"The computer programmer is a creator of universes for which he alone is the lawgiver ... Universes of virtually unlimited complexity can be created in the form of computer programs ... They compliantly obey their laws and vividly exhibit their obedient behavior. No playwright, no stage director, no emperor, however powerful, has ever exercised such absolute authority to arrange a stage or a field of battle and to command such unswervingly dutiful actors of troops."

-- J. Weizenbaum, Computer Power and Human Reason: From Judgment to Calculation

Concepts and Strategies in Parallel Programming

A Brief Introduction

"As soon as an Analytical Engine exists, it will necessarily guide the future course of the science. Whenever any result is sought by its aid, the question will then arise — by what course of calculation can these results be arrived at by the machine in the shortest time?"

~Charles Babbage (1864)

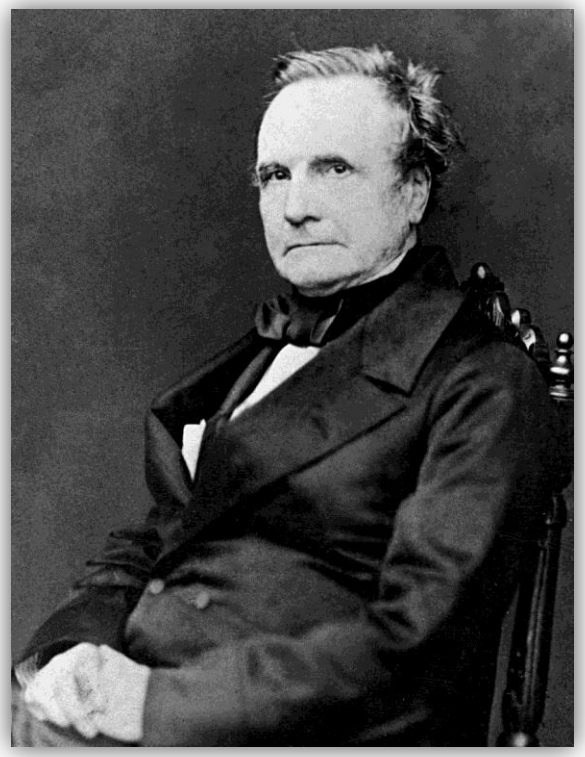

Part I – Basic Concepts

So what does it mean to do Parallel Computing/Programming?

"I decided long ago to stick to what I know best. Other people understand parallel machines much better than I do; programmers should listen to them, not me, for guidance on how to deal with simultaneity."

--**Donald Knuth** (Stanford)

"Redesigning your application to run multithreaded on a multicore machine is a little like learning to swim by jumping into the deep end."

--**Herb Sutter**
(Chair of the ISO C++ standards committee, Microsoft)

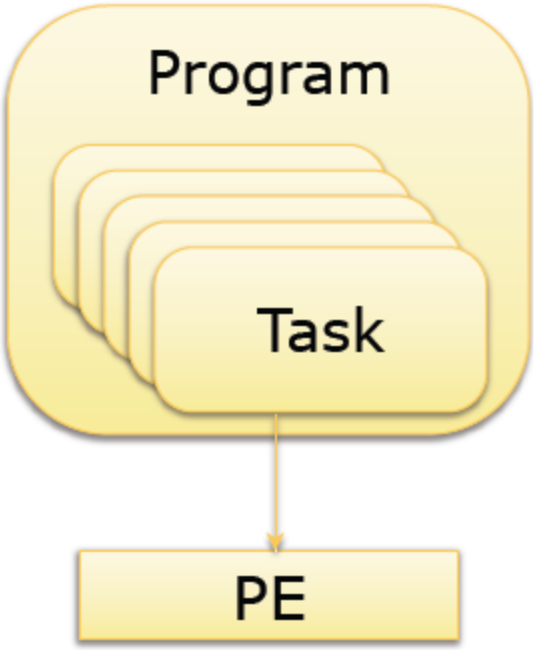
Program
Task
PE

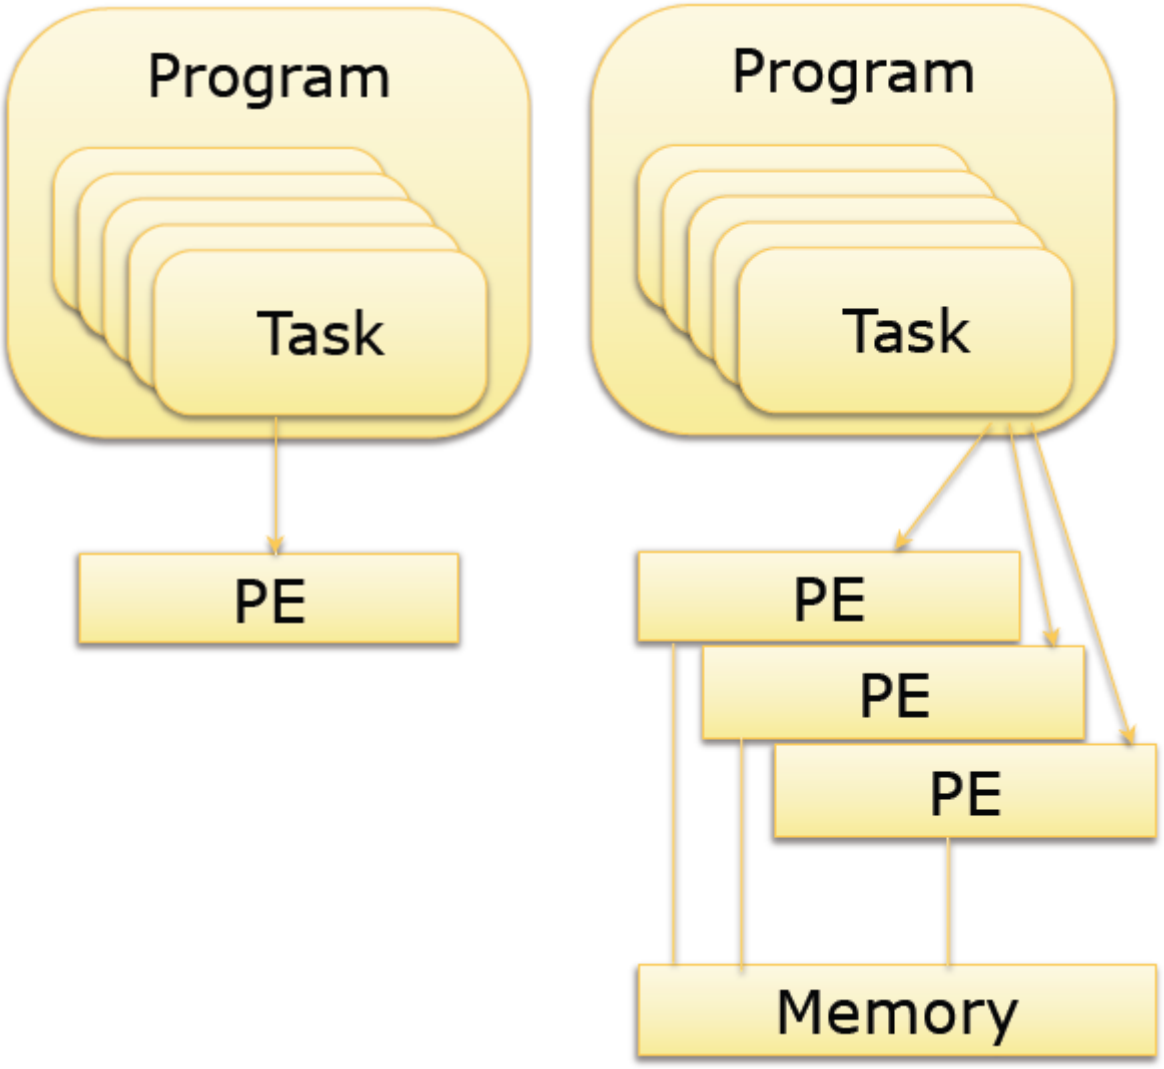
Program
Task
PE
Program
Task
PE
PE
PE
Memory

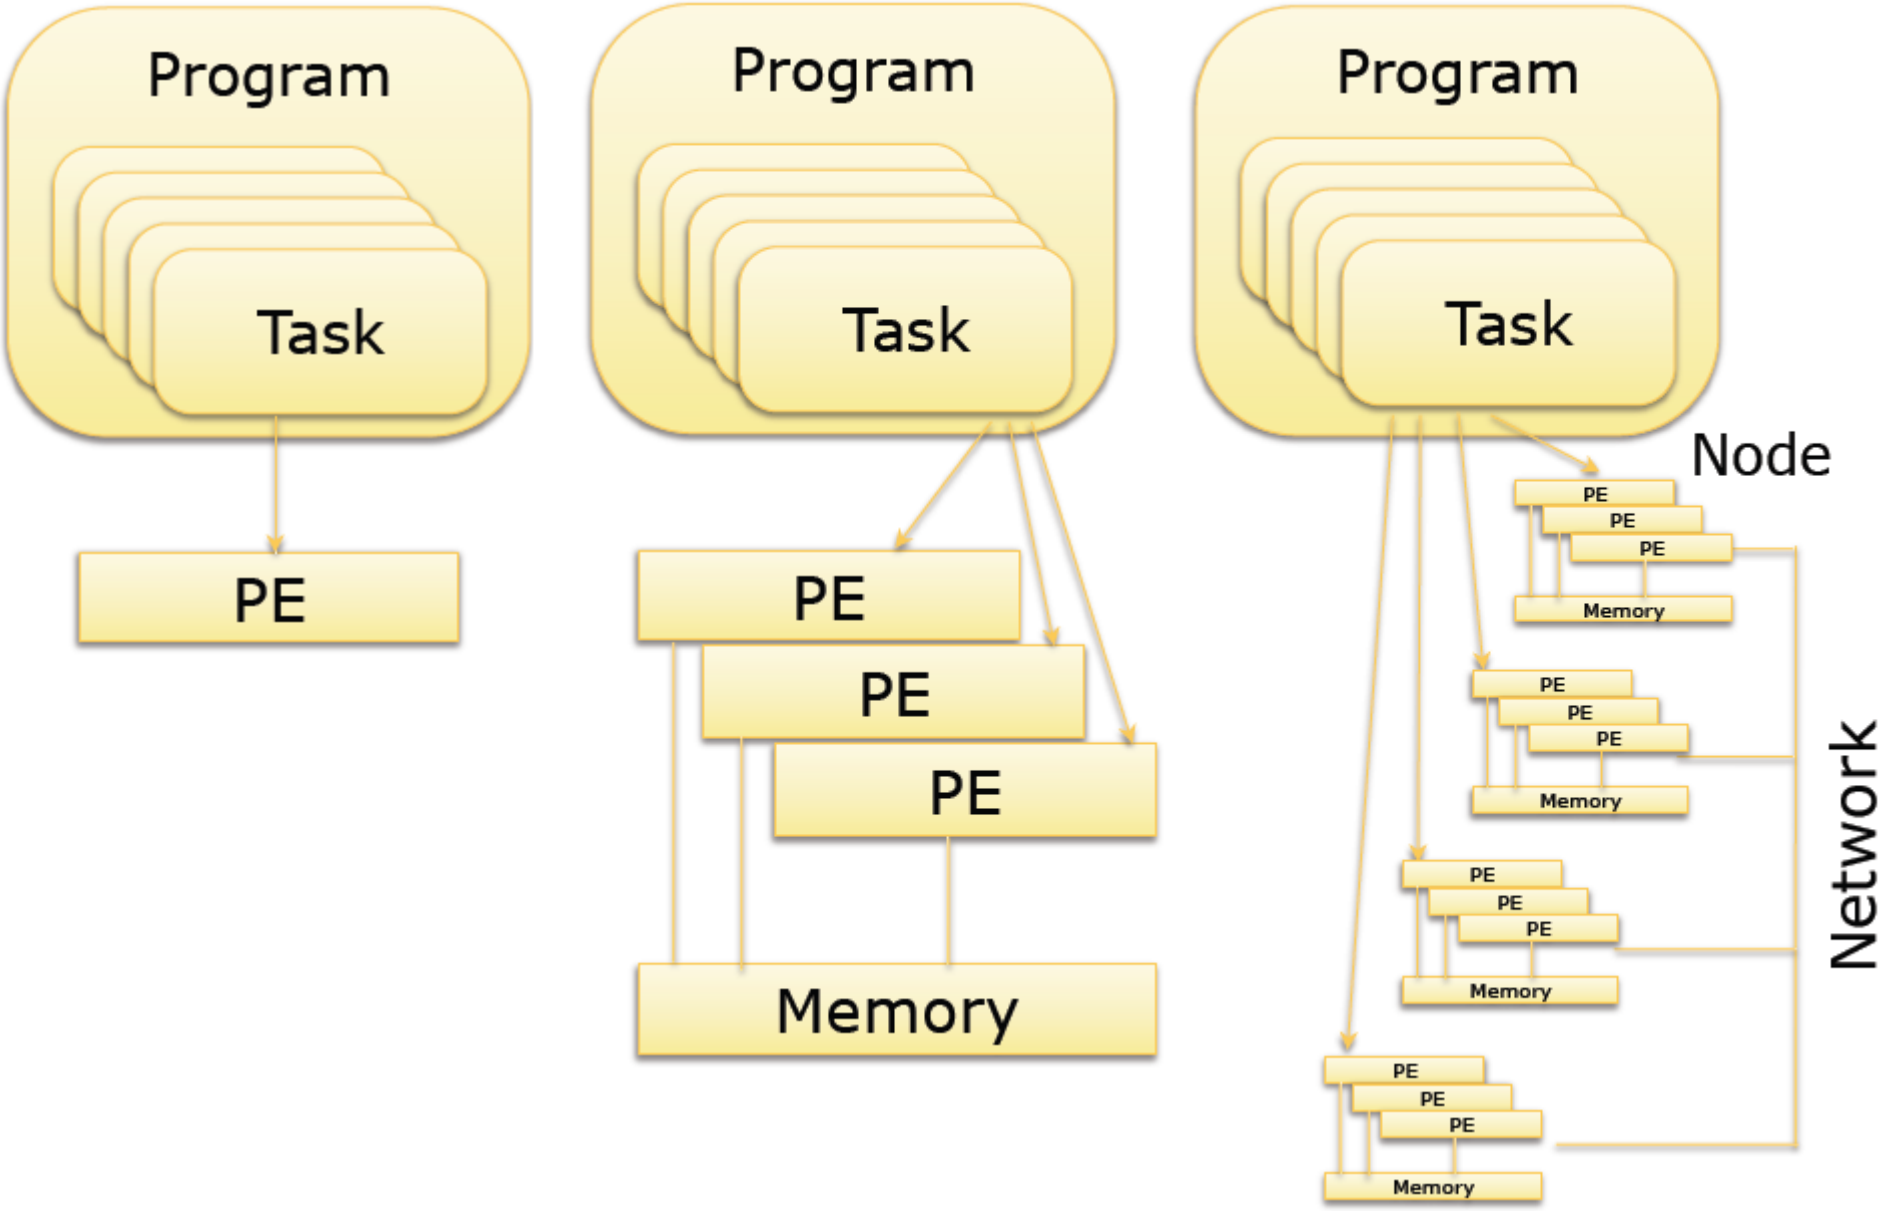
Program
Task
PE
Program
Task
PE
PE
PE
Memory
Program
Task
Node
PE
PE
PE
Memory
PE
PE
PE
Memory
PE
PE
PE
Memory
PE
PE
PE
Memory
Network

Shared-Memory vs. Shared-Nothing

- **Pfister:** *Shared-Memory* vs. *Distributed-Memory*
- **Foster:** *Multi-Processor* vs. *Multi-Computer*
- **Tennenbaum:** *Shared-Memory* vs. *Private-Memory*

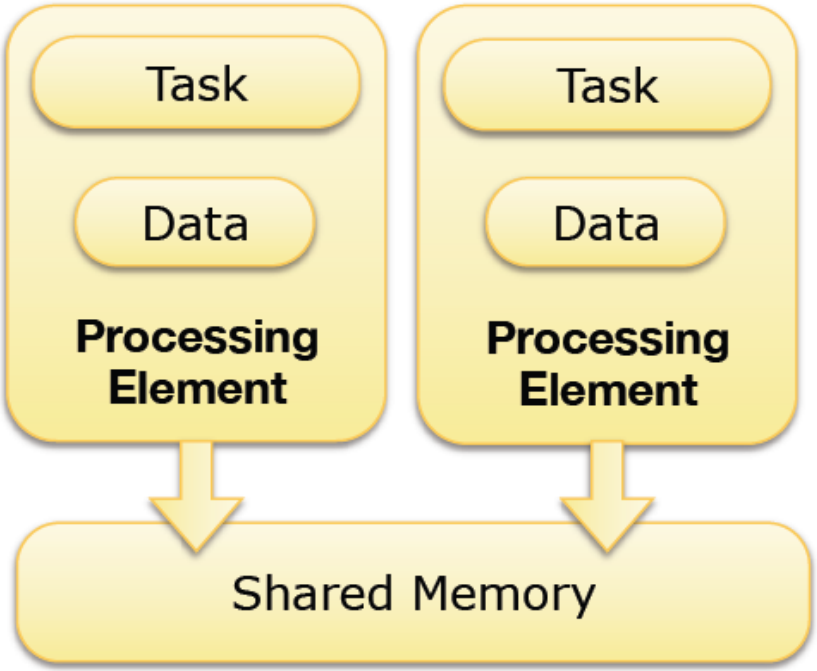

Shared-Memory vs. Shared-Nothing

- **Pfister:** *Shared-Memory* vs. *Distributed-Memory*
- **Foster:** *Multi-Processor* vs. *Multi-Computer*
- **Tennenbaum:** *Shared-Memory* vs. *Private-Memory*

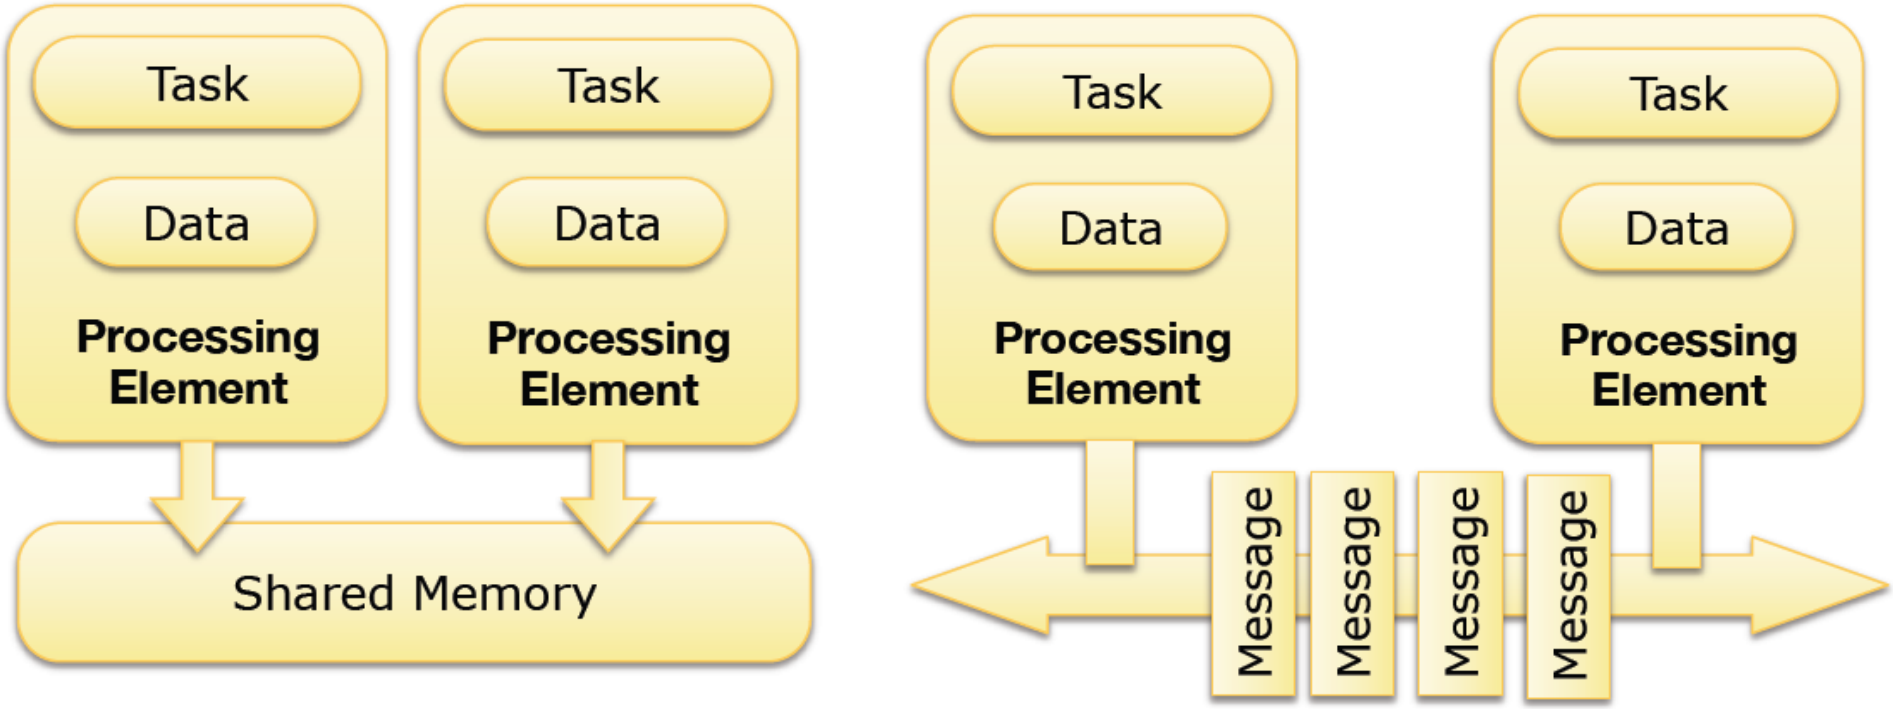

The Essence of Parallel Programming

- *So essentially, what we try to do is,*
  - Look at a serial program and identify the parts that are the most compute intensive → Computational *Hotspots*
    - The *80/20* rule … mehh!

The Essence of Parallel Programming

- *So essentially, what we try to do is,*
  - Look at a serial program and identify the parts that are the most compute intensive → Computational *Hotspots*
    - The *80/20* rule … mehh!
  - Extract *parallelism* from Hotspots

The Essence of Parallel Programming

- *So essentially, what we try to do is,*
  - Look at a serial program and identify the parts that are the most compute intensive → Computational *Hotspots*
    - The *80/20* rule … mehh!
  - Extract *parallelism* from Hotspots
  - Design parallelization strategy

The Essence of Parallel Programming

- *So essentially, what we try to do is,*
  - Look at a serial program and identify the parts that are the most compute intensive → Computational *Hotspots*
    - The *80/20* rule … mehh!
  - Extract *parallelism* from Hotspots
  - Design parallelization strategy
  - Make it possible for the code to exploit the available parallel processing resources

Concurrency vs. Parallelism

- Concurrent
  - "A system is said to be *concurrent* if it can support two or more actions *in progress* at the same time."

Concurrency vs. Parallelism

- Concurrent
  - "A system is said to be *concurrent* if it can support two or more actions *in progress* at the same time."
- Parallel
  - "A system is said to be *parallel* if it can support two or more actions *executing simultaneously*."

Reference: Breshears, Clay. *The art of concurrency: A thread monkey's guide to writing parallel applications*. O'Reilly Media, Inc., 2009.

Geometric Decomposition

Geometric Decomposition

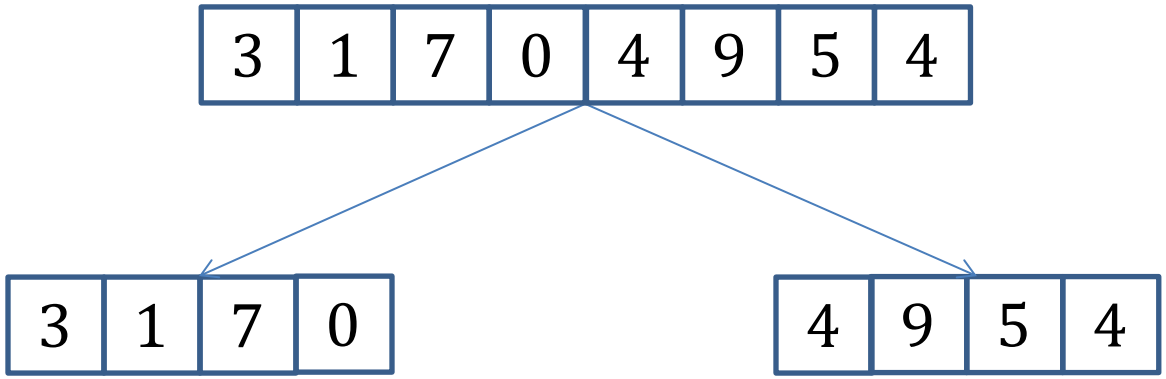

Geometric Decomposition

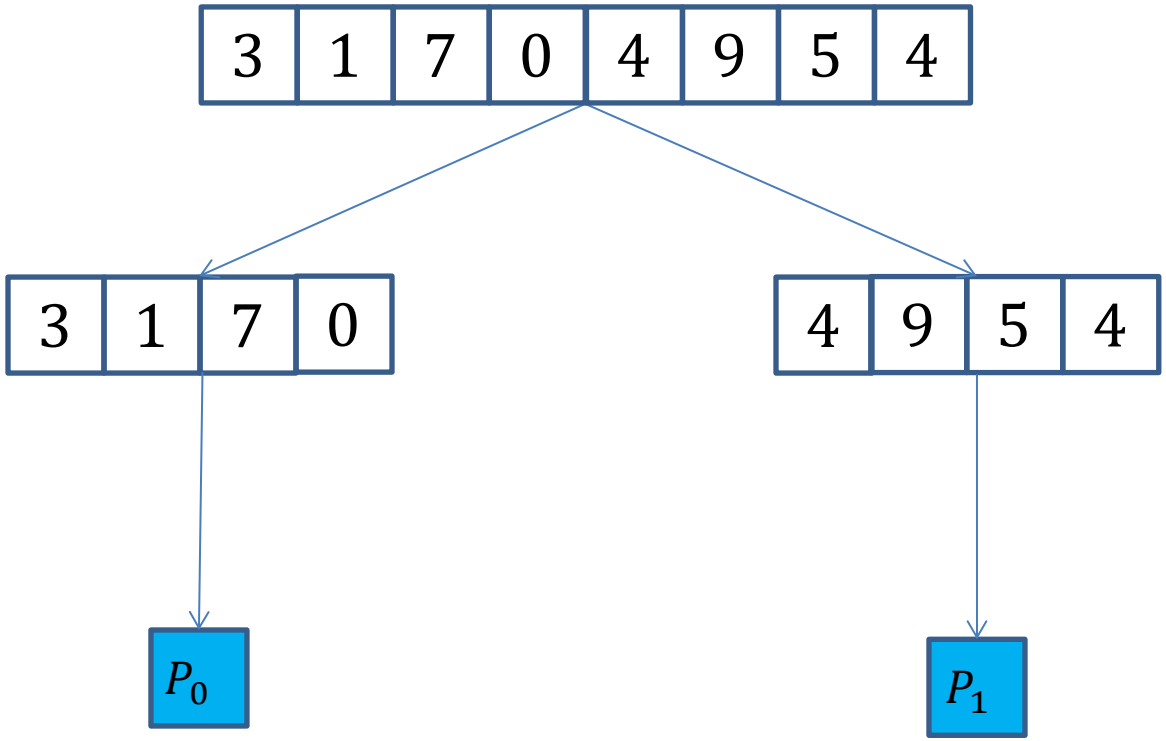

Map Pattern – Independent Loop Iterations

$$w = \frac{b}{8\pi(1-\sigma)} \times \left(\frac{\eta \cos\alpha}{(r-\zeta)} - \frac{y}{r} - \frac{\eta z}{r(r-\zeta)} - (1-2\sigma)\sin\alpha \log(r-\zeta) \right)$$

Elizabeth H. Yoffe (1960): The angular dislocation, *Philosophical Magazine*, 5:50, 161-175

Map Pattern – Independent Loop Iterations

$$w = \frac{b}{8\pi(1-\sigma)} \times \left(\frac{\eta \cos\alpha}{(r-\zeta)} - \frac{y}{r} - \frac{\eta z}{r(r-\zeta)} - (1-2\sigma)\sin\alpha \log(r-\zeta) \right)$$

Elizabeth H. Yoffe (1960): The angular dislocation, *Philosophical Magazine*, 5:50, 161-175

```
// Compute wx
#pragma ivdep
for ( i = 0; i < dimension; i++ )
{
    wx[i] = (bx / ( 8.0 * M_PI * (1 - nu) )) \
          * ( ((eta[i] * cosA) / (r[i] - zeta[i])) \
          - (y[i] / r[i]) \
          - ((eta[i] * z[i]) / ( r[i] * (r[i] - zeta[i]) )) \
          - ((1 - (2 * nu)) * (sinA * log(r[i] - zeta[i]))) );
}
```

Mapping Independent Loop Iterations to Processing Units

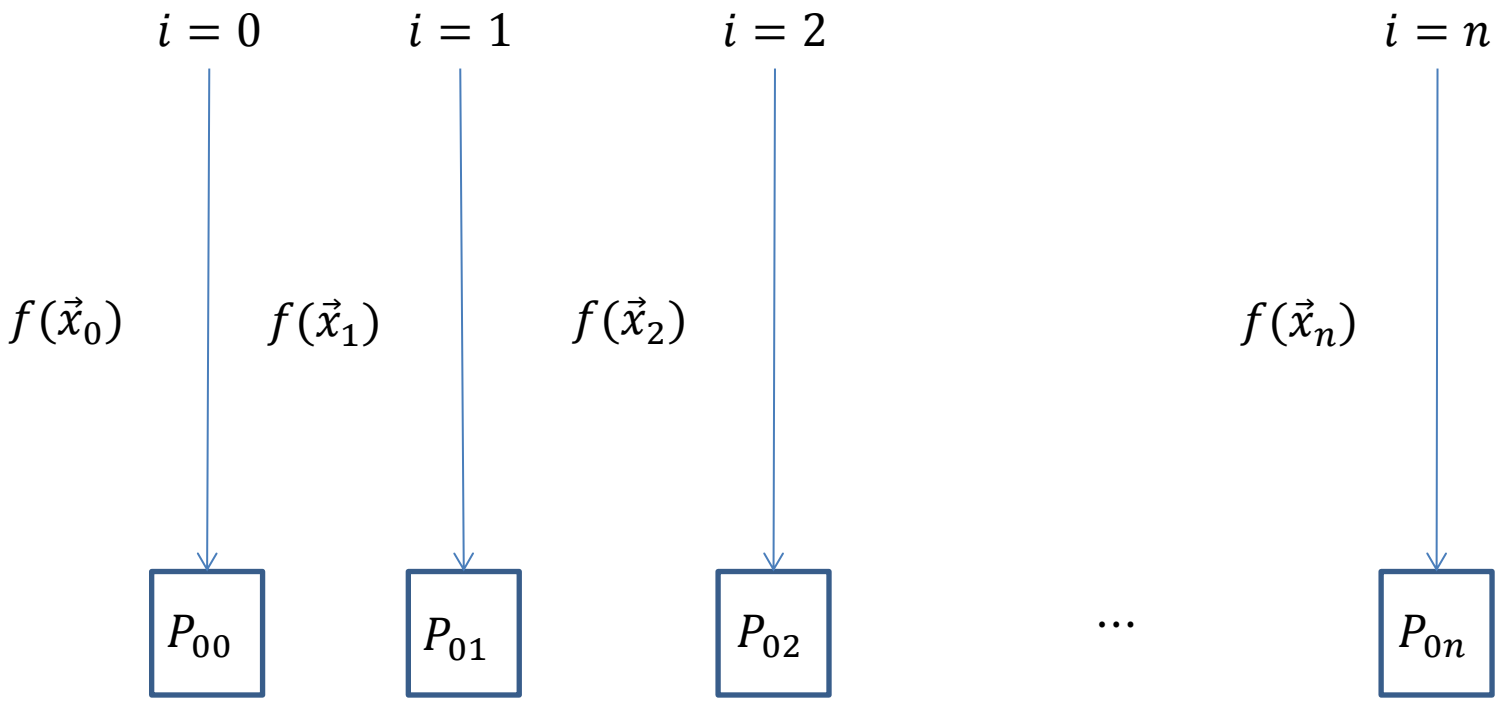

Reduce Pattern – Data Dependency in a Loop

- *Reduction*
  - Fuse all elements in an array into one element, using an associative (and commutative) operation, e.g., sum
- Serial +Reduction

```
for(int i = 0; i < dimension; ++i) {
    sum += input[i];
}
```

Parallelization despite Dependence – Parallel Reduction

Parallelization despite Dependence – Parallel Reduction

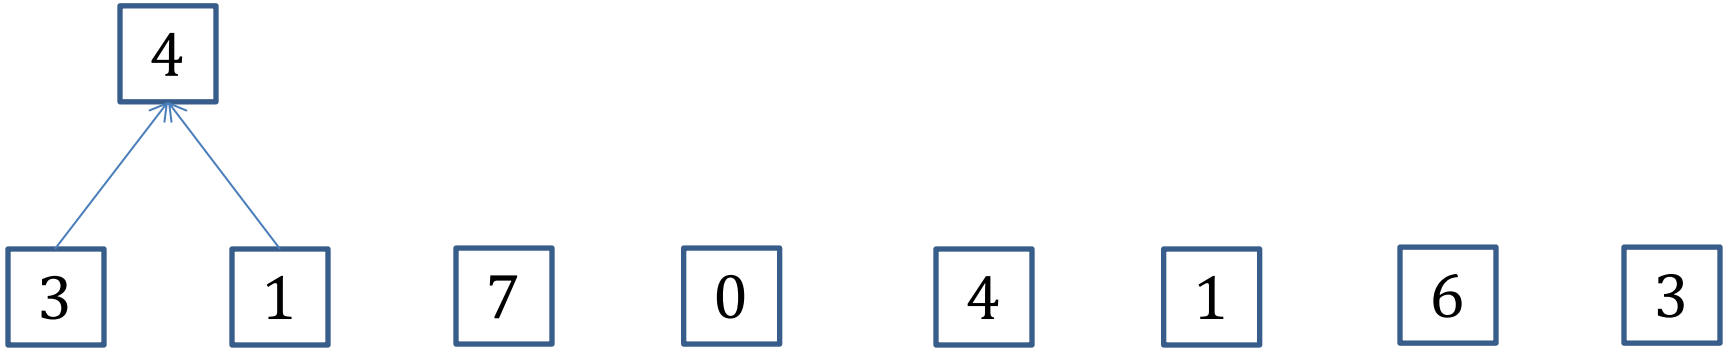

Parallelization despite Dependence – Parallel Reduction

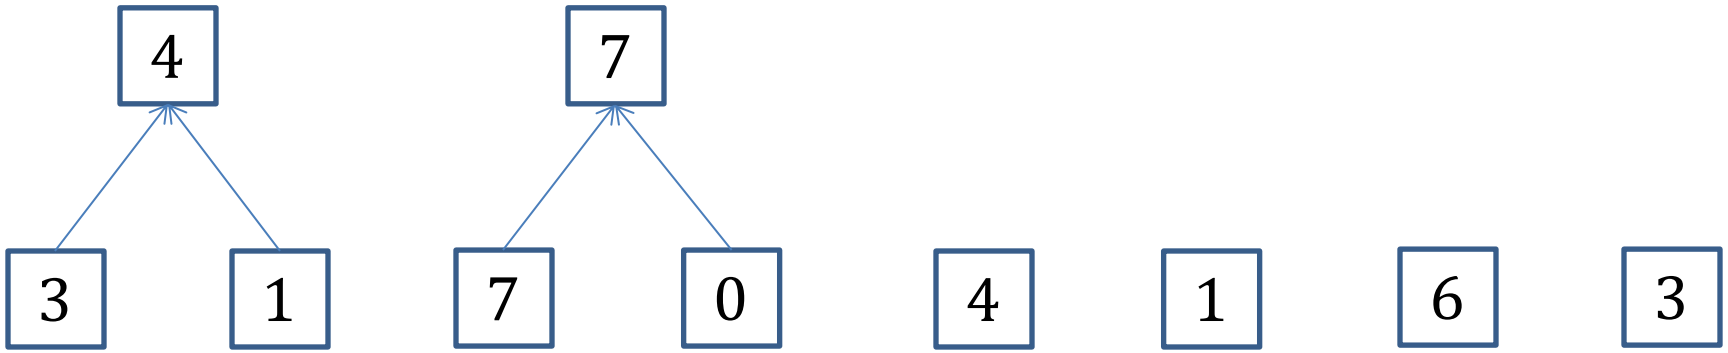

Parallelization despite Dependence – Parallel Reduction

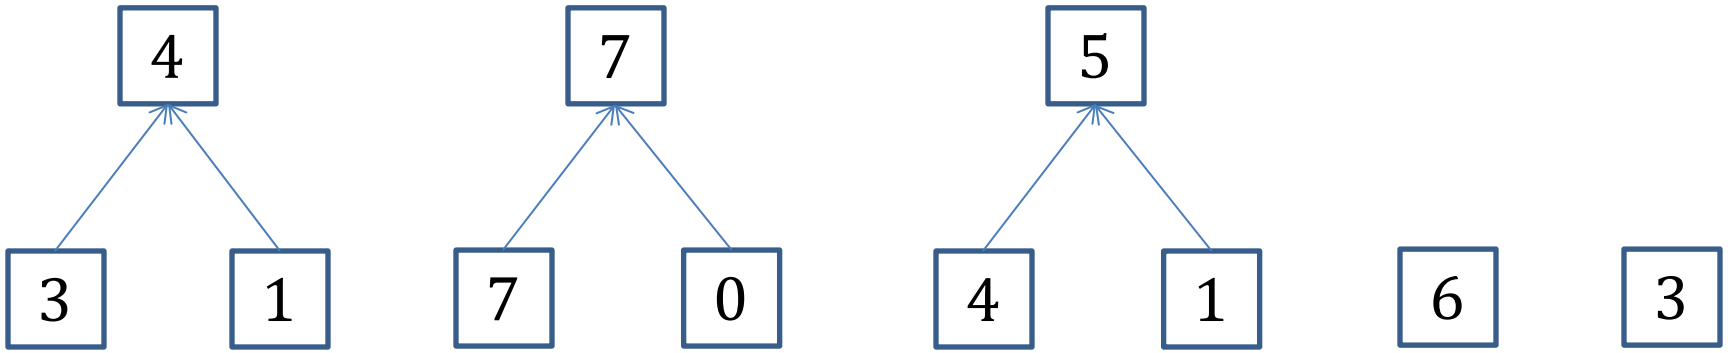

Parallelization despite Dependence – Parallel Reduction

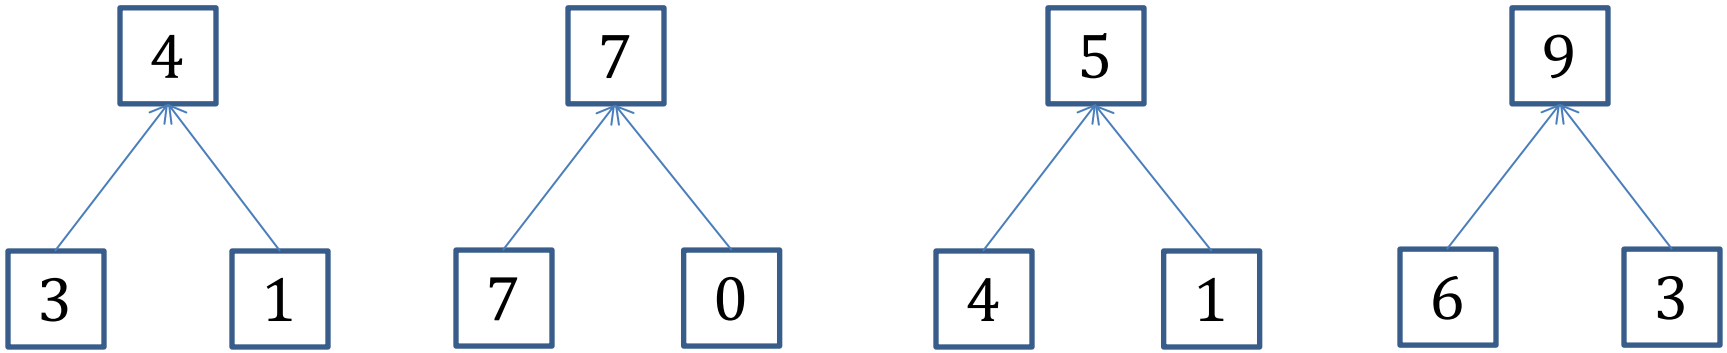

Parallelization despite Dependence – Parallel Reduction

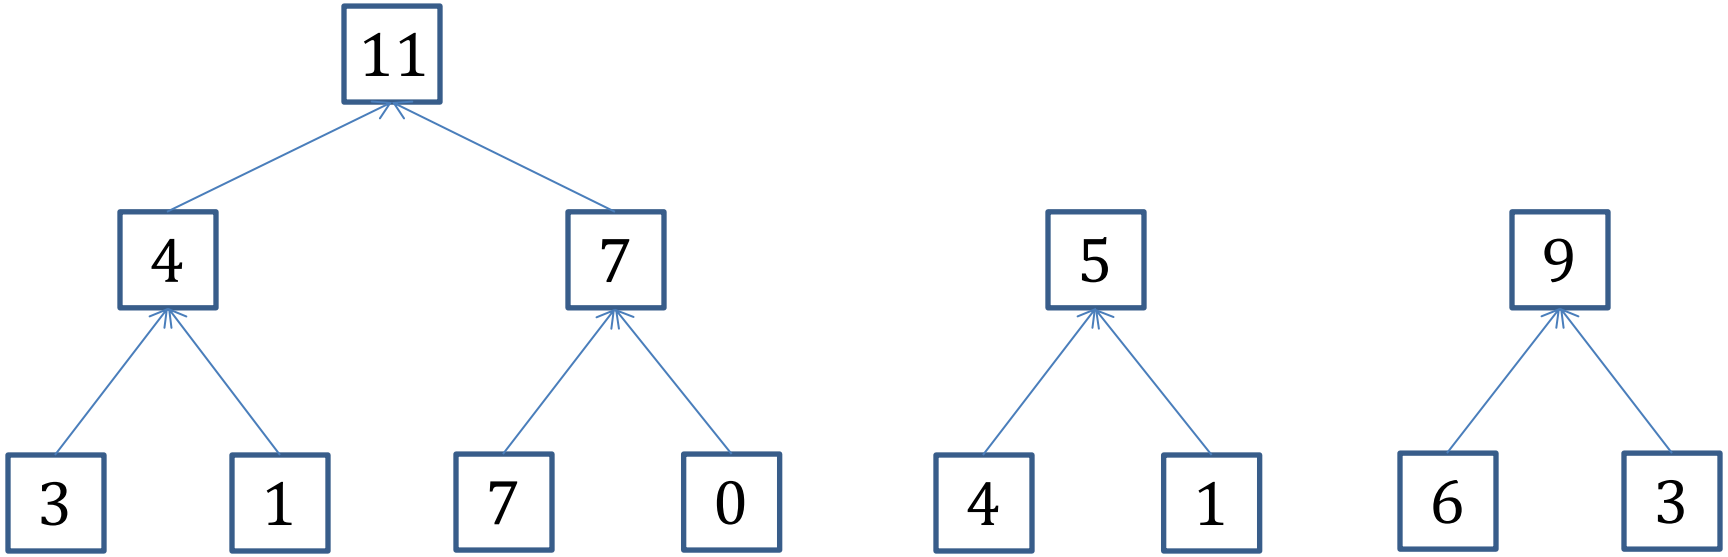

Parallelization despite Dependence – Parallel Reduction

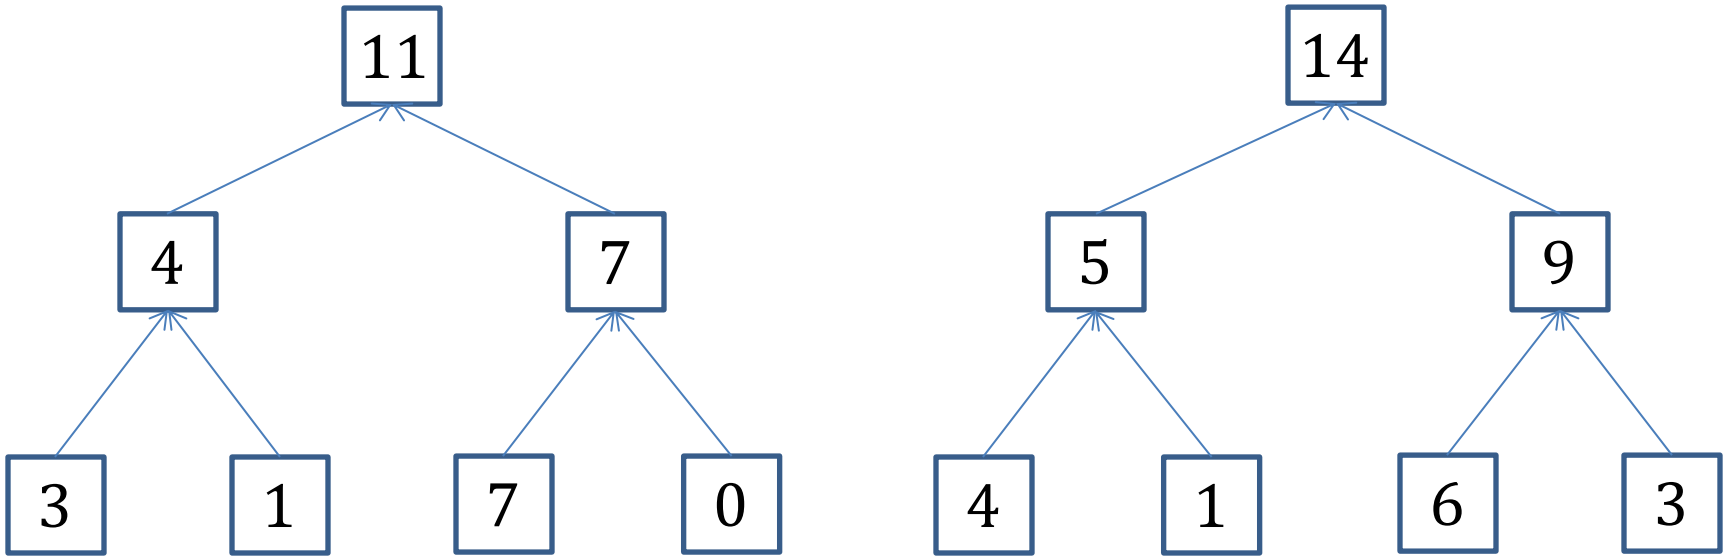

Parallelization despite Dependence – Parallel Reduction

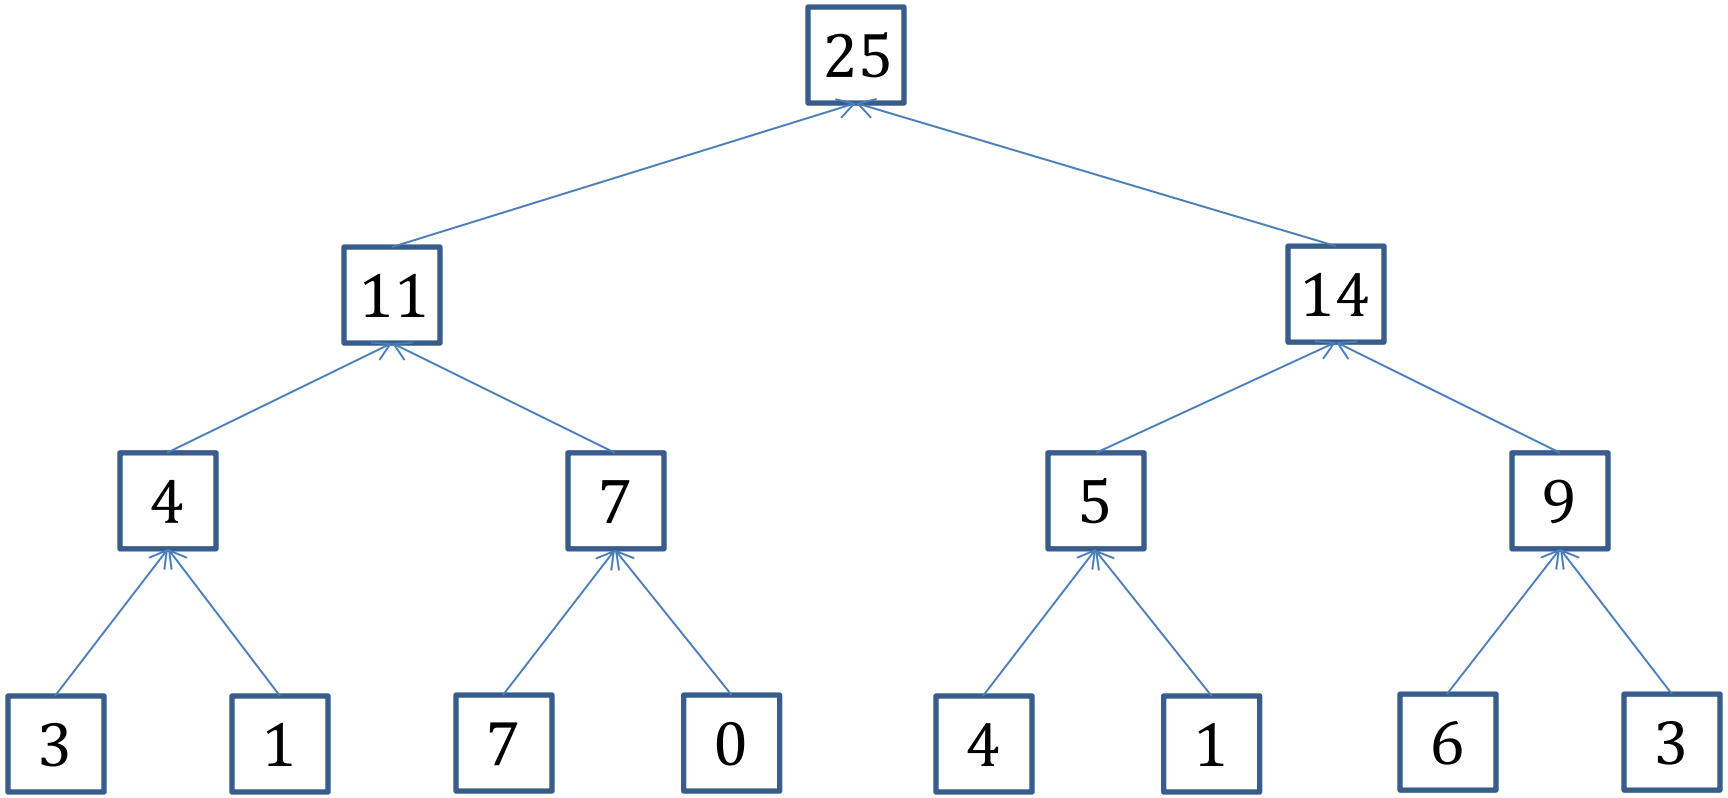

Mapping Tree Reduction to Processing Units

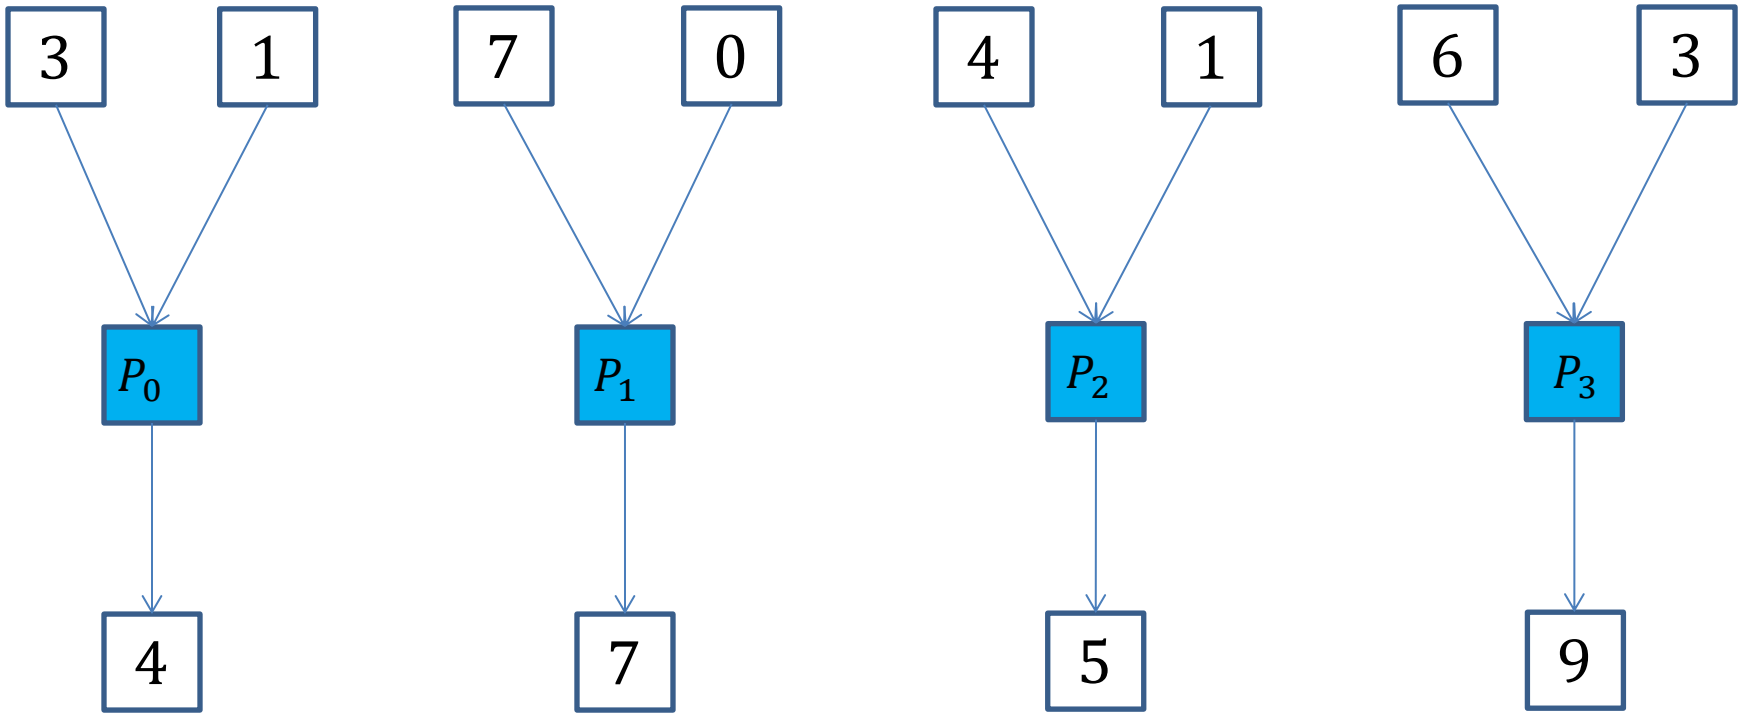

Mapping Tree Reduction to Processing Units

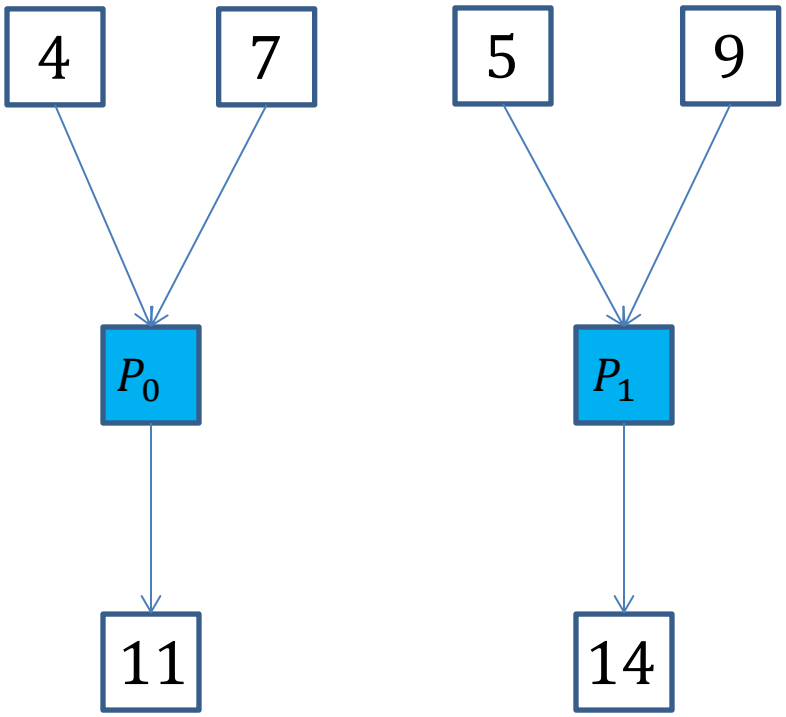

Mapping Tree Reduction to Processing Units

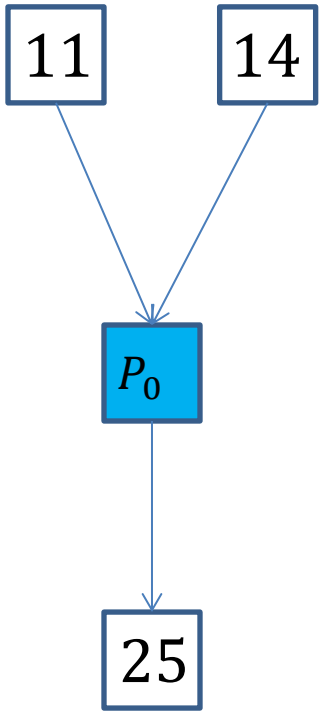

The Stencil Pattern – Neighborhoods

- *Stencil*
  - "A map in which each output depends on a *neighborhood* of inputs specified using a set of fixed offsets relative to the output position."

Ref: McCool, Michael, James Reinders, and Arch Robison. *Structured parallel programming: patterns for efficient computation*. Elsevier, 2012.

The Stencil Pattern – Neighborhoods

- *Stencil*
  - "A map in which each output depends on a *neighborhood* of inputs specified using a set of fixed offsets relative to the output position."

Ref: McCool, Michael, James Reinders, and Arch Robison. *Structured parallel programming: patterns for efficient computation*. Elsevier, 2012.

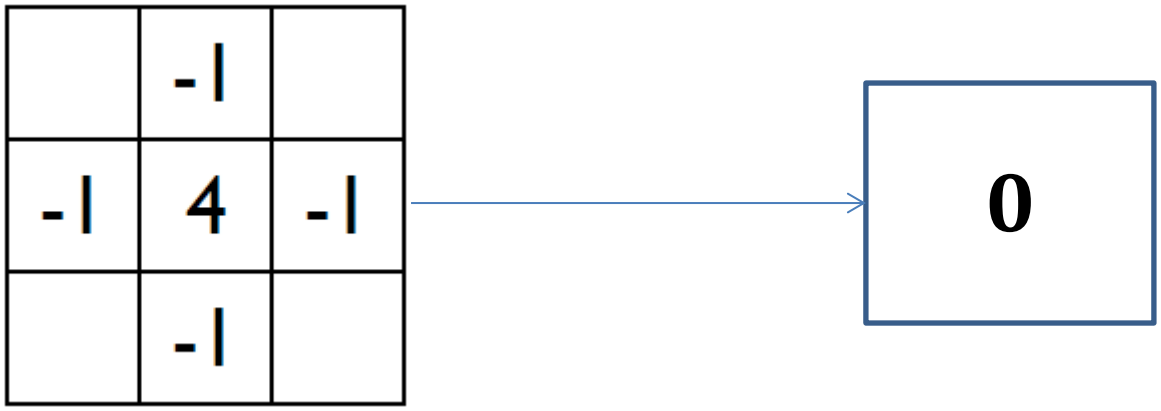

The Stencil Pattern

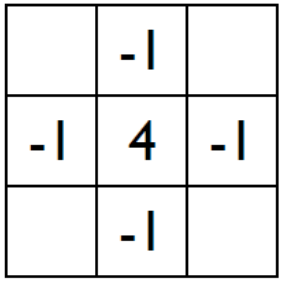

(a) 5-Point Stencil

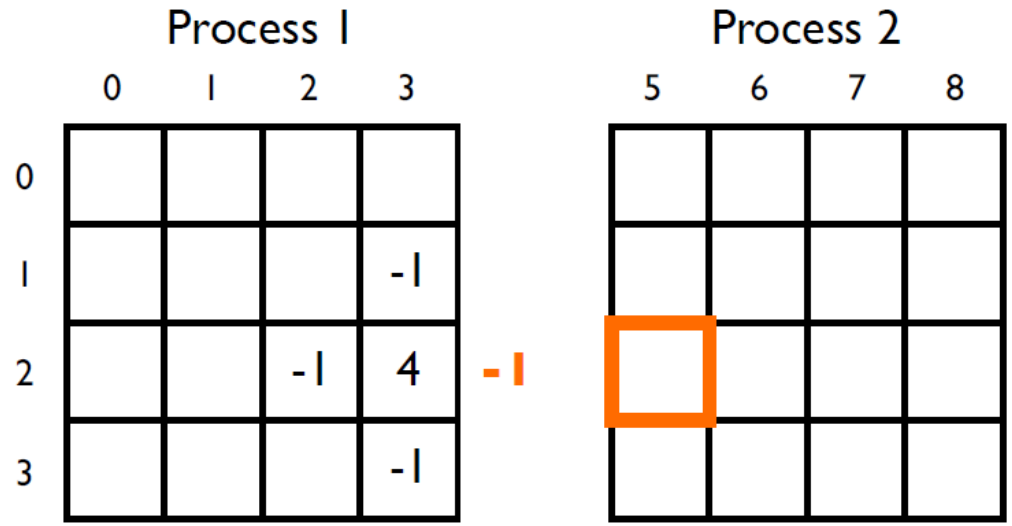

(b) Stencil that needs a cell from its neighbor

Reference: Kjolstad, Fredrik Berg, and Marc Snir. "Ghost cell pattern." *Proceedings of the 2010 Workshop on Parallel Programming Patterns*. ACM, 2010.

The Stencil Pattern – Halo/Ghost Cells

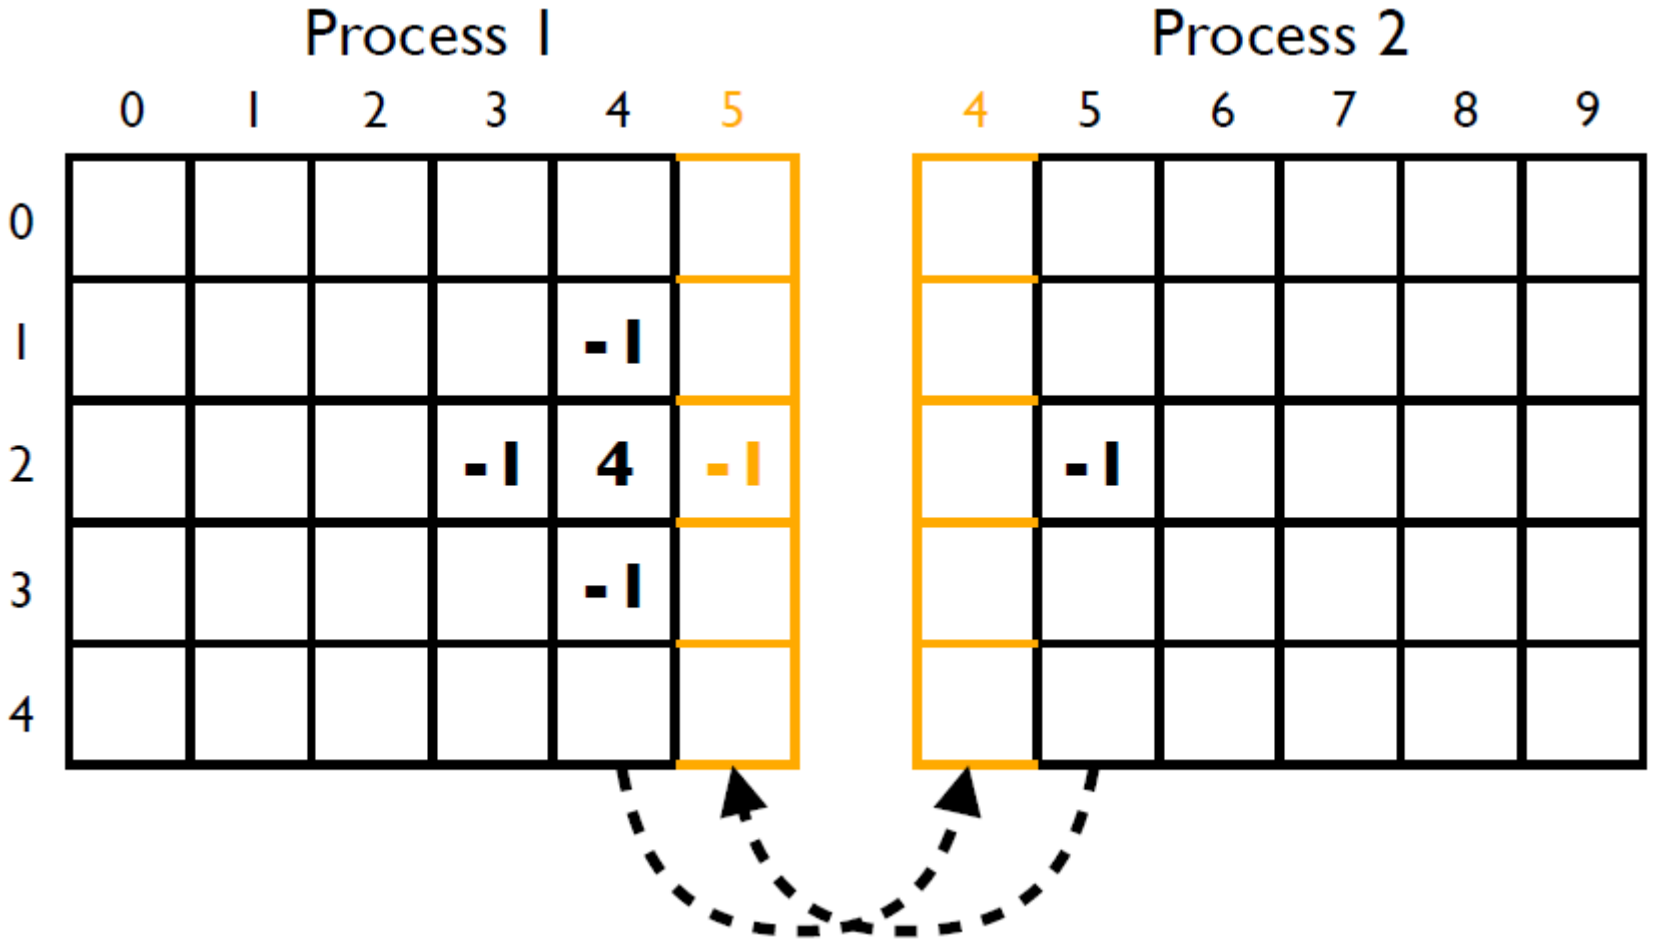

Reference: Kjolstad, Fredrik Berg, and Marc Snir. "Ghost cell pattern." *Proceedings of the 2010 Workshop on Parallel Programming Patterns*. ACM, 2010.

Concurrency and Resource Sharing

- Multiple processes trying to access a shared resource, e.g., concurrent read and write to the same memory location
  - Leads to *Resource Contention*

Concurrency and Resource Sharing

- Multiple processes trying to access a shared resource, e.g., concurrent read and write to the same memory location
  - Leads to *Resource Contention*
- This makes concurrent/parallel programming hard
  - May lead to *Race Condition,* i.e., the final result of an operation depends on the order of execution

Concurrency and Resource Sharing

- Multiple processes trying to access a shared resource, e.g., concurrent read and write to the same memory location
  - Leads to *Resource Contention*
- This makes concurrent/parallel programming hard
  - May lead to *Race Condition,* i.e., the final result of an operation depends on the order of execution
    - *Deadlock:* Two or more concurrent tasks are unable to proceed, because each is waiting for one of the others to do something.

Concurrency and Resource Sharing

- Multiple processes trying to access a shared resource, e.g., concurrent read and write to the same memory location
  - Leads to *Resource Contention*
- This makes concurrent/parallel programming hard
  - May lead to *Race Condition,* i.e., the final result of an operation depends on the order of execution
    - *Deadlock:* Two or more concurrent tasks are unable to proceed, because each is waiting for one of the others to do something.
    - *Starvation:* A runnable task is overlooked indefinitely. Even though it is able to proceed, it is never chosen to run.

The *Minion-Banana* Tutorial on Concurrency and Resource Contention

Race Condition

```
1  #define NUMTHREADS 10000
2
3  int sharedData=0;
4
5  void *doWork(void *parm) {
6      ++sharedData;
7      return NULL;
8  }
9
10 int main(int argc, char **argv) {
11     int i;
12     pthread_t thread[NUMTHREADS];
13     for (i=0; i<NUMTHREADS; ++i) {    // create 3 threads
14       pthread_create(&thread[i], NULL, doWork, NULL);
15     }
16     for (i=0; i <NUMTHREADS; ++i) {
17       pthread_join(thread[i], NULL);
18     }
19     printf("Shared Data: %d\n",sharedData);
20     pthread_exit(NULL);
21     return 0;
22 }
```

Mutex

```
1  #define NUMTHREADS 10000
2
3  int sharedData=0;
4  pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
5
6  void *doWork(void *parm) {
7      pthread_mutex_lock(&mutex);
8      ++sharedData;
9      pthread_mutex_unlock(&mutex);
10     return NULL;
11 }
12
13 int main(int argc, char **argv) {
14     int i;
15     pthread_t thread[NUMTHREADS];
16     for (i=0; i<NUMTHREADS; ++i) {   // create 3 threads
17       pthread_create(&thread[i], NULL, doWork, NULL);
18     }
19     for (i=0; i <NUMTHREADS; ++i) {
20       pthread_join(thread[i], NULL);
21     }
22     printf("Shared Data: %d\n",sharedData);
23     pthread_mutex_destroy(&mutex);
24     pthread_exit(NULL);
25     return 0;
26 }
```

Starvation

```
1  #define NUMTHREADS 2
2  int sharedData = 0;
3  pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
4
5  void *doWork(void *param) {
6      int threadId = *(int*)param;
7      pthread_mutex_lock(&mutex);
8      ++sharedData;
9      return NULL;
10 }
11
12 int main(int argc, char **argv) {
13     int i;
14     pthread_t thread[NUMTHREADS];
15     for (i = 0; i < NUMTHREADS; ++i) {
16       pthread_create(&thread[i], NULL, doWork, &i);
17     }
18     for (i = 0; i < NUMTHREADS; ++i) {
19       pthread_join(thread[i], NULL);
20     }
21     pthread_mutex_destroy(&mutex);
22     pthread_exit(NULL);
23     return 0;
24
25 }
```

The Fork-Join Pattern

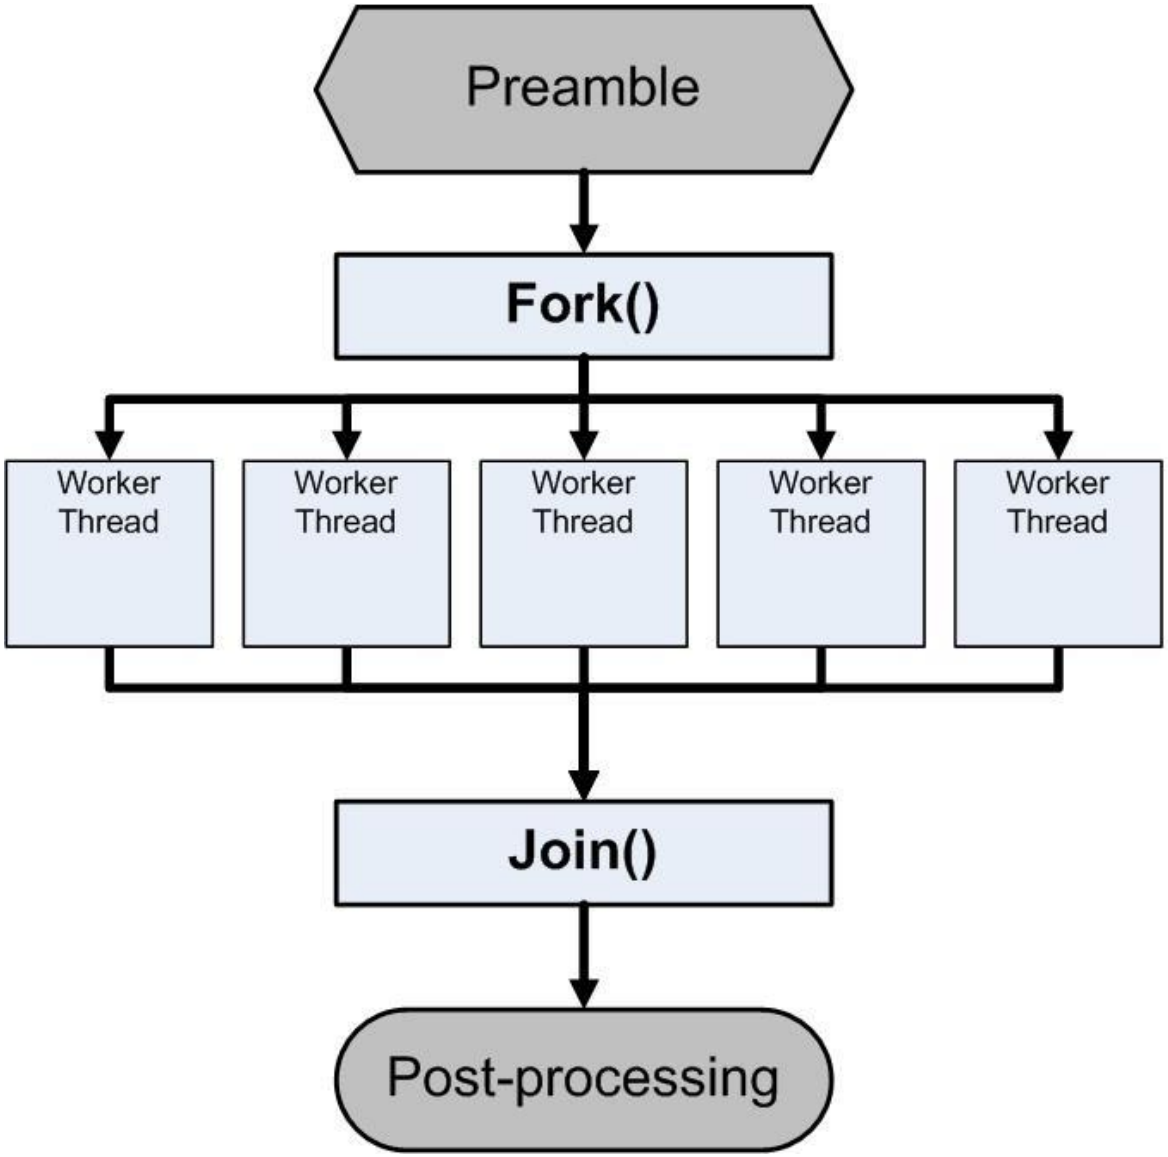

Ref: Mark B Friedman. *Parallel Scalability Isn't Child's Play, Part 2: Amdahl's Law vs. Gunther's Law*. Developer Division Performance Engineering blog, MSDN blogs.

Parallel Pipeline Pattern

Serial processing of stages

Parallel Pipeline Pattern

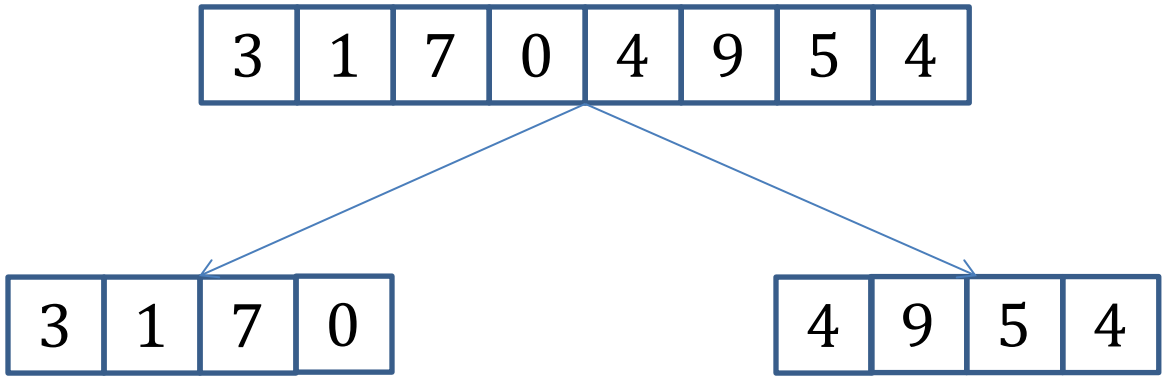

Serial processing of stages

Parallel Pipeline Pattern

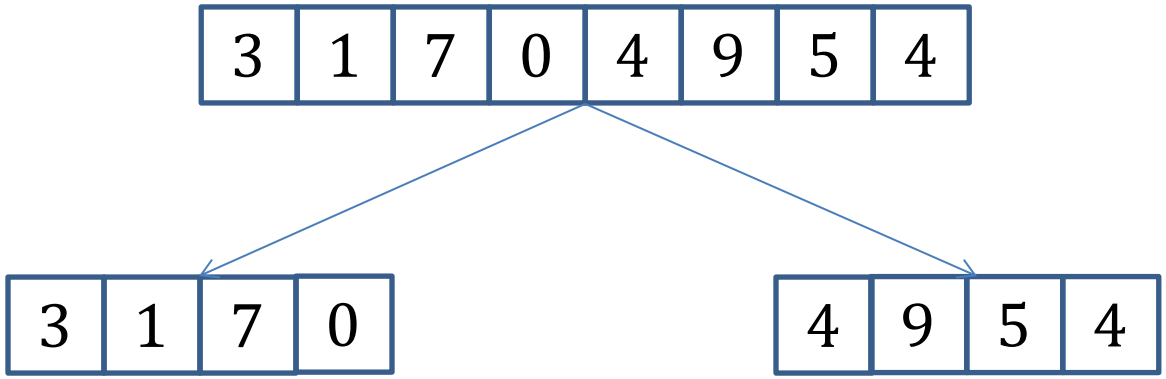

Serial processing of stages

Pipelined processing of stages

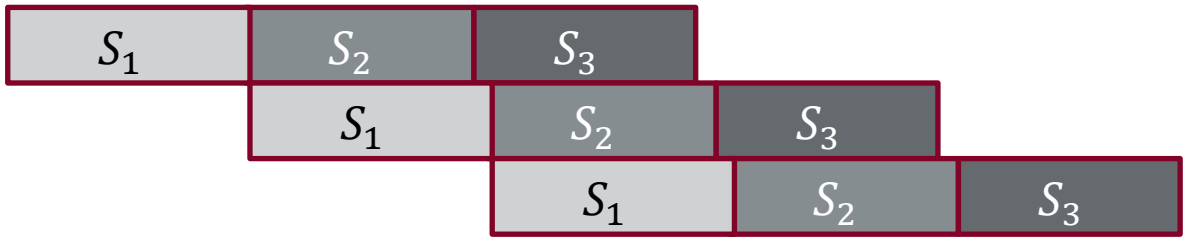

Part - III

Software Performance Optimization

"I really hate this damned machine;
I wish that they would sell it.
It never does quite what I want
but only what I tell it.“

A Programmer’s Lament

Parallel Hardware Architecture and code Optimization

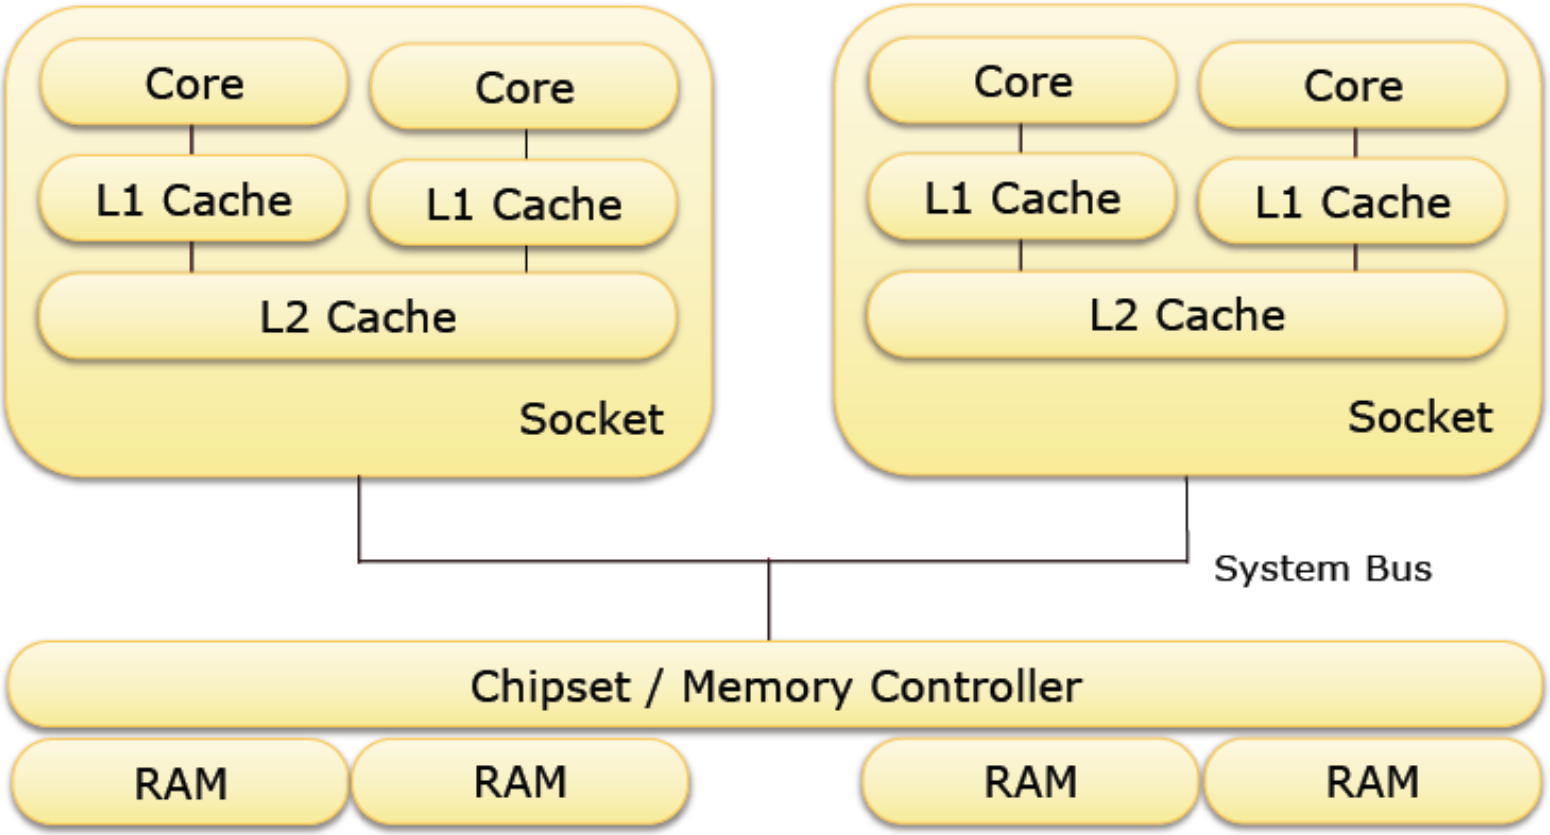

- Code must optimally utilize this complex hierarchy

Parallel Hardware Architecture and code Optimization

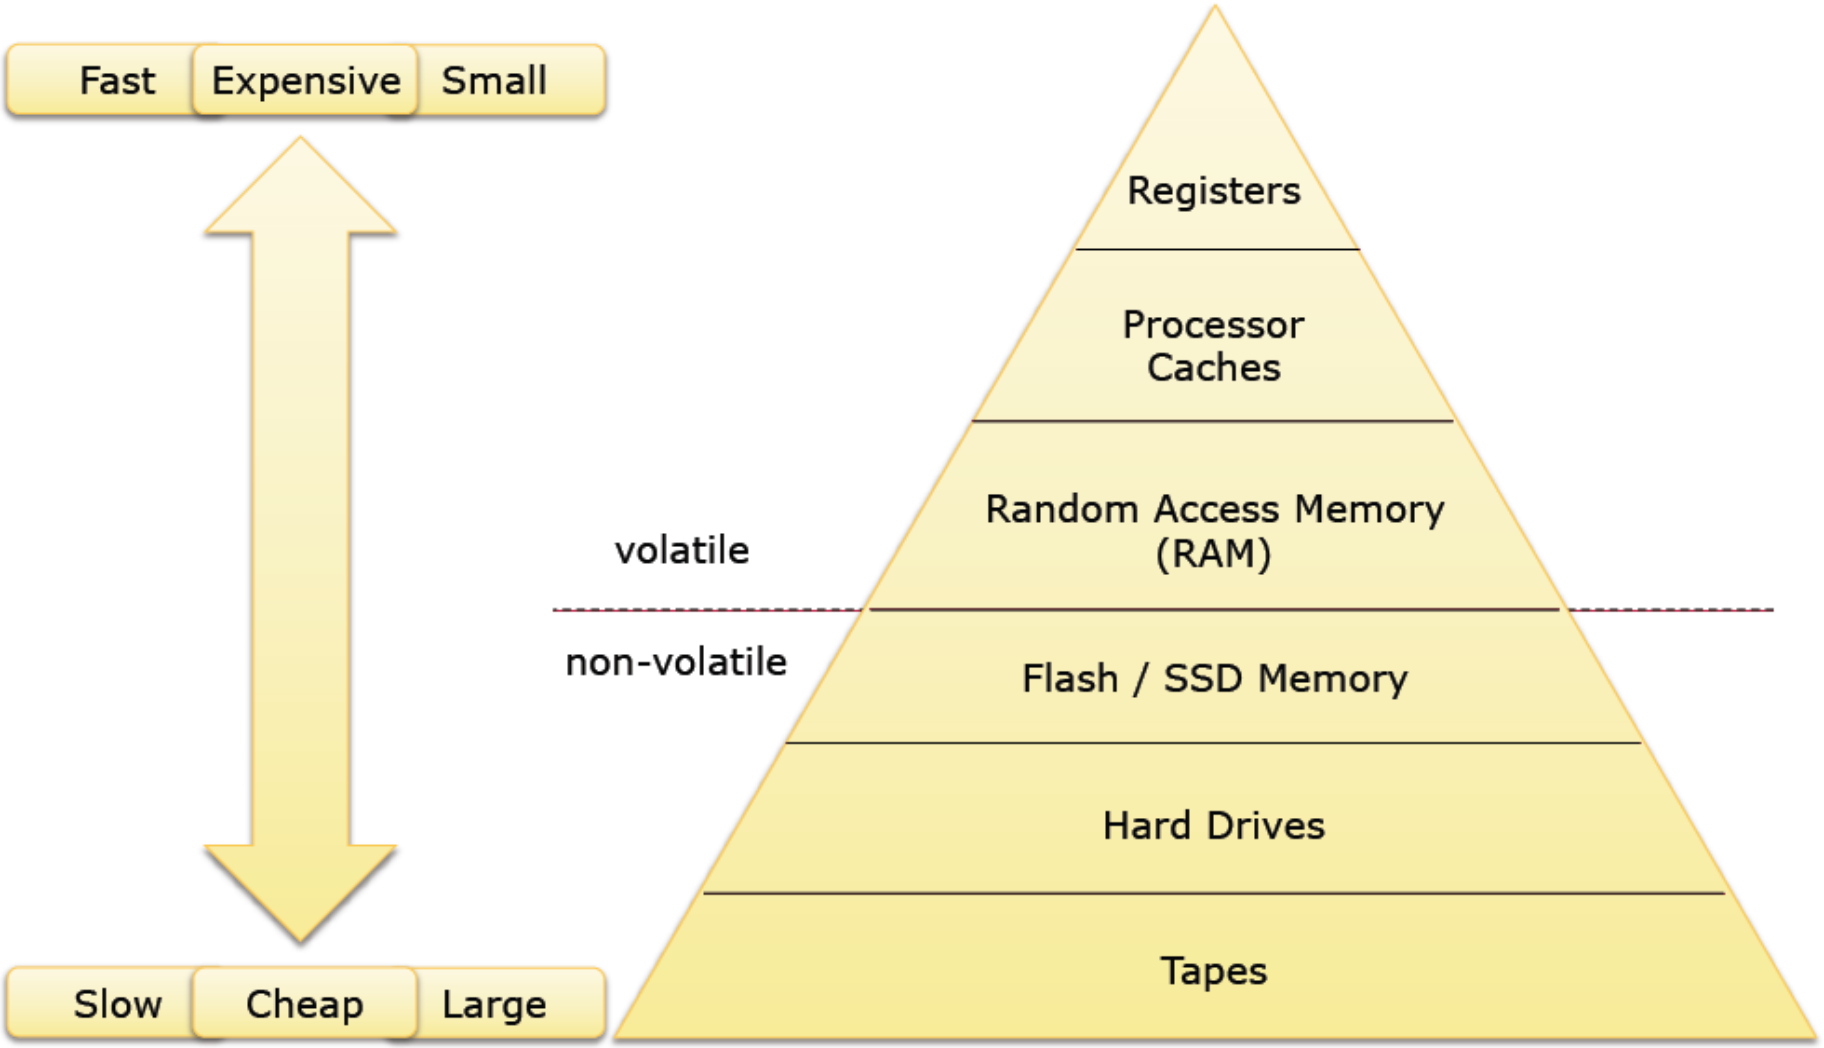

Part II – Motivation

Escaping the Ivory Tower Approach to Parallel Computing

Real Scientist: You suck!
Fahad: I know … I'm conscious, self-aware, …
Real Scientist : No, I mean you *really* suck!
Fahad: Oh, gee … thanks?
Real Scientist: I see all this parallel programming crap, but where are your applications? … How do I use all this stuff?
Fahad: ☹

Simulation of Seismic Events following Earthquakes

- Collaborative project with the *GFZ German Research Center for Geosciences*
  - The *2011 earthquake off the Pacific coast of Tōhoku, Japan*
  - 250 days of data used in the numerical simulation
  - Prediction of the *No. and location of aftershocks*, e.g., *Fukushima*

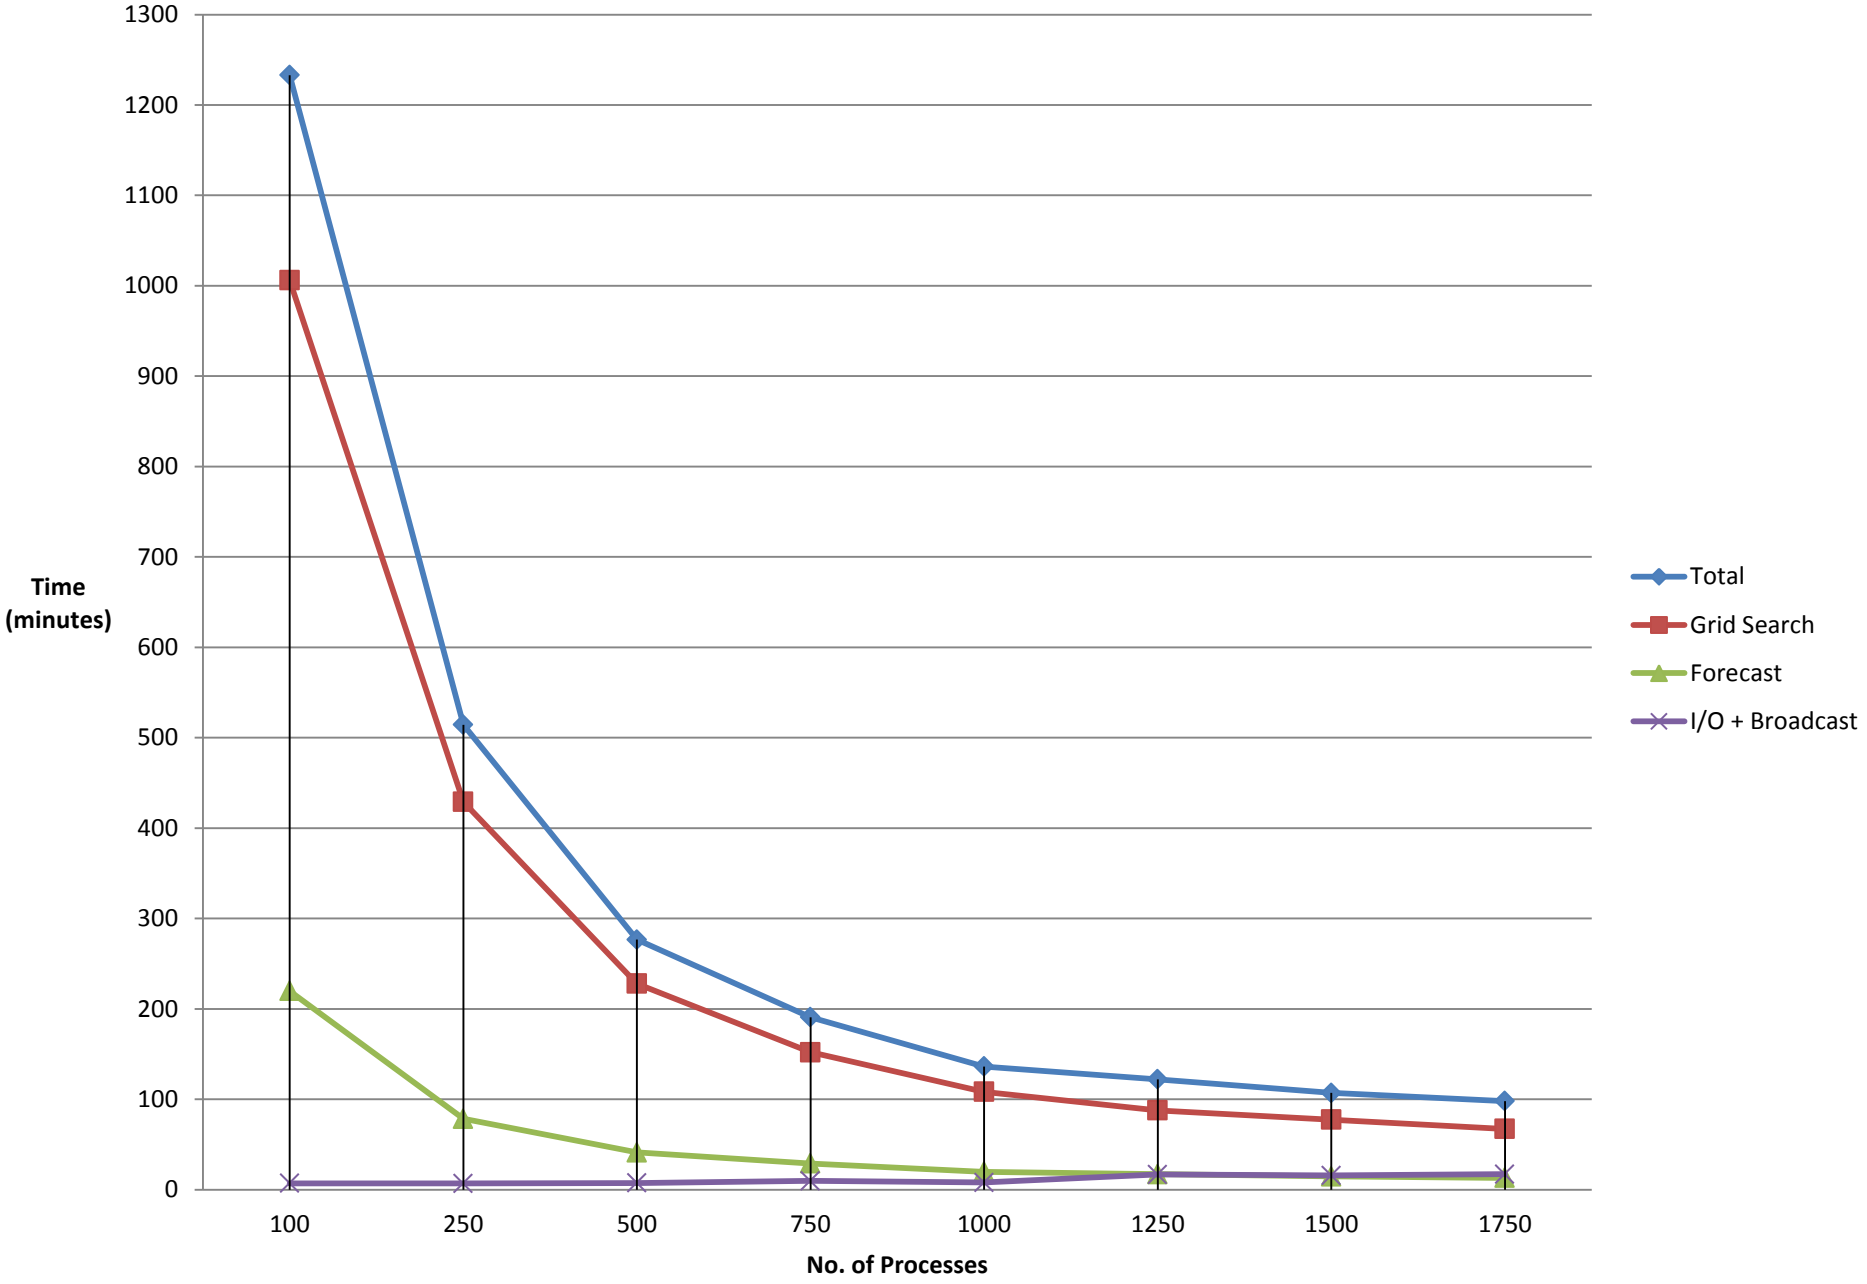
1300
1200
1100
1000
900
800
700
600
500
400
300
200
100
0
Time (minutes)
100
250
500
750
1000
1250
1500
1750
No. of Processes
Total
Grid Search
Forecast
I/O + Broadcast

Simulation of Seismic Events following Earthquakes

- Distributed-memory Parallelization using MPI
  - *1000-core FutureSOC cluster* used for performance evaluation
  - Successful scaling up to *1750* processes
  - Simulation run that would have taken *~6 days* can now be done in *~1.5 hours*

Enumeration of Elementary Flux Modes in Metabolic Networks

Guinea Pig

Enumeration of Elementary Flux Modes in Metabolic Networks

Guinea Pig

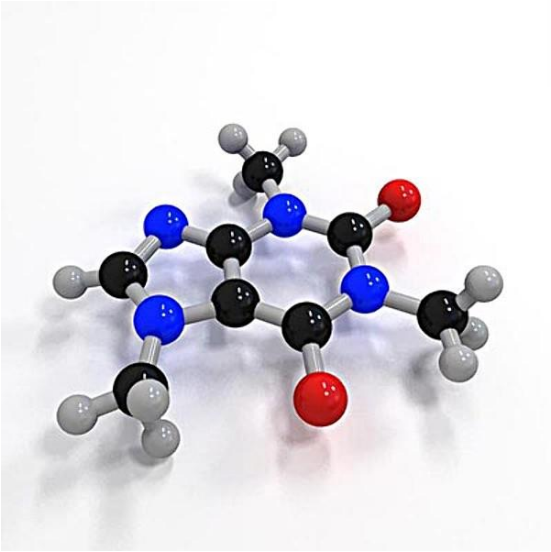

Caffeine Molecule

Enumeration of Elementary Flux Modes in Metabolic Networks

Guinea Pig

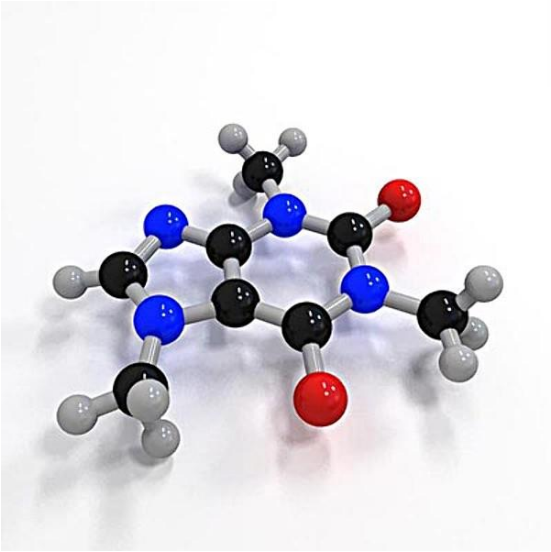

Caffeine Molecule

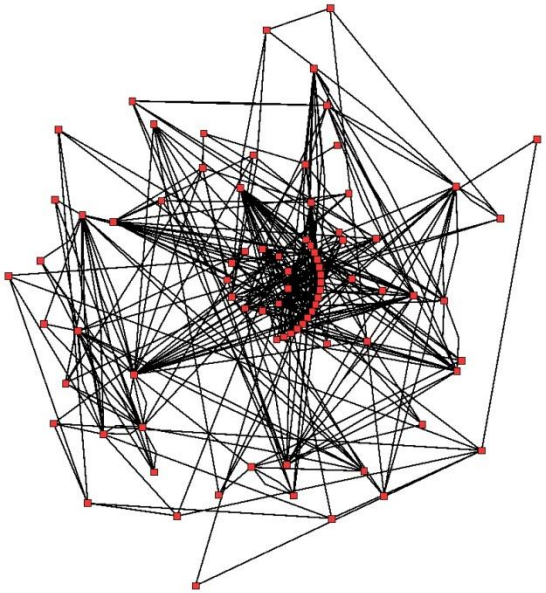

Plant Metabolic Network

Enumeration of Elementary Flux Modes in Metabolic Networks

Guinea Pig

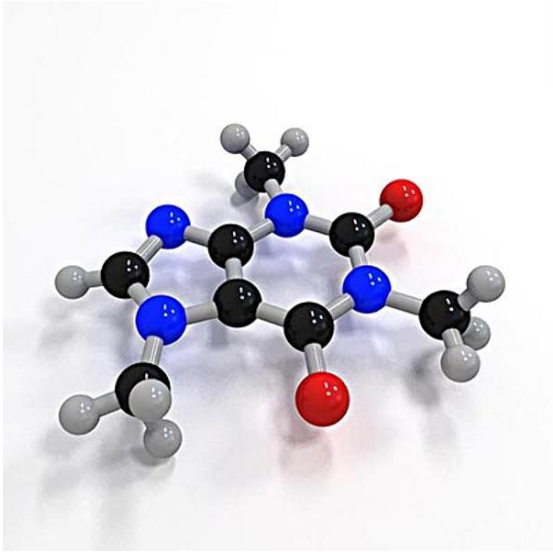

Caffeine Molecule

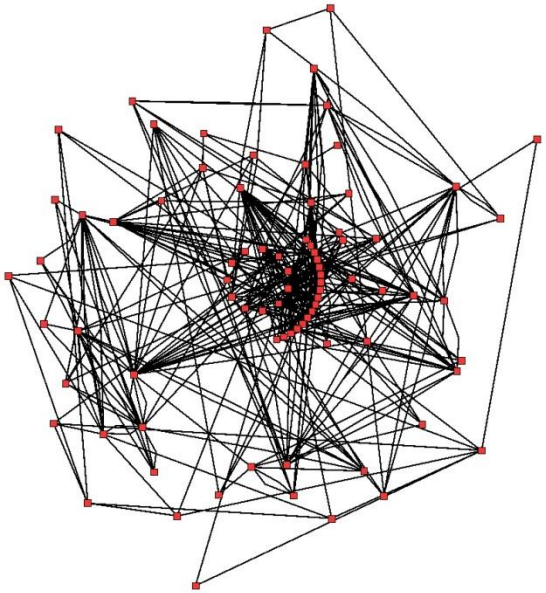

Plant Metabolic Network

Enumeration of Elementary Flux Modes in Metabolic Networks

- Objective
  - To understand structural properties of *Metabolic Networks*
  - Understand the range of *Metabolic Pathways* in the network
  - Gain insight into the *contribution of pathways* to the overall metabolic network behavior
- Potential Applications
  - Treatments for metabolic disorders such as *Diabetes*
  - Therapies for *Cancer* patients
  - Metabolic Engineering: Production of *Biofuels*

Enumeration of Elementary Flux Modes in Metabolic Networks

$$\frac{dX_1}{dt} = 1 \cdot v_1 - 1 \cdot v_2 - 1 \cdot v_4$$

$$\frac{dX_2}{dt} = 2 \cdot v_2 - 1 \cdot v_3$$

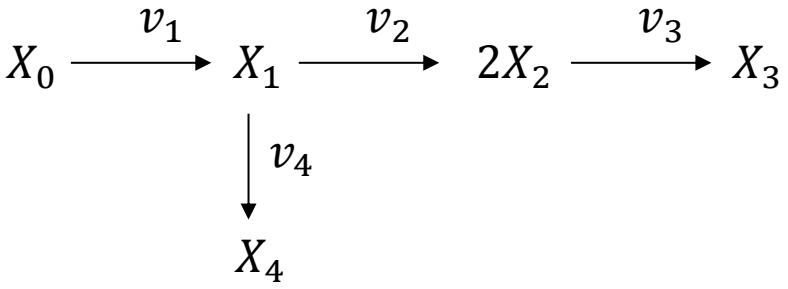

$$\Rightarrow \begin{pmatrix} \frac{dX_1}{dt} \\ \vdots \\ \frac{dX_m}{dt} \end{pmatrix} = \begin{pmatrix} 1 & -1 & \cdots & S_{1n} \\ \vdots & \ddots & & \vdots \\ S_{m1} & \cdots & & S_{mn} \end{pmatrix} \cdot \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$$

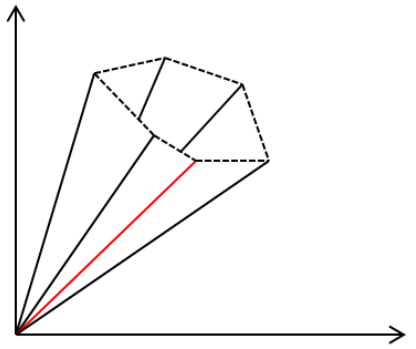

- S is the stoichiometric matrix
- At steady state, $0 = S \cdot v$
- Also, some reactions are irreversible, i.e., $v_i \geq 0$
- This makes the solution space a *convex polyhedral cone*

Enumeration of Elementary Flux Modes in Metabolic Networks

- Elementary Flux Mode (EFM)
  - *is a minimal set of reactions that can operate at steady state*
  - EFMs span the feasible space of *flux distributions*
  - Any feasible flux distribution can be represented as a linear combination of EFMs
- Example

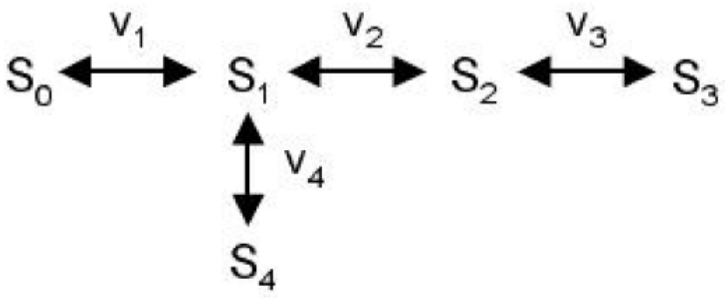

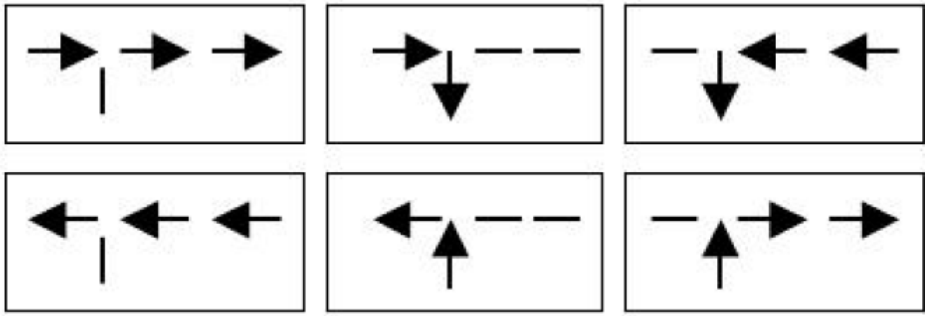

Enumeration of Elementary Flux Modes in Metabolic Networks

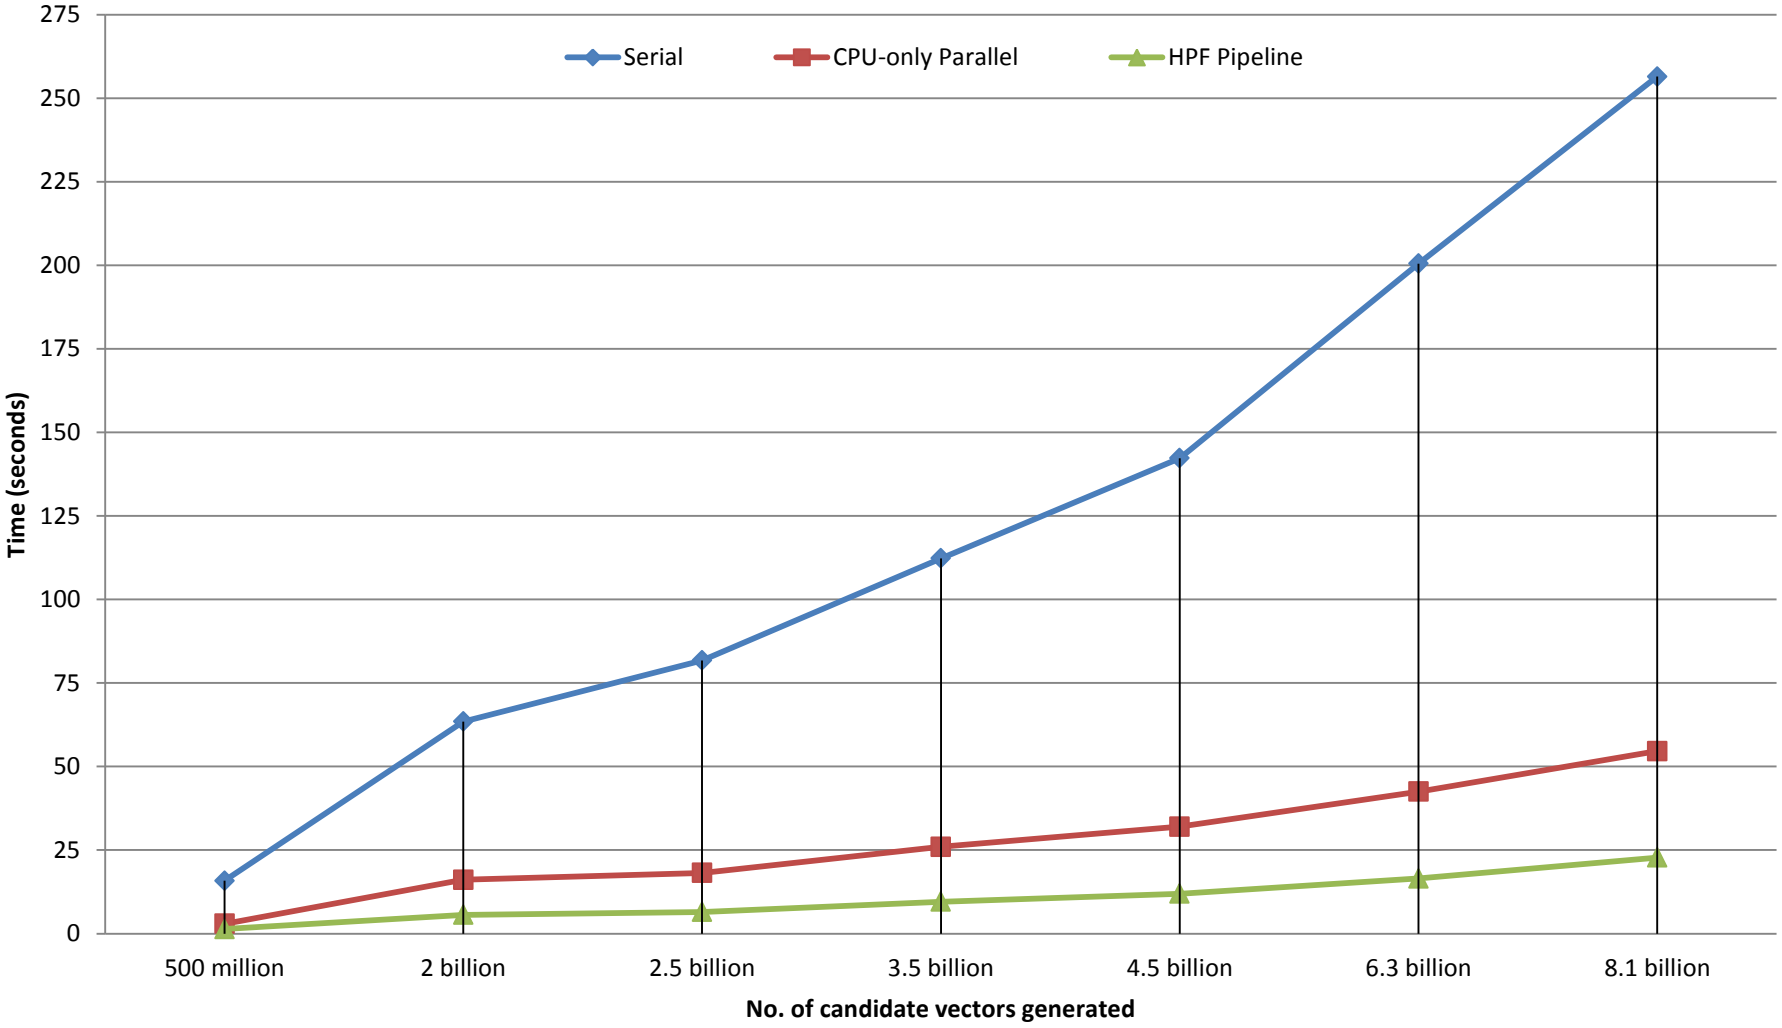

Recap

- *Concurrency* vs. *Parallelism*
- *Shared-memory* vs. *Distributed-memory* Parallelism
- *Hotspots*
- Patterns in Parallel Programming: *Geometric Decomposition*, *Map*, *Reduce*, *Stencil*, *Fork-Join*, *Pipeline*, …
- *Resource contention*, *Race condition*, *Deadlock*, *Starvation*
- *Software Performance Optimization*

Recap

- *Concurrency* vs. *Parallelism*
- *Shared-memory* vs. *Distributed-memory* Parallelism
- *Hotspots*
- Patterns in Parallel Programming: *Geometric Decomposition*, *Map*, *Reduce*, *Stencil*, *Fork-Join*, *Pipeline*, …
- *Resource contention*, *Race condition*, *Deadlock*, *Starvation*
- *Software Performance Optimization*

Thank You!

Literature

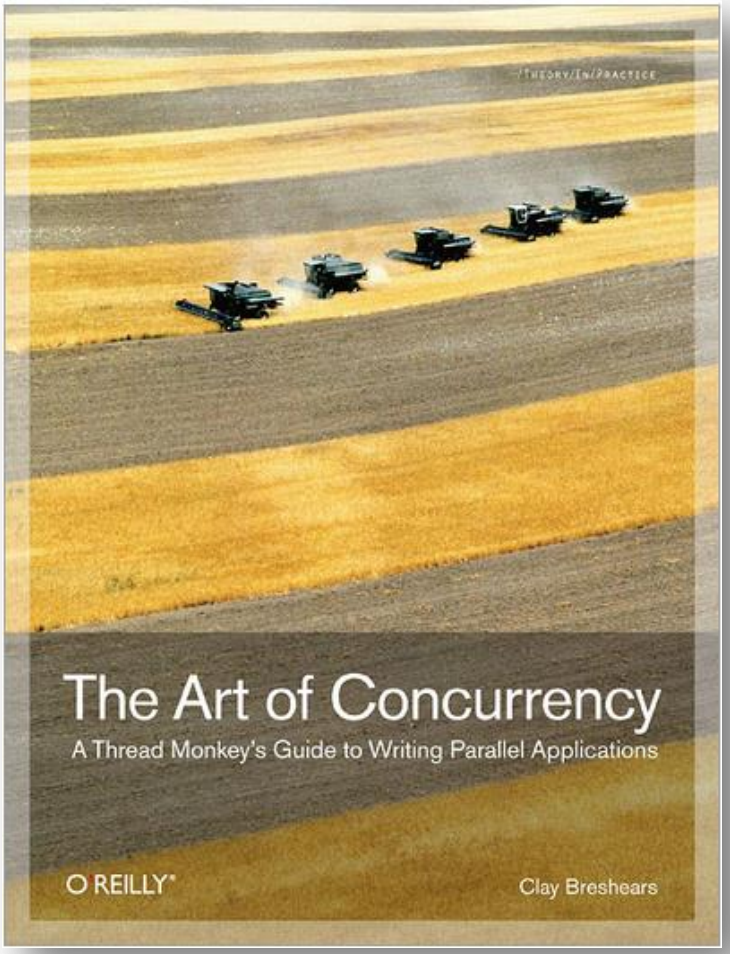

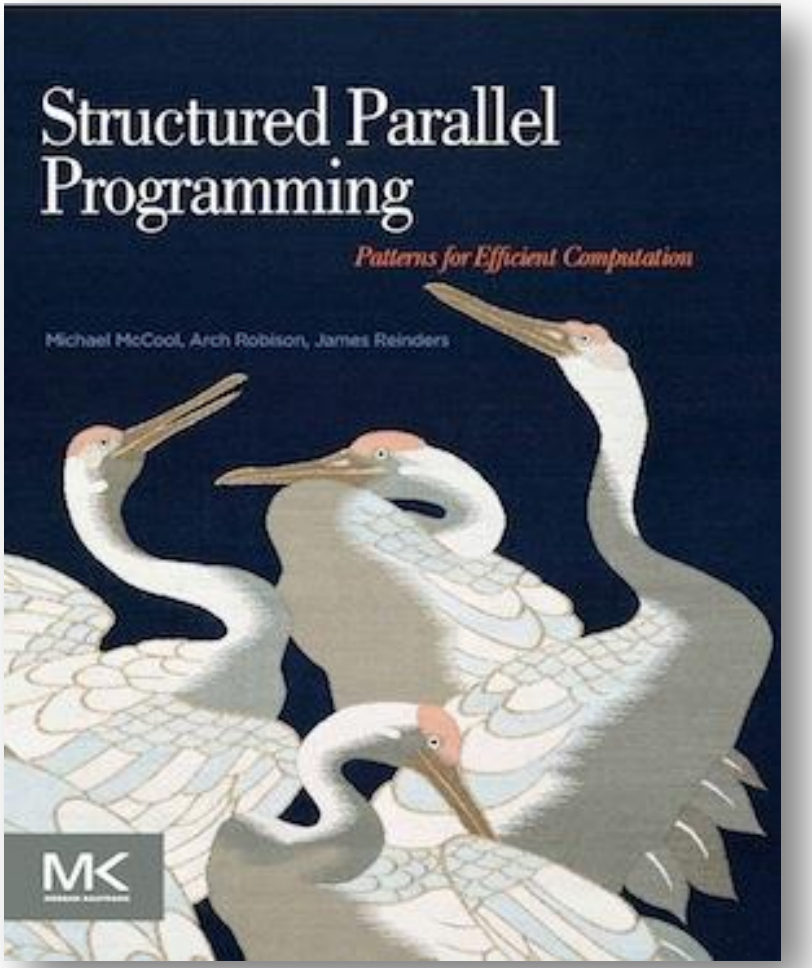

Free MOOC – *Parallel Programming Concepts*, OpenHPI