21st Summer School of PARALLEL COMPUTING

July 2 - 13, 2012 (Italian)
September 3 - 14, 2012 (English)

MPI Exercises

SuperComputing, Applications and Innovation Department

CINECA

Exercise 1

Using MPI, print:

1. Hello world
2. Hello world, I am proc X of total Y (from 1 to total 4 tasks)
3. Submit it as a batch job

Exercise 2:
a) Data exchange

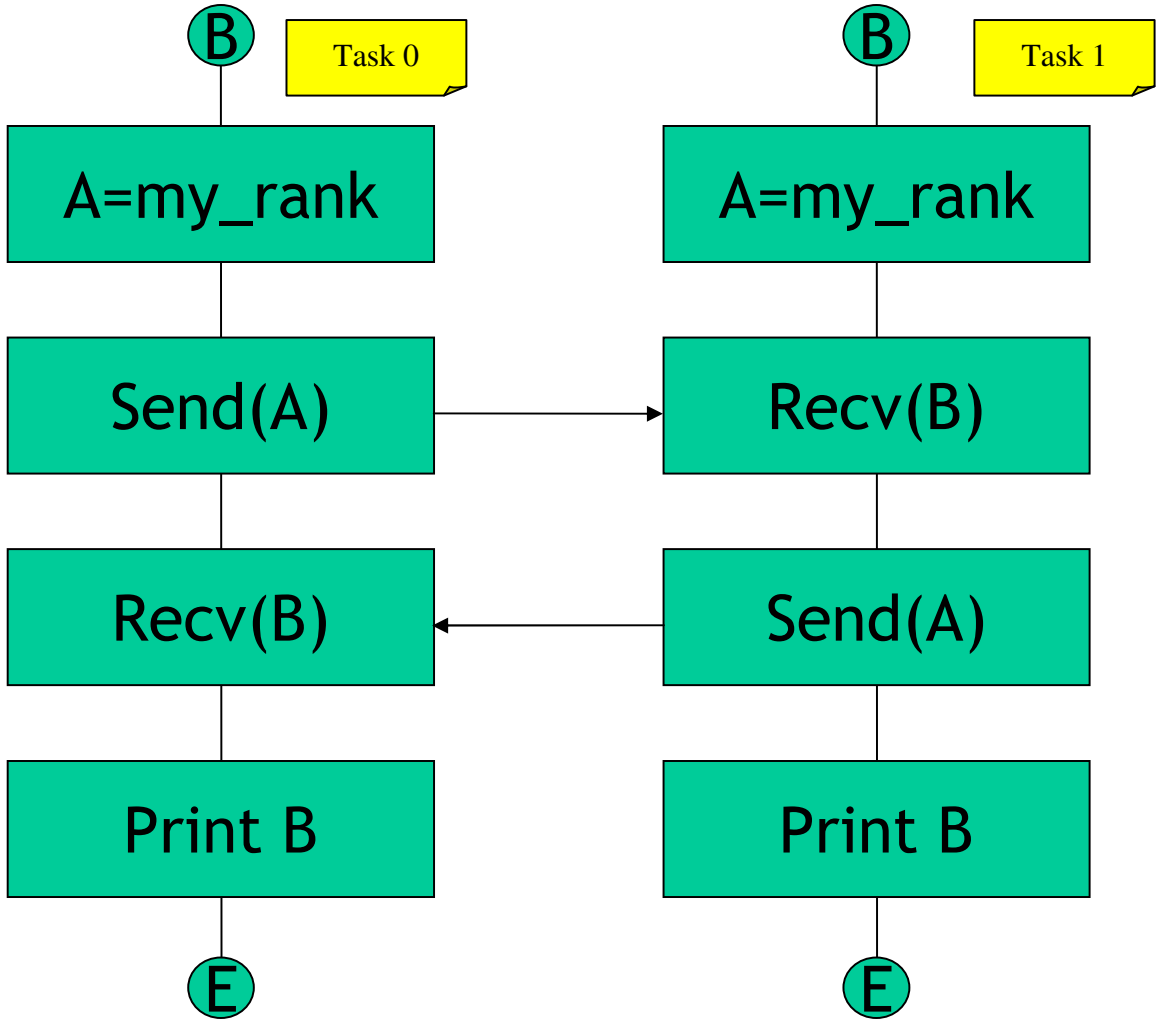

b) try an array[N] to see a deadlock!

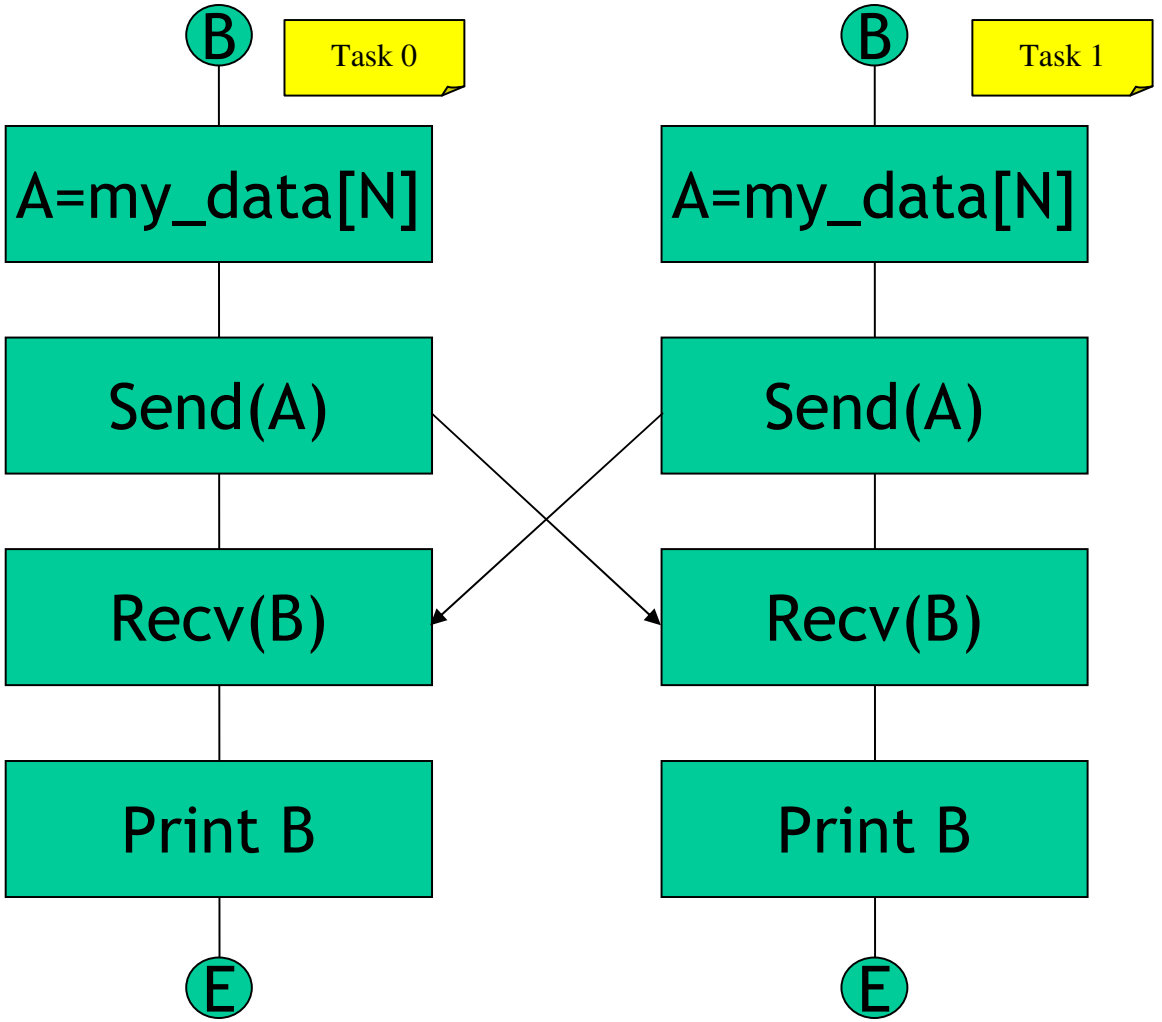

c) without ANY if and without deadlock (using non-blocking send)

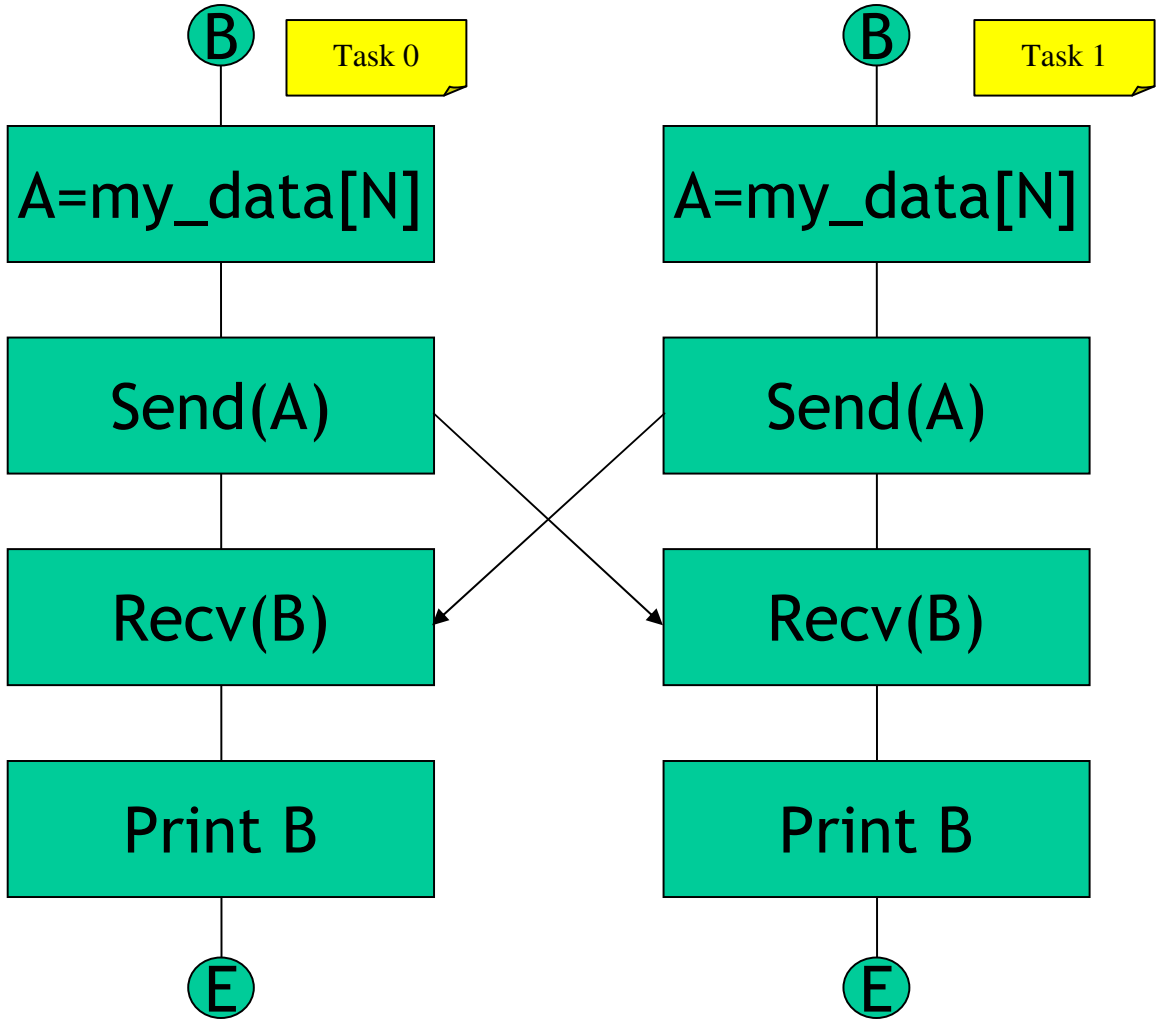

Exercise 3:
as Exercise 2c, but scalable, circular communications

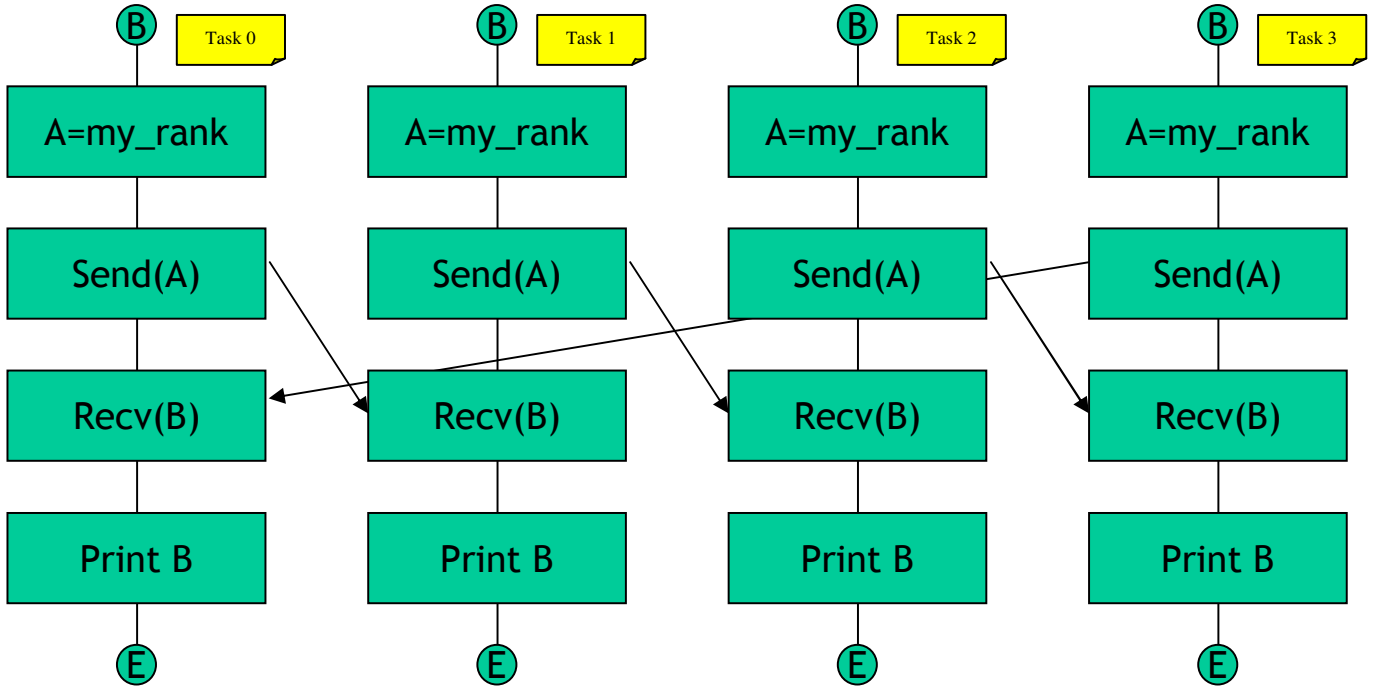

Exercise 4 (optional): a) Sum with circular communications

Each process must print the sum of all ranks. Only communication with the left neighbour process is allowed, as showed in the figure.

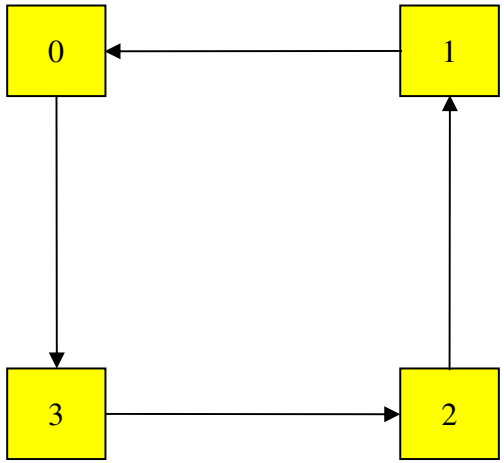

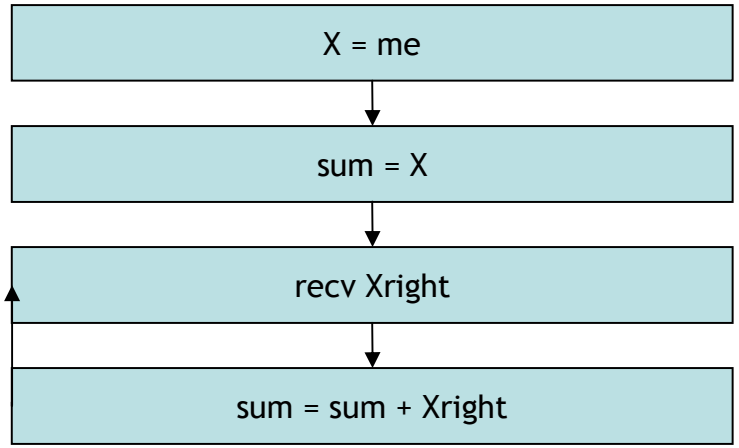

Who/when does the SEND?

What does it send?

How many times?

sum = 6 (on 4 processors)

b) Sum with Binary Tree (optional)

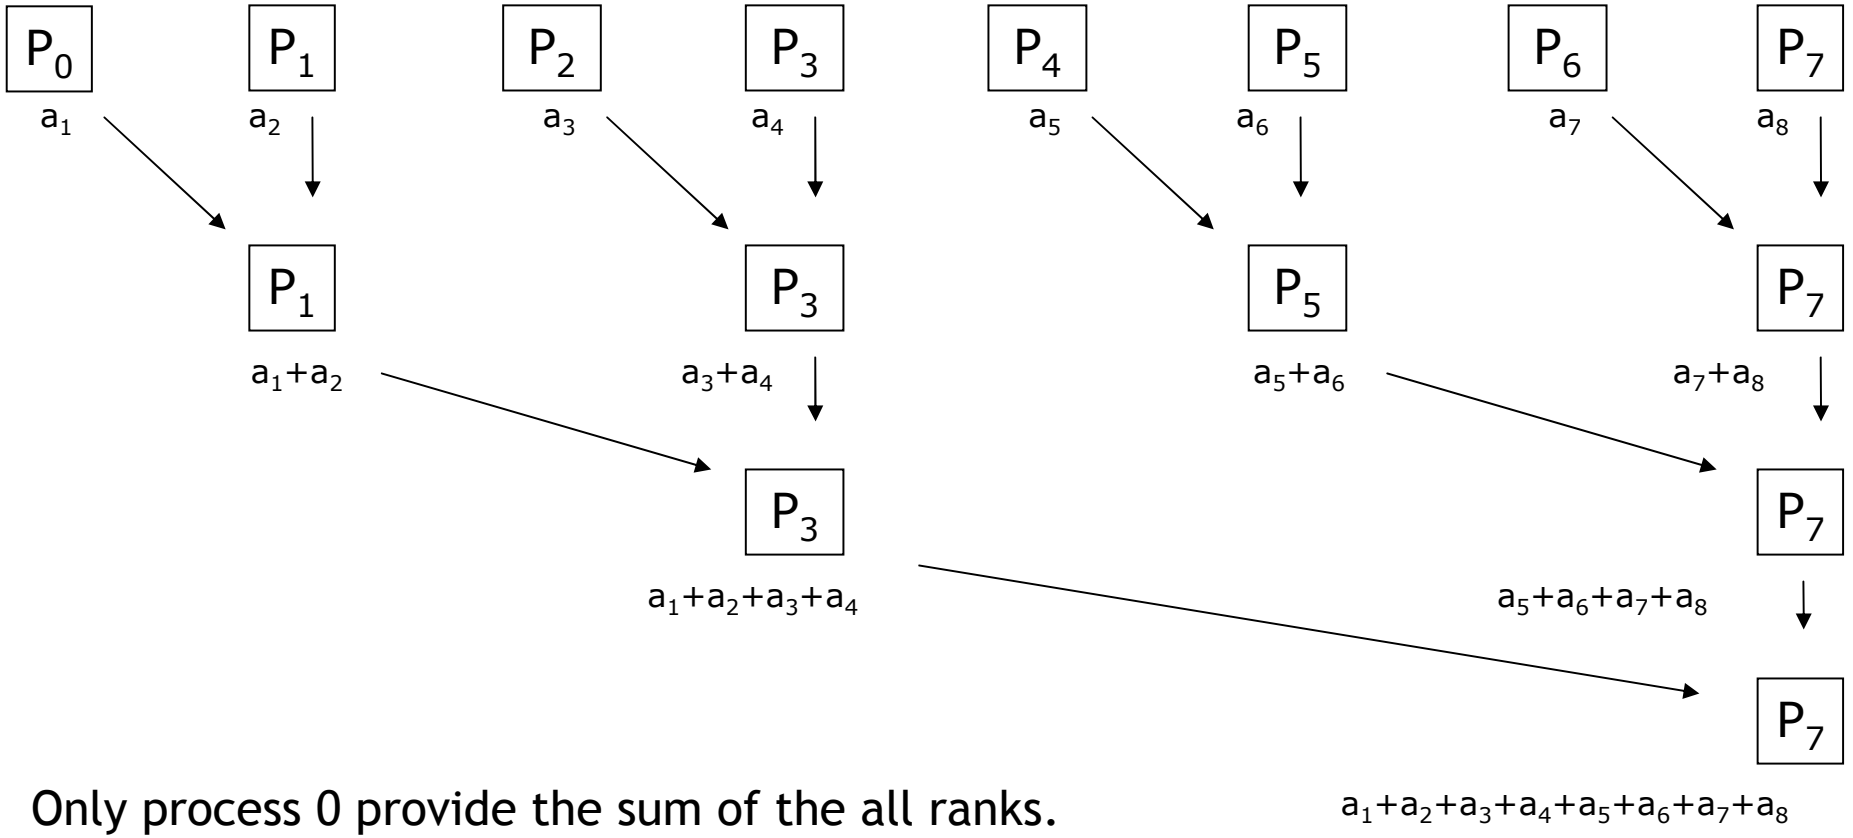

Only process 0 provide the sum of the all ranks.
In this case communication works according to the binary tree showed in figure. This algorithm complete in $\log_2 n$ steps.

Exercise 5 (optional): Data distribution of identity

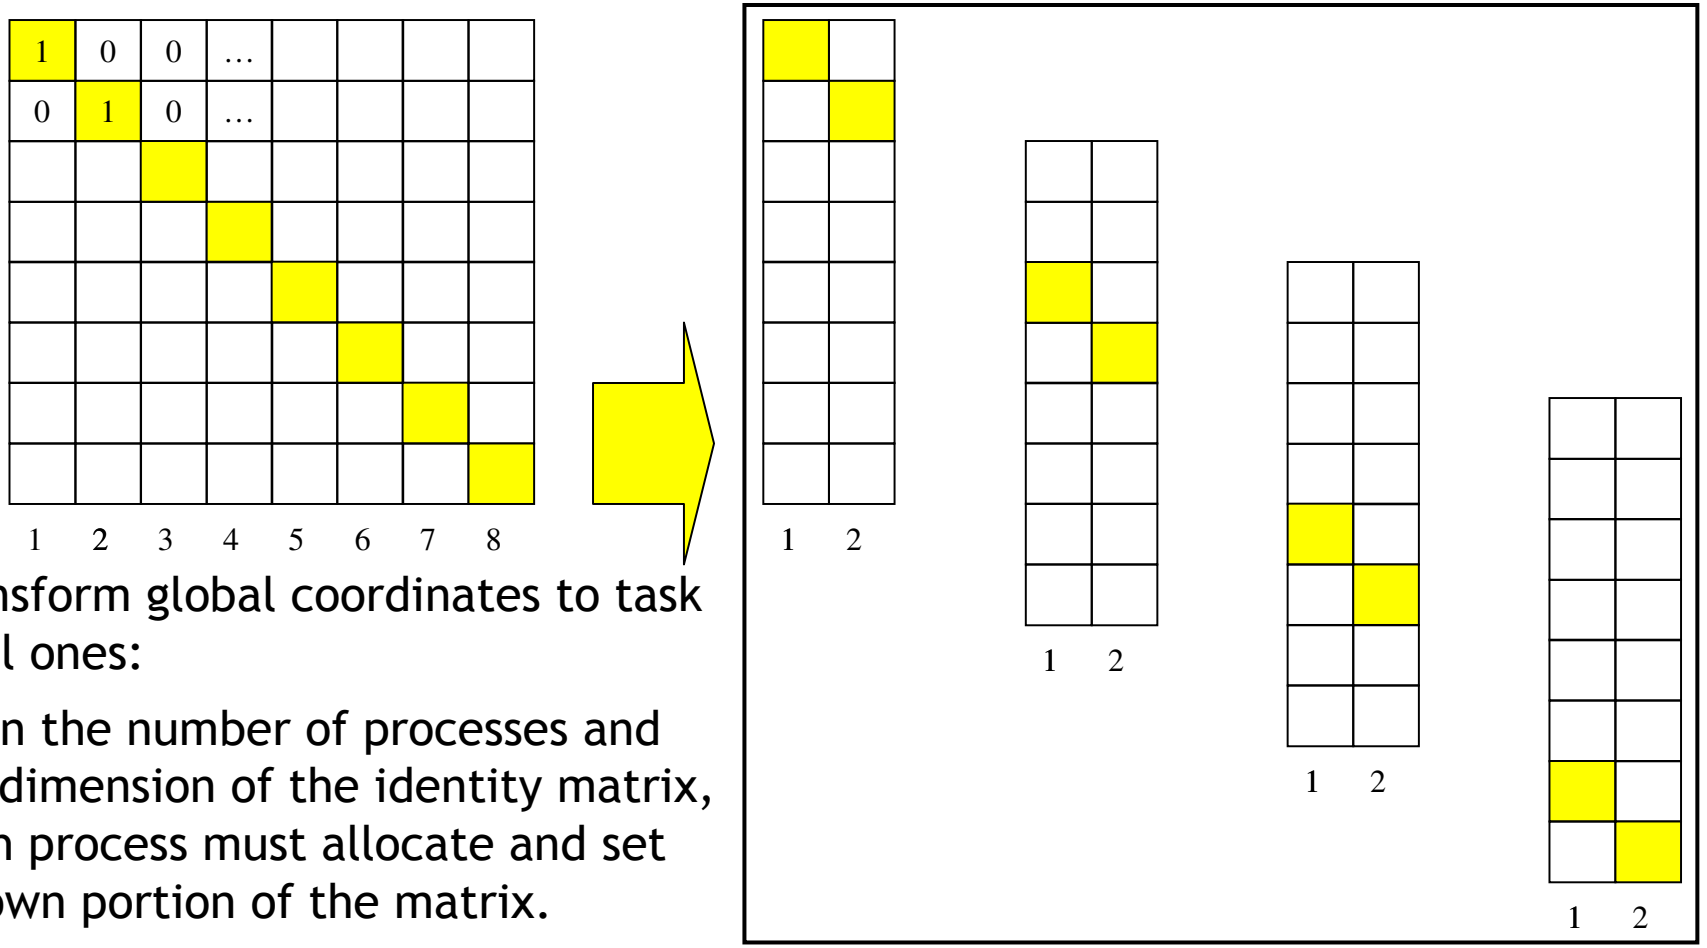

Transform global coordinates to task local ones:

given the number of processes and the dimension of the identity matrix, each process must allocate and set its own portion of the matrix.

Exercise 8: matrix transposition

- Initialize a 8x8 distributed matrix A by setting A_{ij} = 1000*i+j
- Both A and B are distributed by lines over 4 processors
- Print out A
- Evaluate: B = A^T
- Print out B

FORTRAN

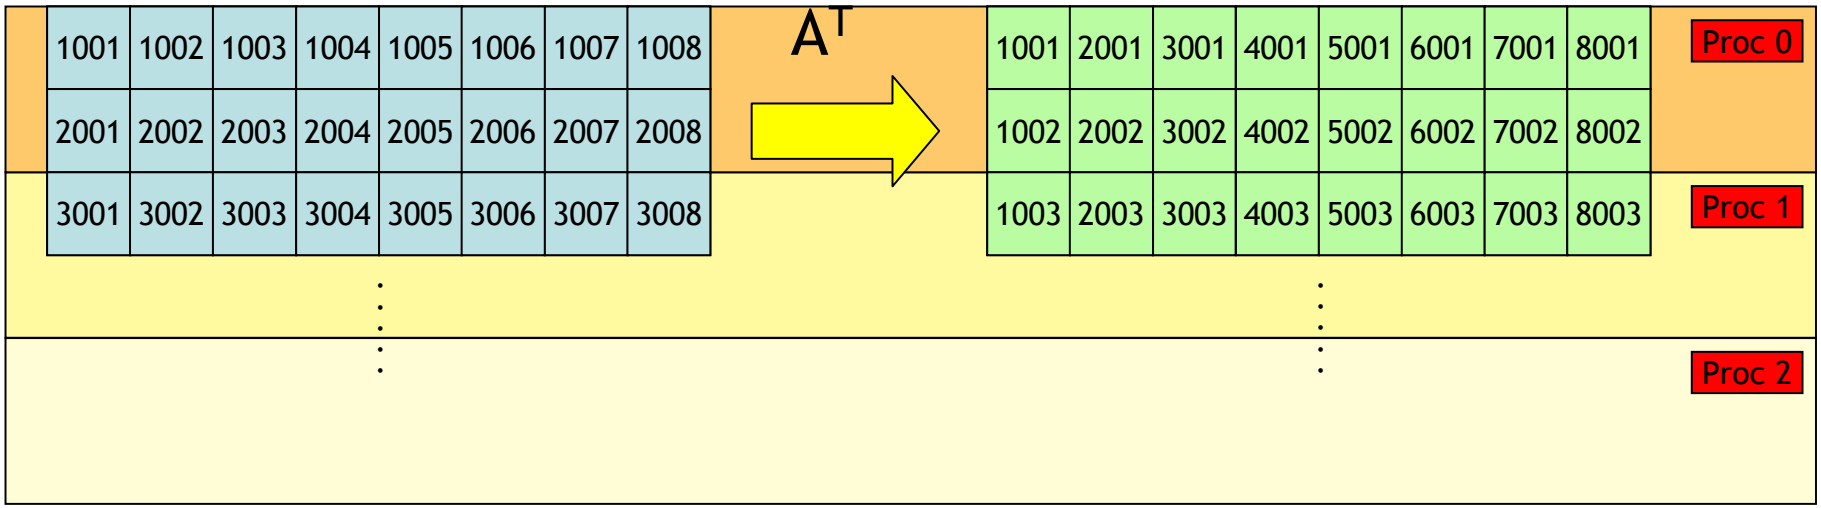

1. Use a collective communication. Which one?
2. Each block passed to the collective MPI function must contain data stored contiguously in memory!!

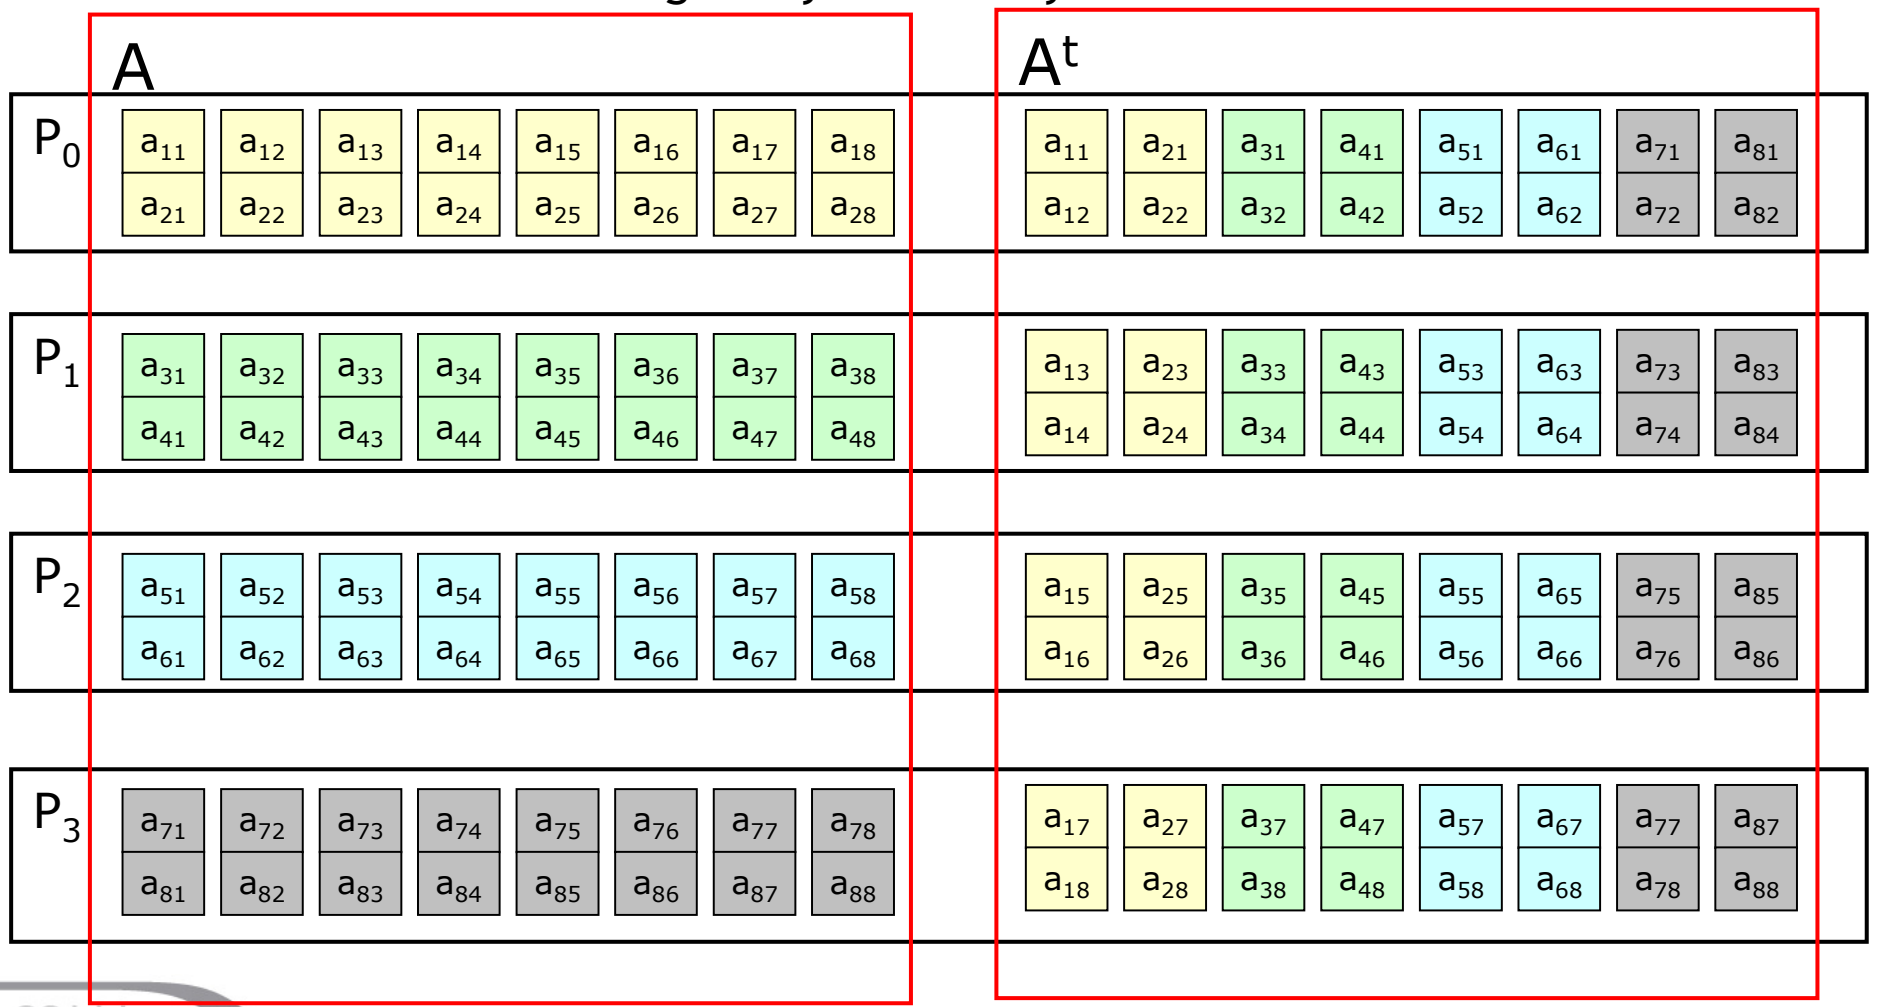

FORTRAN case

Exercise 9: Matrix Multiplication

Write a **subroutine** implementing matrix multiplication and test it.

$$C = A\,B \longrightarrow c_{ij} = \sum_k a_{ik} b_{kj}$$

A, B and C being NxN matrices distributed by row among processes (at least 8x8).
Inizialize A and B matrices respectively as $a_{ij} = i*j$ and $b_{ij}=1/(i*j)$.
Try to minimize memory allocation and the number of MPI calls.

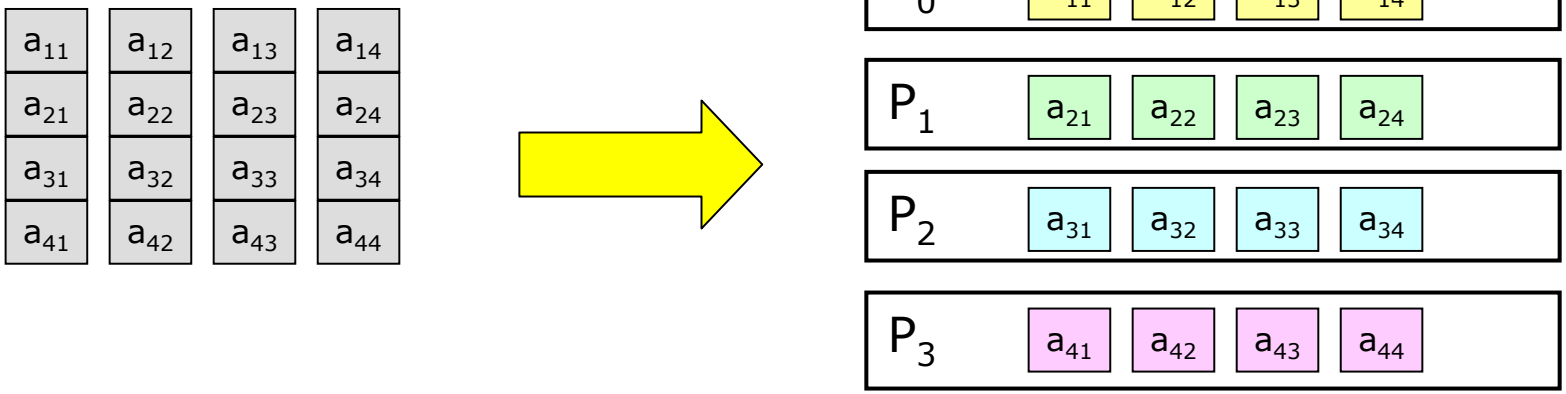

FORTRAN CASE: each process manage a row of elements (or blocks)

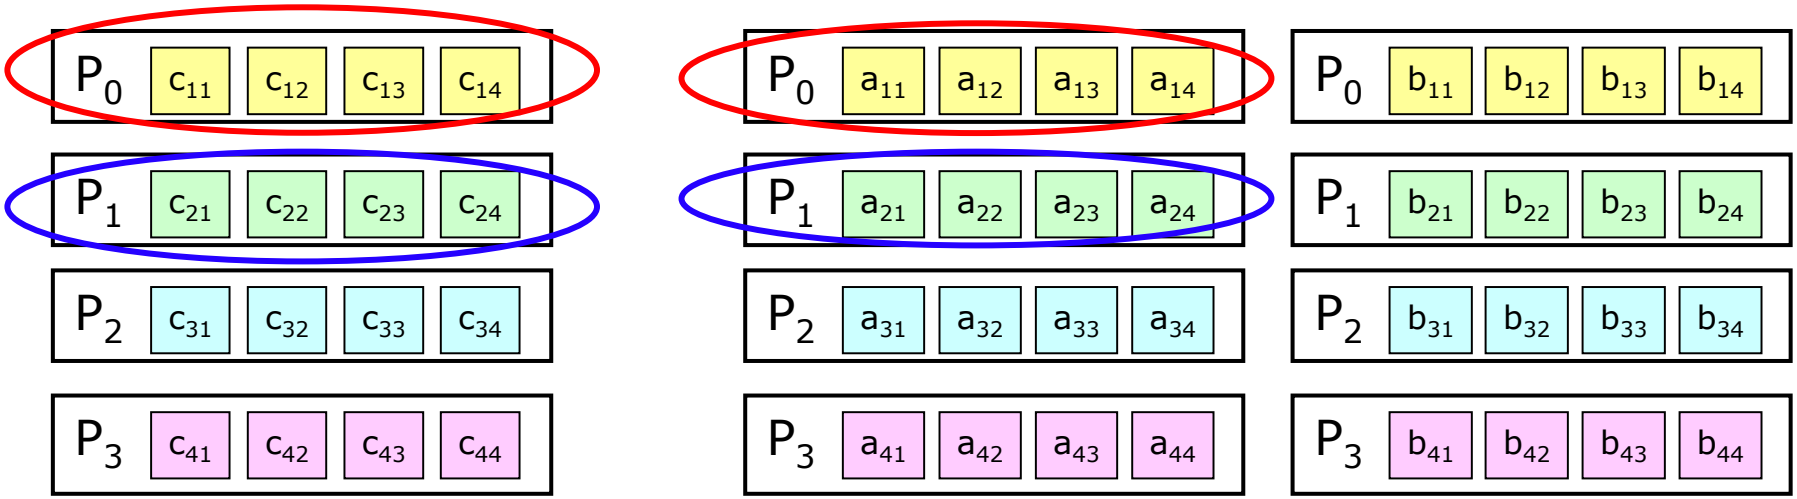

$$c_{11} = a_{11} b_{11} + a_{12} b_{21} + a_{13} b_{31} + a_{14} b_{41}$$

Each process compute the first element (block) of its own row

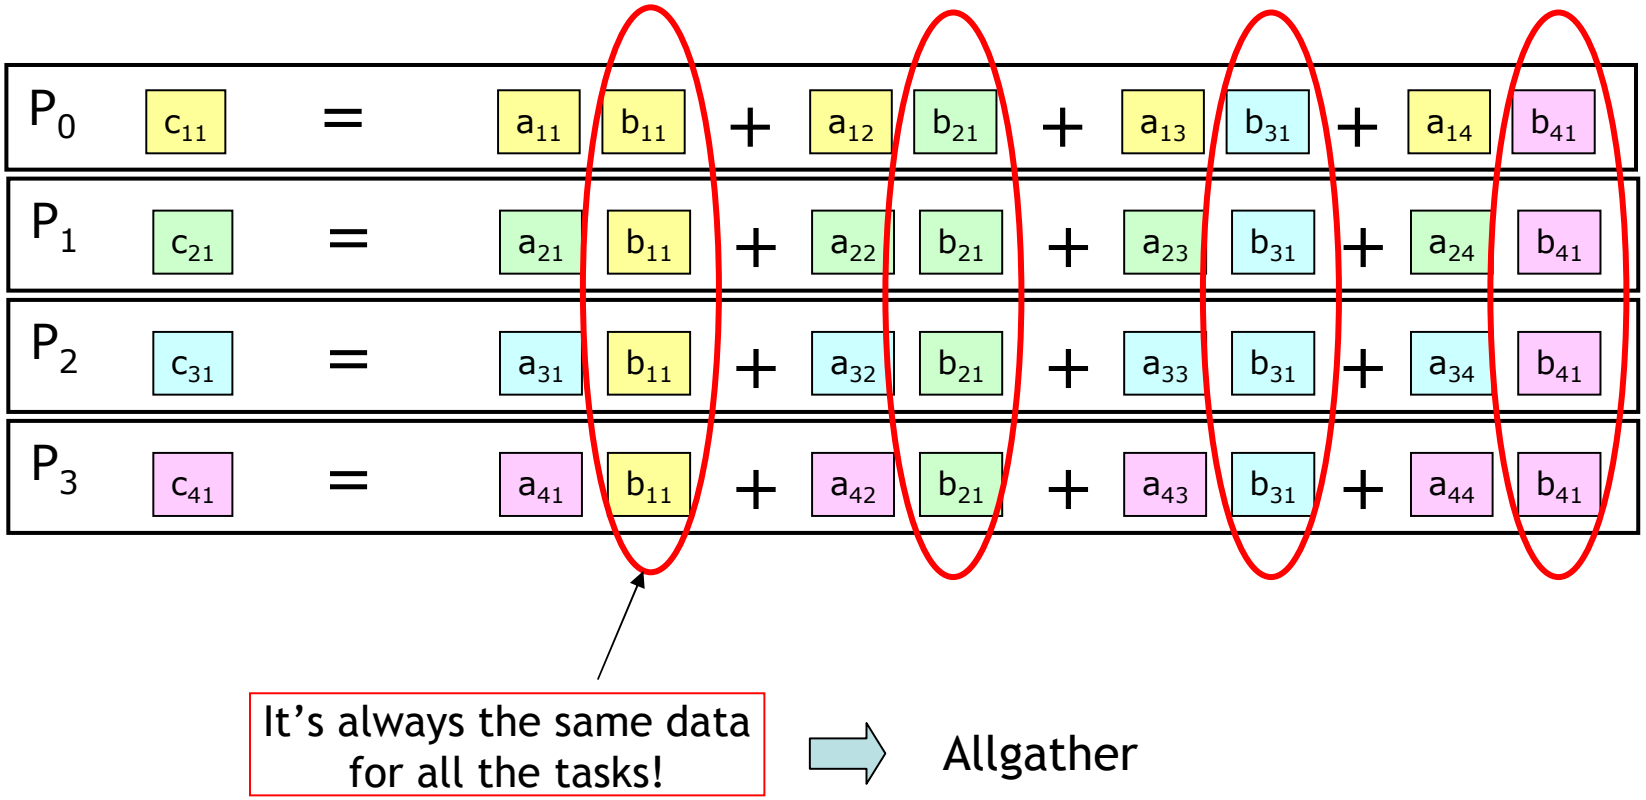

Step 1: allgather

Perform an All gather, of the first colum of blocks

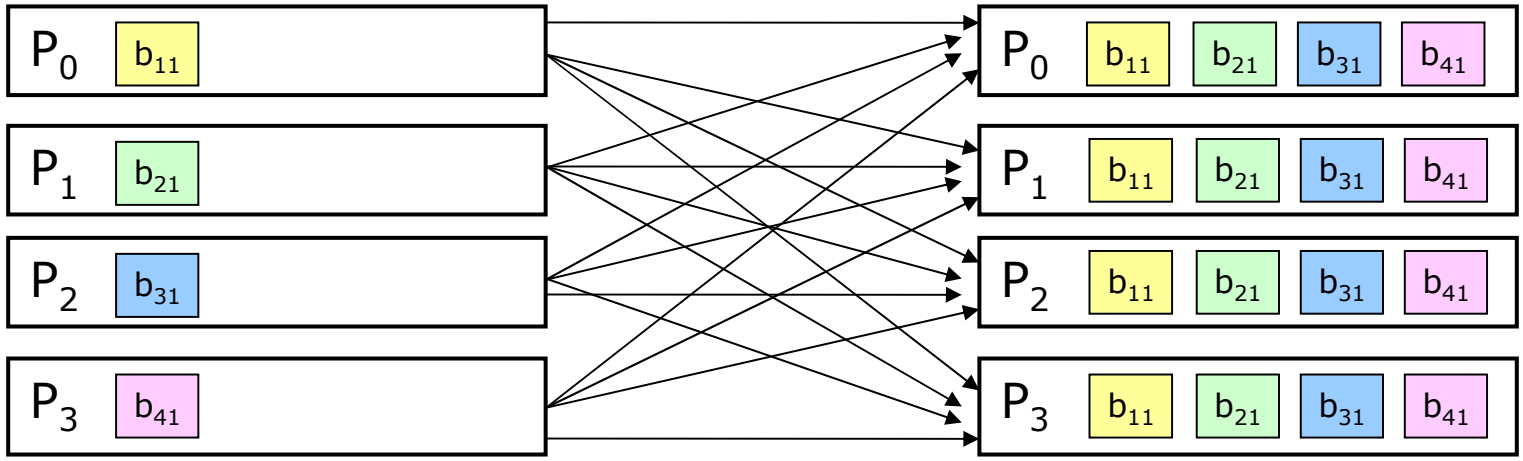

Step 2: local computation

Each processor calculate the first column of the matrix C

P_0 c_{11} = a_{11} b_{11} + a_{12} b_{21} + a_{13} b_{31} + a_{14} b_{41}

P_1 c_{21} = a_{21} b_{11} + a_{22} b_{21} + a_{23} b_{31} + a_{24} b_{41}

P_2 c_{31} = a_{31} b_{11} + a_{32} b_{21} + a_{33} b_{31} + a_{34} b_{41}

P_3 c_{41} = a_{41} b_{11} + a_{42} b_{21} + a_{43} b_{31} + a_{44} b_{41}

Generalize

Repeat Step 1 and Step 2 for each column elements or blocks of matrix C, until matrix C is complete

Exercise 10 (optional): MPI-2 Parallel I/O

Write a code (mpi2io.f90) that reads a binary file (mpi2io.bin) according to the following instructions (see mpi2io_template.f90):

1. read initial data in parallel, using parallel data distribution and standard I/O.
2. write in a "critical" way (one processor after the other)
3. do the same using MPI2 I/O
4. use MPI2 I/O with Views

Exercise 7: Parallelization of a code

transport problem

A code evolving the motion equation:

d/dx + d/dy = -d/dt,

is provided.

Original data have a gaussian distribution and must be moved along Y=X direction, i.e. toward the up-right corner of the system.

Execise: MAKE IT PARALLEL using domain decomposition

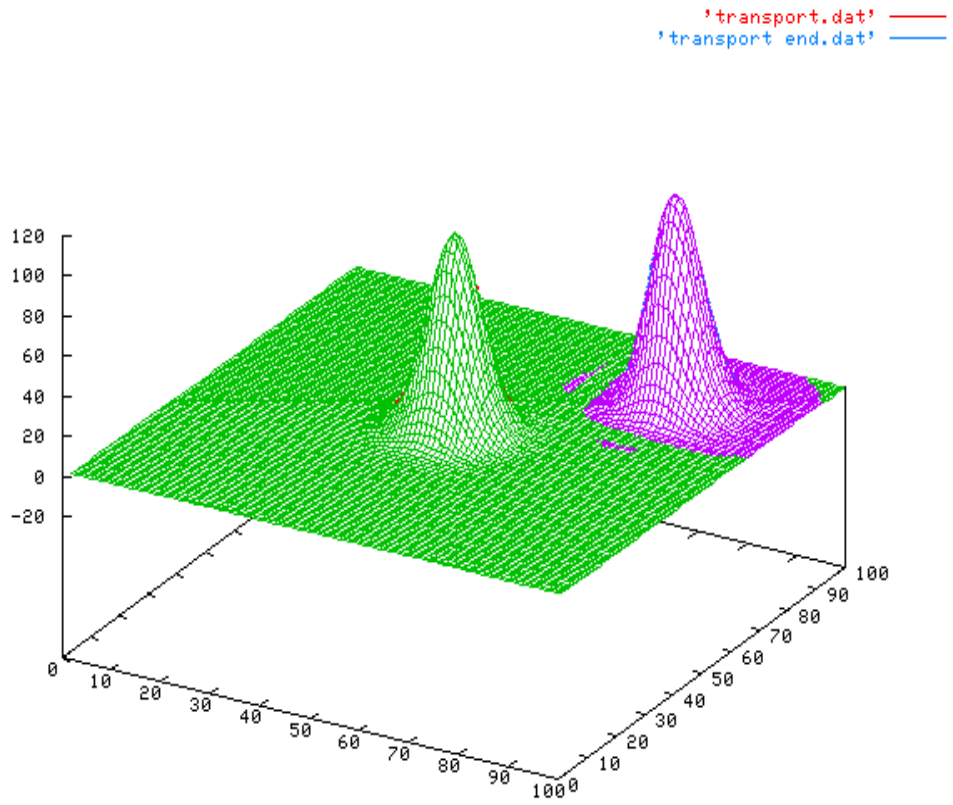

echo "set hidden3d; splot 'transport.dat' w l, 'transport_end.dat' w l" | gnuplot -persist

Exercise 7: Data distribution, ghost-cells

- Distribution by “leading dimension”:

  -FORTRAN: first coordinate

  -C: last coordinate

- The important thing is that the DATA to send/recv are CONTIGUOUS in memory. Otherwise a copy of them to/from a temporary, contiguous buffer is needed.

MORE TASKS: these are easy for the serial version… do it in parallel too.

- Evaluate the average over all the system and check it every 10 steps

- Find the global maximum and print its position every 10 steps

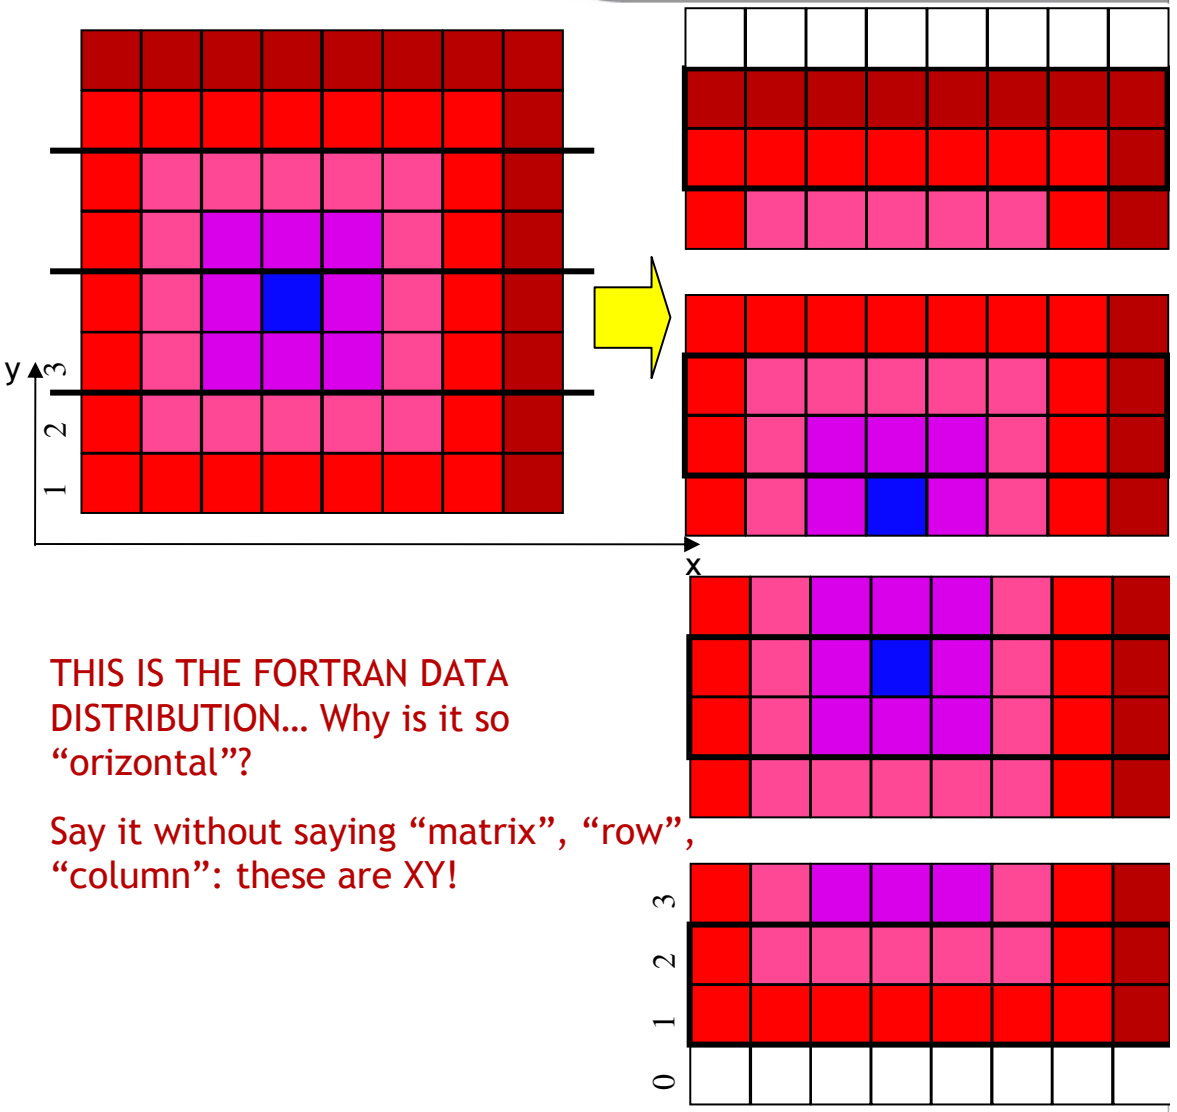

THIS IS THE FORTRAN DATA DISTRIBUTION… Why is it so “orizontal”?

Say it without saying “matrix”, “row”, “column”: these are XY!

Exercise 7: Hybridization

- Add the OpenMP directives to the MPI code to parallelize some loops and to manage the MPI communications.
- Select and check the right MPI level of thread support.
- Print both the process and thread identifiers.
- Compile your code with the OpenMP support.
- Run with different configurations for processes and threads.