



MPI Programing Models

Roman Baranowski

WestGrid UBC

Goals of Parallel Programming




- Serial Computing - One processor executes a series of instructions to produce result
- Parallel Computing - produce the same result using multiple processors
- Solve larger problems
- Ideally a program running on N processors executes N time faster
- In practice performance depends on the manner in which the problem is divided between the processors
- Each processor performs similar amount of work \Leftrightarrow load balancing

**Why do not we have a parallel
compiler ???**

"holy grail" of parallel computing


Parallel Computer Architectures

- Shared Memory
- Distributed Memory
- Hybrid



Models of Parallel Programming

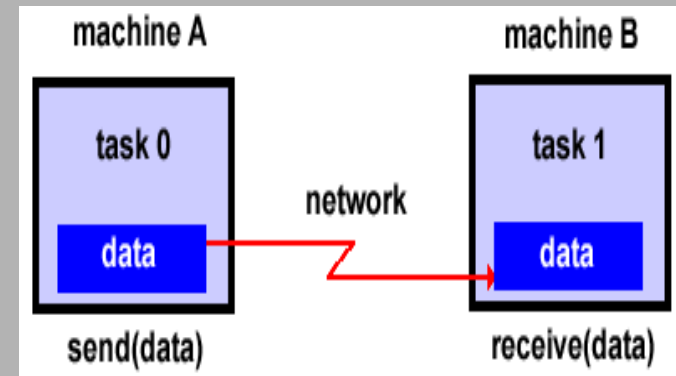
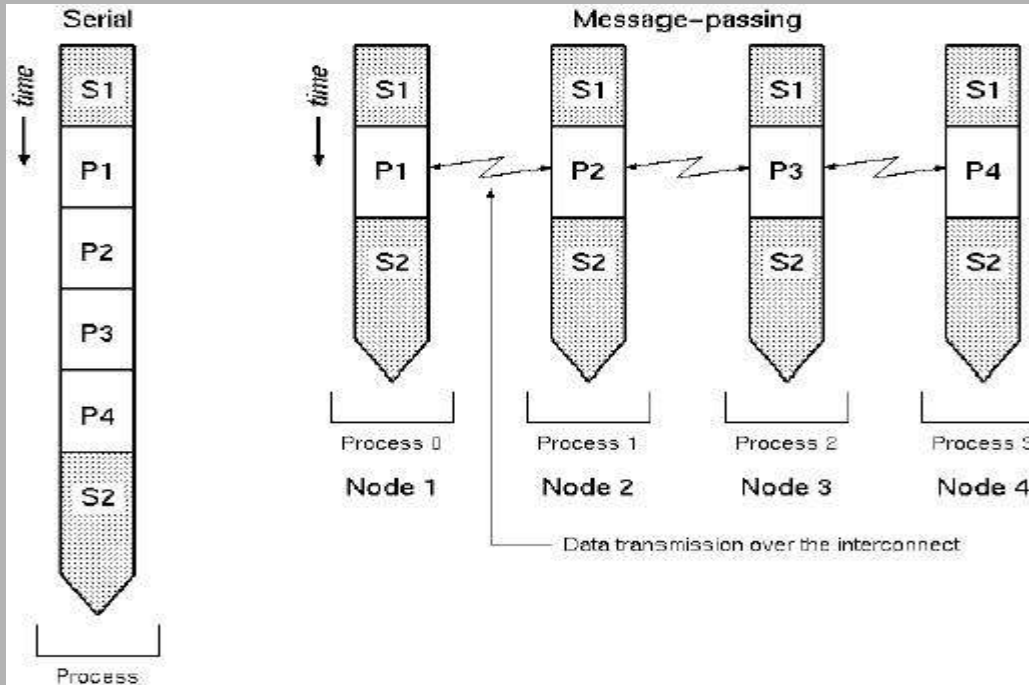
- SMP Based
 - threads
 - Open MP
 - MPI
- Scalability to 100s of processors



Models of Parallel Programming

- Distributed Memory Based (MPP - Massively Parallel Processors)

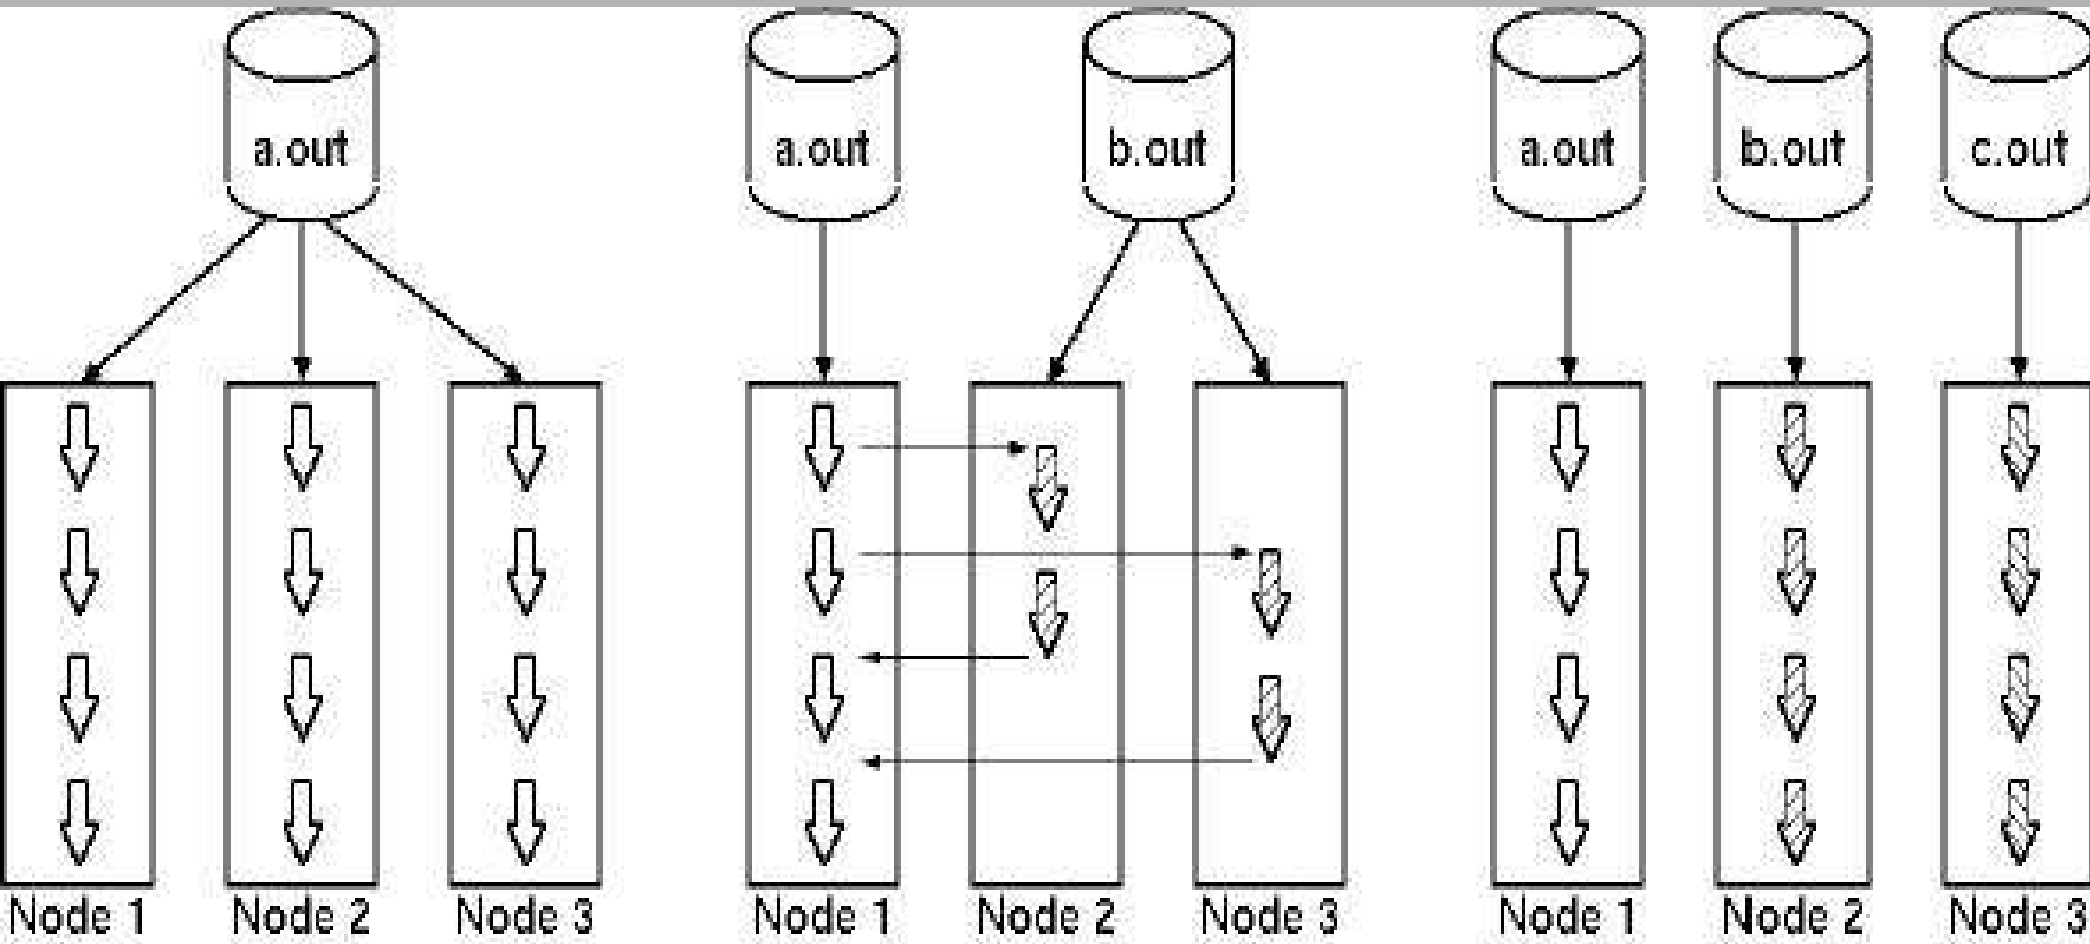
- MPI (Message Passing)



- scalability to 1000 of processors

SMPD

MPMD



(a) SMPD

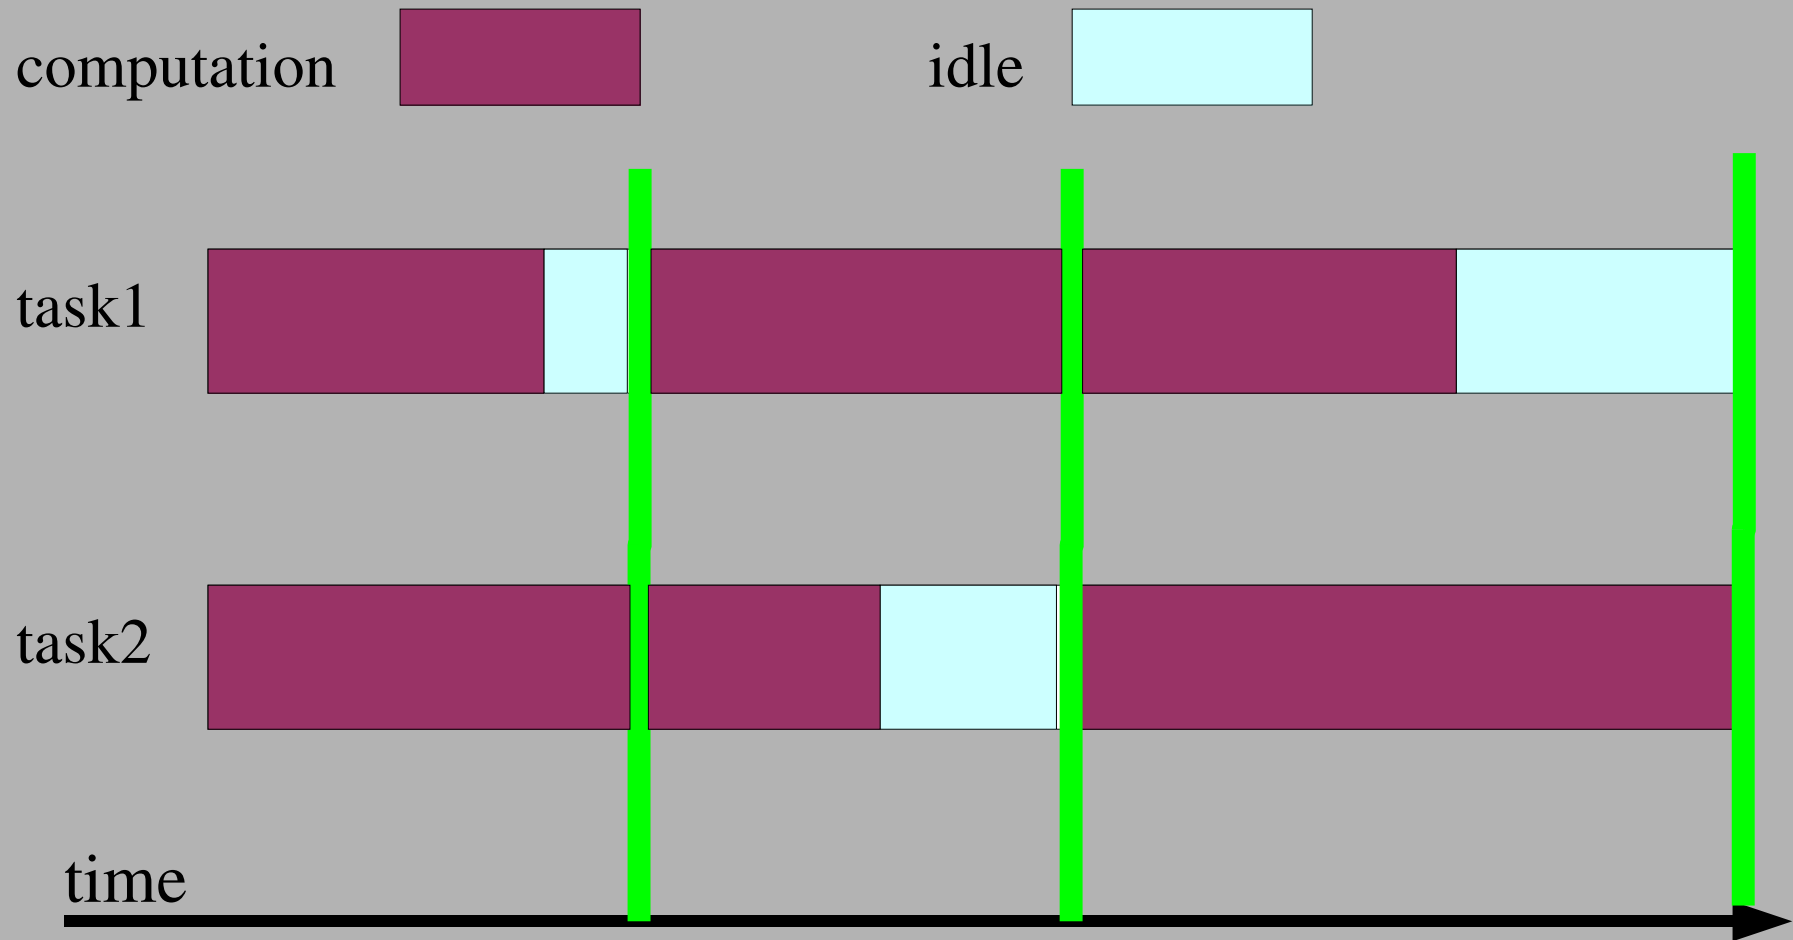
(b) MPMD: Master/Worker

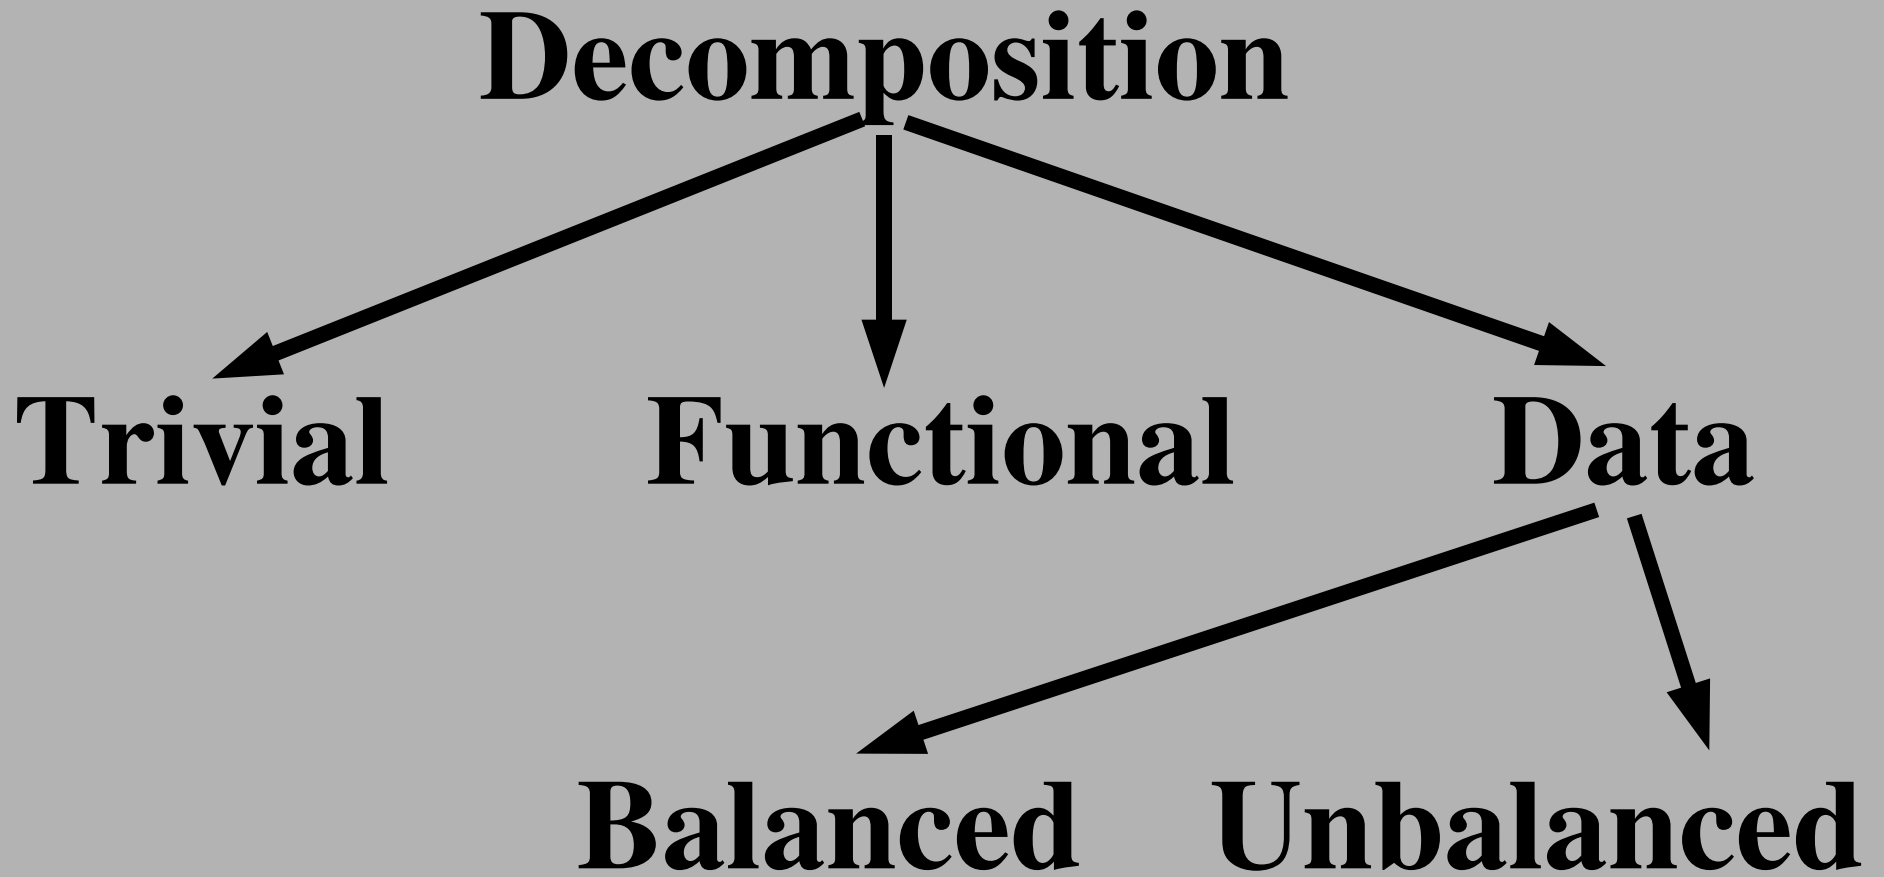
(c) MPMD: Coupled Analysis

Do I need to parallelize my program

- Speed-up $S(n,P) = T(n,1)/T(n,P)$
 - $T(n,P) = 1/P (\sum T_{\text{comp}}^i + \sum T_{\text{comm}}^i + \sum T_{\text{idle}}^i)$
- Efficiency $E(n,P) = S(n,P)/P = T(n,1)/(P T(n,P))$
- For fixed size problems \longrightarrow Amdahl's Law
 - For example, when 10% of the code must be executed sequentially, the maximum speed-up is limited by 10, independent of the number of processors available


Load Balancing & Synchronization





Embarrassingly Parallel or Trivial (SPMD)


- Monte Carlo
- Molecular Dynamics




Trivial (SPMD) II

- PI code:

```
do i = 1, n, 1
  generate random x&y in [0,1]
  if (x^2+y^2) .le. 1 then cc = cc + 1
end do
PI = 4.0*cc/n
```



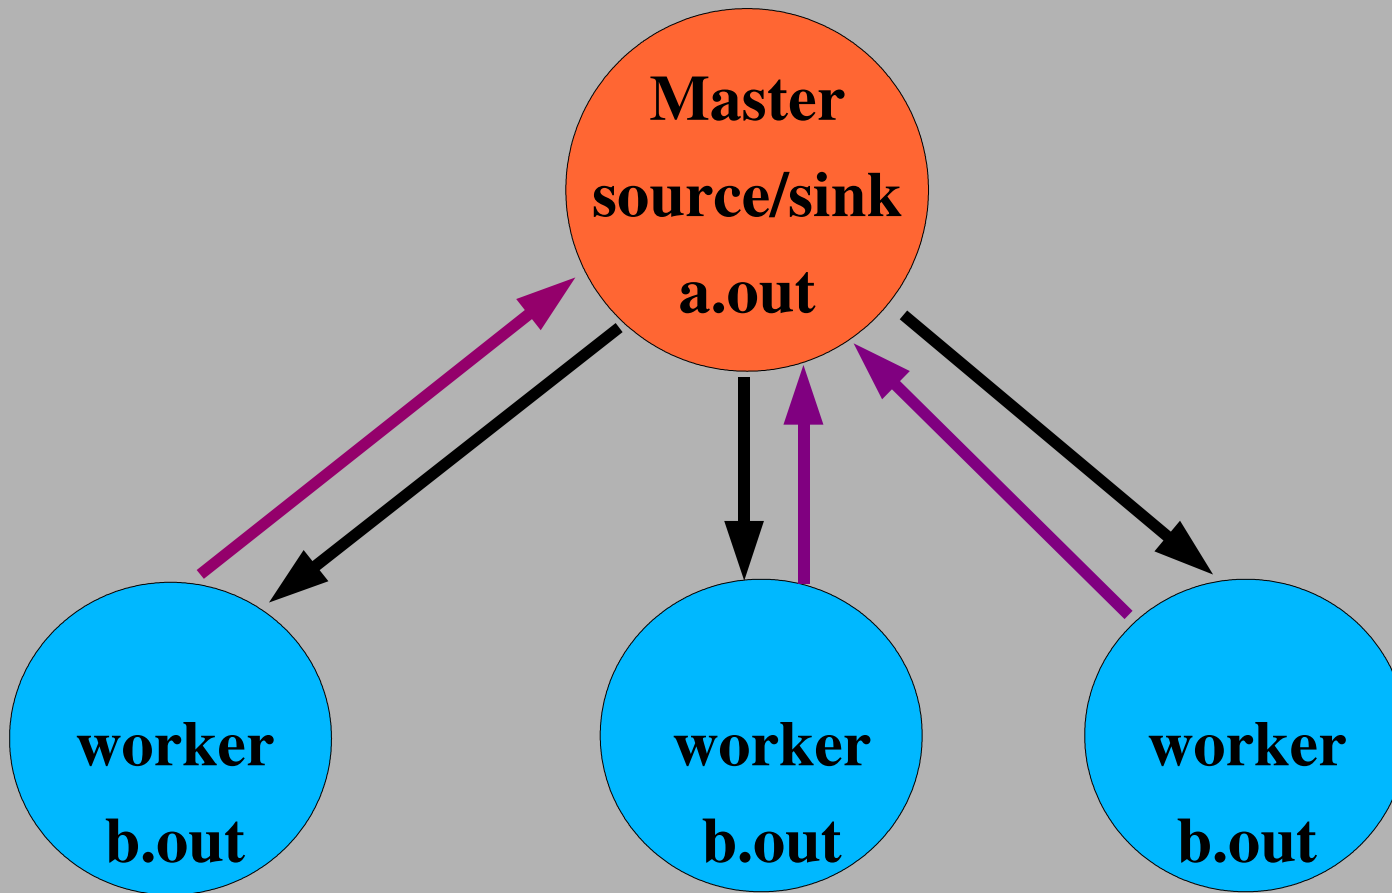
Functional Decomposition pipeline (MPMD)



- startup
- shutdown
- load balancing
- number of tasks

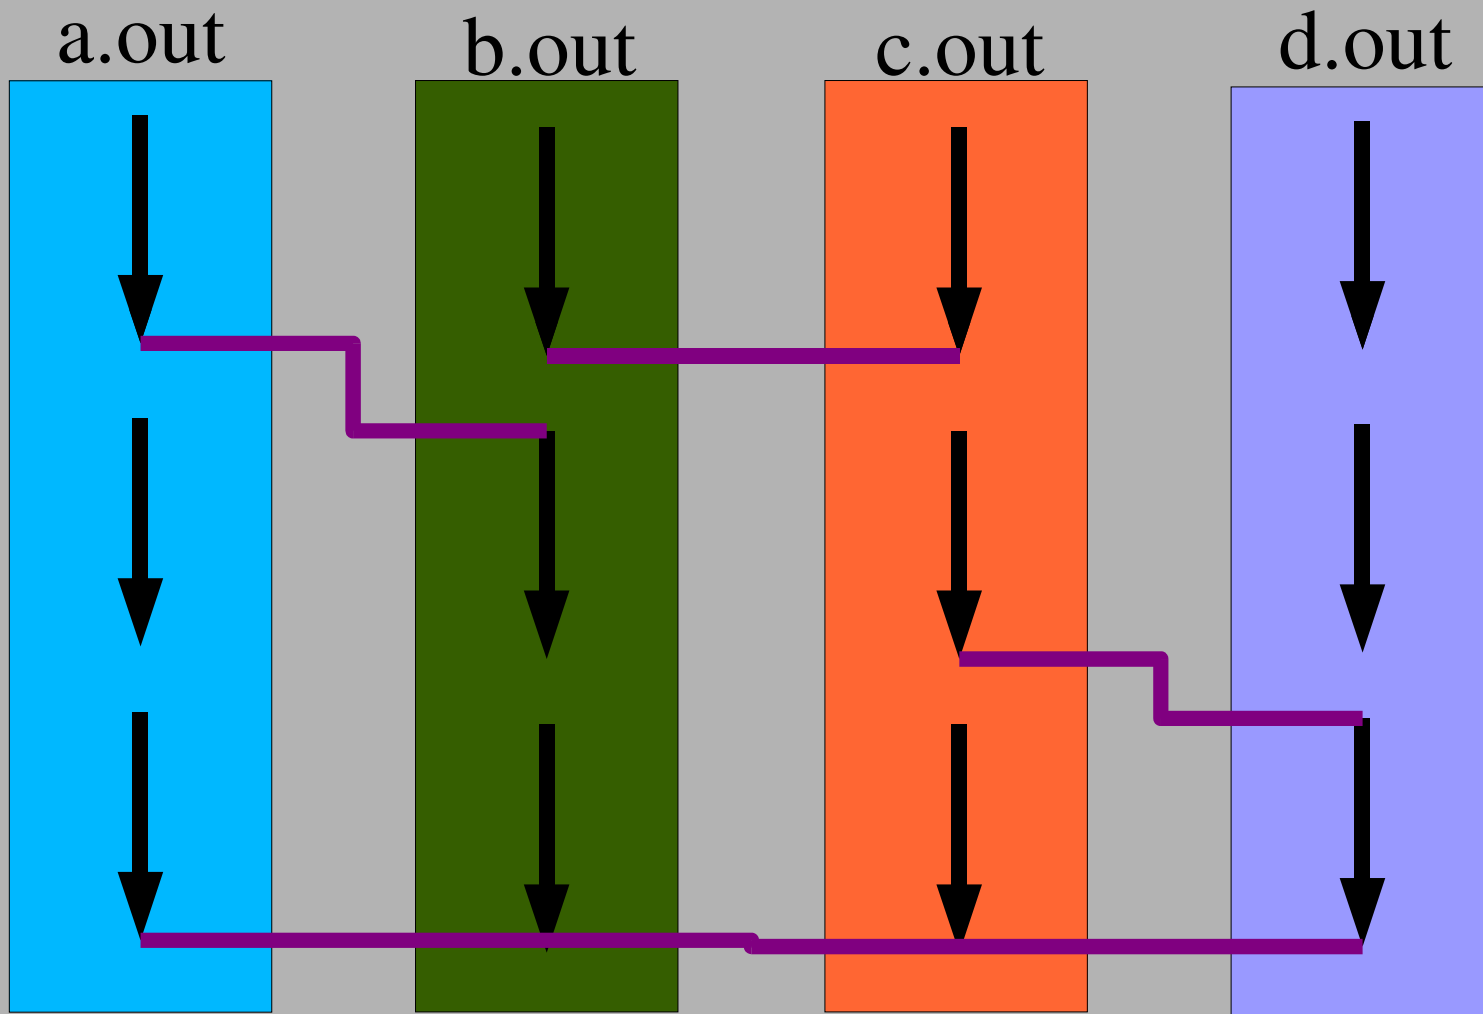
Functional Decomposition

Master - Worker



Functional Decomposition


Coupled Analysis



Data (domain) Decomposition

-

p1-p2-p3-p4



Data (domain) Decomposition

- **SPMD**
 - **load balancing**
 - **synchronization**
 - **communication overhead**
- **MPMD**
 - **synchronization**
 - **load balancing**

What do I do if I have a serial program

- **Do not submit with #PBS -l nodes=100**
- **Is it a software package or your own code ?**
 - software (look for parallel equivalent)
- **Tune up serial code**
- **Consider the scope of parallelization**
 - coarse-grain
 - fine-grain
- **Parallelize**

Tools

- **Portable, Extensible Toolkit for Scientific Computation**
 - <http://www-unix.mcs.anl.gov/petsc/petsc-as>
- **MPI documentation**
 - <http://www-unix.mcs.anl.gov/mpi/>
- **Contact WestGrid**
 - <http://www.westgrid.ca> support@westgrid.ca

