
COMP 322: Fundamentals of Parallel Programming

Lecture 18: Midterm Review

Vivek Sarkar, Shams Imam
Department of Computer Science, Rice University

Contact email: vsarkar@rice.edu, shams.imam@twosigma.com

<http://comp322.rice.edu/>

COMP 322

Lecture 18

22 February 2015



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

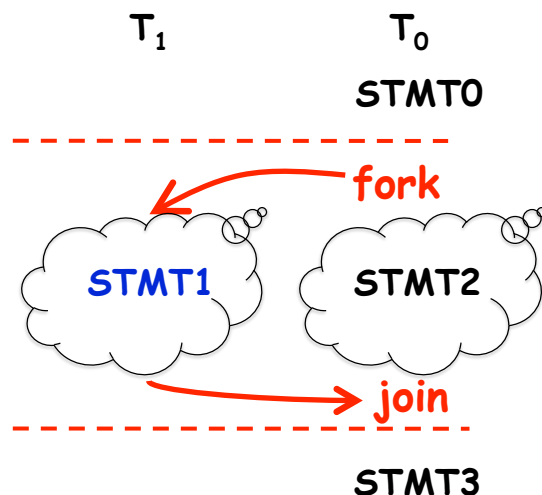
async S

- 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 Problem #2 in 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] = 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)

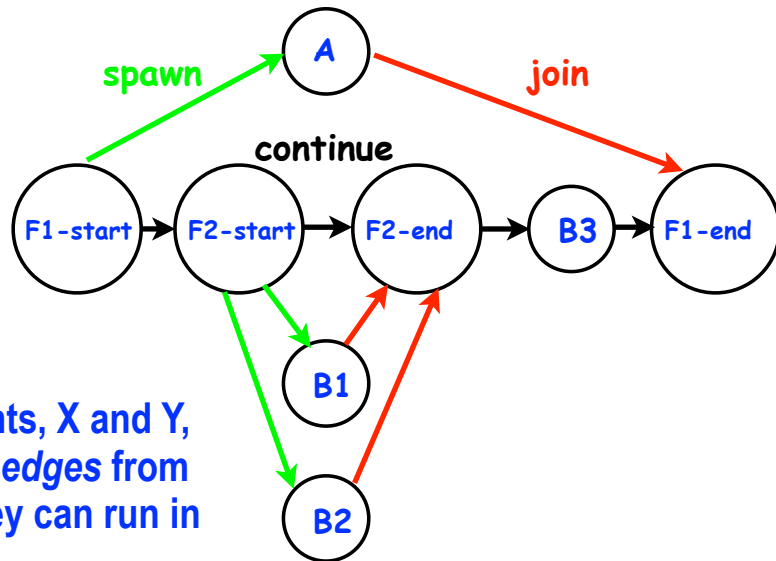
- A Computation Graph (CG) captures the dynamic execution of a parallel program, for a specific input
- CG nodes are “steps” in the program’s execution
 - A step is a sequential subcomputation without any `async`, `begin-finish` and `end-finish` operations
- CG edges represent ordering constraints
 - “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
```

Computation Graph



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

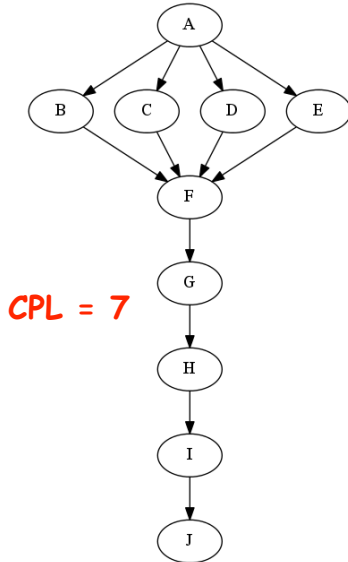
- $\text{TIME}(N)$ = execution time of node N
- $\text{WORK}(G)$ = sum of $\text{TIME}(N)$, for all nodes N in CG G
 - $\text{WORK}(G)$ is the total work to be performed in G
- $\text{CPL}(G)$ = length of a longest path in CG G , when adding up execution times of all nodes in the path
 - Such paths are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



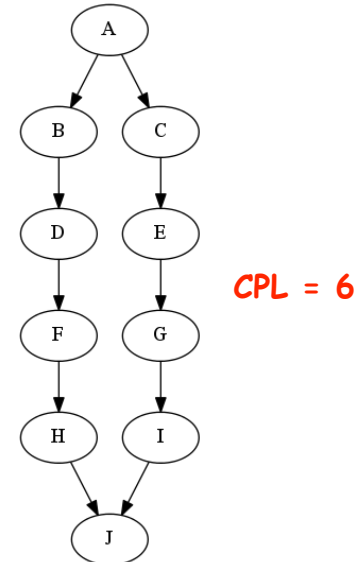
Which Computation Graph has more ideal parallelism?

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

Computation Graph 1



Computation Graph 2



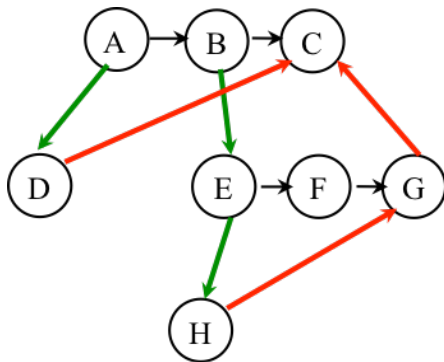
Data Races

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 theoretically undefined if there are two or more writes to the same location (even if the writes write the same value)
 - Above definition includes all “potential” data races i.e., we consider it to be a data race even if S1 and S2 end up executing 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.   {
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

9

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



Abstract Performance Metrics (Lecture 3)

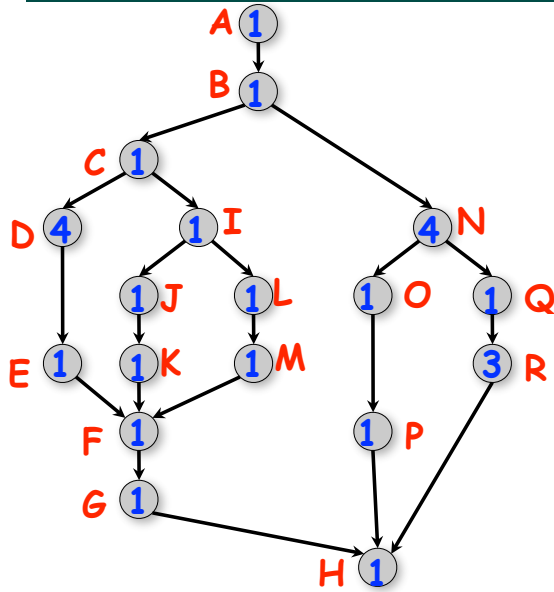
- Basic Idea
 - Count operations of interest, as in big-O analysis
 - Abstraction ignores many overheads that occur on real systems
- Calls to `doWork()`
 - Programmer inserts calls of the form, `doWork(N)`, within a step to indicate abstraction execution of N application-specific abstract operation
 - e.g., in the Homework 1 programming assignment (Parallel Sort), we included one call to `doWork(1)` in each call to `compareTo()`, and ignore the cost of everything else.
- Abstract metrics are enabled by calling
 - `HjSystemProperty.abstractMetrics.set(true)`;
- If an HJ program is executed with this option, abstract metrics can be printed at end of program execution with `WORK(G)`, `CPL(G)`, `Ideal Parallelism = WORK(G) / CPL(G)` after retrieving the metrics using call to `abstractMetrics()`

10

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



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 $\max(\text{CPL}, \text{WORK}/2) \leq T_2 < \text{CPL} + \text{WORK}/2$

11

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



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, $1/P \leq \text{Efficiency}(P) \leq 1$
- 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

12

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



Solution to Worksheet 4

- Estimate $T(S,P) \sim \text{WORK}(G,S)/P + \text{CPL}(G,S) = (S-1)/P + \log_2(S)$ for the parallel array sum computation shown in Lecture 4 Slide 4.
- Assume $S = 1024 \implies \log_2(S) = 10$
- Compute for 10, 100, 1000 processors
 - $T(P) = 1023/P + 10$
 - $\text{Speedup}(10) = T(1)/T(10) = 1033/112.3 \sim 9.2$
 - $\text{Speedup}(100) = T(1)/T(100) = 1033/20.2 \sim 51.1$
 - $\text{Speedup}(1000) = T(1)/T(1000) = 1033/11.0 \sim 93.7$
- Why does the speedup not increase linearly in proportion to the number of processors?
 - Because of the critical path length, $\log_2(S)$, is a bottleneck

13

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



Functional Parallelism: Adding Return Values to Async Tasks (Lecture 5)

Example Scenario (PseudoCode)

```
// Parent task creates child async task
future<int> 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**

14

COMP 322, Spring 2016 (V.Sarkar, S.Imam)



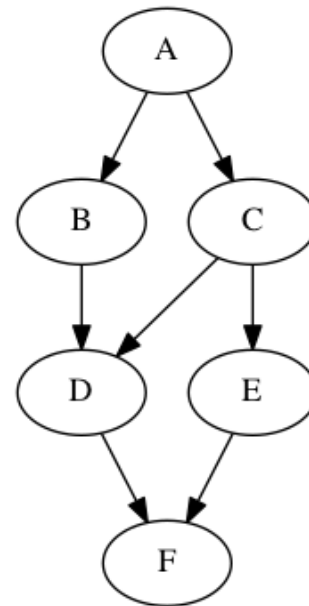
Worksheet #5: Computation Graphs for Async-Finish and Future Constructs

1) Can you write pseudocode with async-finish 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.

No. finish cannot be used to ensure that D waits for both B and C, while E waits only for C.

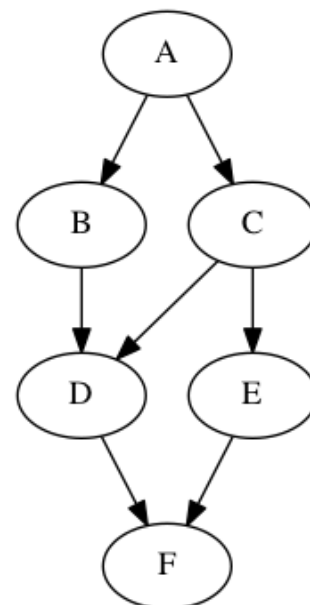
2) Can you write pseudocode with future async-get constructs that generates a Computation Graph with the same ordering constraints as the graph on the right? If so, provide a sketch of the program.

Yes, see program sketch with void futures. A dummy return value can also be used.



Worksheet #5 solution (contd)

```
1. HjFuture<String> A = future(() -> {
2.     return "A"; });
3. HjFuture<String> B = future(() -> {
4.     A.get(); return "B"; });
5. HjFuture<String> C = future(() -> {
6.     A.get(); return "C"; });
7. HjFuture<String> D = future(() -> {
8.     // Order of B.get() & C.get() doesn't matter
9.     B.get(); C.get(); return "D"; });
10. HjFuture<String> E = future(() -> {
11.     C.get(); return "E"; });
12. HjFuture<String> F = future(() -> {
13.     D.get(); E.get(); return "F"; });
14. F.get();
```



Memoization (Lecture 6)

- Memoization - saving and reusing previously computed values of a function rather than recomputing them
 - A optimization technique with 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

Helpful Link: <http://en.wikipedia.org/wiki/Memoization>

17

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



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

```
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.         future<int> left = future { return choose (N-1, K-1); }
9.         future<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

18

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



Worksheet #6 solution: Parallelizing Pascal's Triangle with Futures and Memoization

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

- Sequential Recursive without Memoization (chooseRecursiveSeq())
- Parallel Recursive without Memoization (chooseRecursivePar())
- Sequential Recursive with Memoization (chooseMemoizedSeq())
- 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:

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



Use of Finish Accumulators to count solutions in Parallel NQueens (Lecture 7)

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



Worksheet #7 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 e.g., an accumulator can be implemented using one “partial accumulator” per processor core.

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**

⇒ *Incorrect answers found in some worksheets: Associative / Both / Neither*

c) $h(s1,s2) = \text{concat}(s1, s2)$ for strings $s1, s2$, e.g., $h(\text{“ab”}, \text{“cd”}) = \text{“abcd”}$, is **associative but not commutative**

⇒ *Incorrect answers found in some worksheets: Commutative / Neither*



Functional vs. Structural Determinism (Lecture 8)

- 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*



Worksheet #8: Classifying different versions of parallel search algorithm

Example: String Search variation	Data Race Free?	Functionally Deterministic?	Structurally Deterministic?
V1: Count of all occurrences	YES	YES	YES
V2: Existence of an occurrence	NO	YES	YES
V3: Index of any occurrence	NO	NO	YES
V4: Optimized existence of an occurrence: do not create more async tasks after occurrence is found	NO	YES	NO
V5: Optimized index of any occurrence: do not create more async tasks after occurrence is found	NO	NO	NO

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 Summary (Lecture 9)

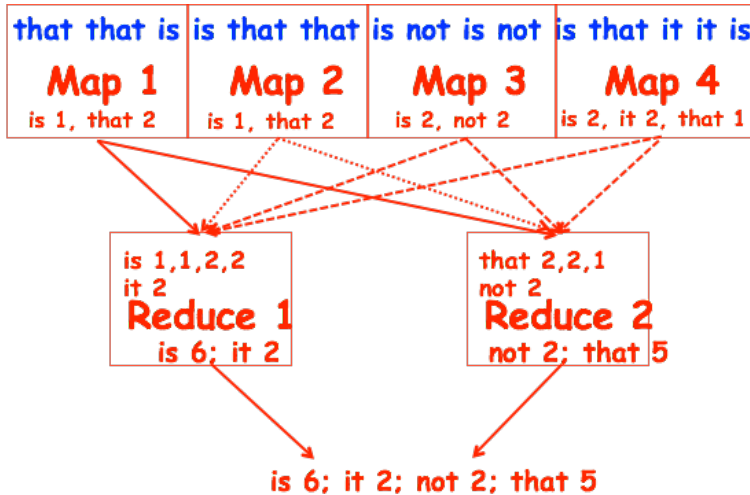
- Input set is of the form $\{(k_1, v_1), \dots, (k_n, v_n)\}$, where (k_i, v_i) consists of a key, k_i , and a value, v_i .
 - Assume that the key and value objects are immutable, and that equality comparison is well defined on all key objects.
- Map function f generates sets of intermediate key-value pairs, $f(k_i, v_i) = \{(k_1', v_1'), \dots, (k_{m_i}', v_{m_i}')\}$. The k_j' keys can be different from k_i key in the in 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



Worksheet #9: Analysis of Map Reduce Example

Analyze the total WORK and CPL for the Map reduce example:

- 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 = $\log_2(\# \text{ occurrences for input word with largest } \# \text{ pairs})$



WORK/CPL for all Map steps:

- WORK = 15
- CPL = 1

WORK/CPL for Reduce 1 step:

- WORK = 5
- CPL = $\log_2(4) = 2$

WORK/CPL for Reduce 2 step:

- WORK = 4
- CPL = $\log_2(3) = 1.58$

Total WORK and CPL

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



One-Dimensional Iterative Averaging Example (Lecture 11)

- 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 $1..n$

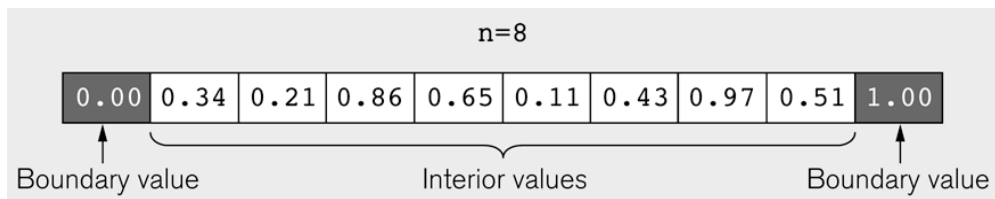


Illustration of an intermediate step for $n = 8$ (source: Figure 6.19 in Lin-Snyder book)



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
```

27

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



Solution to Worksheet #11: 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) Write the formula for the final value of `myNew[i]` as a function of i and n . In general, this is the value that we will get if m (= #iterations in sequential for-iter loop) is large enough.

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

28

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



Barriers (Lecture 12)

- **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
} Phase 1

- **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)



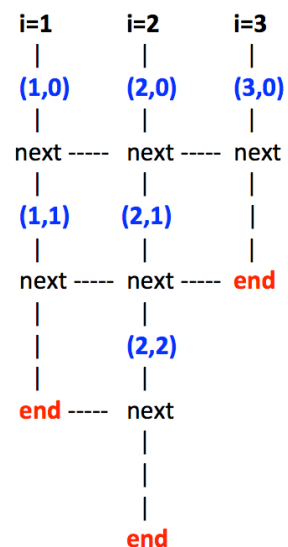
Worksheet #12: Forall Loops and Barriers

Draw a “barrier matching” figure similar to lecture 12 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. });
    
```

Solution



Converting forseq-forall version (Slide 5) into a forall-foreach version with barriers

```

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 iteration of iter loop
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

```

31

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



Worksheet #13 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)?	$m*n$	$m*nc$ Incorrect: $n * nc$	n Incorrect: $n * m$	nc Incorrect: $n*m, m*nc$
How many barrier operations (calls to next per task) are performed?	0 Incorrect: m	0 Incorrect: m	m Incorrect: $m*n$	m Incorrect: $m*nc, nc$

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.)

32

COMP 322, Spring 2016 (V. Sarkar, S. Imam)



Extending HJ Futures for Macro-Dataflow: Data-Driven Futures (DDFs, Lecture 14)

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

- Allocate an instance of a data-driven-future object (container)
- Object in container must be of type T1, and can only be assigned once via put() operations
- HjDataDrivenFuture extends the HjFuture interface

```
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



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

```
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”)
- Await clause can be used to implement “nodes” and “edges” in a computation graph

```
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



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.

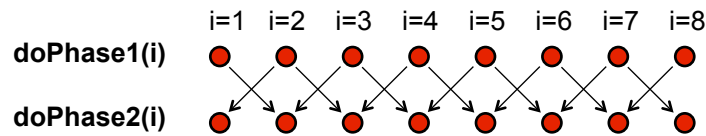


Summary of Phaser Construct (Lecture 15)

- Phaser allocation
 - `HjPhaser ph = new Phaser(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
 - Semantics depends on registration mode
 - Barrier is a special case of phaser, which is why `next` is used for both



Solution to Worksheet #15: 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

```

NOTE: Task-to-
phaser mappings can be
many-to-many in general. In
general, it is important to
understand the difference between
computation tasks (async's) and
synchronization objects
(phasers).



Solution to Worksheet #16: 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

```



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



Midterm exam (Exam 1)

- **Midterm exam (Exam 1) will be held during COMP 322 lab time at 7pm on Wednesday, February 24, 2016 in lab rooms (DH 1064, DH 1070)**
 - Closed-notes, closed-book, closed computer, written exam scheduled for 3 hours during 7pm — 10pm (but you can leave early if you're done early!)
 - Scope of exam is limited to Lectures 1 - 16 (all topics in Module 1 handout)
 - “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.”

