
COMP 322: Fundamentals of Parallel Programming

Lecture 26: Parallel Programming Patterns

Vivek Sarkar

Department of Computer Science, Rice University

vsarkar@rice.edu

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



Acknowledgments for Today's Lecture

- “A Pattern language for Parallel Programming”
 - Presentation by Beverly Sanders, U. Florida
 - <http://www.cise.ufl.edu/research/ParallelPatterns/sasplas.ppt>
- “Parallel Programming with Microsoft .NET”
 - Book published by Microsoft; free download available at <http://parallelpatterns.codeplex.com>
- “Introduction to Concurrency in Programming Languages”, Matthew J. Sottile, Timothy G. Mattson, and Craig E Rasmussen
 - Book published by CRC press; slides available from book web site, <http://www.parlang.com/>
- “Parallel Programming Patterns”, ME964 lecture, Oct 2008, U. Wisconsin
 - <http://sbel.wisc.edu/Courses/ME964/2008/.../me964Oct16.ppt>




Outline

- Introduction to Design Patterns
- Algorithm Design Patterns from "Parallel Programming with Microsoft .NET", illustrated using HJ
- Supporting Design Patterns from "Introduction to Concurrency in Programming Languages", illustrated using HJ




Design Patterns = formal discipline of design



- Christopher Alexander's approach to (civil) architecture:
 - A design pattern “describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice.” Page x, *A Pattern Language*, Christopher Alexander, 1977
- A pattern language is an organized way of tackling an architectural problem using patterns

- The gang of 4 used patterns to bring order to the chaos of object oriented design.
- The book “overnight” turned object oriented design from “an art” to a systematic design discipline.



GOF Pattern Example

- Behavioral Pattern: Visitor
 - Separate the structure of an object collection from the operations performed on that collection.
 - Example: Abstract syntax tree in a compiler
 - Multiple node types (declaration, command, expression, etc.)
 - Action during traversal depends on both type of node and compiler pass (type checking, code generation)
 - Can add new functionality by implementing new visitor without modifying AST code.



Parallel Programming

- Can a pattern language/taxonomy providing guidance for the entire development process make parallel programming easier?
 - Need to identify basic patterns, along with refinements (usually for efficiency)
- By relating HJ constructs to parallel programming patterns, you can apply HJ concepts to any parallel programming model you encounter in the future



Algorithm Patterns vs. Supporting Patterns

- **Algorithm Patterns**
 - Selection of task and data decompositions to solve a given problem in parallel
 - Task decomposition = identification of parallel steps
 - Data decomposition = partitioning of data into task-local vs. shared storage classes (with ownership specified for local data)
 - Examples: Parallel Loops, Parallel Tasks, Reductions, Dataflow, Pipeline
- **Supporting Patterns**
 - Selection of execution model for a given algorithm pattern
 - Execution model specifies scheduling and synchronization of tasks
 - May be implemented using code and/or configuration parameters
 - Examples of supporting patterns for task scheduling: Master-worker, Single Instruction Multiple Data (SIMD), Single Program Multiple Data (SPMD)
 - Examples of supporting patterns for mutual exclusion: Isolated, Object-based isolation, Atomic variables, Concurrent collections, Actors
- **Boundary between Algorithm & Supporting Patterns can sometimes be blurred**



Outline

- Introduction to Design Patterns
- Algorithm Design Patterns from “Parallel Programming with Microsoft .NET”, illustrated using HJ
- Supporting Design Patterns from “Introduction to Concurrency in Programming Languages”, illustrated using HJ



Selecting the Right Pattern

(adapted from page 9, Parallel Programming w/ Microsoft .Net)

Application characteristics	Algorithmic pattern	Relevant HJ constructs
Sequential loop with independent iterations	1) Parallel Loop	forall, forasync
Independent operations with well-defined control flow	2) Parallel Task	async, finish
Aggregating data from independent tasks/iterations	3) Parallel Aggregation (reductions)	finish accumulators, atomic variables
Ordering of steps based on data flow constraints	4) Futures	futures, data-driven tasks
Divide-and-conquer algorithms with recursive data structures	5) Dynamic Task Parallelism	async, finish
Repetitive operations on data streams	6) Pipelines	streaming phasers (deterministic), actors (non-deterministic)



1) Example of Parallel Loop Pattern using HJ forall statement (Lecture 10)

```
1. double[] myVal = new double[n+2];
2. . . . // Initialize myVal
3. for (point [iter] : [0:iterations-1]) {
4.     // Compute MyNew array as function of input array, MyVal
5.     myNew = new double[n+2];
6.     // PARALLEL LOOP PATTERN
7.     forall (point [j] : [1:n]) { // Create n tasks
8.         myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
9.     } // forall
10.    // myNew becomes input array for next iteration
11.    myVal=myNew;
12.} // for
```



Partitioning/Chunking refinement of Parallel Loops for Efficiency (Lecture 10)

```
1. double[] myVal = new double[n+2];
2. . . . // Initialize myVal
3. for (point [iter] : [0:iterations-1]) {
4.     // Compute MyNew array as function of input array, MyVal
5.     myNew = new double[n+2];
6.     // PARALLEL LOOP PATTERN WITH CHUNKING
7.     forall (point [jj]:[0:Cj-1]) { // Create Cj tasks
8.         for (point [j]:getChunk([1:n],Cj,jj)) // Execute chunk jj
9.             // Loop body is unchanged from before
10.            myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
11.     } // forall
12.     // myNew becomes input array for next iteration
13.     myVal=myNew;
14.} // for
```



Double Buffering for Further Efficiency (Lecture 10)

```
1. double[] myVal = new double[n+2];
2. double[] myNew = new double[n+2];
3. . . . // Initialize myVal
4. for (point [iter] : [0:iterations-1]) {
5.     // Compute MyNew array as function of input array, MyVal
6.     // PARALLEL LOOP PATTERN WITH CHUNKING AND DOUBLE BUFFERING
7.     forall (point [jj]:[0:Cj-1]) { // Create Cj tasks
8.         for (point [j]:getChunk([1:n],Cj,jj)) // Execute chunk jj
9.             // Loop body is unchanged from before
10.            myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
11.     } // forall
12.     // myNew becomes input array for next iteration
13.     // DOUBLE BUFFERING: swap myVal & myNew
14.     temp=myVal; myVal=myNew; myNew=temp;
15.} // for
```



2) Example of Parallel Task Pattern using HJ async and finish statements (Lecture 1)

```
1. // Start of Task T0 (main program)
2. sum1 = 0; sum2 = 0; // sum1 & sum2 are static fields
3. finish {
4.     async { // Task T1 computes sum of upper half of array
5.         for(int i=X.length/2; i < X.length; i++) sum2 += X[i];
6.     }
7.     // Continue in T0 and compute sum of lower half of array
8.     for(int i=0; i < X.length/2; i++) sum1 += X[i];
9. } // finish
10. // Task T0 waits for Task T1 (join)
11. return sum1 + sum2;
```



Sequential clause refinement to Parallel Task Pattern for Efficiency (Lecture 9)

```
1. // Start of Task T0 (main program)
2. sum1 = 0; sum2 = 0; // sum1 & sum2 are static fields
3. finish {
4.     async seq(X.length < 100000) {
5.         // Only create async if array length is >= 100000
6.         for(int i=X.length/2; i < X.length; i++) sum2 += X[i];
7.     } // async
8.     // Continue in T0 and compute sum of lower half of array
9.     for(int i=0; i < X.length/2; i++) sum1 += X[i];
10. } // finish
11. // Task T0 waits for Task T1 (join)
12. return sum1 + sum2;
```



3) Example of Parallel Aggregation Pattern using HJ finish accumulators (Lecture 12)

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



Refinement of Parallel Aggregation Pattern using Atomic Variables for Efficiency (Lecture 6)

```
1. import java.util.concurrent.atomic.AtomicInteger;
2. . . .
3. AtomicInteger count = new AtomicInteger();
4. finish nqueens_kernel(new int[0], 0);
5. . . .
6. void nqueens_kernel(int [] a, int depth) {
7.     if (size == depth) count.addAndGet(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()
```

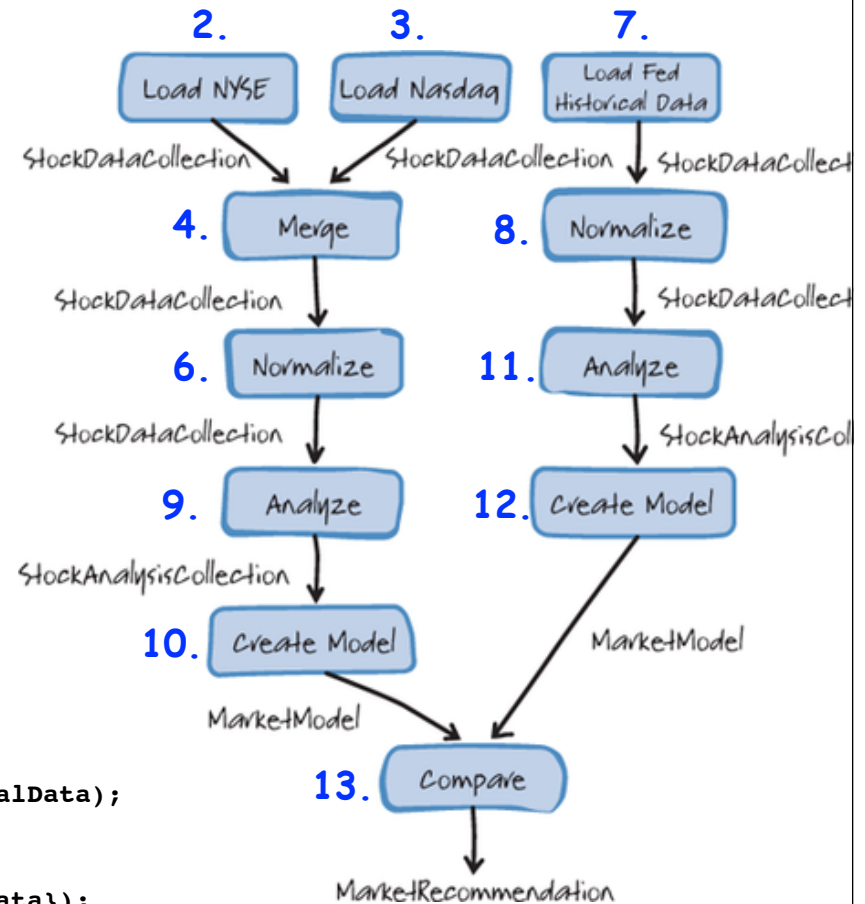


“Adatum Dashboard” Example: Sequential Version (Lecture 8)

```

1. public MarketRecommendation DoAnalysisSequential() {
2.   StockDataCollection nyseData = LoadNyseData();
3.   StockDataCollection nasdaqData = LoadNasdaqData();
4.   StockDataCollection mergedMarketData =
5.     MergeMarketData(new[] {nyseData, nasdaqData});
6.   StockDataCollection normalizedMarketData =
7.     NormalizeData(mergedMarketData);
8.   StockDataCollection fedHistoricalData =
9.     LoadFedHistoricalData();
10.  StockDataCollection normalizedHistoricalData =
11.    NormalizeData(fedHistoricalData);
12.  StockAnalysisCollection analyzedStockData =
13.    AnalyzeData(normalizedMarketData);
14.  MarketModel modeledMarketData = RunModel(analyzedStockData);
15.  StockAnalysisCollection analyzedHistoricalData =
16.    AnalyzeData(normalizedHistoricalData);
17.  MarketModel modeledHistoricalData = RunModel(analyzedHistoricalData);
18.  MarketRecommendation recommendation =
19.    CompareModels(new[] {modeledMarketData, modeledHistoricalData});
20.  return recommendation;
21. }

```



Source: <http://programming4.us/enterprise/3004.aspx>



4) Example of Future pattern using HJ future tasks

```
1. public MarketRecommendation DoAnalysisParallelDDT() {
2.     future<StockDataCollection> nyseData = async<StockDataCollection> {return LoadNyseData();};
3.     future<StockDataCollection> nasdaqData = async<StockDataCollection> {return LoadNasdaqData();};
4.     future<StockDataCollection> mergedMarketData =
        async<StockDataCollection> {return MergeMarketData(new[] {nyseData.get(), nasdaqData.get()});};
5.     future<StockDataCollection> normalizedMarketData =
        async<StockDataCollection> {return NormalizeData(mergedMarketData.get());};
6.     future<StockDataCollection> fedHistoricalData = async<StockDataCollection> {return LoadFedHistoricalData();};
7.     future<StockDataCollection> normalizedHistoricalData =
        async<StockDataCollection> {return NormalizeData(fedHistoricalData.get());};
8.     future<StockAnalysisCollection> analyzedStockData =
        async<StockAnalysisCollection> {return AnalyzeData(normalizedMarketData.get());};
9.     future<MarketModel> modeledMarketData = async<MarketModel> {return RunModel(analyzedStockData.get());};
10.    future<StockAnalysisCollection> analyzedHistoricalData =
        async<StockAnalysisCollection> {return AnalyzeData(normalizedHistoricalData.get());};
11.    future<MarketModel> modeledHistoricalData = async<MarketModel> {return RunModel(analyzedHistoricalData.get());};
12.    MarketRecommendation recommendation =
        CompareModels(new[] {modeledMarketData.get(), modeledHistoricalData.get()});
13.    return recommendation;
14. }
```



Refinement of Future Pattern using Data-Driven Tasks for Efficiency (Lecture 8)

```
1. public MarketRecommendation DoAnalysisParallelDDT() {
2.     async nyseData.put(LoadNyseData());
3.     async nasdaqData.put(LoadNasdaqData());
4.     async await(nyseData, nasdaqData)
           mergedMarketData.put(MergeMarketData(new[] {nyseData.get(), nasdaqData.get()}));
5.     async await(mergedMarketData) normalizedMarketData.put(NormalizeData(mergedMarketData.get()));
6.     async fedHistoricalData.put(LoadFedHistoricalData());
7.     async await(fedHistoricalData) normalizedHistoricalData.put(NormalizeData(fedHistoricalData.get()));
8.     async await(normalizedMarketData) analyzedStockData.put(AnalyzeData(normalizedMarketData.get()));
9.     async await(analyzedStockData) modeledMarketData.put(RunModel(analyzedStockData.get()));
10.    async await(normalizedHistoricalData) analyzedHistoricalData.put(AnalyzeData(normalizedHistoricalData.get()));
11.    async await(analyzedHistoricalData) modeledHistoricalData.put(RunModel(analyzedHistoricalData.get()));
12.    MarketRecommendation recommendation =
           CompareModels(new[] {modeledMarketData.get(), modeledHistoricalData.get()});
13.    return recommendation;
14. }
```



5) Example of Dynamic Task Pattern using HJ async and finish constructs (Lecture 11)

```
1. static void quicksort(int[] A, int M, int N) {
2.   if (M < N) {
3.     // partition() selects a pivot element in A[M..N]
4.     // to partition A[M..N] into A[M..J] and A[I..N]
5.     point p = partition(A, M, N);
6.     int I=p.get(0); int J=p.get(1);
7.     async quicksort(A, M, J);
8.     async quicksort(A, I, N);
9.   }
10.} //quicksort
11. . . .
12. finish quicksort(A, 0, A.length-1);
```



Sequential clause refinement to Dynamic Task Pattern for Efficiency

```
1. static void quicksort(int[] A, int M, int N) {
2.   if (M < N) {
3.     // partition() selects a pivot element in A[M..N]
4.     // to partition A[M..N] into A[M..J] and A[I..N]
5.     point p = partition(A, M, N);
6.     int I=p.get(0); int J=p.get(1);
7.     async seq(J-M < 10000) quicksort(A, M, J);
8.     async seq(N-I < 10000) quicksort(A, I, N);
9.   }
10.} //quicksort
11. . . .
12. finish quicksort(A, 0, A.length-1);
```



6) Example of Pipeline Pattern using HJ bounded phasers and accumulators (Lecture 15)

```
void Pipeline() {
    phaser phI      = new phaser(SIG_WAIT, bnd);
    accumulator I  = new accumulator(phI, accumulator.ANY);
    phaser phM      = new phaser(SIG_WAIT, bnd);
    accumulator M  = new accumulator(phM, accumulator.ANY);
    phaser phO      = new phaser(SIG_WAIT, bnd);
    accumulator O  = new accumulator(phO, accumulator.ANY);
    async phased (phI<SIG>)          source(I);
    async phased (phI<WAIT>, phM<SIG>) avg(I,M);
    async phased (phM<WAIT>, phO<SIG>) abs(M,O);
    async phased (phO<WAIT>)        sink(O);
}

void avg(accumulator I, accumulator M) {
    while(...) {
        wait; wait;           // wait for two elements on I
        v1 = I.get(0);        // read first element
        v2 = I.get(-1);       // read second element (offset = -1)
        M.put((v1+v2)/2);     // put result on M
        signal;
    }
}
```



Outline

- Introduction to Design Patterns
- Algorithm Design Patterns from "Parallel Programming with Microsoft .NET", illustrated using HJ
- Supporting Design Patterns from "Introduction to Concurrency in Programming Languages", illustrated using HJ



Supporting Patterns

1) Master-worker

- A process or thread (the master) sets up a task queue and manages other threads (the workers) as they grab a task from the queue, carry out the computation, and then return for their next task. This continues until the master detects that a termination condition has been met, at which point the master ends the computation.

2) Single Instruction Multiple Data (SIMD)

- A supporting pattern for data parallelism, in which a single instruction stream is applied to multiple data elements in parallel.

3) Single Program Multiple Data (SPMD)

- Multiple copies of a single program are launched typically with their own view of the data. The path through the program for each copy is determined in part based on a unique ID (a rank).



1) Master-Worker pattern

- A master manages a collection of tasks (often with a queue) and schedules them for execution on by a collection of workers.
- Workers request more tasks as they finish their assigned work,
- The result ... the computational work is automatically balanced between a collection of workers.



Master Worker: Pseudo-code for the master

```
Global SharedQueue tasks ;
Global Results R;

master {
int i, taskCount ;

// First , generate the set of tasks and
// place them in the shared queue .

    for (i=0;i< taskCount ;i++) {
        Task t = generateSubTask ();
        tasks .add(t);
    }
// Next , consume the set of results produced
// by the workers .

    for (i=0;i< taskCount ;i++) {
        Result res = R.pop ();
        consumeResult (res );
    }
    shut_down();
}
```

Create data structures shared between workers and the master:

- A Queue to hold the tasks
- A Queue to hold the results

Fill the queue with a set of tasks

Pull results off the results-queue and consume them.

Shut down system results from all tasks have been consumed



Master Worker: Pseudo-code for the worker

Global SharedQueue tasks ;

Global Results R;

worker {

while (true) {

 // Get a task

 task = tasks .pop ();

 // Perform the work

 Result res = performWork (task);

 // Put the result in the result queue

 R. push (res);

}

}

Access key shared data
structures:

- A Queue to hold the tasks
- A Queue to hold the results

Grab next available
task

Do the work

Store the result

This version of the worker loops until the master
shuts down the entire computation

Pi Calculation Example (Lecture 22) --- Master Actor

```
1. class Master extends Actor<Object> {
2.     private double result = 0; private int nrMsgsReceived = 0;
3.     private Worker[] workers;
4.     Master(nrWrkrs, nrEls, nrMsgs) {...} // constructor
5.     void start() {
6.         super.start(); // Starts the master actor
7.         // Create and start workers
8.         workers = new Worker[nrWrkrs];
9.         for (int i = 0; i < nrwrkrs; i++) {
10.            workers[i] = new Worker();
11.            workers[i].start();
12.        }
13.        // Send messages to workers
14.        for (int j = 0; j < nrMsgs; j++) {
15.            someWrkr = ... ; // Select worker for message j
16.            someWrkr.send(new Work(...));
17.        }
18.    } // start()
```



Pi Calculation --- Master Actor (contd)

```
19. void exit() {
20.     for (int i = 0; i < nrWrkrs; i++) workers[i].send(new Stop());
21.     super.exit(); // Terminates the actor
22. } // exit()
23. void process(final Object msg) {
24.     if (msg instanceof Result) {
25.         result += ((Result) msg).result;
26.         nrMsgsReceived += 1;
27.         if (nrMsgsReceived == nrMsgs) exit();
28.     }
29.     // Handle other message cases here
30. } // process()
31.} // Master
32. . . .
33. // Main program
34. Master master = new Master(w, e, m);
35. finish master.start();
36. println("PI = " + master.getResult());
```




Pi Calculation --- Worker Actor

```
1. class Worker extends Actor<Object> {
2.     void process(Object msg) {
3.         if (msg instanceof Stop) exit();
4.         else if (msg instanceof Work) {
5.             Work wm = (Work) msg;
6.             double result = calculatePiFor(wm.start, wm.end)
7.             master.send(new ResultMessage(result));    }
8.     } // process()
9.
10.    private double calculatePiFor(int start, int end) {
11.        double acc = 0.0;
12.        for (int i = start; i < end; i++) {
13.            acc += 4.0 * (1 - (i % 2) * 2) / (2 * i + 1);
14.        }
15.        return acc;
16.    }
17. } // Worker
```




2) SIMD pattern: Single Instruction Multiple Data

- A supporting pattern for data parallelism algorithmic pattern.
- Definition: A single instruction stream is applied to multiple data elements.
 - One program text
 - One instruction counter
 - Distinct data streams per Processing Element (PE)



Data Parallelism Algorithmic Pattern

- **Use when:**
 - Your problem is defined in terms of collections of data elements operated on by a similar (if not identical) sequence of instructions; i.e. the concurrency is in the data.
- **Solution**
 - Define collections of data elements that can be updated in parallel.
 - Define computation as a sequence of collective operations applied together to each data element.



Matrix Multiplication (serial)

```
double[][] a, b, c; // three 2D arrays : a,b,c
int n; // Assume that all arrays are of size n*n
```

```
for(point[i,j] : [0:n-1,0:n-1] {
```

```
    c[i][j] = 0.0
```

```
        for(int k=0; k<n; k++) {
            c[i][j] += a[i][k] * b[k][j];
        }
```

```
    }
```

This operation is called a dot product of the two vectors a and b to produce a scalar value for c(i,j)



Matrix Multiplication (with multiple sources of parallelism)

```
double[][] a, b, c; // three 2D arrays : a,b,c  
int n; // Assume that all arrays are of size n*n
```

```
forall(point[i,j] : [0:n-1,0:n-1] {
```


Loop parallelism



```
c[i][j] = sum(a[i][0:n-1] * b[0:n-1][j]);
```


```
}
```

Dot product is expressed as SIMD parallelism
(This is pseudocode, not real HJ code)




Array slice notation


- Designating different slices of an array.




$A[:, :]$




$A[3][:]$




$A[:, 3]$



$A[2][3]$



$A[1:3][:]$



$A[1:3][2:4]$



Flynn's Taxonomy for Parallel Computers

	Single Instruction	Multiple Instructions
Single Data	SISD	MISD
Multiple Data	SIMD	MIMD

Single Instruction, Single Data stream (SISD)

A sequential computer which exploits no parallelism in either the instruction or data streams. e.g., old single processor PC

Single Instruction, Multiple Data streams (SIMD)

A computer which exploits multiple data streams against a single instruction stream to perform operations which may be naturally parallelized. e.g. graphics processing unit

Multiple Instruction, Single Data stream (MISD)

Multiple instructions operate on a single data stream. Uncommon architecture which is generally used for fault tolerance. Heterogeneous systems operate on the same data stream and must agree on the result. e.g. the Space Shuttle flight control computer.

Multiple Instruction, Multiple Data streams (MIMD)

Multiple autonomous processors simultaneously executing different instructions on different data. e.g. a PC cluster memory space.



3) SPMD Supporting Pattern

- SPMD: Single Program Multiple Data
- Run the same program on P processing elements (PEs)
- Use the “rank” ... an ID ranging from 0 to $(P-1)$... to determine what computation is performed on what data by a given PE
- Different PEs can follow different paths through the same code (unlike the SIMD pattern)
- Convenient pattern for hardware platforms that are not amenable to efficient forms of dynamic task parallelism
 - General-Purpose Graphics Processing Units (GPGPUs)
 - Distributed-memory parallel machines
- Key design decisions --- what data and computation should be replicated or partitioned across PEs?



Typical SPMD Program Phases

- **Initialize**
 - Establish localized data structure and communication channels
- **Obtain a unique identifier**
 - Each thread acquires a unique identifier, typically range from 0 to $N-1$, where N is the number of threads.
 - Both OpenMP and CUDA have built-in support for this.
- **Distribute Data**
 - Decompose global data into chunks and localize them, or
 - Sharing/replicating major data structure using thread ID to associate subset of the data to threads
- **Run the core computation**
 - More details in next slide...
- **Finalize**
 - Reconcile global data structure, prepare for the next major iteration

SPMD Example #1

- **Assign a chunk of iterations to each thread**
 - The last thread also finishes up the remainder iterations

```
num_steps = 1000000;  
  
i_start = my_id * (num_steps / num_threads);  
i_end = i_start + (num_steps / num_threads);  
if (my_id == (num_threads-1)) i_end = num_steps;  
  
for (i = i_start; i < i_end; i++) {  
    ....  
}  
Reconciliation of results across threads if necessary.
```

SPMD Example #2: Iterative Averaging Example (Slide 9, Lecture 13)

```
1. double[] gVal=new double[n+2]; double[] gNew=new double[n+2];
2. gVal[n+1] = 1; // Boundary condition
3. int Cj = Runtime.getNumOfWorkers();
4. forall (point [jj]:[0:Cj-1]) { // SPMD computation
5.     double[] myVal = gVal; double[] myNew = gNew; // Local copy
6.     for (point [iter] : [0:numIters-1]) {
7.         // Compute MyNew as function of input array MyVal
8.         for (point [j]:getChunk([1:n],[Cj],[jj]))
9.             myNew[j] = (myVal[j-1] + myVal[j+1])/2.0;
10.        next; // Barrier before executing next iteration of iter loop
11.        // Swap myVal and myNew (replicated computation)
12.        double[] temp=myVal; myVal=myNew; myNew=temp;
13.        // myNew becomes input array for next iter
14.    } // for
15.} // forall
```

