

PARALLEL SORTING

- where the i/p & o/p sequences are stored
- how comparisons are performed
- 1 element/processor :
compare-exchange step: $(t_s + t_w)$. As $t_s \gg t_w$, poor performance
- multiple elements/processor :
compare split step: $(t_s + t_w \frac{n}{p})$
For large blocks $\Theta(n/p)$ to merge 2 sorted blocks




Figure 6.1 A parallel compare-exchange operation. Processors P_i and P_j send their elements to each other. Processor P_i keeps $\min\{a_i, a_j\}$, and P_j keeps $\max\{a_i, a_j\}$.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

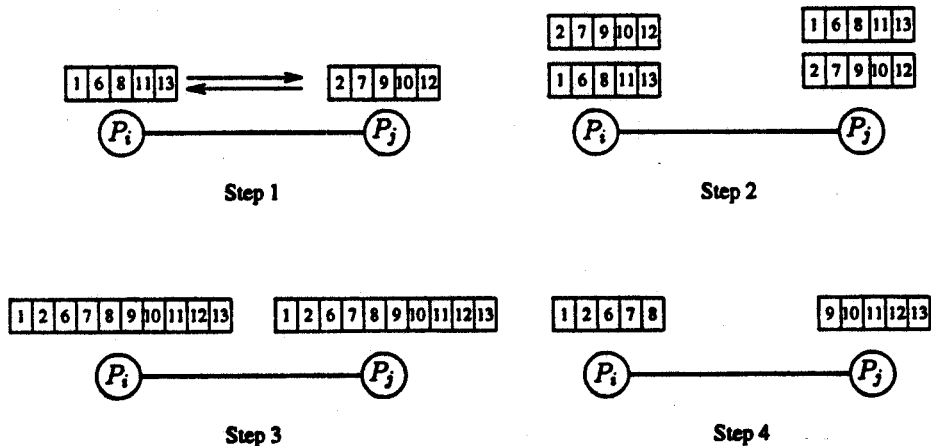


Figure 6.2 A compare-split operation. Each processor sends its block of size n/p to the other processor. Each processor merges the received block with its own block and retains only the appropriate half of the merged block. In this example, processor P_i retains the smaller elements and processor P_j retains the larger elements.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

SORTING NETWORKS

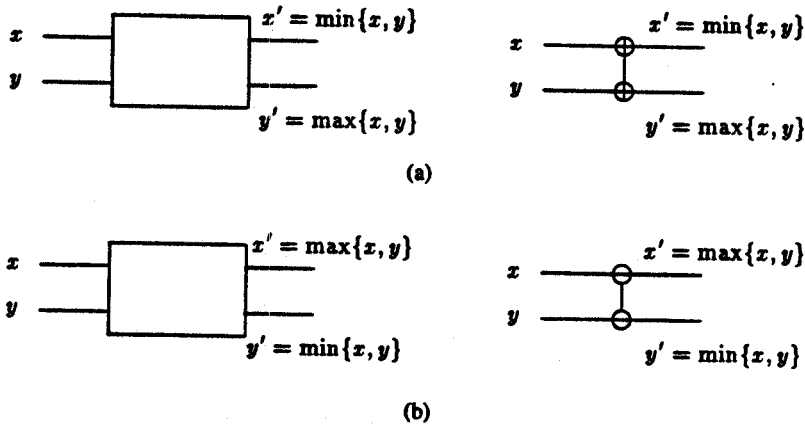


Figure 6.3 A schematic representation of comparators: (a) an increasing comparator, and (b) a decreasing comparator.
 Copyright (r) 1994 Benjamin/Cummings Publishing Co.

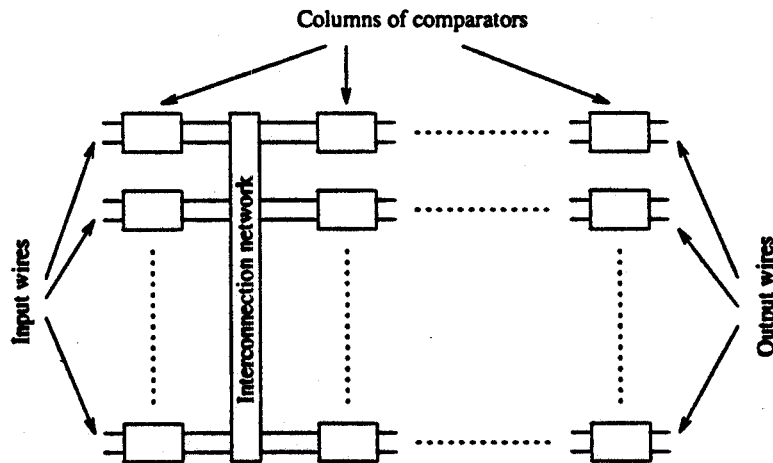


Figure 6.4 A typical sorting network. Every sorting network is made up of a series of columns, and each column contains a number of comparators connected in parallel.
 Copyright (r) 1994 Benjamin/Cummings Publishing Co.

BITONIC SORT : $\Theta(\log^2 n)$

• bitonic sequence : $\uparrow\downarrow$ or $\downarrow\uparrow$ or a cyclic shift thereof

• let $s = \langle \underbrace{a_0, a_1, \dots, a_{n/2-1}}_{\uparrow}, \underbrace{a_{n/2}, \dots, a_{n-1}}_{\downarrow} \rangle$

$s_1 = \langle \min\{a_0, a_{n/2}\}, \min\{a_1, a_{n/2+1}\}, \dots, \min\{a_{n/2-1}, a_{n-1}\} \rangle$


$s_2 = \langle \max\{ \quad \}, \max\{ \quad \}, \dots, \max\{ \quad \} \rangle$

Claim: s_1 & s_2 are bitonic & each element of s_1 is $<$ each element of s_2

• BITONIC SPLIT : split bitonic seq of size n into 2 bitonic subseqs

• BITONIC MERGE : sort bitonic seq using bitonic splits. RECURSION
 $\log n$ steps

Use bitonic merging network



Original																
sequence	3	5	8	9	10	12	14	20	95	90	60	40	35	23	18	0
1st Split	3	5	8	9	10	12	14	0	95	90	60	40	35	23	18	20
2nd Split	3	5	8	0	10	12	14	9	35	23	18	20	95	90	60	40
3rd Split	3	0	8	5	10	9	14	12	18	20	35	23	60	40	95	90
4th Split	0	3	5	8	9	10	12	14	18	20	23	35	40	60	90	95

Figure 6.5 Merging a 16-element bitonic sequence through a series of $\log 16$ bitonic splits.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

• To sort n numbers, repeatedly merge bitonic sequences of \uparrow length (seq of 2 elements is bitonic)

→ Each stage of network merges adjacent bitonic seqs in \uparrow or \downarrow order
⇒ seq. obtained by concatenating these seqs is bitonic

• last stage has depth $\log n$ [for bitonic merge with n inputs]

• Depth of network $d(n) = d\left(\frac{n}{2}\right) + \log n = \sum_{i=1}^{\log n} (\log^2 n + \log n) / 2$




Figure 6.6 A bitonic merging network for $n = 16$. The input wires are numbered $0, 1, \dots, n - 1$, and the binary representation of these numbers is shown. Each column of comparators is drawn separately; the entire figure represents a $\oplus\text{BM}[16]$ bitonic merging network. The network takes a bitonic sequence and outputs it in sorted order.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

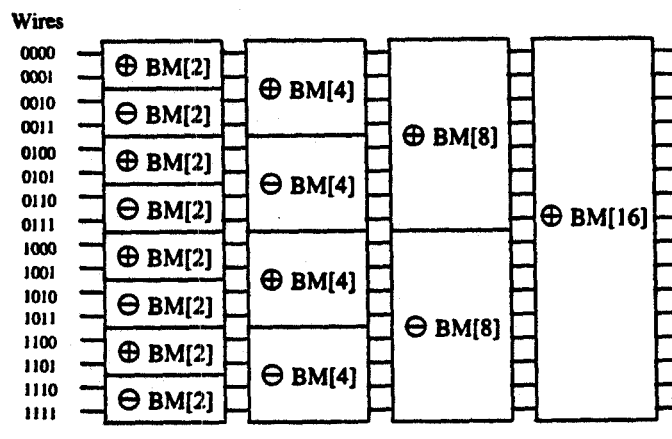


Figure 6.7 A schematic representation of a network that converts an input sequence into a bitonic sequence. In this example, $\oplus\text{BM}[k]$ and $\ominus\text{BM}[k]$ denote bitonic merging networks of input size k that use \oplus and \ominus comparators, respectively. The last merging network ($\oplus\text{BM}[16]$) sorts the input. In this example, $n = 16$.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

First 3 stages in Fig 6.8

Last stage in Fig 6.6

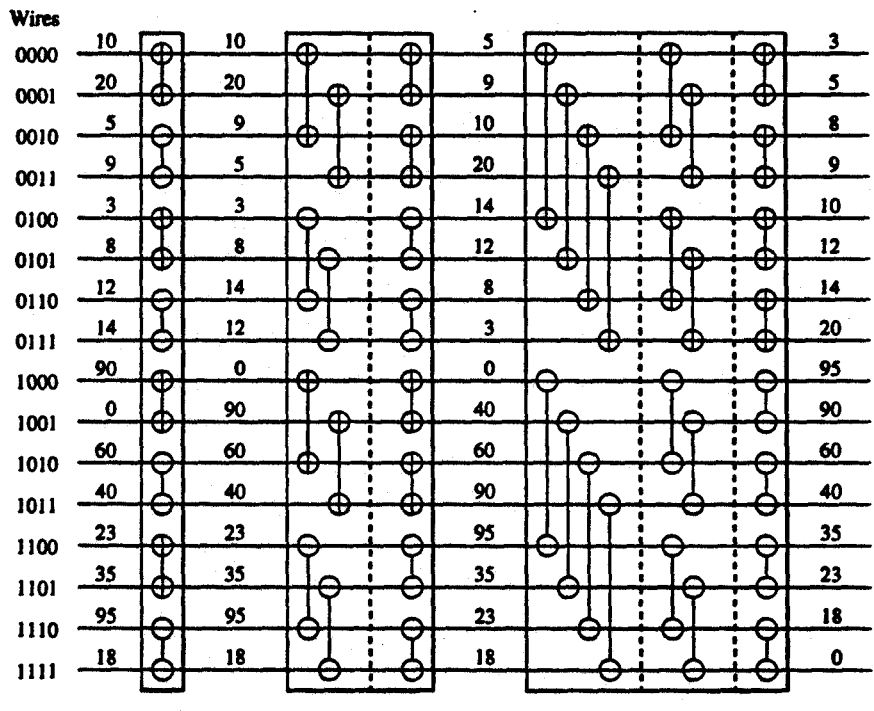


Figure 6.8 The comparator network that transforms an input sequence of 16 unordered numbers into a bitonic sequence. In contrast to Figure 6.6, the columns of comparators in each bitonic merging network are drawn in a single box, separated by a dashed line.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

HC: one element/processor

→ In any step, compare-exchange performed bet. 2 wires if labels differ in 1 bit

→ Last step of each stage: labels differing in LSB

→ Last 3 stages, 2nd-last step: labels differing in 2nd LSB.

⇒ Labels differing in i^{th} -LSB: perform compare-exchange $(\log n - i + 1)$ times

• On HC: i^{th} step of final stage: communicate along $(d - (i - 1))^{\text{st}}$ dimension

• $(1 + \log n)(\log n)/2$ steps ⇒ cost-optimal w.r.t. sequential implementation of bitonic sort




Figure 6.9 Communication during the last stage of bitonic sort. Each wire is mapped to a hypercube processor; each connection represents a compare-exchange between processors.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

```

1 procedure BITONIC_SORT (label, d)
2 begin
3   for i := 0 to d-1 do
4     for j := i downto 0 do
5       if  $(i+1)^{\text{st}}$  bit of label  $\neq j^{\text{th}}$  bit of label then
6         comp_exchange_max(j);
7       else
8         comp_exchange_min(j);
9   end BITONIC_SORT

```

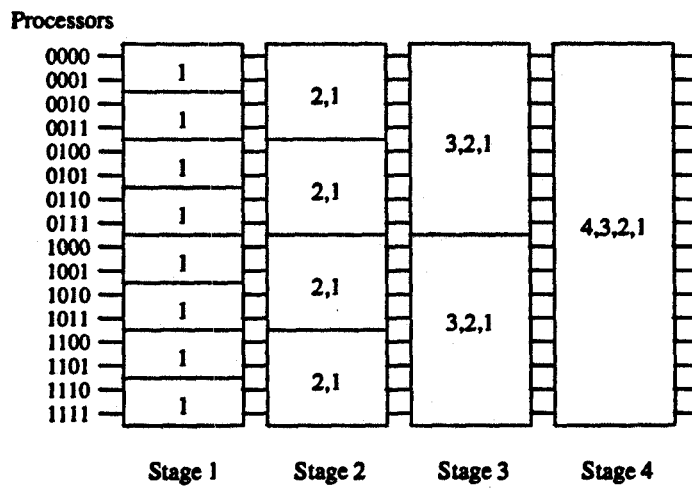


Figure 6.10 Communication characteristics of bitonic sort on a hypercube. During each stage of the algorithm, processors communicate along the dimensions shown.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

- Mesh: impossible to map such that each compare-exchange between neighbors
- Row-major shuffled mapping: processors that do c-e are on square subsections whose size is inversely related to frequency of c-e's
- Wires that differ in i^{th} LSB mapped to mesh processors
 $2^{\lfloor (i-1)/2 \rfloor}$ links away
- Communication overhead performed per processor = $\sum_{i=1}^{\log n} \sum_{j=1}^i 2^{\lfloor (j-1)/2 \rfloor} \approx 7\sqrt{n} = \Theta(\sqrt{n})$

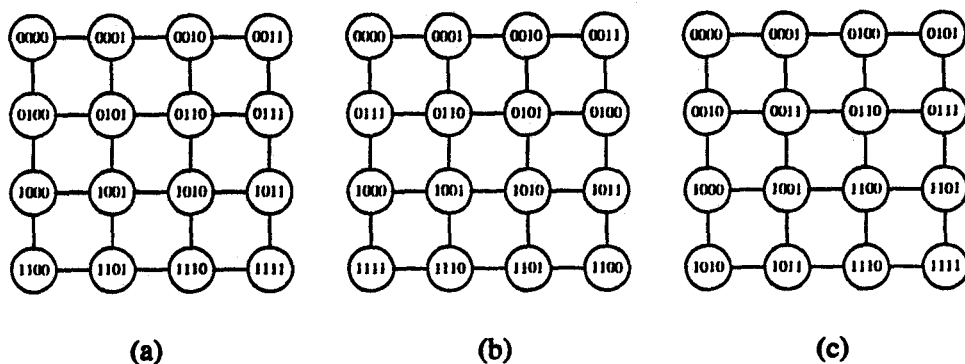


Figure 6.11 Different ways of mapping the input wires of the bitonic sorting network to a mesh of processors: (a) row-major mapping, (b) row-major snakelike mapping, and (c) row-major shuffled mapping.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

- In each of $\Theta(\log^2 n)$ steps, each processor performs max. 1 comparison
 \Rightarrow Parallel run time $T_p = \Theta(\log^2 n) + \Theta(\sqrt{n})$
- \Rightarrow Processor-time product = $\Theta(n\sqrt{n}) \Rightarrow$ not cost-optimal!
- Run time of sorting on mesh bounded by $\Omega(\sqrt{n})$. Why?
 \therefore parallel formulation is asymptotically optimal for mesh architecture

Block of elements per processor

Hypercube: $T_p = \theta\left(\frac{n}{p} \log \frac{n}{p}\right) + \theta\left(\frac{n}{p} \log^2 p\right) + \theta\left(\frac{n}{p} \log^2 p\right)$

local sort

comparisons

communication

Mesh: $T_p = \theta\left(\frac{n}{p} \log \frac{n}{p}\right) + \theta\left(\frac{n}{p} \log^2 p\right) + \theta\left(\frac{n}{p} \sqrt{p}\right)$




Figure 6.12 The last stage of the bitonic sort algorithm for $n = 16$ on a mesh, using the row-major shuffled mapping. During each step, processor pairs compare-exchange their elements. Arrows indicate the pairs of processors that perform compare-exchange operations.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

HC: $S = \frac{\theta(n \log n)}{\theta\left(\left(\frac{n}{p} \log \frac{n}{p}\right) + \theta\left(\frac{n}{p} \log^2 p\right)\right)}$

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

For cost optimality

$\left. \begin{aligned} &(\log^2 p) / \log n = O(1) \\ &\Rightarrow P = \theta\left(2^{\sqrt{\log n}}\right) \text{ procs} \end{aligned} \right\}$

Isoefficiency function = $\theta\left(p^{\log p} \log^2 p\right)$

Mesh: $S = \frac{\theta(n \log n)}{\theta\left(\frac{n}{p} \log \frac{n}{p}\right) + \theta\left(\frac{n}{p} \log^2 p\right) + \theta\left(\frac{n}{\sqrt{p}}\right)}$

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

$\left. \begin{aligned} &\text{For cost optimality,} \\ &\sqrt{p} / \log n = O(1) \Rightarrow \\ &P = \theta(\log^2 n) \end{aligned} \right\}$
 Isoefficiency function = $\theta\left(2^{\sqrt{p}} \sqrt{p}\right)$

- Bubble-sort $\Theta(n^2)$ inherently sequential \therefore use odd-even X position sort which has n phases, each w/ $\Theta(n)$ comparisons
Use RING. Processor-time product is $\Theta(n^2) \Rightarrow$ NOT cost-optimal
- For cost optimality, each proc has n/p elements

$$T_p = \underbrace{\Theta\left(\frac{n}{p} \log \frac{n}{p}\right)}_{\text{local sort}} + \underbrace{\Theta(n)}_{\text{comparisons}} + \underbrace{\Theta(n)}_{\text{communication}}$$




Figure 6.13 Sorting $n = 8$ elements, using the odd-even transposition sort algorithm. During each phase, $n = 8$ elements are compared.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

$$S = \frac{\Theta(n \log n)}{\Theta\left(\frac{n}{p} \log \frac{n}{p}\right) + \Theta(n)}$$

$$E = \frac{1}{1 - \Theta\left(\frac{\log p}{\log n}\right) + \Theta\left(\frac{p}{\log n}\right)}$$

- cost optimal when $p = O(\log n)$
- Isoefficiency function is $\Theta(p \log p)$

QUICKSORT: Worst case $T(n) = T(n-1) + \theta(n) \Rightarrow T(n) = \theta(n^2)$

$T(n) = 2T(n/2) + \theta(n) \Rightarrow T(n) = \theta(n \log n) \Rightarrow$ depends on pivot selection

QUICKSORT(A, q, r)

if $q < r$ then

$x := A[q]$

$s := q$

for $i := q+1$ to r do

if $A[i] \leq x$ then

$s := s+1$

swap(A[s], A[i])

swap(A[q], A[s])

QUICKSORT(A, q, s)

QUICKSORT(A, s+1, r)

Naive parallel formulation: give a subproblem to another processor

Drawback: partitioning done by single processor $\Rightarrow \Omega(n)$

\therefore important to do partitioning in parallel




Figure 6.15 Example of the quicksort algorithm sorting a sequence of size $n = 8$.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

Parallel Formulation for CRCW PRAM

→ construct a binary tree; sorted sequence is its in-order traversal




Figure 6.16 A binary tree generated by the execution of the quicksort algorithm. Each level of the tree represents a different array-partitioning iteration. If pivot selection is optimal, then the height of the tree is $\Theta(\log n)$, which is also the number of iterations.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

```

1. procedure BUILD.TREE (A[1...n])
2. begin
3.   for each processor  $i$  do
4.     begin
5.        $root := i$ ;
6.        $parent_i := root$ ;
7.        $leftchild[i] := rightchild[i] := n + 1$ ;
8.     end for
9.     repeat for each processor  $i \neq root$  do
10.      begin
11.        if ( $A[i] < A[parent_i]$ ) or
12.           ( $A[i] = A[parent_i]$  and  $i < parent_i$ ) then
13.          begin
14.             $leftchild[parent_i] := i$ ;
15.            if  $i = leftchild[parent_i]$  then exit
16.            else  $parent_i := leftchild[parent_i]$ ;
17.          end for
18.        else
19.          begin
20.             $rightchild[parent_i] := i$ ;
21.            if  $i = rightchild[parent_i]$  then exit
22.            else  $parent_i := rightchild[parent_i]$ ;
23.          end else
24.        end repeat
25.      end BUILD.TREE

```

Program 6.6 The binary tree construction procedure for the CRCW PRAM parallel quicksort formulation.

$A[root]$ is used as first pivot
 \rightarrow copied into local $parent_i$

Processors with elements $< A[parent_i]$ write their labels into $LC[parent_i]$
 etc.

- Algorithm continues until n pivot elements are selected
- Processor exits when its element becomes pivot
- In each iteration, 1 level of tree is constructed in $\Theta(1)$ time
 Avg complexity of binary tree building = $\Theta(\log n)$
- To traverse tree $\Theta(\log n)$ on n -processor PRAM

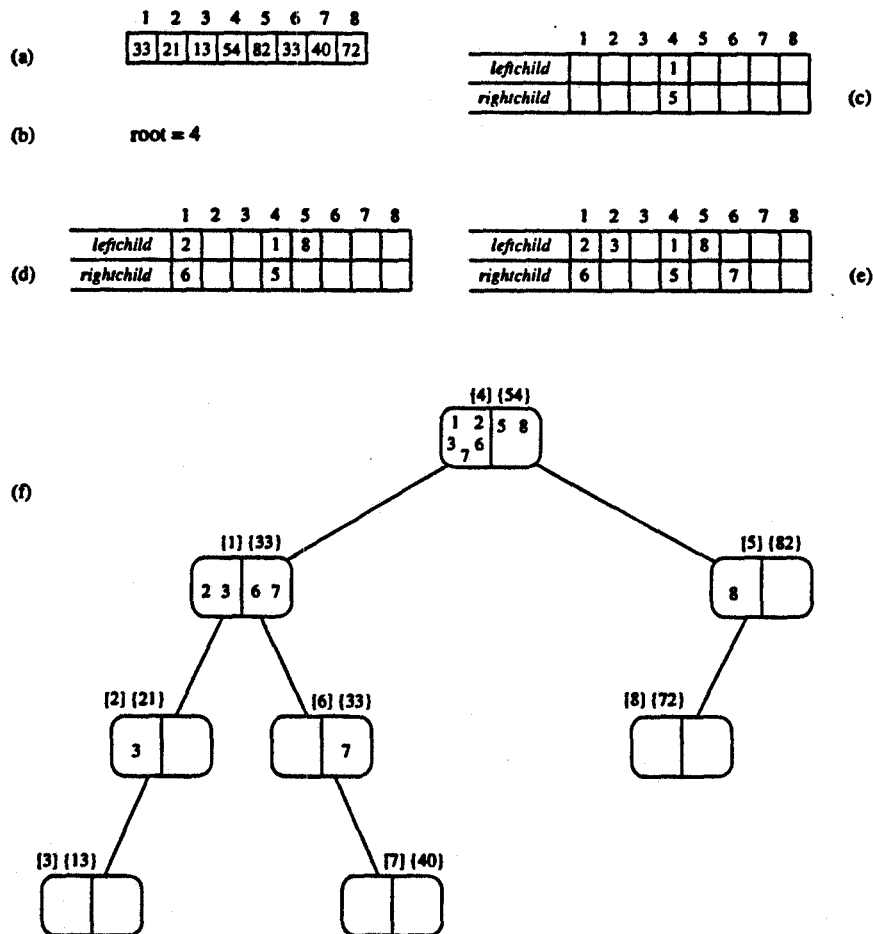


Figure 6.17 The execution of the PRAM algorithm on the array shown in (a). The arrays *leftchild* and *rightchild* are shown in (c), (d), and (e) as the algorithm progresses. Figure (f) shows the binary tree constructed by the algorithm. Each node is labeled by the processor (in square brackets), and the element is stored at that processor (in curly brackets). The element is the pivot. In each node, processors with smaller elements than the pivot are grouped on the left side of the node, and those with larger elements are grouped on the right side. These two groups form the two partitions of the original array. For each partition, a pivot element is selected at random from the two groups that form the children of the node.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

Quicksort on HC:

d steps: $1 \leq i \leq d$ } pivot selection is critical!

select a pivot
 BC the pivot
 Xchange lower/higher elements along dimension i

$$T_p = \underbrace{\Theta\left(\frac{n}{p} \log \frac{n}{p}\right)}_{\text{local sort}} + \underbrace{\Theta\left(\frac{n}{p} \log p\right)}_{\text{communication}} + \underbrace{\Theta(\log^2 p)}_{\text{pivot broadcasting}}$$

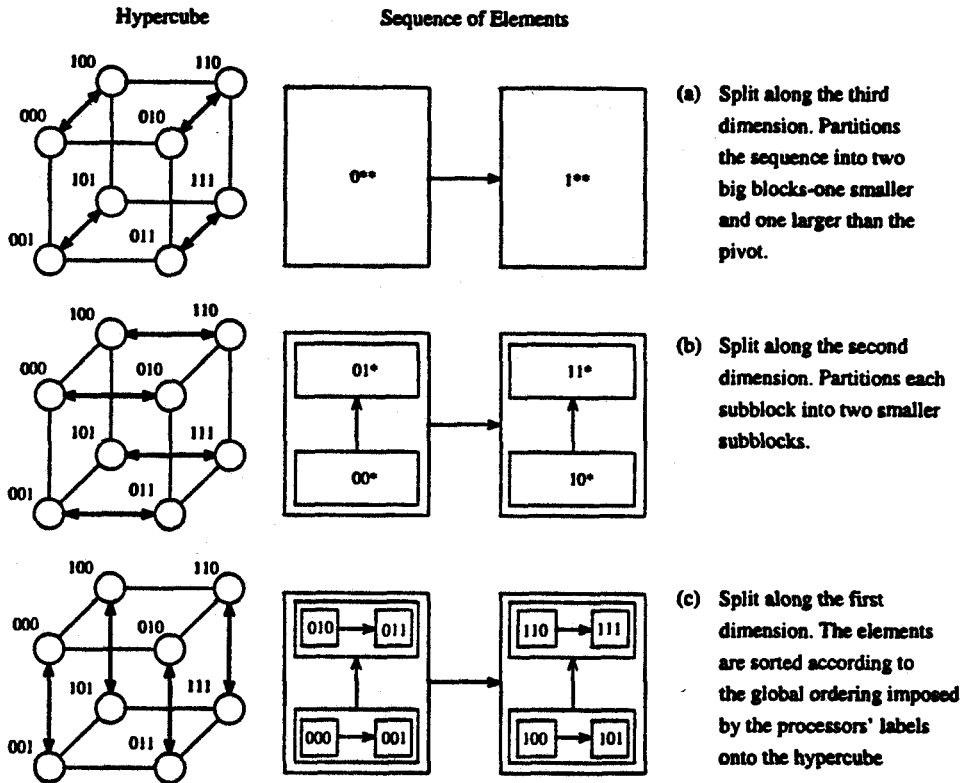


Figure 6.18 The execution of the hypercube formulation of quicksort for $d = 3$. The three splits—one along each communication link—are shown in (a), (b), and (c). The second column represents the partitioning of the n -element sequence into subcubes. The arrows between subcubes indicate the movement of larger elements. Each box is marked by the binary representation of the processor labels in that subcube. A * denotes that all the binary combinations are included.

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

$$S = \frac{\Theta(n \log n)}{\Theta\left(\frac{n}{p} \log \frac{n}{p}\right) + \Theta\left(\frac{n}{p} \log p\right) + \Theta(\log^2 p)}$$

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

For cost optimality
 $(p \log^2 p) / (n \log n) = O(1)$
 $\Rightarrow p = \Theta(n / \log n)$
 Isoefficiency = $\Theta(p \log^2 p)$