


# Introduction to Parallel Programming

Center for Institutional Research  
Computing




*Slides for the book "An introduction to Parallel Programming", by Peter Pacheco (available from the publisher*

*website):*  
[742605/](http://booksite.elsevier.com/9780123742605/)

<http://booksite.elsevier.com/9780123742605/>

# Serial hardware and software



# Why we need to write parallel programs

- Running multiple instances of a serial program often is not very useful.
  - Have the same program run 100 times
  - Have 100 computers run the same program 1 time

# Why we need to write parallel programs

- Running multiple instances of a serial program often is not very useful.
  - Have the same program run 100 times
  - Have 100 computers run the same program 1 time
- What you really want is to make the overall process finish faster.

# How do we write parallel programs?

- Partition the workload and let CPU cores work in parallel
  - Task parallelism
    - Partition various tasks used in solving the problem among the cores.

# How do we write parallel programs?


- Partition the workload and let CPU cores work in parallel
  - Task parallelism
    - Partition various tasks used in solving the problem among the cores.
  - Data parallelism
    - Partition the data used in solving the problem among the cores.

# How do we write parallel programs?

- Partition the workload and let CPU cores work in parallel
  - Task parallelism
    - Partition various tasks used in solving the problem among the cores.
  - Data parallelism
    - Partition the data used in solving the problem among the cores.
    - Each core carries out similar operations on it's part of the data.

# Professor P

Grade an exam:  
300 exam papers  
15 questions each






# Professor P' s grading assistants




# Division of work – data parallelism

TA#1




100 exams



TA#3

100 exams



TA#2

100 exams

# Division of work – task parallelism

TA#1



Questions 1 - 5

or

Questions 1 - 7



TA#3

Questions 11 - 15

or

Questions 12 - 15



TA#2

Questions 6 - 10

or

Questions 8 - 11

Partitioning strategy:  
- either by number  
- Or by workload

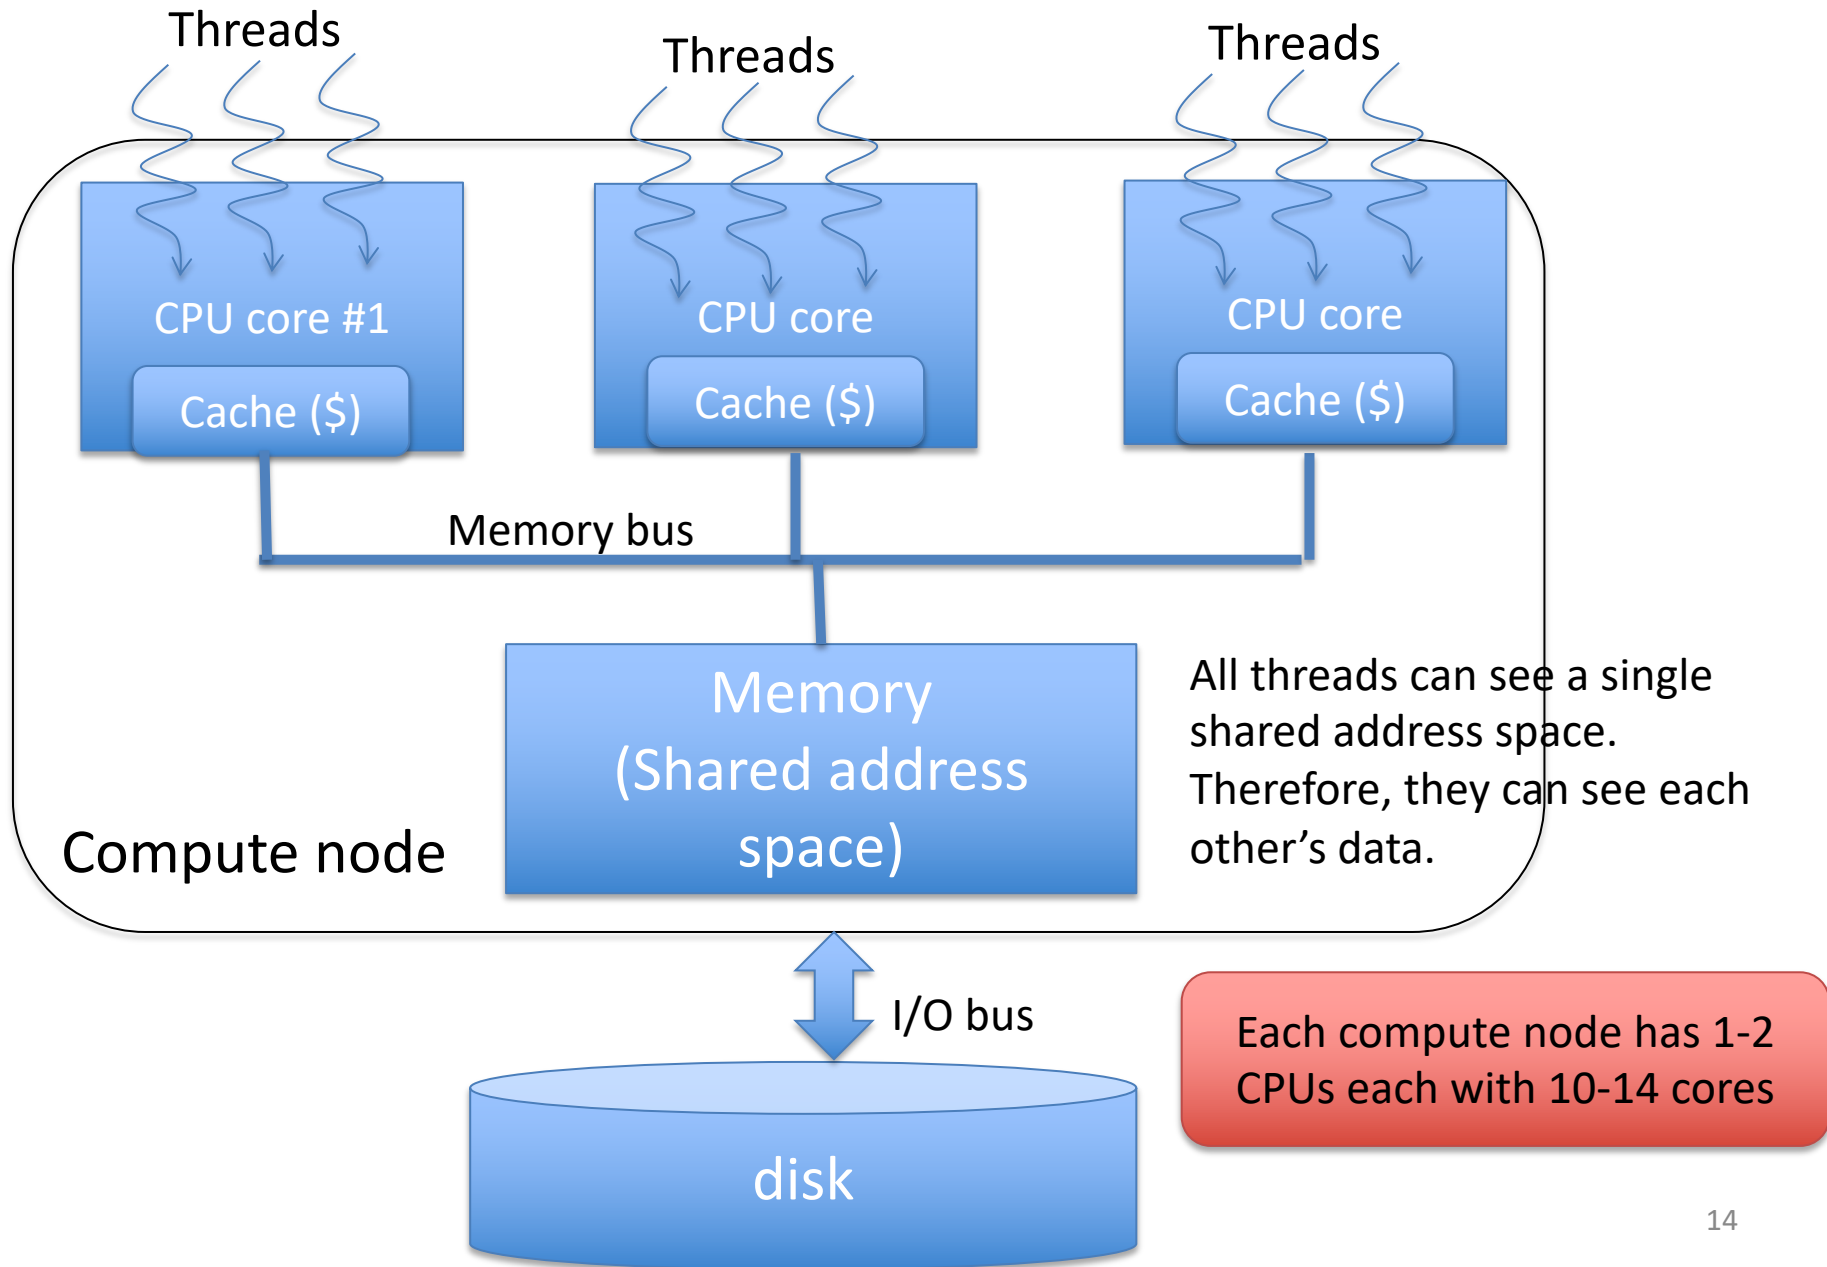
# Coordination

- Cores usually need to coordinate their work.
- **Communication** – one or more cores send their current partial sums to another core.
  - E.g., ML algorithms, PageRank
- **Load balancing** – share the work evenly among the cores so that one is not heavily loaded.
- **Synchronization** – because each core works at its own pace, make sure cores do not get too far ahead of the rest.

# Memory


- Two major classes of parallel programming models:
  - Shared Memory
  - Distributed Memory

# Shared Memory Architecture

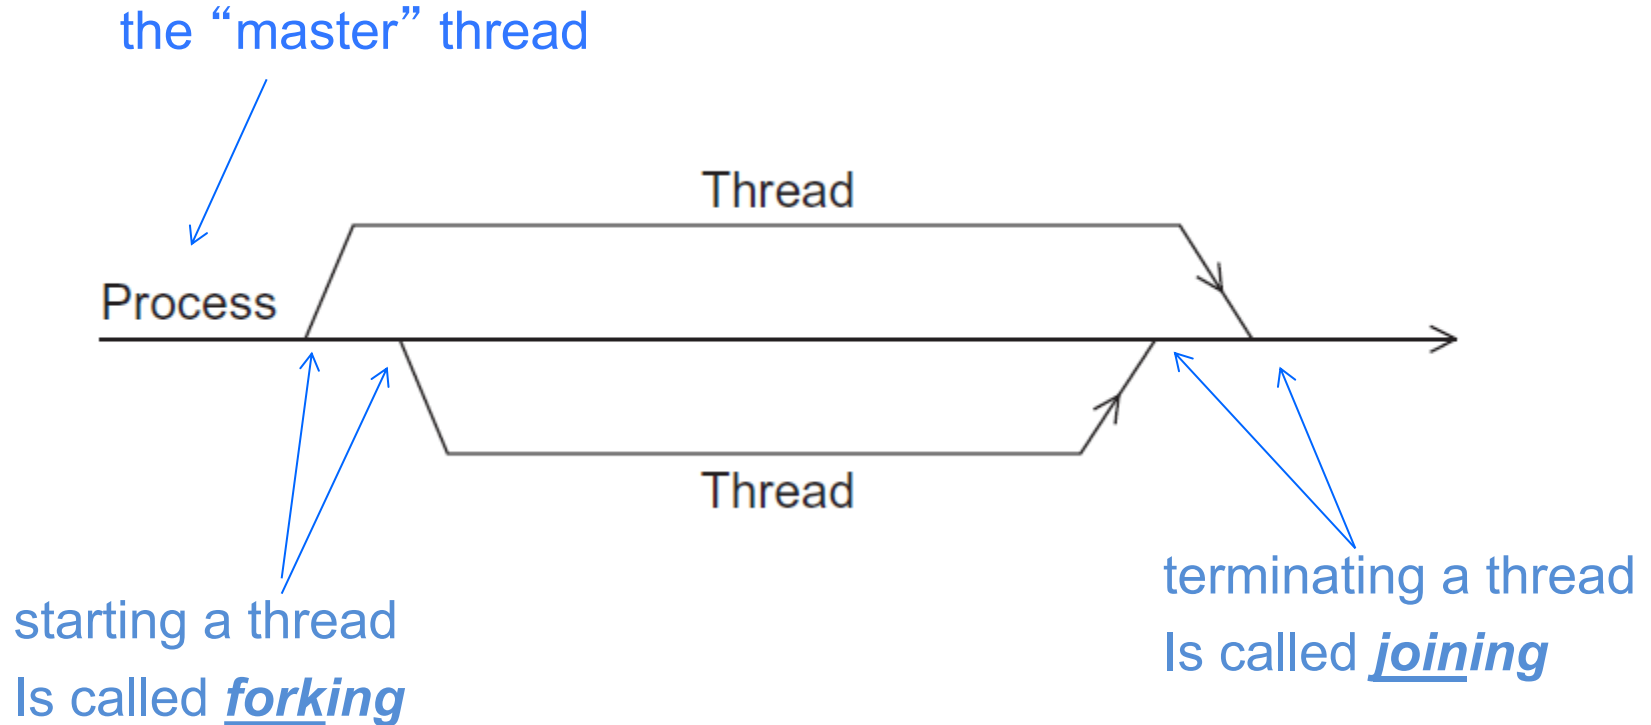


# Multi-Threading (for shared memory architectures)

- Threads are contained within processes
  - One process => multiple threads
- All threads of a process share the same address space (in memory).
- Threads have the capability to run concurrently (executing different instructions and accessing different pieces of data at the same time)
- But if the resource is occupied by another thread, they form a queue and wait.
  - For maximum throughput, it is ideal to map each thread to a unique/distinct core

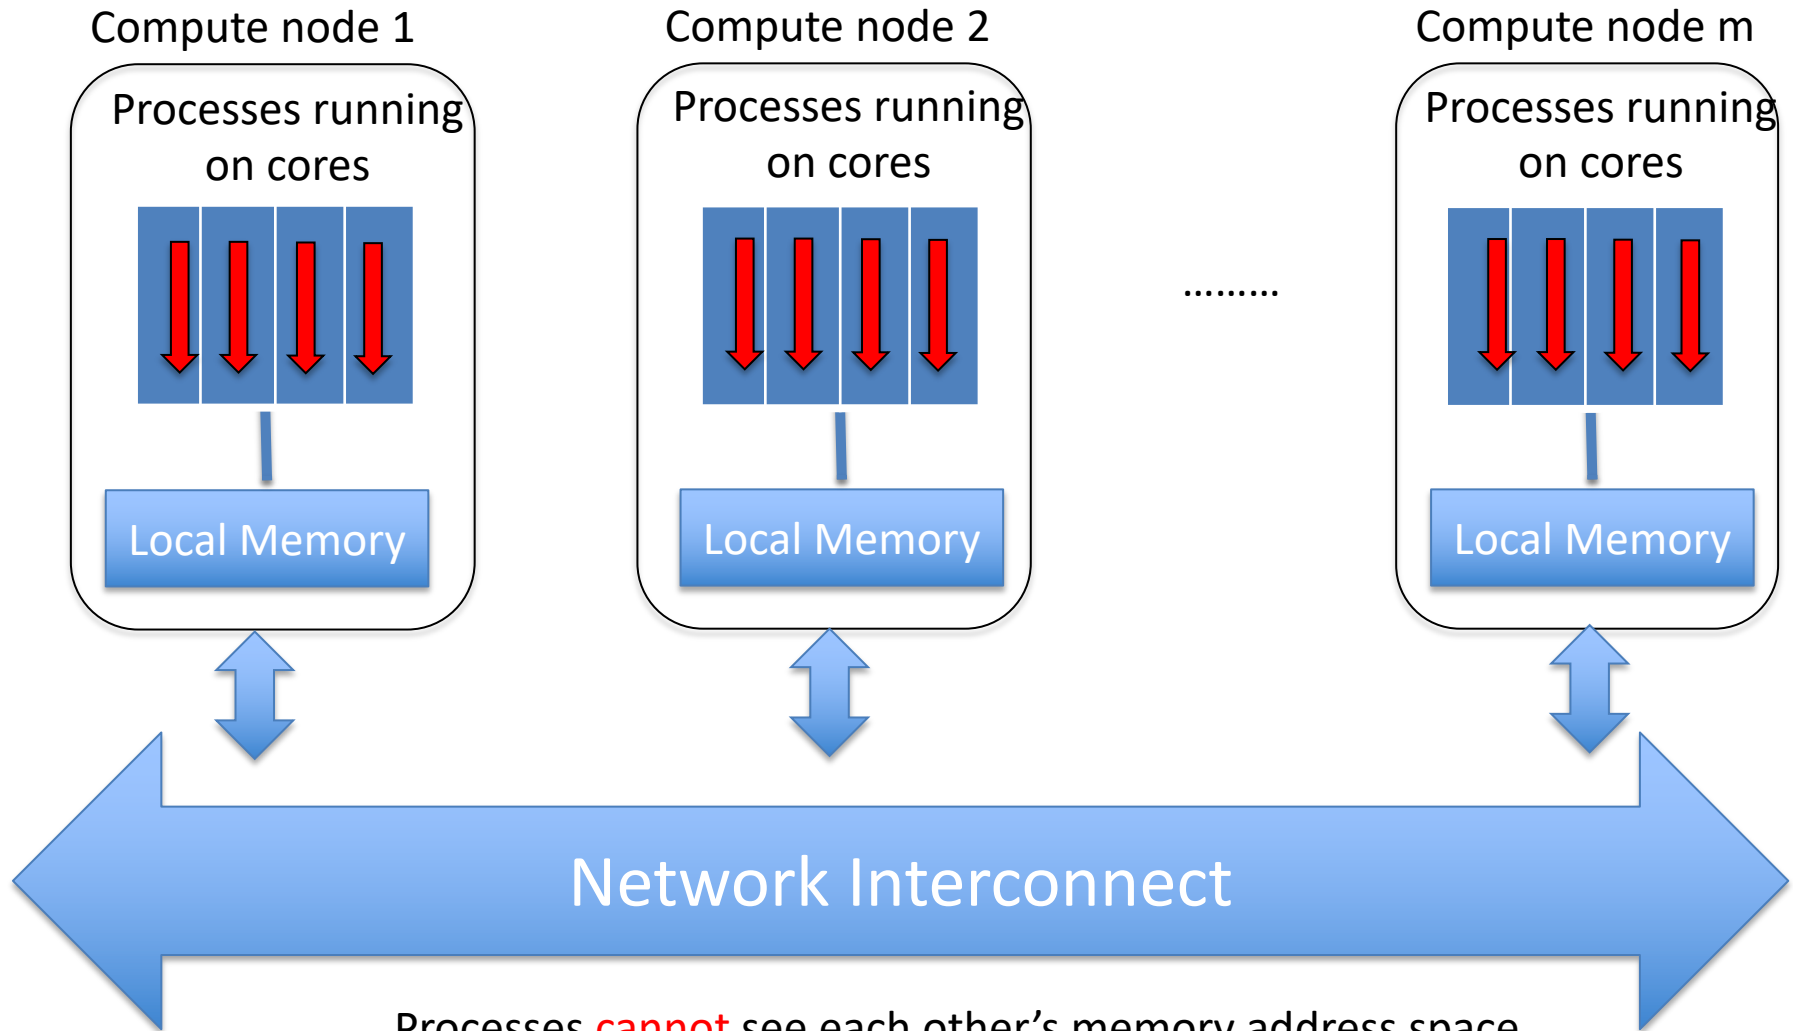


# A process and two threads





# Distributed Memory Architecture



Processes **cannot** see each other's memory address space.  
They have to send inter-process messages (using MPI).

# Distributed Memory System

- **Clusters** (most popular)
  - A collection of commodity systems.
  - Connected by a commodity interconnection network.
- **Nodes** of a cluster are individual computers joined by a communication network.

# How to change your program to a parallel program?

## Foster's methodology

1. **Partitioning**: divide the **computation** to be performed and the **data** operated on by the computation into small tasks.

The focus here should be on identifying tasks that can be executed in parallel.

# Foster' s methodology

**2. Communication:** determine what communication needs to be carried out among the tasks identified in the previous step.



# Foster' s methodology

3. **Aggregation**: combine tasks and communications identified in the first step into larger tasks.

For example, if task A must be executed before task B can be executed, it may make sense to aggregate them into a single composite task.

# Foster' s methodology

4. **Mapping**: assign the composite tasks identified in the previous step to processes/threads.

This should be done so that communication is minimized, and each process/thread gets roughly the same amount of work.