function of i and n.

After a large number of iterations, the iterated averaging code will converge with myNew[i] = myVal[i] = i / (n+1)

Barriers (Lecture 11)

- **Question: how can we transform this code so as to ensure that all tasks say hello before any tasks say goodbye, without having to change local ?**
- **Approach 2: insert a "barrier" ("next" statement) between the hello's and goodbye's**

```
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 (lines 3–5)

Phase 1 (lines 6–7)

- **next ➔ each forall iteration waits at barrier until all iterations arrive (previous phase is completed), after which the next phase can start**
  - **Scope of next is the closest enclosing forall statement**
  - **If a forall iteration terminates before executing "next", then the other iterations don't wait for it**
  - **Special case of "phaser" construct (will be discussed later in class)**

Worksheet #11: Forall Loops and Barriers

Draw a "barrier matching" figure similar to slide 13 for the code fragment below.

```
1. String[] a = { "ab", "cde", "f" };
2. . . . int m = a.length; . . .
3. forallPhased (0, m-1, (i) -> {
4.   for (int j = 0; j < a[i].length(); j++) {
5.     // forall iteration i is executing phase j
6.     System.out.println("(" + i + "," + j + ")");
7.     next();
8.   }
9. });
```

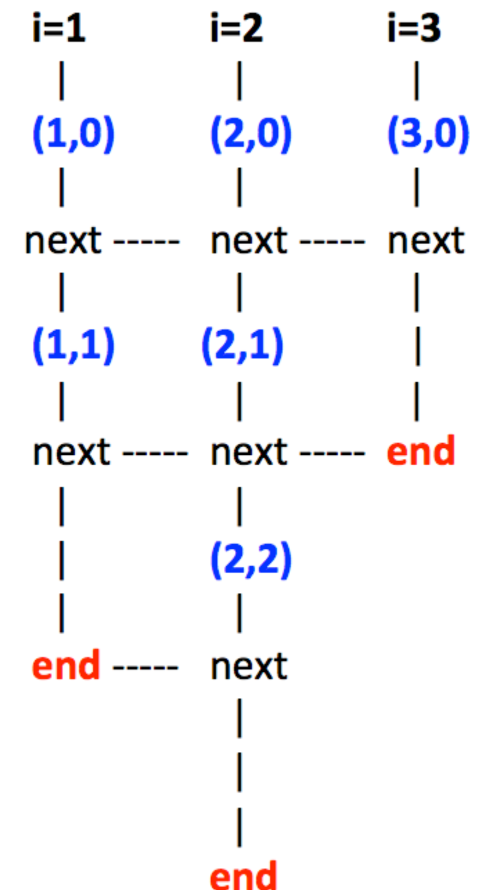

HJ code for One-Dimensional Iterative Averaging with forall-forseq structure and barriers (Lecture 12)

```
1.  double[] gVal=new double[n+2]; gVal[n+1] = 1;
2.  double[] gNew=new double[n+2];
3.  forallPhased(1, n, (j) -> { // Create n tasks
4.    // Initialize myVal and myNew as local pointers
5.    double[] myVal = gVal; double[] myNew = gNew;
6.    forseq(0, m-1, (iter) -> {
7.      // Compute MyNew as function of input array MyVal
8.      myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
9.      next(); // Barrier before executing next iter
10.     // Swap local pointers, myVal and myNew
11.     double[] temp=myVal; myVal=myNew; myNew=temp;
12.     // myNew becomes input array for next iteration
13.   }); // forseq
14. }); // forall
```

Worksheet #12 Solution: Iterative Averaging Revisited

Answer the questions in the table below for the versions of the Iterative Averaging code shown in slides 5, 7, 8, 10. Write in your answers as functions of m, n, and nc.

	Slide 5	Slide 7	Slide 8	Slide 10
How many tasks are created (excluding the main program task)?	m*n	m*nc	n	nc
How many barrier operations (calls to next per task) are performed?	0	0	m	m

The SPMD version in slide 10 is the most efficient because it only creates nc tasks. (Task creation is more expensive than a barrier operation.)

Using Java's Fork/Join Library (Lecture 13)

- Thus far, your introduction to parallel programming has been through HJlib's API with lambda parameters
- Today, we will look at popular library for task parallelism available since Java 7 (pre-dates Java 8 lambdas)
- We can perform recursive subdivision using the Fork/Join libraries provided in Java 7 as follows:

```
public abstract class RecursiveAction extends ForkJoinTask<Void>
{
   protected abstract void compute();
   …
}
```

Solution to Worksheet #13: RecursiveAction Computation Graph

1) Consider the compute() method on slide 8 in Lecture 13 (also in the next slide). Let us suppose we supply it with an 8 element array with values [0,1,2,3,4,5,6,7] and THRESHOLD value of 2. Draw a computation graph corresponding to a call to compute().

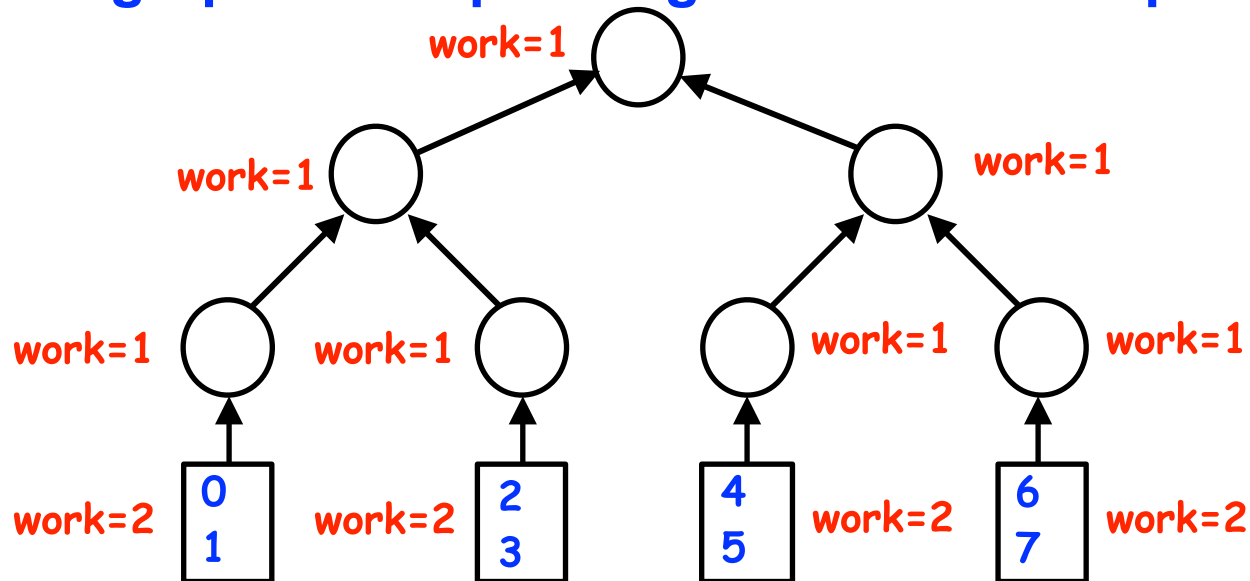

2) Define each leaf computation as 2 units of work and each recursive subdivision as one unit of work.

What is the total work? 4*2 + 1*7 = 15

What is the critical path length? 1 + 1 + 1 + 2 = 5

Example of Compute function for ForkJoin framework (Worksheet #13)

```
1.   protected void compute() {
2.     if (hi - lo < THRESHOLD) {
3.       for (int i = lo; i < hi; ++i)
4.         array[i] = array[i] / (i + 1);
5.     }
6.     else {
7.       int mid = (lo + hi)/2;
8.       invokeAll(new DivideTask(array,lo,mid),
9.                 new DivideTask(array,mid+1,hi));
10.     }
11.  }
```

Macro-Dataflow Programming (Lecture 14)

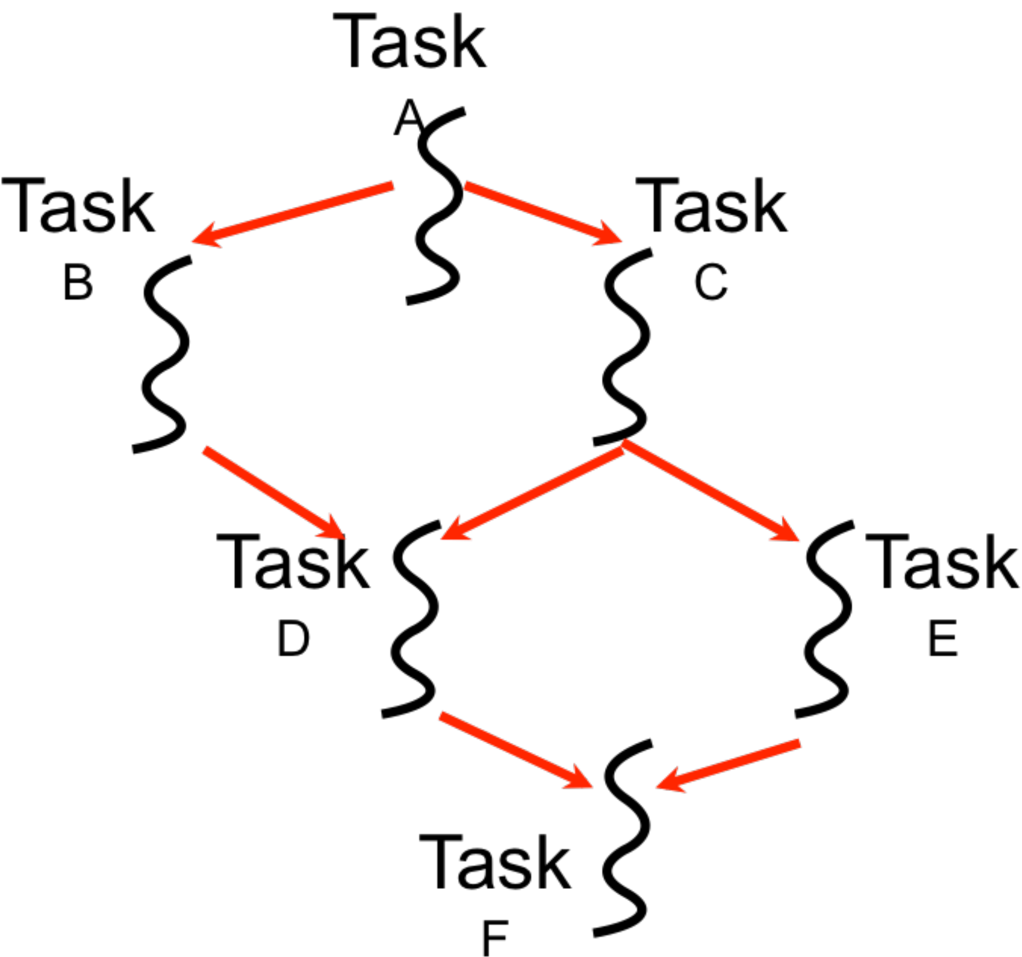

- "Macro-dataflow" = extension of dataflow model from instruction-level to task-level operations
- General idea: build an arbitrary task graph, but restrict all inter-task communications to single-assignment variables
  - Static dataflow ==> graph fixed when program execution starts
  - Dynamic dataflow ==> 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
- Used to implement “edges” in a computation graph

```
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”)
- Used to implement “nodes” in a computation graph

```
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
- Throws an exception if put() has not been performed
  - **Should be performed by async’s that contain ddfA in their await clause, or if there’s some other synchronization to guarantee that the put() was performed**

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
  - Consumer tasks can be created before the producer tasks
- DDTs and DDFs can be more implemented more efficiently than futures
  - An "asyncAwait" statement does not block the worker, unlike a future.get()

Worksheet #14 solution: Data-Driven Tasks

For the example below, will reordering the five async statements change the meaning of the program (assuming that the semantics of the reader/writer methods depends only on their parameters)? If so, show two orderings that exhibit different behaviors. If not, explain why not. (You can use the space below this slide for your answer.)

```
1. DataDrivenFuture left = new DataDrivenFuture();
2. DataDrivenFuture right = new DataDrivenFuture();
3. finish {
4.   async await(left) leftReader(left); // Task3
5.   async await(right) rightReader(right); // Task5
6.   async await(left,right)
7.         bothReader(left,right); // Task4
8.   async left.put(leftWriter()); // Task1
9.   async right.put(rightWriter());// Task2
10. }
```

No, reordering consecutive async's will never change the meaning of the program, whether or not the async's have await clauses.

Summary: Abstract vs. Real Performance in HJlib (Lecture 15)

- **Abstract Performance**
  - Abstract metrics focus on operation counts for WORK and CPL, regardless of actual execution time
- **Real Performance**
  - HJlib uses ForkJoinPool implementation of Java Executor interface with Blocking or Cooperative Runtime (option-controlled)

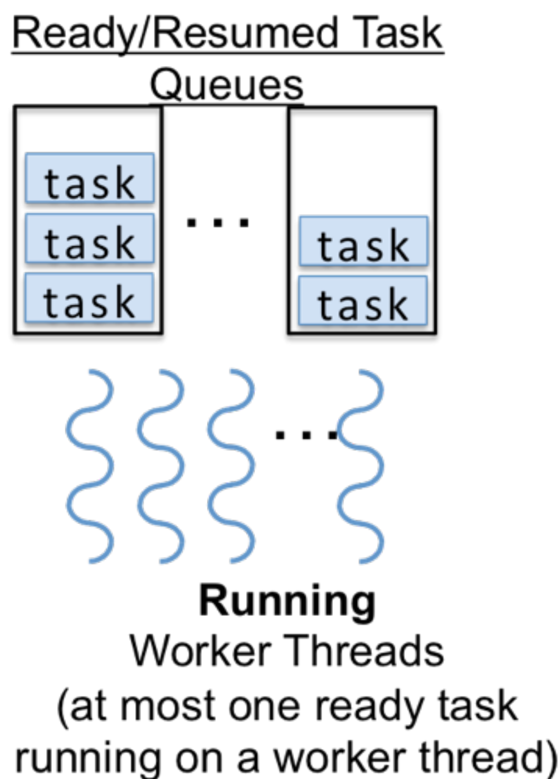

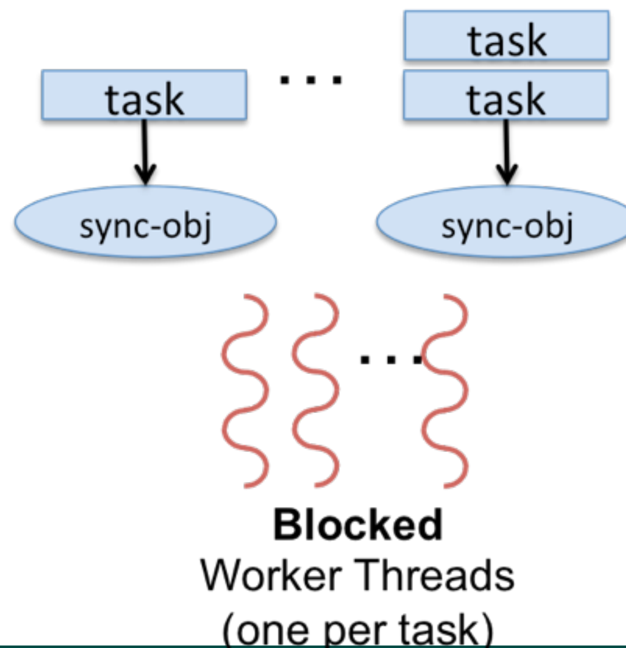

We'll study ForkJoinPool and other Java libraries in detail later in the course --- they manage parallelism at a lower level than HJ

Worksheet #15 solution: Work-First vs. Help-First Work-Stealing Policies

For each of the continuations below, label it as “WF” if a work-first worker can switch from one task to another at that point and as “HF” if a help-first worker can switch from one task to another at that point. Some continuations may have both labels.

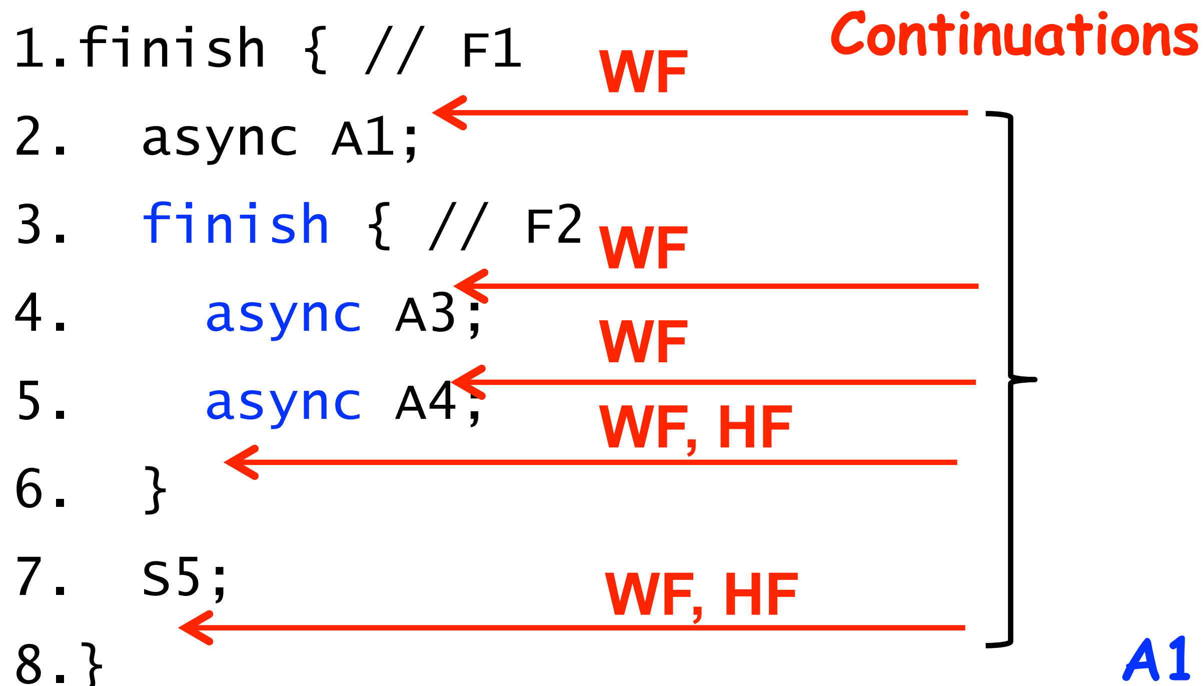

fork
activity
fork
fork
A3
A4
activity
activity
A1
activity
join
F2
join
F1

SPAWN
JOIN
FUTURE
ISOLATED
CONTINUATION
SIGNAL
WAIT
PHASE

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 (ph_1.inMode(<$mode_1$>), ph_2.inMode(<$mode_2$>), … () -> <stmt>)
    - Spawned task is registered with ph_1 in $mode_1$, ph_2 in $mode_2$, …
    - 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
    - Semantics depends on registration mode
    - Barrier is a special case of phaser, which is why *next* is used for both

Solution to Worksheet #16: Left-Right Neighbor Synchronization using Phasers

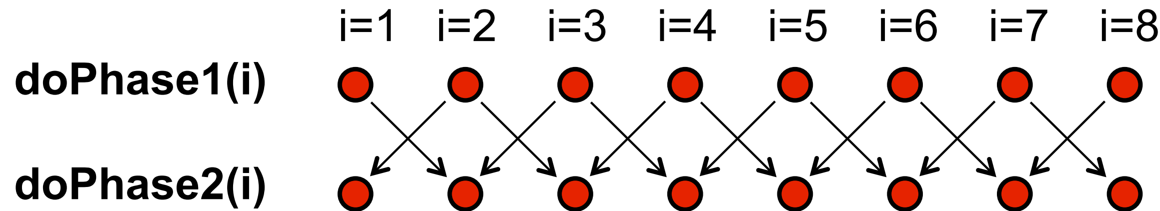

Complete the phased clause below to implement the left-right neighbor synchronization shown above.

```
1. finish (() -> {
2.   final HjPhaser[] ph =
         new HjPhaser[m+2]; // array of phaser objects
3.   forseq(0, m+1, (i) -> { ph[i] = newPhaser(SIG_WAIT) });
4.   forseq(1, m, (i) -> {
5.     asyncPhased(
          ph[i-1].inMode(WAIT),
          ph[i].inMode(SIG),
          ph[i+1].inMode(WAIT), () -> {
6.      doPhase1(i);
7.      next();
8.      doPhase2(i); }); // asyncPhased
9.   }); // forseq
10.}); // finish
```

Complexity Analysis of One-Dimensional Pipeline (Lecture 17)

- Assume
  - n = number of items in input sequence
  - p = number of pipeline stages
  - each stage takes 1 unit of time to process a single data item
- WORK = $n \times p$ is the total work for all data items
- CPL = $n + p - 1$ is the critical path length of the pipeline
- Ideal parallelism, PAR = WORK/CPL = $np/(n + p - 1)$
- Boundary cases
  - $p = 1 \rightarrow$ PAR $= n/(n + 1 - 1) = 1$
  - $n = 1 \rightarrow$ PAR $= p/(1 + p - 1) = 1$
  - $n = p \rightarrow$ PAR $= p/(2 - 1/p) \approx p/2$
  - $n \gg p \rightarrow$ PAR $\approx p$

Solution to Worksheet #17: Critical Path Length for Computation with Signal Statement

Compute the WORK and CPL values for the program shown below. (WORK = 204, CPL = 102). How would they be different if the signal() statement was removed? (CPL would increase to 202.)

```
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
```

Implementing a pipeline stage with phaser-specific doNext() operations (Lecture 18)

```
1. asyncPhased(ph1.inMode(WAIT),
2.              ph2.inMode(SIG), () -> {
3.   for (int i = 0; i < rounds; i++) {
4.       // Wait-only operation on ph1
5.       ph1.doNext();
6.       x = buffer1.remove();
7.       y = f(x);
8.        buffer2.insert();
9.       // Signal-only operation on ph2
10.       ph2.doNext();
11.   }
12. });
```

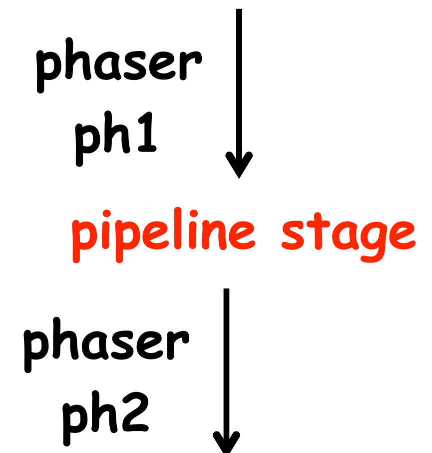

Solution to Worksheet #18: Semantic Classification of Parallel Programs

For each case below, either sketch a parallel program that satisfies the given properties, or explain why no such program exists:

1. A program that satisfies the Serial Elision property, but not the Deadlock Freedom Property
   - No such program exists. If a parallel program satisfies the Serial Elision property, the sequential version can run to completion without deadlock.
2. A program that satisfies the Structural Determinism property, but not the Deadlock Freedom property
   - The program with two DDFs, left and right, and two DDTs such that each produces one DDF and awaits the other is one such example (slide 13, Lecture 18)
3. A program with data races that satisfies the Functional and Structural Determinism properties
   - As an example, take any data-race-free program, and add (say) the statements X=1 and X=2 in any two parallel steps, but don't use the value of X for anything else.

Summary of Parallel Programming Constructs you've learned so far

- Task Parallelism (Unit 1)
  - —Async (task creation)
  - —Finish (structured task termination)
- Functional Parallelism (Unit 2)
  - —Future (task creation)
  - —Future get() (task termination with return value)
  - —Accumulators (functional reduction)
  - —Map-Reduce (functional parallelism & reduction on key-value pairs)
- Loop Parallelism (Unit 3)
  - —Forall (parallel loops)
  - —Barriers (all-to-all synchronization)
- Dataflow Parallelism (Unit 4)
  - —Data-Driven Tasks (dataflow parallelism)
  - —Phasers (point-to-point synchronization)
  - —Phaser-specific next operations

Announcements

- Take-home midterm exam (Exam 1) will be given today after lecture
- Closed-book, closed computer, written exam that can be taken in any 2-hour duration during that period
  - Will need to be returned to Bel Martinez (Duncan Hall 3122) by 4pm on Friday, February 27, 2015
    - Exam can also be picked up from Bel Martinez starting 2pm today if you're unable to attend lecture.
  - No lecture on Friday, Feb 27th
  - Today's lab is optional
    - Unix command line
    - Exercise with phasers
- Homework 3 is due by by 5:00pm on Friday, March 13, 2015
  - Programming assignment is more challenging than in previous homeworks --- start early!

Scope of Midterm Exam

- Midterm exam will cover material from Lectures 1 - 18
  - —Lecture 19 (TODAY) is 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."
- Good luck!