



# Introduction to Parallel Computing

George Karypis

Dense Matrix Algorithms



# Outline

- Focus on numerical algorithms involving dense matrices:
  - Matrix-Vector Multiplication
  - Matrix-Matrix Multiplication
  - Gaussian Elimination
- Decompositions & Scalability

# Review

**Table 4.1** Summary of communication times of various operations discussed in Sections 4.1–4.7 on a hypercube interconnection network. The message size for each operation is  $m$  and the number of nodes is  $p$ .

Operation	Hypercube Time	B/W Requirement
One-to-all broadcast, All-to-one reduction	$\min((t_s + t_w m) \log p, 2(t_s \log p + t_w m))$	$\Theta(1)$
All-to-all broadcast, All-to-all reduction	$t_s \log p + t_w m(p - 1)$	$\Theta(1)$
All-reduce	$\min((t_s + t_w m) \log p, 2(t_s \log p + t_w m))$	$\Theta(1)$
Scatter, Gather	$t_s \log p + t_w m(p - 1)$	$\Theta(1)$
All-to-all personalized	$(t_s + t_w m)(p - 1)$	$\Theta(p)$
Circular shift	$t_s + t_w m$	$\Theta(p)$

# Matrix-Vector Multiplication


- Compute:  $y = Ax$ 
  - $y, x$  are  $n \times 1$  vectors
  - $A$  is an  $n \times n$  dense matrix
- Serial complexity:  $W = O(n^2)$ .
- We will consider:
  - 1D & 2D partitioning.

---

```
1.  procedure MAT_VECT ( $A, x, y$ )
2.  begin
3.      for  $i := 0$  to  $n - 1$  do
4.          begin
5.               $y[i] := 0$ ;
6.              for  $j := 0$  to  $n - 1$  do
7.                   $y[i] := y[i] + A[i, j] \times x[j]$ ;
8.              endfor;
9.          end MAT_VECT
```

---

# Row-wise 1D Partitioning




How do we perform the operation?

# Row-wise 1D Partitioning


Each processor needs to have the entire  $x$  vector.

All-to-all broadcast



(b) Distribution of the full vector among all the processes by all-to-all broadcast

Local computations



(c) Entire vector distributed to each process after the broadcast

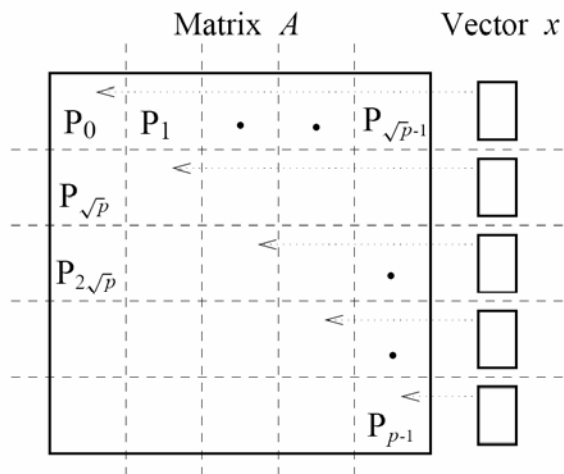
Analysis?

$$T_P = \frac{n^2}{p} + t_s \log p + t_w n.$$



$$T_o = t_s p \log p + t_w np.$$

$$W = \Theta(p^2)$$

# Block 2D Partitioning



(a) Initial data distribution and communication steps to align the vector along the diagonal




(d) Final distribution of the result vector

How do we perform the operation?

# Block 2D Partitioning

Each processor needs to have the portion of the x vector that corresponds to the set of columns that it stores.



(b) One-to-all broadcast of portions of the vector along process columns

(c) All-to-one reduction of partial results

Analysis?

$$\begin{aligned}
 T_P &= \underbrace{\frac{n^2}{p}}_{\text{computation}} + \underbrace{t_s + t_w n / \sqrt{p}}_{\text{aligning the vector}} + \\
 &\quad \underbrace{(t_s + t_w n / \sqrt{p}) \log(\sqrt{p})}_{\text{columnwise one-to-all broadcast}} + \underbrace{(t_s + t_w n / \sqrt{p}) \log(\sqrt{p})}_{\text{all-to-one reduction}} \\
 &\approx \frac{n^2}{p} + t_s \log p + t_w \frac{n}{\sqrt{p}} \log p
 \end{aligned}$$

$$T_o = t_s p \log p + t_w n \sqrt{p} \log p.$$

$$W = \Theta(p \log^2 p).$$





# 1D vs 2D Formulation

- Which one is better?

# Matrix-Matrix Multiplication

- Compute:  $C = AB$ 
  - $A$ ,  $B$ , &  $C$  are  $n \times n$  dense matrices.
- Serial complexity:  
 $W = O(n^3)$ .
- We will consider:
  - 2D & 3D partitioning.

---

```
1. procedure MAT_MULT (A, B, C)
2. begin
3.   for i := 0 to n - 1 do
4.     for j := 0 to n - 1 do
5.       begin
6.         C[i, j] := 0;
7.         for k := 0 to n - 1 do
8.           C[i, j] := C[i, j] + A[i, k] × B[k, j];
9.         endfor;
10.      end MAT_MULT
```

---

# Simple 2D Algorithm

- Processors are arranged in a logical  $\sqrt{p} \times \sqrt{p}$  2D topology.
- Each processor gets a block of  $(n/\sqrt{p}) \times (n/\sqrt{p})$  block of A, B, & C.
- It is responsible for computing the entries of C that it has been assigned to.
- Analysis?

$$T_P = \frac{n^3}{p} + t_s \log p + 2t_w \frac{n^2}{\sqrt{p}}.$$

$$W = \Theta(p^{3/2}).$$

How about the  
memory  
complexity?

# Cannon's Algorithm

- Memory efficient variant of the simple algorithm.
- Key idea:

- Replace traditional loop:

$$C_{i,j} = \sum_{k=0}^{\sqrt{p}-1} A_{i,k} * B_{k,j}$$

- With the following loop:

$$C_{i,j} = \sum_{k=0}^{\sqrt{p}-1} A_{i,(i+j+k)\% \sqrt{p}} * B_{(i+j+k)\% \sqrt{p},j}$$

- During each step, processors operate on different blocks of  $A$  and  $B$ .

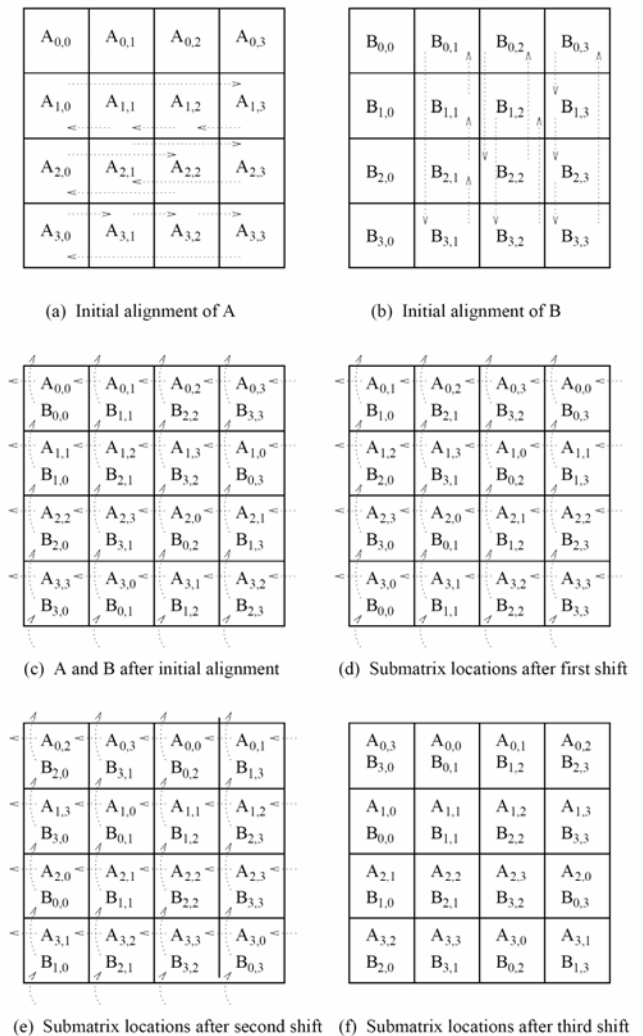


Figure 8.3 The communication steps in Cannon's algorithm on 16 processes.

$$T_P = \frac{n^3}{p} + 2\sqrt{p}t_s + 2t_w \frac{n^2}{\sqrt{p}}$$

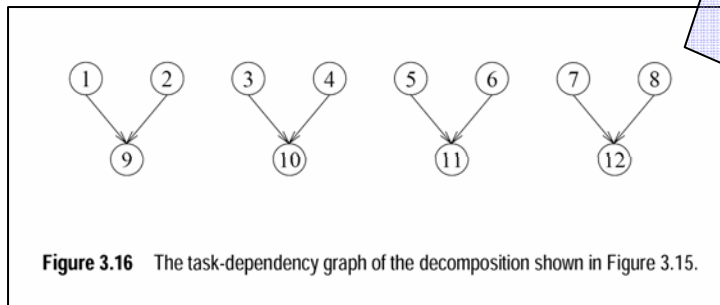
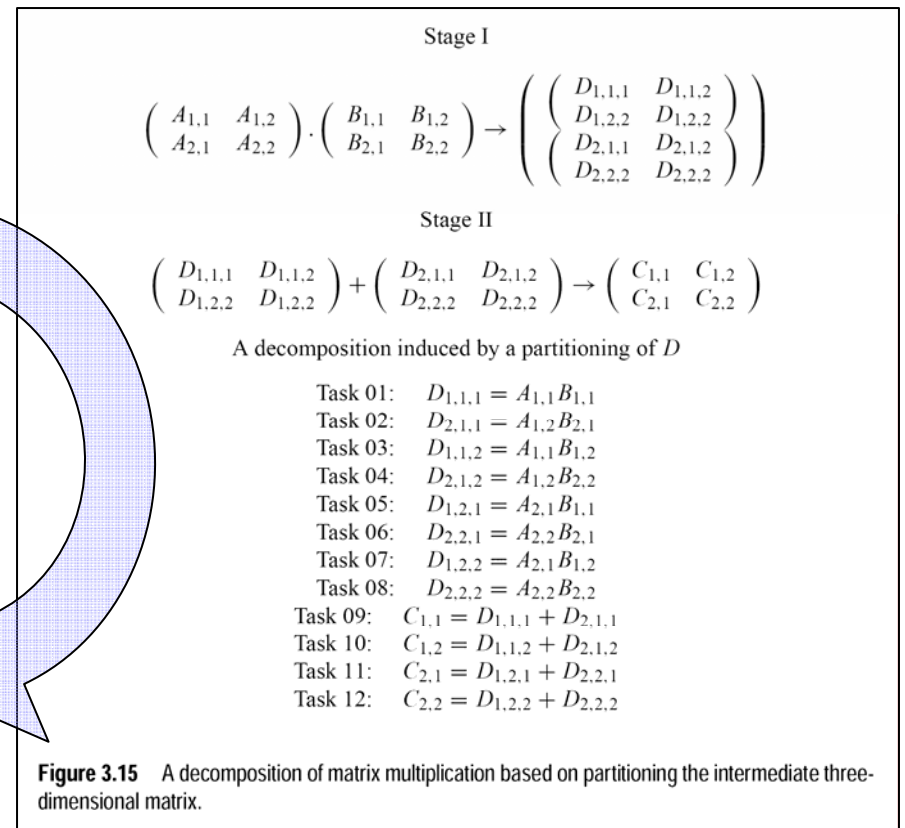
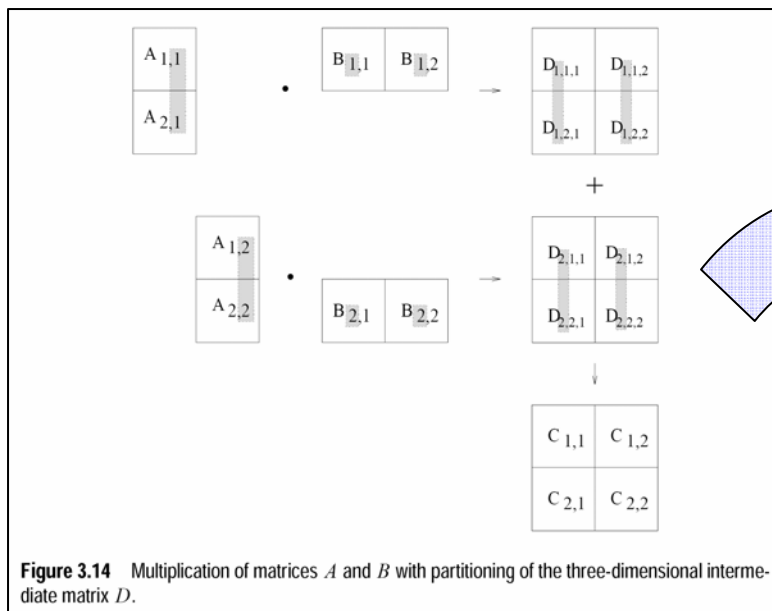


# Can we do better?

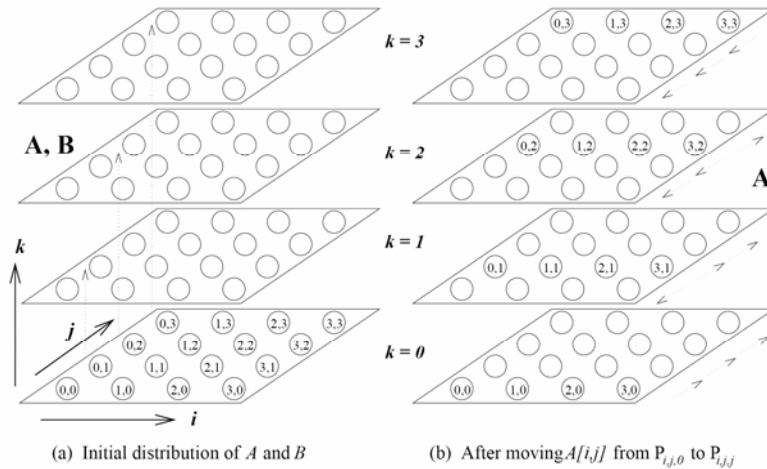
- Can we use more than  $O(n^2)$  processors?
- So far the task corresponded to the dot-product of two vectors
  - i.e.,  $C_{i,j} = A_{i,*} \cdot B_{*,j}$
- How about performing this dot-product in parallel?
- What is the maximum concurrency that we can extract?

# 3D Algorithm—DNS Algorithm

## ■ Partitioning the intermediate data



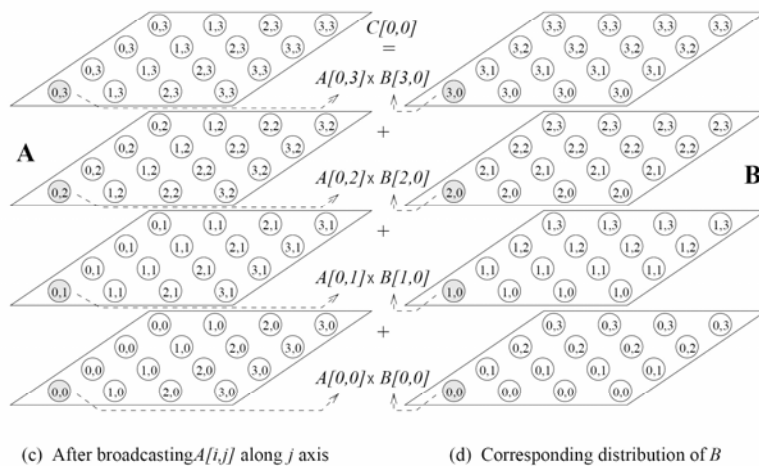
# 3D Algorithm—DNS Algorithm



$$q = p^{1/3}$$

$$T_P \approx \left(\frac{n}{q}\right)^3 + 3t_s \log q + 3t_w \left(\frac{n}{q}\right)^2 \log q$$

$$T_P = \frac{n^3}{p} + t_s \log p + t_w \frac{n^2}{p^{2/3}} \log p.$$



$$W = \Theta(p(\log p)^3)$$

**Figure 8.4** The communication steps in the DNS algorithm while multiplying  $4 \times 4$  matrices  $A$  and  $B$  on 64 processes. The shaded processes in part (c) store elements of the first row of  $A$  and the shaded processes in part (d) store elements of the first column of  $B$ .



Which one is better?



# Gaussian Elimination


- Solve  $Ax=b$ 
  - $A$  is an  $n \times n$  dense matrix.
  - $x$  and  $b$  are dense vectors
- Serial complexity:  
 $W = O(n^3)$ .
- There are two key steps in each iteration:
  - Division step
  - Rank-1 update
- We will consider:
  - 1D & 2D partitioning, and introduce the notion of pipelining.

```

1. procedure GAUSSIAN_ELIMINATION ( $A, b, y$ )
2. begin
3.   for  $k := 0$  to  $n - 1$  do           /* Outer loop */
4.     begin
5.       for  $j := k + 1$  to  $n - 1$  do
6.          $A[k, j] := A[k, j]/A[k, k];$  /* Division step */
7.        $y[k] := b[k]/A[k, k];$ 
8.        $A[k, k] := 1;$ 
9.       for  $i := k + 1$  to  $n - 1$  do
10.        begin
11.          for  $j := k + 1$  to  $n - 1$  do
12.             $A[i, j] := A[i, j] - A[i, k] \times A[k, j];$  /* Elimination step */
13.           $b[i] := b[i] - A[i, k] \times y[k];$ 
14.           $A[i, k] := 0;$ 
15.        endfor;                       /* Line 9 */
16.      endfor;                           /* Line 3 */
17.    end GAUSSIAN_ELIMINATION

```

**Algorithm 8.4** A serial Gaussian elimination algorithm that converts the system of linear equations  $Ax = b$  to a unit upper-triangular system  $Ux = y$ . The matrix  $U$  occupies the upper-triangular locations of  $A$ . This algorithm assumes that  $A[k, k] \neq 0$  when it is used as a divisor on lines 6 and



**Figure 8.5** A typical computation in Gaussian elimination.

# 1D Partitioning

- Assign  $n/p$  rows of  $A$  to each processor.
- During the  $i^{\text{th}}$  iteration:
  - Divide operation is performed by the processor who stores row  $i$ .
  - Result is broadcasted to the rest of the processors.
  - Each processor performs the rank-1 update for its local rows.
- Analysis?

$$T_P = \frac{3}{2}n(n-1) + t_s n \log n + \frac{1}{2}t_w n(n-1) \log n.$$

(one element per processor)

P <sub>0</sub>	1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
P <sub>1</sub>	0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
P <sub>2</sub>	0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
P <sub>3</sub>	0	0	0	(3,3)	(3,4)	(3,5)	(3,6)	(3,7)
P <sub>4</sub>	0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
P <sub>5</sub>	0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
P <sub>6</sub>	0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
P <sub>7</sub>	0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(a) Computation:

- (i)  $A[k,j] := A[k,j]/A[k,k]$  for  $k < j < n$
- (ii)  $A[k,k] := 1$

P <sub>0</sub>	1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
P <sub>1</sub>	0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
P <sub>2</sub>	0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
P <sub>3</sub>	0	0	0	1	(3,4)	(3,5)	(3,6)	(3,7)
P <sub>4</sub>	0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
P <sub>5</sub>	0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
P <sub>6</sub>	0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
P <sub>7</sub>	0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(b) Communication:

One-to-all broadcast of row  $A[k,*]$

P <sub>0</sub>	1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
P <sub>1</sub>	0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
P <sub>2</sub>	0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
P <sub>3</sub>	0	0	0	1	(3,4)	(3,5)	(3,6)	(3,7)
P <sub>4</sub>	0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
P <sub>5</sub>	0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
P <sub>6</sub>	0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
P <sub>7</sub>	0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(c) Computation:

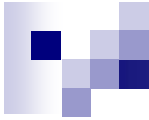
- (i)  $A[i,j] := A[i,j] - A[i,k] \times A[k,j]$   
for  $k < i < n$  and  $k < j < n$
- (ii)  $A[i,k] := 0$  for  $k < i < n$

Figure 8.6 Gaussian elimination steps during the iteration corresponding to  $k = 3$  for an  $8 \times 8$  matrix partitioned rowwise among eight processes.



# 1D Pipelined Formulation

- Existing Algorithm:  
Next iteration starts only when the previous iteration has finished.
- Key Idea:  
The next iteration can start as soon as the rank-1 update involving the next row has finished.
  - Essentially multiple iterations are performed simultaneously!



(0,0)	(0,1)	(0,2)	(0,3)	(0,4)
(1,0)	(1,1)	(1,2)	(1,3)	(1,4)
(2,0)	(2,1)	(2,2)	(2,3)	(2,4)
(3,0)	(3,1)	(3,2)	(3,3)	(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(a) Iteration k = 0 starts

1	(0,1)	(0,2)	(0,3)	(0,4)
(1,0)	√(1,1)	√(1,2)	√(1,3)	√(1,4)
(2,0)	(2,1)	(2,2)	(2,3)	(2,4)
(3,0)	(3,1)	(3,2)	(3,3)	(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(b)

1	(0,1)	(0,2)	(0,3)	(0,4)
(1,0)	(1,1)	(1,2)	(1,3)	(1,4)
(2,0)	√(2,1)	√(2,2)	√(2,3)	√(2,4)
(3,0)	(3,1)	(3,2)	(3,3)	(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(c)

1	(0,1)	(0,2)	(0,3)	(0,4)
(1,0)	(1,1)	(1,2)	(1,3)	(1,4)
(2,0)	(2,1)	(2,2)	(2,3)	(2,4)
(3,0)	√(3,1)	√(3,2)	√(3,3)	√(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(d)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	(1,1)	(1,2)	(1,3)	(1,4)
(2,0)	(2,1)	(2,2)	(2,3)	(2,4)
(3,0)	(3,1)	(3,2)	(3,3)	(3,4)
(4,0)	√(4,1)	√(4,2)	√(4,3)	√(4,4)

(e) Iteration k = 1 starts

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	(2,1)	√(2,2)	√(2,3)	√(2,4)
(3,0)	(3,1)	(3,2)	(3,3)	(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(f)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	(1,1)	(1,2)	(1,3)	(1,4)
0	(2,1)	(2,2)	(2,3)	(2,4)
0	(3,1)	√(3,2)	√(3,3)	√(3,4)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)

(g) Iteration k = 0 ends

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	(2,1)	(2,2)	(2,3)	(2,4)
0	(3,1)	(3,2)	(3,3)	(3,4)
0	(4,1)	√(4,2)	√(4,3)	√(4,4)

(h)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	(2,2)	(2,3)	(2,4)
0	(3,1)	(3,2)	(3,3)	(3,4)
0	(4,1)	(4,2)	(4,3)	(4,4)

(i) Iteration k = 2 starts

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	(3,2)	√(3,3)	√(3,4)
0	(4,1)	(4,2)	(4,3)	(4,4)

(j) Iteration k = 1 ends

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	(3,2)	(3,3)	(3,4)
0	0	(4,2)	√(4,3)	√(4,4)

(k)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	(3,2)	(3,3)	(3,4)
0	0	(4,2)	(4,3)	(4,4)

(l)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	0	(3,3)	(3,4)
0	0	(4,2)	(4,3)	(4,4)

(m) Iteration k = 3 starts

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	0	1	(3,4)
0	0	0	(4,3)	√(4,4)

(n)

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	0	1	(3,4)
0	0	0	(4,3)	(4,4)

(o) Iteration k = 3 ends

1	(0,1)	(0,2)	(0,3)	(0,4)
0	1	(1,2)	(1,3)	(1,4)
0	0	1	(2,3)	(2,4)
0	0	0	1	(3,4)
0	0	0	(4,3)	(4,4)

(p) Iteration k = 4

Cost-optimal with  $n$  processors

---> Communication for k = 0, 3

—> Communication for k = 1

- -> Communication for k = 2

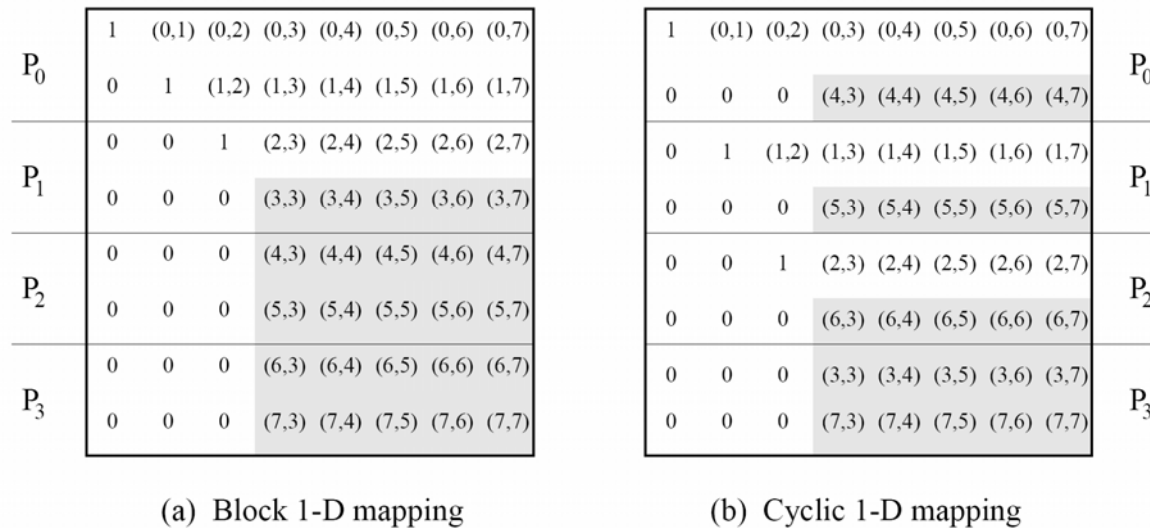
□ Computation for k = 0, 3

□ Computation for k = 1, 4

□ Computation for k = 2

# 1D Partitioning

- Is the block mapping a good idea?



**Figure 8.9** Computation load on different processes in block and cyclic 1-D partitioning of an  $8 \times 8$  matrix on four processes during the Gaussian elimination iteration corresponding to  $k = 3$ .

# 2D Mapping

- Each processor gets a 2D block of the matrix.
- Steps:
  - Broadcast of the “active” column along the rows.
  - Divide step in parallel by the processors who own portions of the row.
  - Broadcast along the columns.
  - Rank-1 update.
- Analysis?

1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
0	0	0	(3,3)	(3,4)	(3,5)	(3,6)	(3,7)
0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(a) Rowwise broadcast of  $A[i,k]$  for  $(k - 1) < i < n$

1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
0	0	0	(3,3)	(3,4)	(3,5)	(3,6)	(3,7)
0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(b)  $A[k,j] := A[k,j]/A[k,k]$  for  $k < j < n$

1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
0	0	0	1	(3,4)	(3,5)	(3,6)	(3,7)
0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

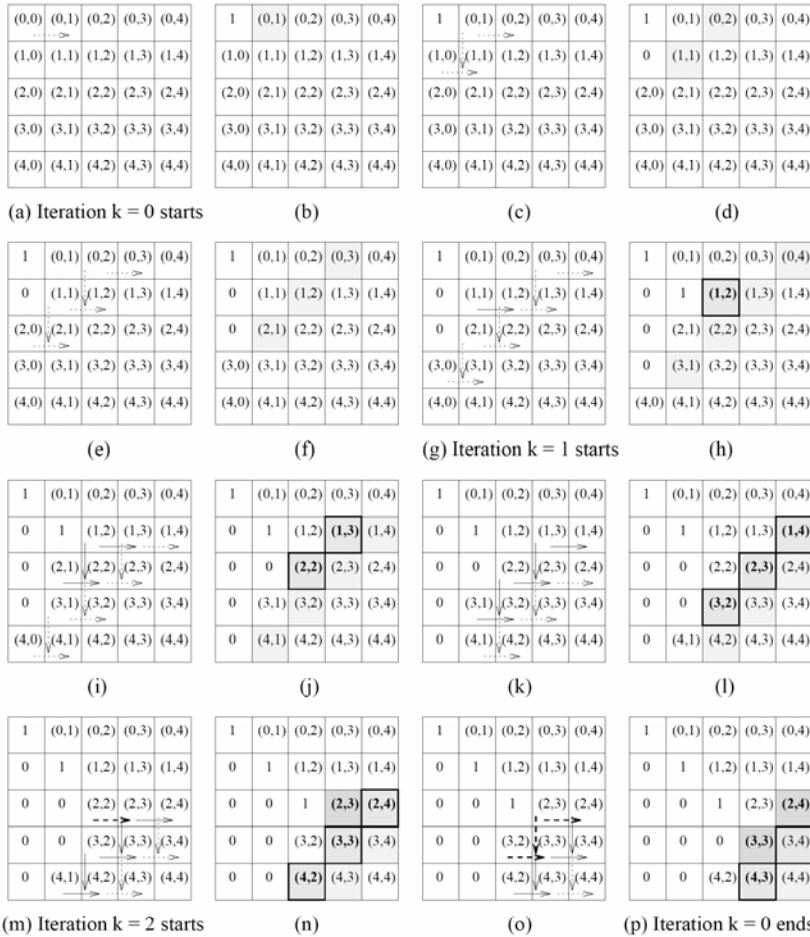
(c) Columnwise broadcast of  $A[k,j]$  for  $k < j < n$

1	(0,1)	(0,2)	(0,3)	(0,4)	(0,5)	(0,6)	(0,7)
0	1	(1,2)	(1,3)	(1,4)	(1,5)	(1,6)	(1,7)
0	0	1	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)
0	0	0	1	(3,4)	(3,5)	(3,6)	(3,7)
0	0	0	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)
0	0	0	(5,3)	(5,4)	(5,5)	(5,6)	(5,7)
0	0	0	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)
0	0	0	(7,3)	(7,4)	(7,5)	(7,6)	(7,7)

(d)  $A[i,j] := A[i,j] - A[i,k] \times A[k,j]$  for  $k < i < n$  and  $k < j < n$

**Figure 8.10** Various steps in the Gaussian elimination iteration corresponding to  $k = 3$  for an  $8 \times 8$  matrix on 64 processes arranged in a logical two-dimensional mesh.

# 2D Pipelined



Cost-optimal with  $n^2$  processors

.....> Communication for  $k = 0$

—> Communication for  $k = 1$

---> Communication for  $k = 2$

□ Computation for  $k = 0$

■ Computation for  $k = 1$

■ Computation for  $k = 2$