Array Distribution Schemes

Ned Nedialkov

McMaster University
Canada

SE 3F03
February 2014

Outline

Block distributions

- Each process receives a contiguous part of the data
- Suitable when there is a locality in the interactions
- E.g. distributing an $n \times n$ matrix
  - row-wise distribution: each process receives $\approx n/p$ rows of the matrix
  - column-wise distribution: each process receives $\approx n/p$ columns of the matrix
  - suitable e.g. for matrix-vector multiplication

2D distribution

- Assume a process grid of size $p_1 \times p_2$
  Number of processes is $p = p_1 p_2$
- Consider distributing an $n \times n$ matrix
  - We can split it into $n/p_1 \times n/p_2$ sub matrices

4×4 process grid

P_0	P_1	P_2	P_3
P_4	P_5	P_6	P_7
P_8	P_9	P_{10}	P_{11}
P_{12}	P_{13}	P_{14}	P_{15}

2×4 process grid

P_0	P_1	P_2	P_3
P_4	P_5	P_6	P_7

Cyclic and block-cyclic distributions

- If the amount of work is different for different parts of a matrix, this could lead to load imbalances
- One way to avoid this is to use cyclic or block-cyclic distributions

1D cyclic distribution on 4 processes

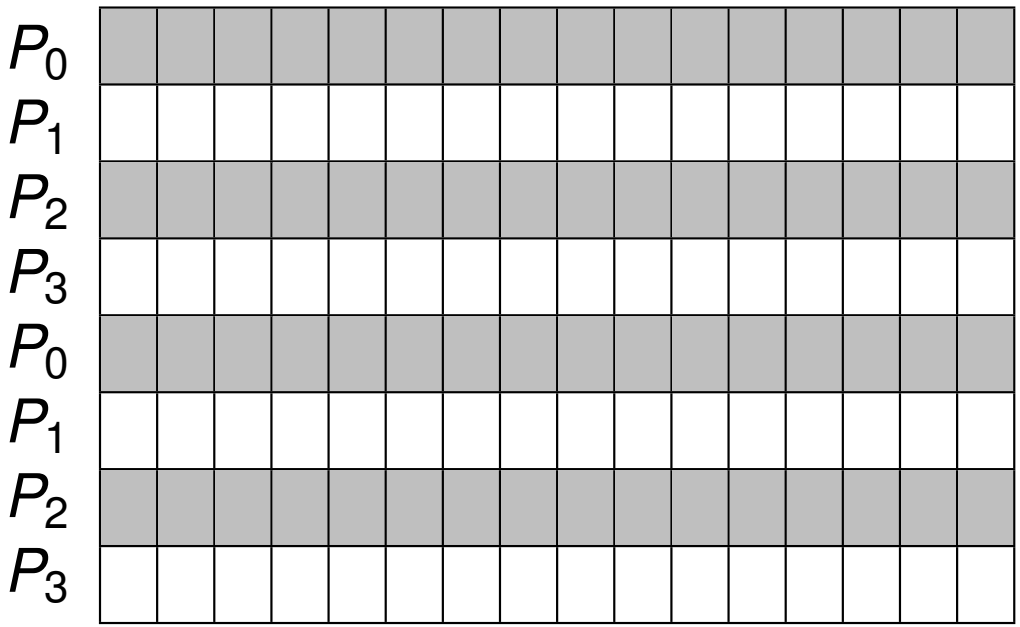

1D block-cyclic distribution of 4 processes

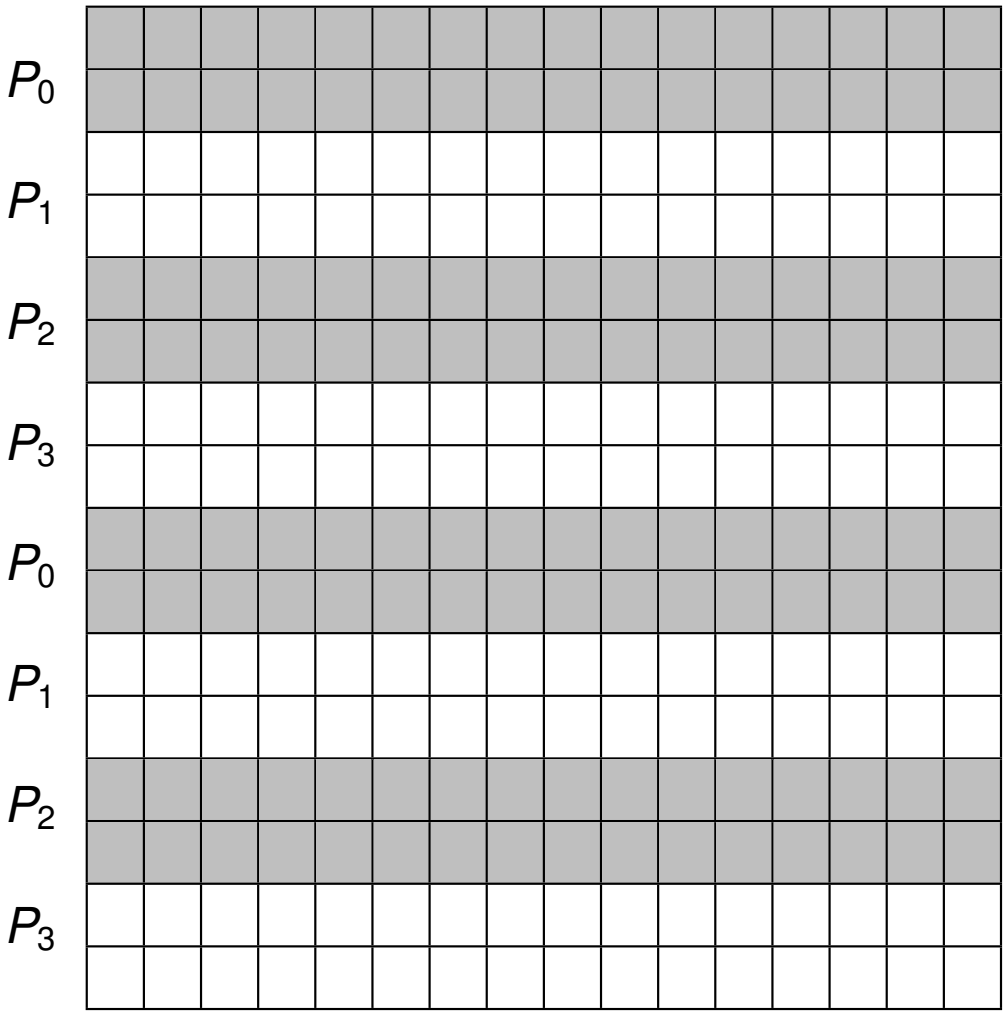

2D distribution on 4 processes

P_0	P_1	P_0	P_1
P_2	P_3	P_2	P_3
P_0	P_1	P_0	P_1
P_2	P_3	P_2	P_3