



The Abdus Salam
International Centre
for Theoretical Physics



IAEA
International Atomic Energy Agency

Basic Principles of Parallelism

Ivan Girotto – igirotto@ictp.it

Information & Communication Technology Section (ICTS)

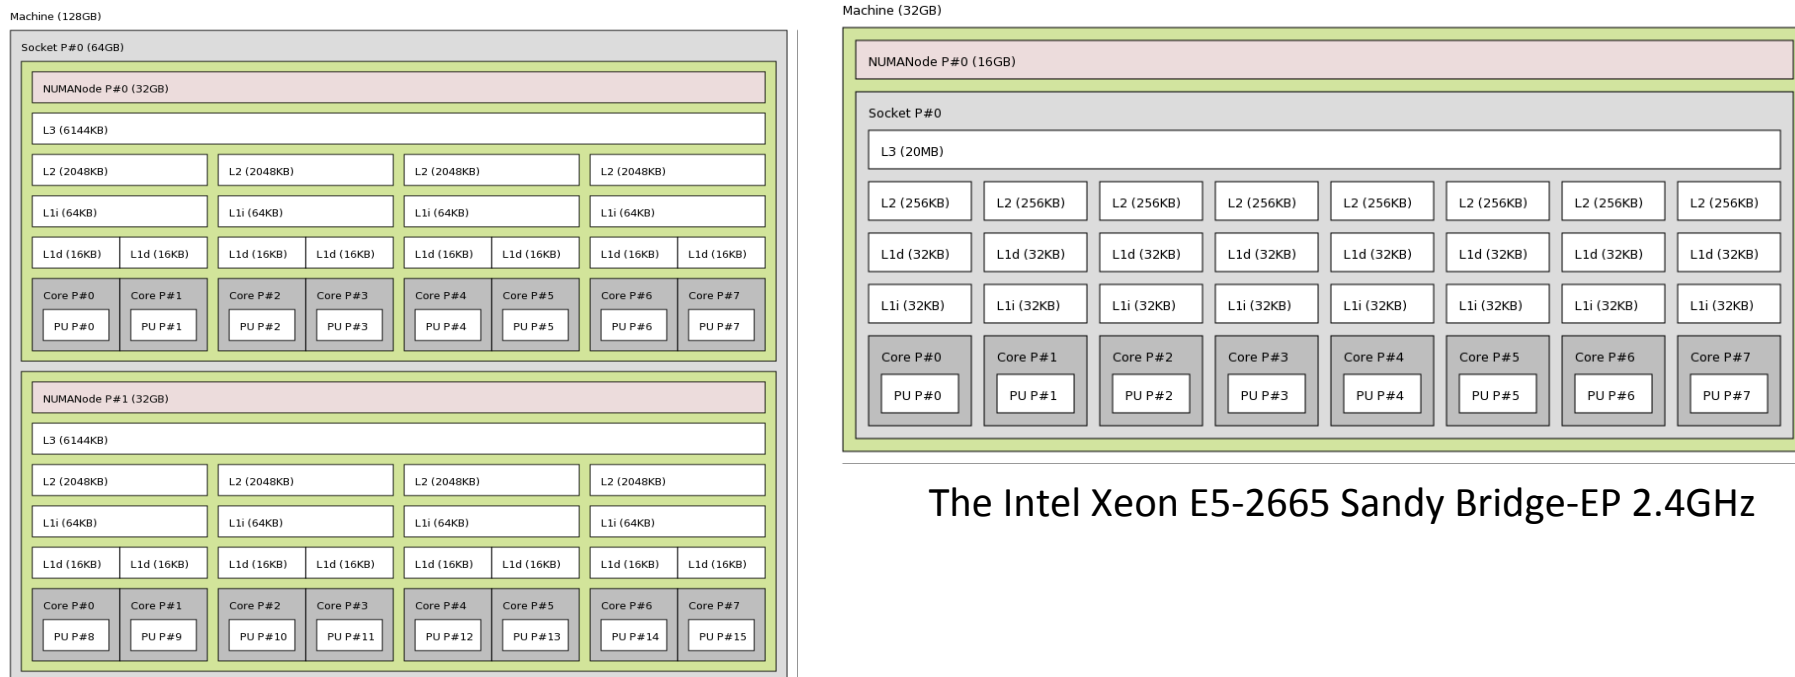
International Centre for Theoretical Physics (ICTP)



Why Parallel Programming

- Solve larger problems
- Run memory demanding codes
- Solve problems with greater speed
- Today?!
 - Processing of massive amount of data
 - Inescapable to exploit modern CPU systems

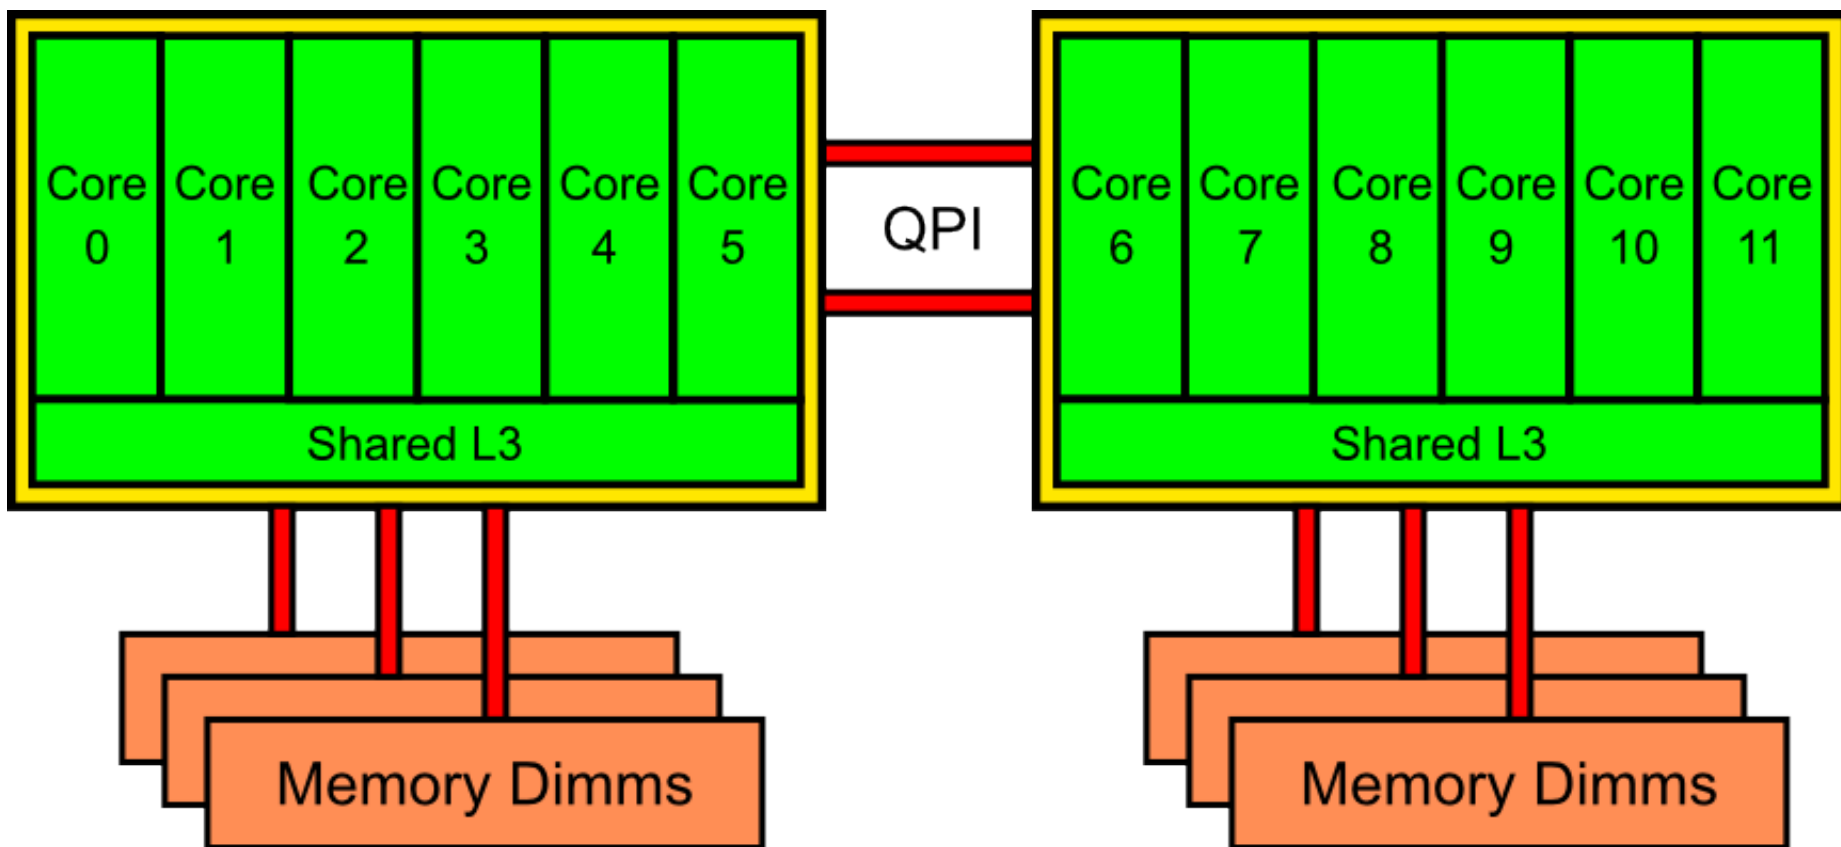
Modern CPUs Models



The Intel Xeon E5-2665 Sandy Bridge-EP 2.4GHz


The AMD Opteron 6380 Abu Dhabi 2.5GHz

Multiple Socket CPUs





Parallel Architectures

- Distributed Memory




- Shared Memory







Serial Programming



A problem is broken into a discrete series of instructions.
Instructions are executed one after another.
Only one instruction may execute at any moment in time.



Parallel Programming



Logical Machine Organization


- Developers can have a different view of the hardware architecture due to the abstraction of the logical organization (how access to memory).
- It is easy to logically partition a Shared Memory system as a Distributed Memory system.
- The opposite is not trivial.

Design of Parallel Algorithm /1

- A serial algorithm is a sequence of basic steps for solving a given problem using a single serial computer
- Similarly, a parallel algorithm is a set of instruction that describe how to solve a given problem using multiple (≥ 1) parallel processors
- The parallelism add the dimension of concurrency. Designer must define a set of steps that can be executed simultaneously!!!

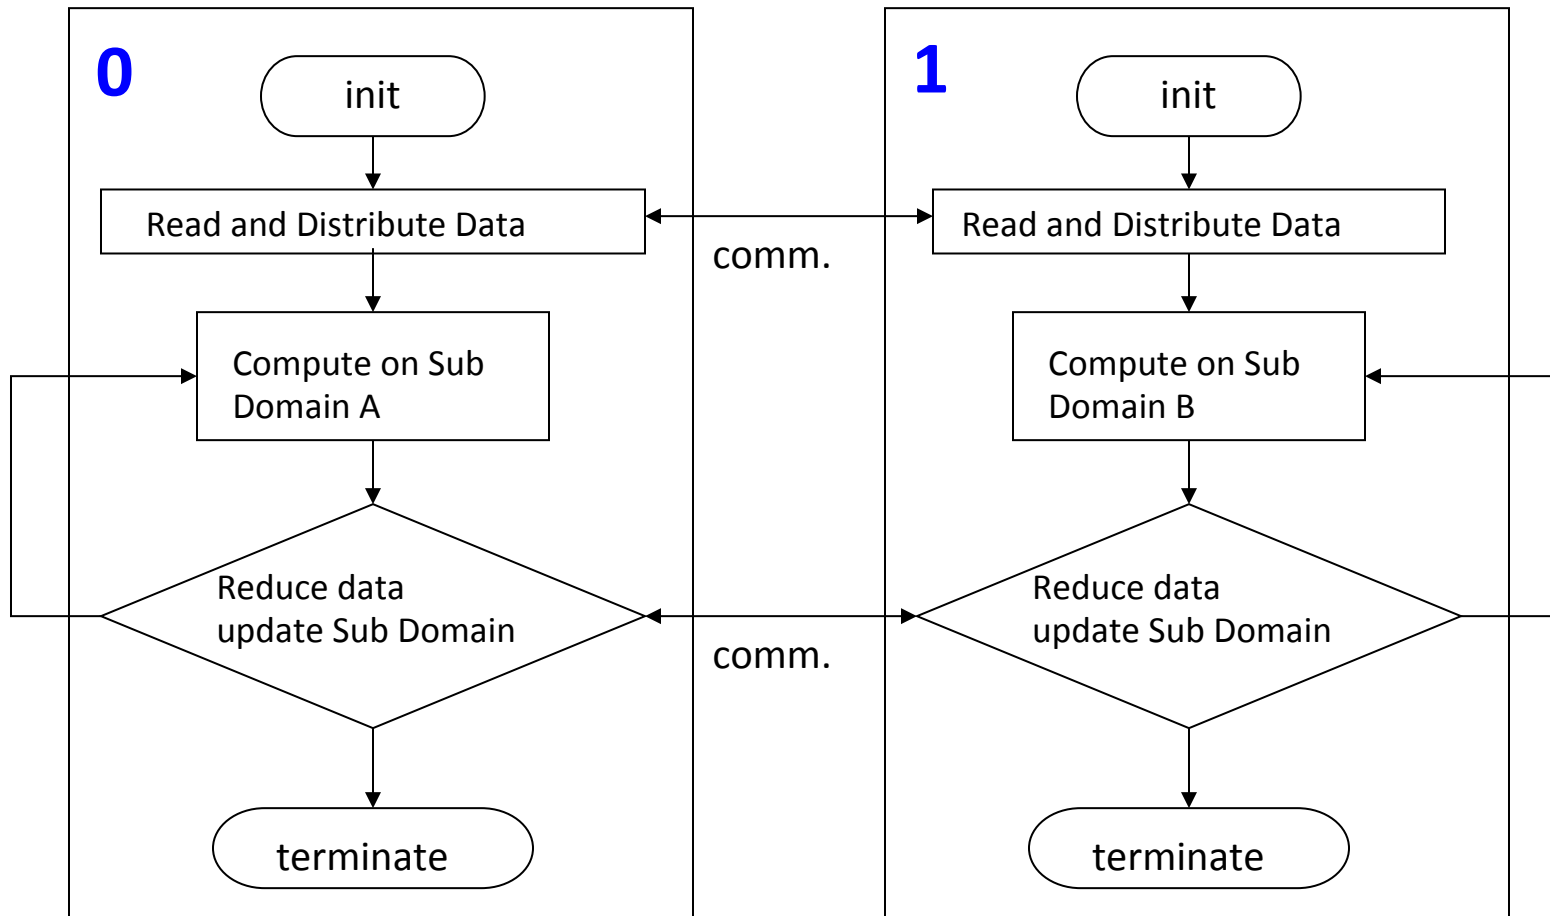
Concurrency

The first step in developing a parallel algorithm is to decompose the problem into tasks that can be executed concurrently



- A problem is broken into discrete parts that can be solved concurrently
- Each part is further broken down to a series of instructions
- Instructions from each part execute simultaneously on different processors
- An overall control / coordination mechanism is used

What is a Parallel Program



Design of Parallel Algorithm /2

- Identify portions of the work that can be performed concurrently
- Mapping the concurrent pieces of work onto multiple processes running in parallel
- Distributing the input, output and intermediate data associated within the program
- Managing accesses to data shared by multiple processors
- Synchronizing the processors at various stages of the parallel program execution


Type of Parallelism

- **Functional (or task) parallelism:**
different people are performing different task at the same time
- **Data Parallelism:**
different people are performing the same task, but on different equivalent and independent objects



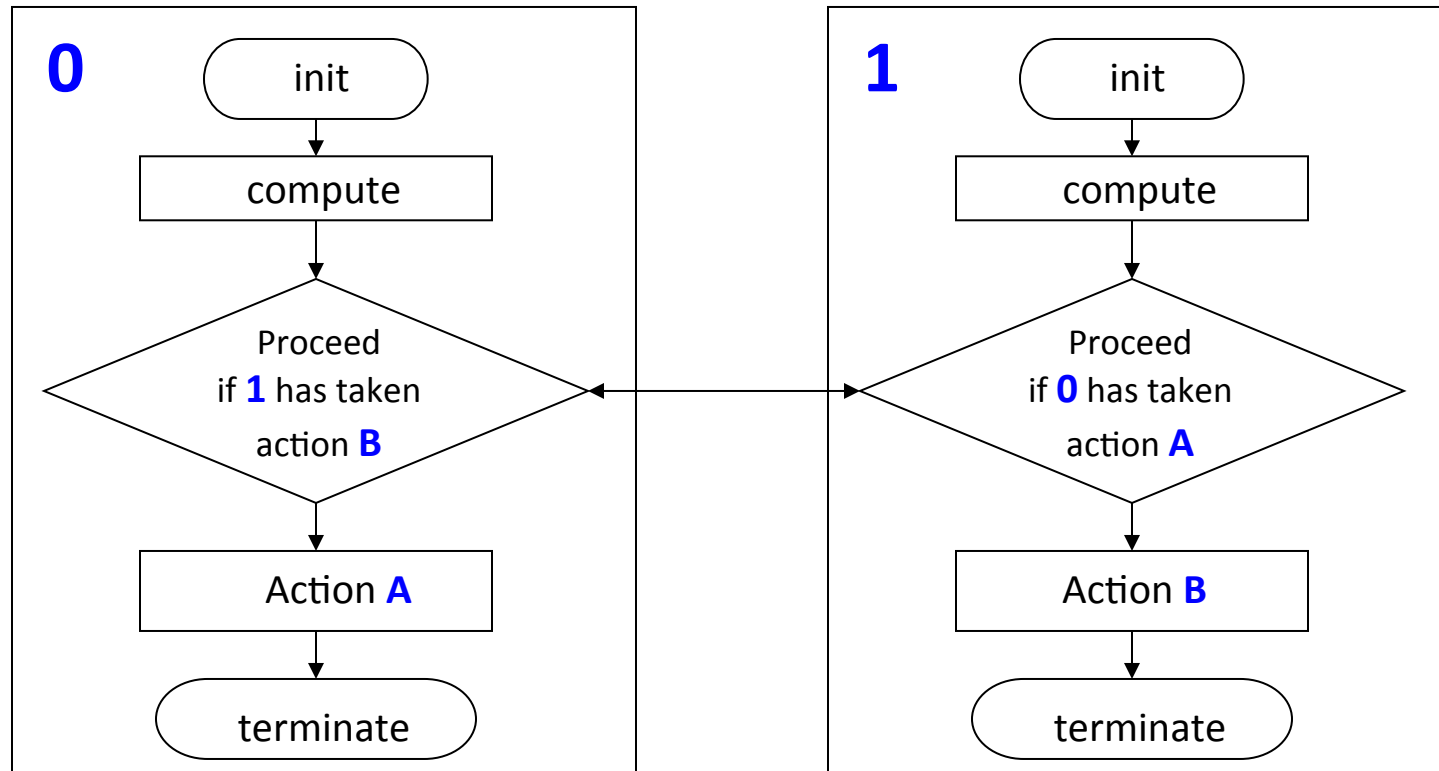
Process Interactions

- The effective speed-up obtained by the parallelization depend by the amount of overhead we introduce making the algorithm parallel
- There are mainly two key sources of overhead:
 1. Time spent in inter-process interactions (communication)
 2. Time some process may spent being idle (synchronization)

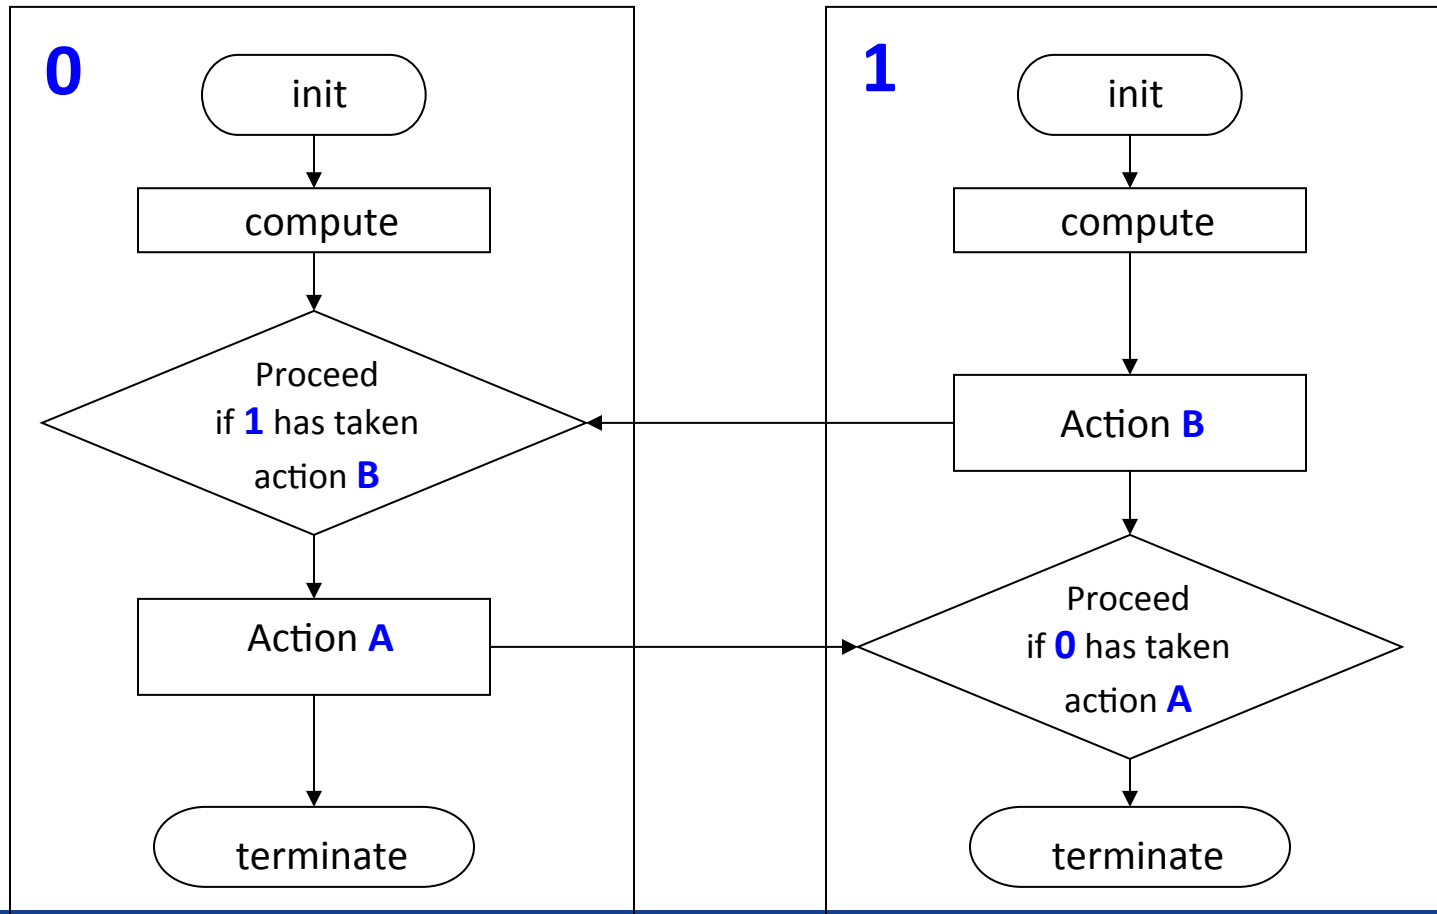


DEADLOCK

- occurs when 2⁺ processes are blocked and each is waiting for the other to progress.



Avoiding DEADLOCK

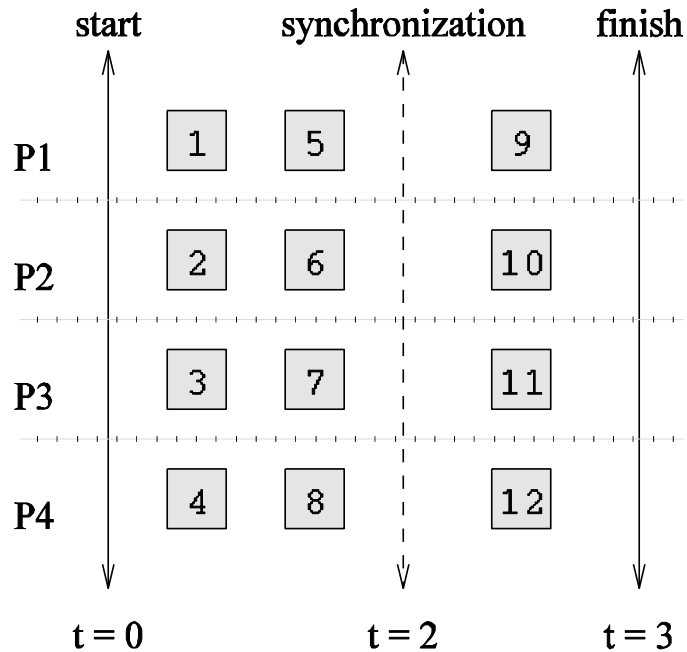





Load Balancing

- Equally divide the work among the available resource: processors, memory, network bandwidth, I/O, ...
- This is usually a simple task for the problem decomposition model
- It is a difficult task for the functional decomposition model

Mapping and Synchronization

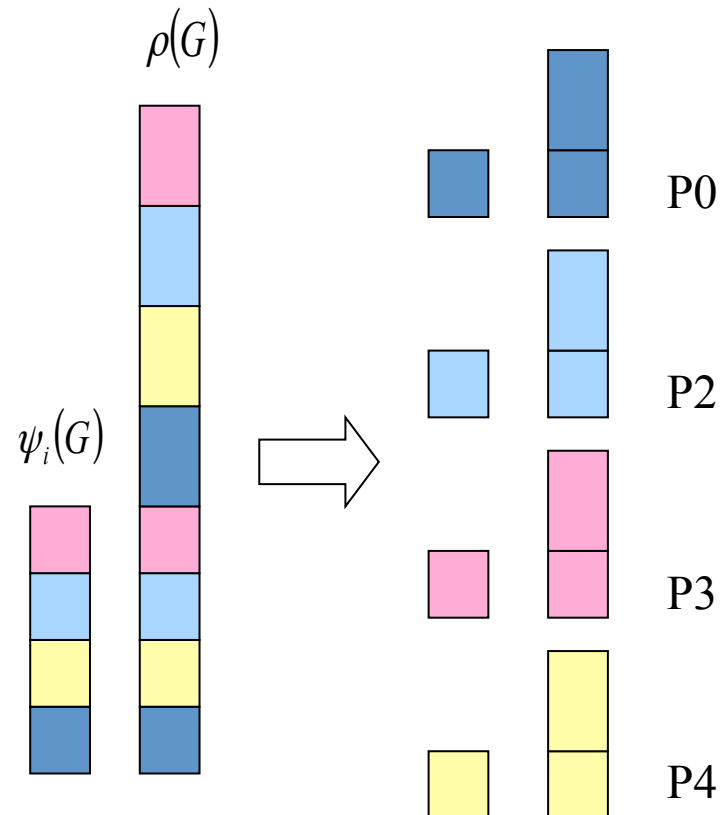


(a)




(b)

Reciprocal Space distribution




Effect of load-unbalancing



Amdahl's law

In a massively parallel context, an upper limit for the scalability of parallel applications is determined by the fraction of the overall execution time spent in non-scalable operations (Amdahl's law).



maximum speedup tends to
 $1 / (1 - P)$

$P =$ parallel fraction

1000000 core

$P = 0.999999$

serial fraction = 0.000001

Minimizing Communication


- When possible reduce the communication events:
 - Group lots of small communications into large one
 - Eliminate synchronizations as much as possible. Each synchronization level off the performance to that of the slowest process



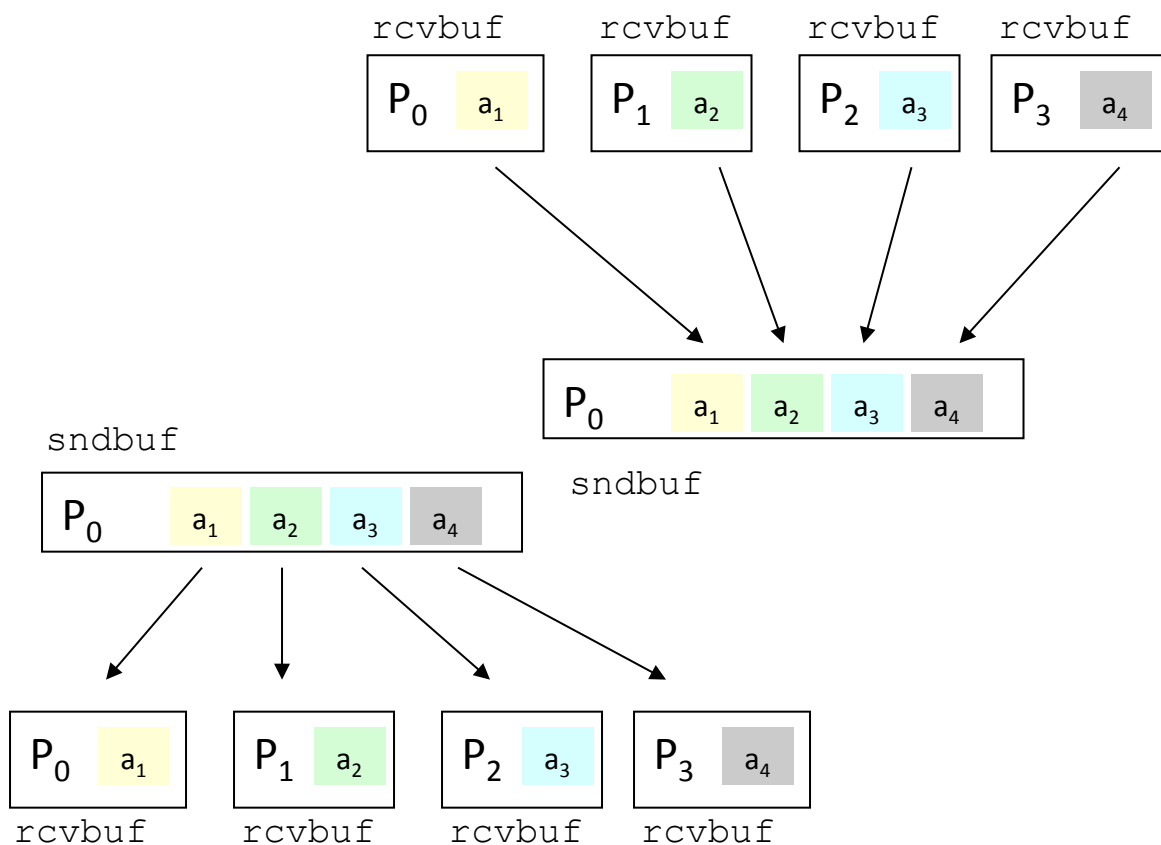
Overlap Communication and Computation

- When possible code your program in such a way that processes continue to do useful work while communicating
- This is usually a non trivial task and is afforded in the very last phase of parallelization.
- A successful implementation can bring a relevant improvement


Broadcast



Scatter/Gather




All to all



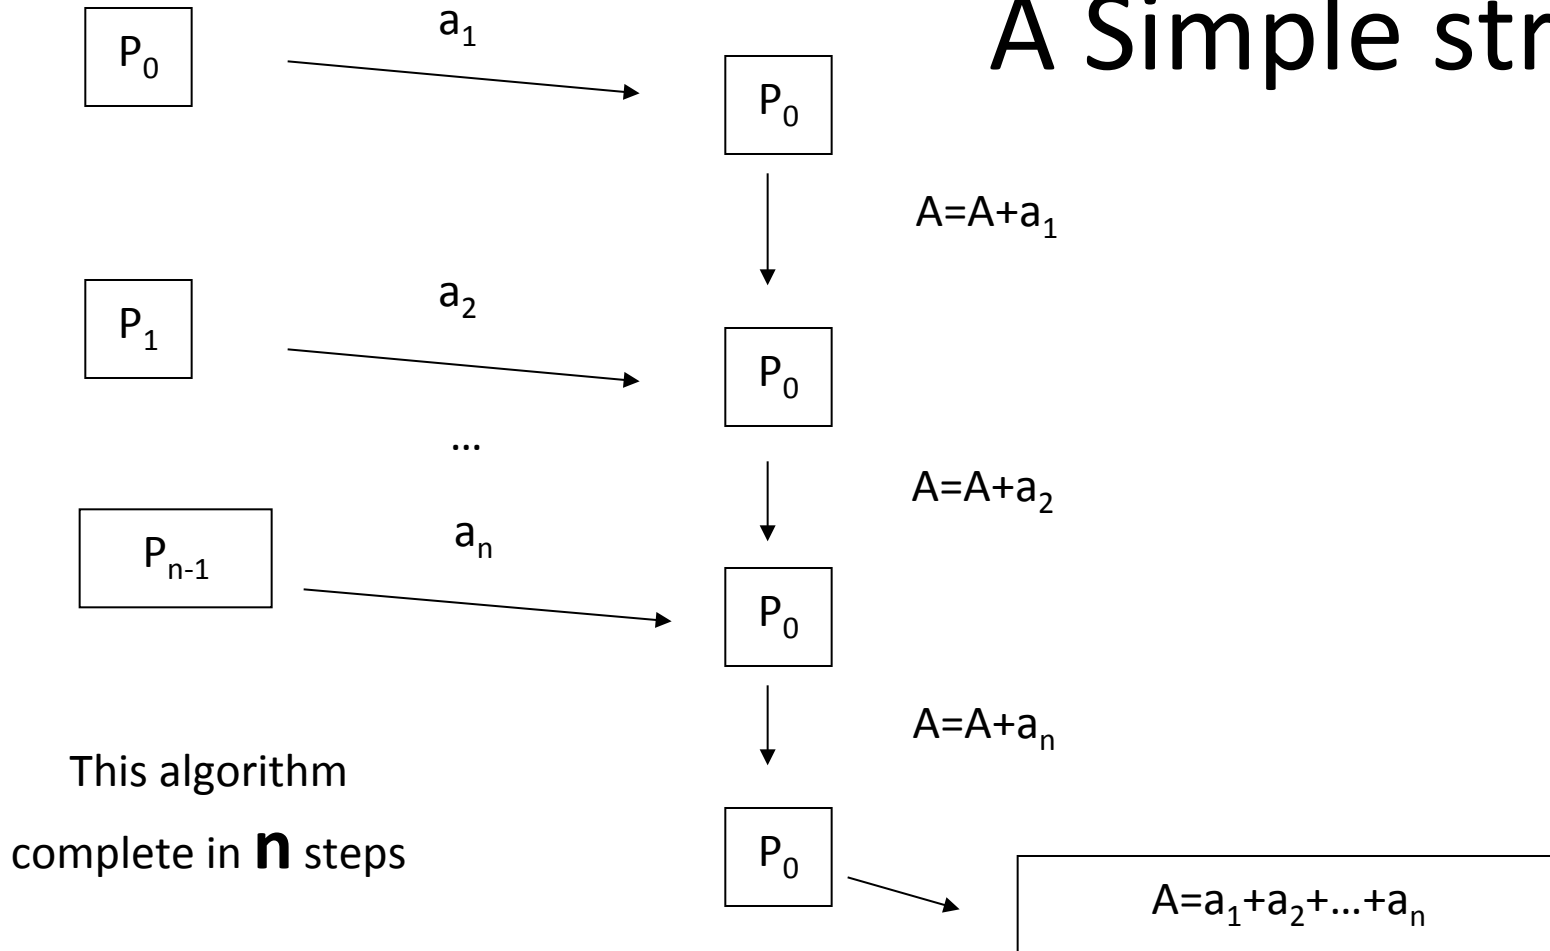
Very useful to implement data transposition

Reduction

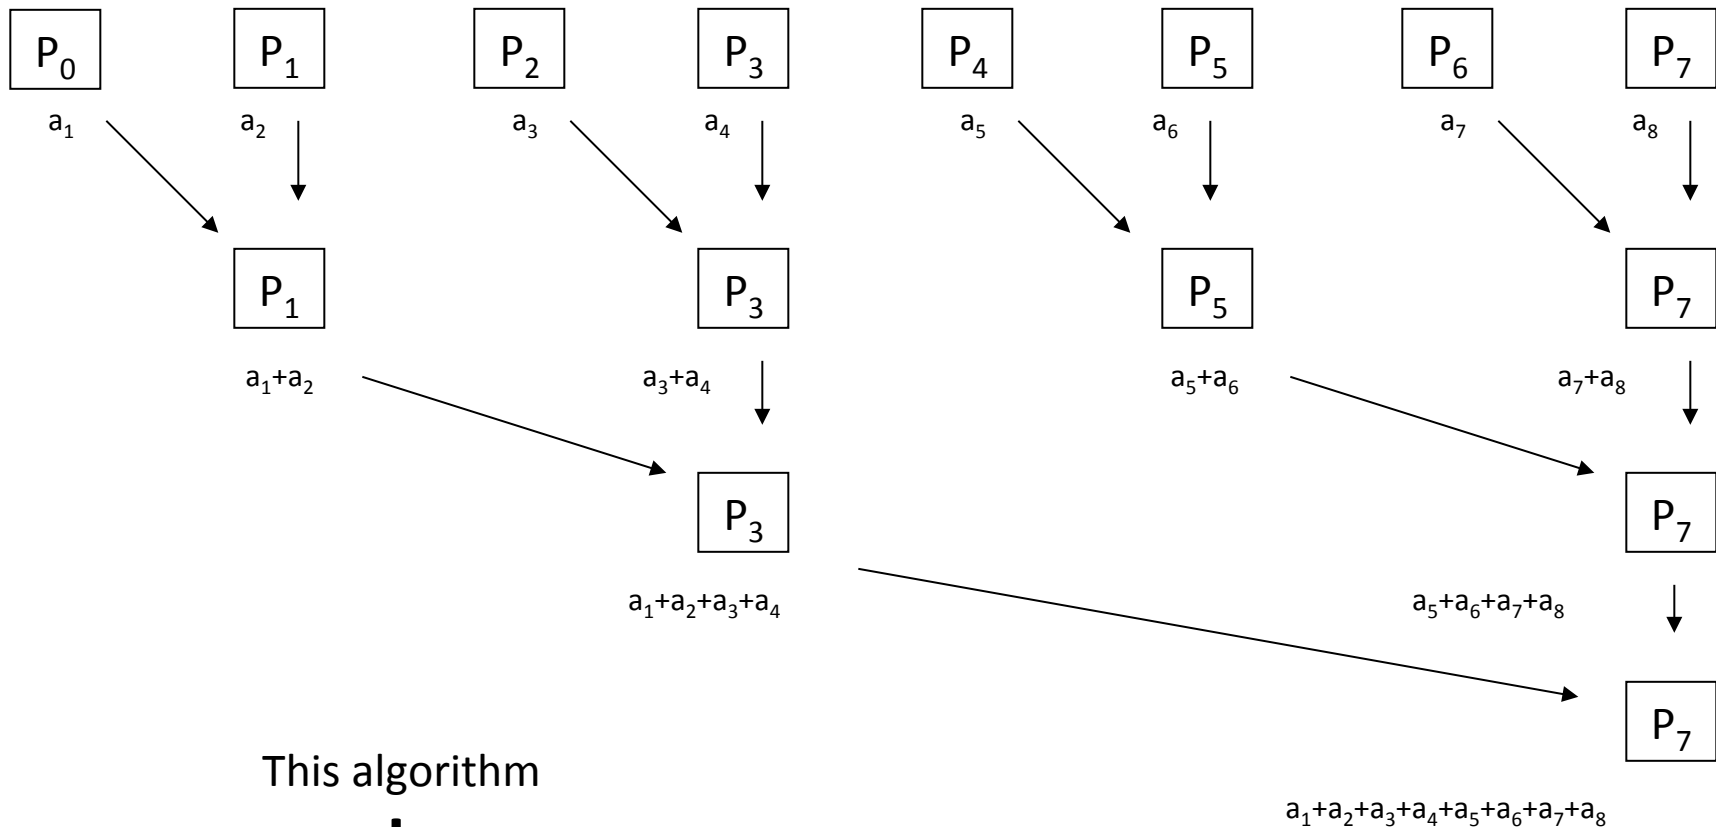
- Reduction: sum up the partial results of different process
(maybe the most common parallel operation required in a parallel program)



A Simple strategy



Binary Tree



This algorithm
complete in **$\log_2 n$** steps

Binary Tree

- Finding the right partner:

- | | Sender | Receiver |
|----------|--------------------------------------|--------------------------------------|
| • Step 1 | $\text{MOD}(\text{myid}, 2) = 0$ | $\text{MOD}(\text{myid}, 2) = 1$ |
| • Step 2 | $\text{MOD}(\text{myid}, 4) = 1$ | $\text{MOD}(\text{myid}, 4) = 3$ |
| • Step 3 | $\text{MOD}(\text{myid}, 8) = 3$ | $\text{MOD}(\text{myid}, 8) = 7$ |
| • ... | | |
| • Step n | $\text{MOD}(\text{myid}, 2^{**n}) =$ | $\text{MOD}(\text{myid}, 2^{**n}) =$ |
| • | $2^{** (n-1)} - 1$ | $2^{**n} - 1$ |
- myid: processor index


How do we evaluate the improvement?

- We want estimate the amount of the introduced overhead $\Rightarrow T_o = n_{pes} T_P - T_S$
- But to quantify the improvement we use the term **Speedup**:


$$S_P = \frac{T_S}{T_P}$$

Speedup

Caspian Test Case 210 x 192 x 18 - 1 Month Simulation



Caspian Test Case 210 x 192 x 18 - 1 Month Simulation




Efficiency

- Only embarrassing parallel algorithm can obtain an ideal Speedup
- The **Efficiency** is a measure of the fraction of time for which a processing element is usefully employed:


$$E_p = \frac{S_p}{p}$$

Efficiency

Caspian Test Case 210 x 192 x 18 - 1 Month Simulation



Caspian Test Case 210 x 192 x 18 - 1 Month Simulation





The Abdus Salam
International Centre
for Theoretical Physics



IAEA
International Atomic Energy Agency

Thanks for your attention!!

