

• $G(V, E, w, z)$

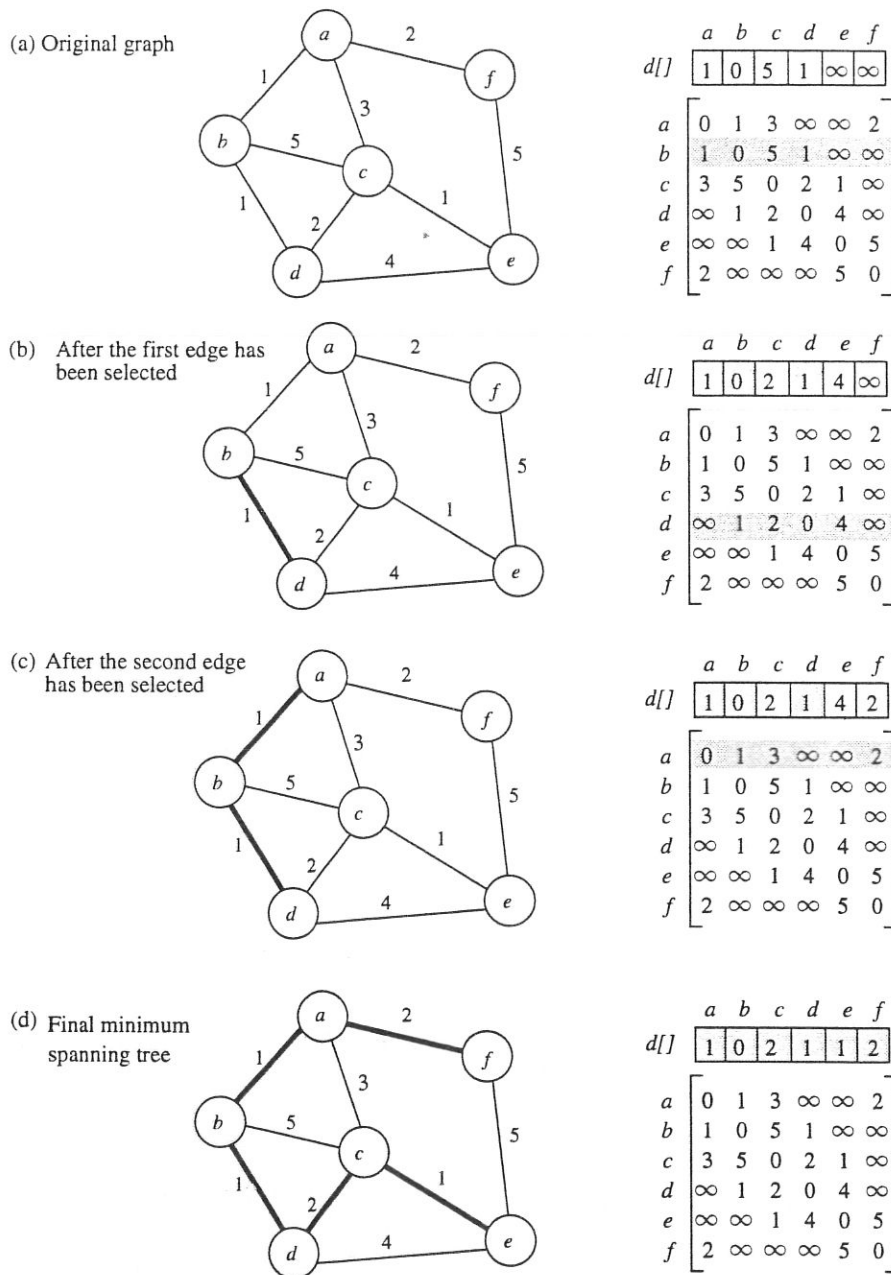


Figure 7.5 Prim's minimum spanning tree algorithm. The MST is rooted at vertex b . During each iteration, the minimum cost edge connecting a vertex in V_T to a vertex in $V - V_T$ is selected and the corresponding vertex is added to V_T (shown shaded in the distance array d). The $d[v]$ values of the vertices in $V - V_T$ are then updated.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

• $\Theta(n^2)$

procedure PRIM_MST (V, E, w, r)

$V_T := \{r\}$

$d[r] := 0$

for all $v \in (V - V_T)$ do

if edge (r, v) exists set $d[v] := w(r, v)$

else set $d[v] := \infty$

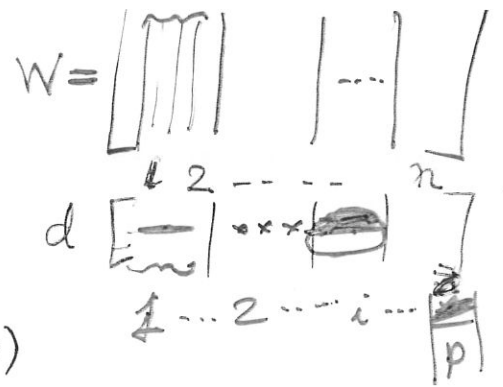
while $V_T \neq V$ do

→ find vertex u such that $d[u] = \min \{d[v] \mid v \in (V - V_T)\}$

$V_T := V_T \cup \{u\}$

for all $v \in (V - V_T)$ do

$d[v] = \min \{d[v], w(u, v)\}$



Dijkstra's single-source shortest paths algorithm is identical to Prim's MST algo, except for each $v \in (V - V_T)$

→ Dijkstra's stores $l[u]$, the min cost to reach u from s via V_T
 ie $l[v] = \min \{l[v], l[u] + w(u, v)\}$ is the update step

→ Prim's stores $d[u]$, cost of min-cost edge connecting vertex in V_T to u

• Parallel formulation of both algorithms are the same

• Dijkstra's all-pairs shortest paths: $\theta(n^3)$

→ Source-partitioned parallel formulation: uses n processors only

$$T_p = \theta(n^2) ; S = \theta(n^3) / \theta(n^2) ; E = \theta(1) ; \text{isoefficiency}_{\text{comm, concurr}} = \theta(p^3)$$

very poor scalability

$$\#p \approx \theta(n^{1/3})$$

→ Use source-parallel formulation: $p (> n)$ processors

$\frac{p}{n}$ procs/partition; each partition runs the 1-source shortest path algo

Can use $\theta(n^2)$ processors efficiently

$$T_p = \theta\left(\frac{n^3}{p}\right) + \theta(n \log p)$$

$$\text{Isoefficiency (communication)} = \theta((p \log p)^{1.5})$$

$$E = \frac{1}{1 + \theta(p \log p / n^2)}$$

$$\text{Isoefficiency (concurrency)} = \theta(p^{1.5})$$

Parallel Formulation: block-striped partitioning

- In each iteration, P_i computes $d_i[u] = \min\{d_i[v] \mid v \in (V - V_T) \cap V_i\}$
- Then do single-node accumulation at P_0 to find global minimum $d_i[u]$
- P_0 does $1 \rightarrow$ all BC of u
- Each P_i updates values of $d[v]$ for its local vertices
- (Each P_i needs to store the columns of the weighted adjacency matrix for V_i assigned to it.)
- Space = $\Theta(n^2/p)$; computation in each iteration = $\Theta(n/p)$

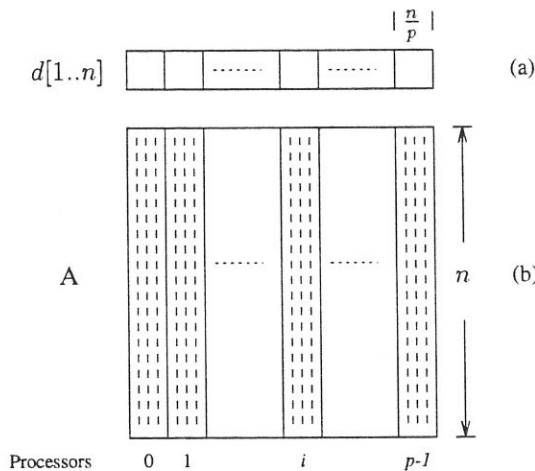


Figure 7.6 The partitioning of the distance array d and the adjacency matrix A among p processors.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

computation communication

HC: $T_p = \Theta\left(\frac{n^2}{p}\right) + \Theta(n \log p)$ For cost-optimality, $(p \log p)/n = O(1)$

$S = \frac{\Theta(n^2)}{\Theta(n^2/p) + \Theta(n \log p)}$ Isoefficiency (due to communication) is $\Theta(p^2 \log^2 p)$

$E = \frac{1}{1 + \Theta((p \log p)/n)}$

Mesh: $T_p = \Theta(n^2/p) + \Theta(n\sqrt{p})$ For cost-optimality, $p^{1.5}/n = O(1)$

$S = \frac{\Theta(n^2)}{\Theta(n^2/p) + \Theta(n\sqrt{p})}$ $\Rightarrow O(n^{2/3})$ processors can be used efficiently

$E = \frac{1}{1 + \Theta(p^{1.5}/n)}$ Isoefficiency (due to communication) = $\Theta(p^3)$

Source-parallel formulation of Dijkstra's algorithm

→ single-source algo. on each submesh: $T_p = \underbrace{\theta(n^3/p)}_{\text{computation}} + \underbrace{\theta(\sqrt{np})}_{\text{communic}}$

$$S = \frac{\theta(n^3)}{\theta(n^3/p) + \theta(\sqrt{np})}, \quad E = \frac{1}{1 + \theta(p^{1.5}/n^{2.5})}$$

For cost-optimality, $p^{1.5}/n^{2.5} = O(1) \Rightarrow$ can use $O(n^{1.66})$ processors efficiently

$$\text{Isoefficiency}_{\text{communic}} = \theta(p^{1.8})$$

$$\text{Isoefficiency}_{\text{concurrency}} = \theta(p^{1.5})$$

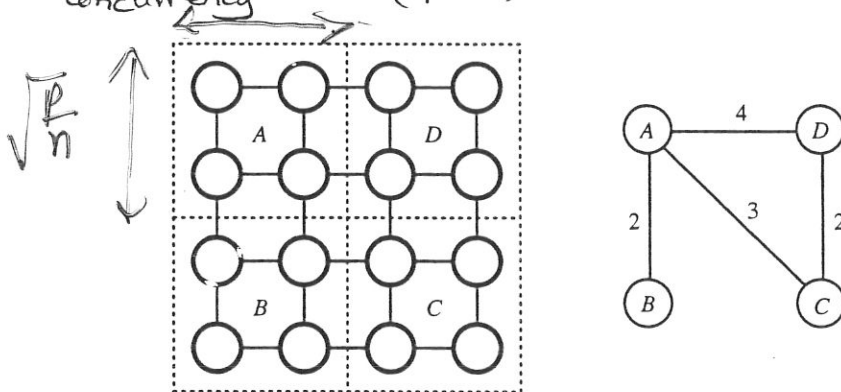


Figure 7.8 Partitioning a $\sqrt{p} \times \sqrt{p}$ mesh into n submeshes, each of size $\sqrt{p/n} \times \sqrt{p/n}$. In this example, $p = 16$ and $n = 4$. Each of the $\sqrt{p/n} \times \sqrt{p/n} = 2 \times 2$ meshes solves the single-source shortest paths problem for a given source vertex. In this example, each submesh is marked by the source vertex of its single-source algorithm.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

Hypercube:

// $p > n$, each ~~processor~~ ^{partition} handles $\frac{p}{n}$ processor

$$T_p = \theta\left(\frac{n^3}{p}\right) + \theta(n \log p)$$

$$S = \frac{\theta(n^3)}{\theta\left(\frac{n^3}{p}\right) + \theta(n \log p)}$$

$$E = \frac{1}{1 + \theta\left(\frac{p \log p}{n^2}\right)}$$

For cost optimality: $\frac{p \log p}{n^2} = \theta(1)$

$$\text{Iso-efficiency}_{\text{communic}} = \theta\left(\frac{p \log p}{n^2}\right)^{3/2}$$

$$\text{Isoefficiency}_{\text{concurr}} = \theta\left(\frac{p^3}{n^2}\right)$$

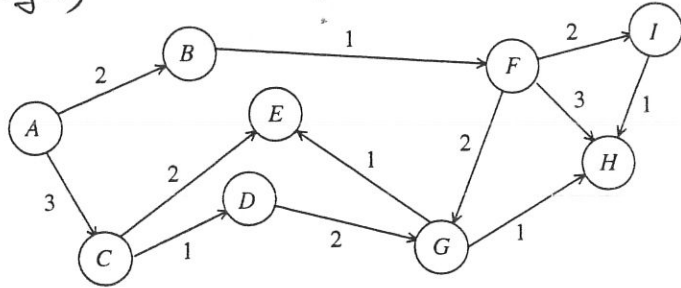
$$d_{ij}^{(k)} = \min_{1 \leq m \leq n} \{ d_{im}^{(k-1)} + w(v_m, v_j) \}$$

• $D^{(k)}$ computed from $D^{(k-1)}$ & A using modified matrix multiplication

$$c_{ij} = \min_{k=1}^n \{ a_{i,k} + b_{k,j} \}, \text{ instead of } \sum_{k=1}^n a_{i,k} \times b_{k,j}$$

• Total $\lceil \log(n-1) \rceil$ steps, ie, A^2, A^4, A^8, \dots

• Complexity = $\Theta(n^3 \log n) \Rightarrow$ not optimal, but has a high degree of parallelism



eg: $n = 20$

$$D^{(0)} \leftarrow A$$

$$A^1 = \begin{pmatrix} 0 & 2 & 3 & \infty & \infty & \infty & \infty & \infty & \infty \\ \infty & 0 & \infty & \infty & \infty & 1 & \infty & \infty & \infty \\ \infty & \infty & 0 & 1 & 2 & \infty & \infty & \infty & \infty \\ \infty & \infty & \infty & 0 & \infty & \infty & 2 & \infty & \infty \\ \infty & \infty & \infty & \infty & 0 & \infty & \infty & \infty & \infty \\ \infty & \infty & \infty & \infty & \infty & 0 & 2 & 3 & 2 \\ \infty & \infty & \infty & \infty & 1 & \infty & 0 & 1 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & 0 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & \infty & 1 & 0 \end{pmatrix}$$

$$A^2 = \begin{pmatrix} 0 & 2 & 3 & 4 & 5 & 3 & \infty & \infty & \infty \\ \infty & 0 & \infty & \infty & \infty & 1 & 3 & 4 & 3 \\ \infty & \infty & 0 & 1 & 2 & \infty & 3 & \infty & \infty \\ \infty & \infty & \infty & 0 & 3 & \infty & 2 & 3 & \infty \\ \infty & \infty & \infty & \infty & 0 & \infty & \infty & \infty & \infty \\ \infty & \infty & \infty & \infty & 3 & 0 & 2 & 3 & 2 \\ \infty & \infty & \infty & \infty & 1 & \infty & 0 & 1 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & 0 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & \infty & 1 & 0 \end{pmatrix}$$

$$A^4 = \begin{pmatrix} 0 & 2 & 3 & 4 & 5 & 3 & 5 & 6 & 5 \\ \infty & 0 & \infty & \infty & 4 & 1 & 3 & 4 & 3 \\ \infty & \infty & 0 & 1 & 2 & \infty & 3 & 4 & \infty \\ \infty & \infty & \infty & 0 & 3 & \infty & 2 & 3 & \infty \\ \infty & \infty & \infty & \infty & 0 & \infty & \infty & \infty & \infty \\ \infty & \infty & \infty & \infty & 3 & 0 & 2 & 3 & 2 \\ \infty & \infty & \infty & \infty & 1 & \infty & 0 & 1 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & 0 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & \infty & 1 & 0 \end{pmatrix}$$

$$A^8 = \begin{pmatrix} 0 & 2 & 3 & 4 & 5 & 3 & 5 & 6 & 5 \\ \infty & 0 & \infty & \infty & 4 & 1 & 3 & 4 & 3 \\ \infty & \infty & 0 & 1 & 2 & \infty & 3 & 4 & \infty \\ \infty & \infty & \infty & 0 & 3 & \infty & 2 & 3 & \infty \\ \infty & \infty & \infty & \infty & 0 & \infty & \infty & \infty & \infty \\ \infty & \infty & \infty & \infty & 3 & 0 & 2 & 3 & 2 \\ \infty & \infty & \infty & \infty & 1 & \infty & 0 & 1 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & 0 & \infty \\ \infty & \infty & \infty & \infty & \infty & \infty & \infty & \infty & 1 & 0 \end{pmatrix}$$

Figure 7.7 An example of the matrix-multiplication-based all-pairs shortest paths algorithm.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

• On a HC w/ n^3 processors, $n \times n$ matrix mult in $\Theta(\log n)$ [see MM chapt.]

$$S = \frac{\Theta(n^3)}{\Theta(\log^2 n)} ; E = \frac{1}{\Theta(\log^2 n)}$$

FLOYD's algorithm

$$d_{ij}^{(k)} = \begin{cases} w(v_i, v_j) & \text{if } k=0 \\ \min \{ d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)} \} & \text{if } k \geq 1 \end{cases}$$

Solve bottom-up in the order of increasing values of k ; $\Theta(n^3)$
 space = $\Theta(n^2)$

$D^{(0)} = A$
 for $k=1$ to n do
 for $i=1$ to n do
 for $j=1$ to n do

$$d_{ij}^{(k)} := \min (d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)})$$

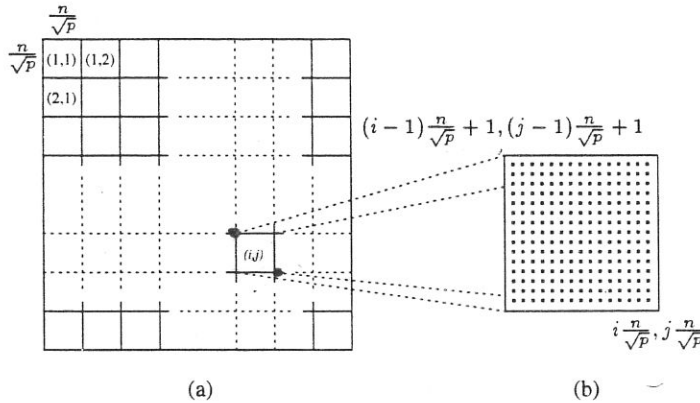
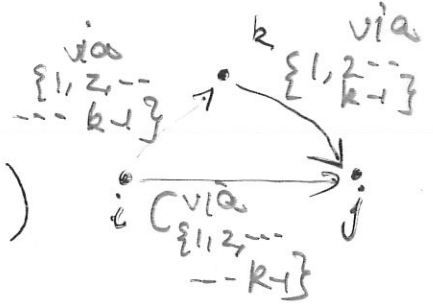


Figure 7.9 (a) Matrix $D^{(k)}$ partitioned by block checkerboarding into $\sqrt{p} \times \sqrt{p}$ subblocks, and (b) the square subblock of $D^{(k)}$ assigned to processor $P_{i,j}$.
 Copyright (r) 1994 Benjamin/Cummings Publishing Co.

Total of p blocks

Parallel formulations: to compute portion of $D^{(k)}$, need corresp. segments of k^{th} row and column of $D^{(k-1)}$

- k^{th} iteration: each of the \sqrt{P} processors containing part of the k^{th} row (column) send it to the $\sqrt{P}-1$ processors in the same column (row)
 - embed mesh into p-proc HC with cut-through routing
 - $\sqrt{P} \times \sqrt{P}$ virtual mesh \leftrightarrow p-proc HC
 - each row or column of mesh \equiv HC of \sqrt{P} procs
 - in each iteration, k^{th} row & k^{th} column do $1 \rightarrow$ all BC along a column or row
(each such proc has n/\sqrt{P} elements of k^{th} row or col)
- \Rightarrow BC requires $\Theta((n \log P)/\sqrt{P})$ time

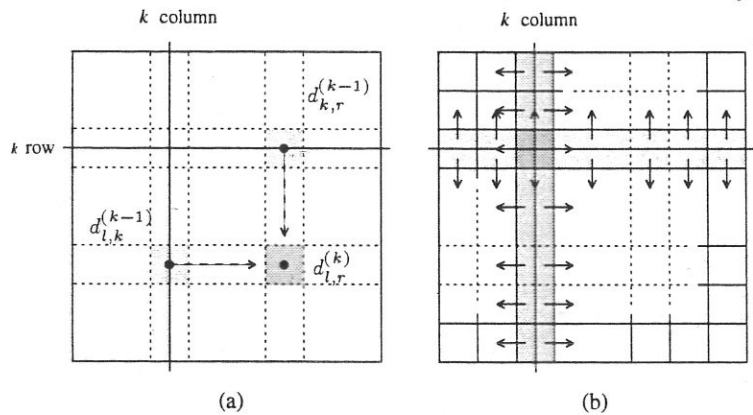


Figure 7.10 (a) Communication patterns used in the block-checkerboard partitioning. When computing $d_{i,j}^{(k)}$, information must be sent to the highlighted processor from two other processors along the same row and column. (b) The row and column of \sqrt{P} processors that contain the k^{th} row and column send them along processor columns and rows.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

→ $\Theta(n^2/P)$ to compute values of $D^{(k)}$ in block

$$T_p = \Theta(n^3/P) + \Theta\left(\frac{n^2}{\sqrt{P}} \log P\right)$$

$$S = \frac{\Theta(n^3)}{\Theta(n^3/P) + \Theta((n^2 \log P)/\sqrt{P})}$$

$$E = \frac{1}{1 + \Theta((\sqrt{P} \log P)/n)}$$

with pipelining, $T_p = \Theta\left(\frac{n^3}{P}\right) + \Theta(n)$

For cost-optimality, $\sqrt{P} \log P/n = O(1)$

$$\text{Isoefficiency}_{\text{comm}} = \Theta(P^{1.5} \log^3 P)$$

$$\text{Isoefficiency}_{\text{compute concurrency}} = \Theta(P^{1.5})$$

$$W = \Theta(n^3); \# \text{ procs} = O(n^2) = P$$