

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

SORTING ON A MESH-CONNECTED PARALLEL COMPUTER

C. D. Thompson and H. T. Kung

March 1976

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, PA 15213

ABSTRACT

Two algorithms are presented for sorting n^2 elements on an $n \times n$ mesh-connected processor array that require $O(n)$ routing and comparison steps. The best previous algorithm takes time $O(n \log n)$. The algorithms of this paper are shown to be optimal in time within small constant factors. Extensions to higher-dimensional mesh-connected processor arrays are also given.

This research was supported in part by the National Science Foundation under Grant MCS75-222-55 and the Office of Naval Research under Contract N00014-76-C-0370, NR 044-422.

1. Introduction

In the course of a parallel computation, individual processors will need to distribute their results to other processors and complicated data flow problems may arise. One way to handle this problem is by sorting "destination tags" attached to each data element, as discussed in Batcher [1968]. Hence efficient sorting algorithms for parallel machines with some fixed processor interconnection pattern are relevant to almost any use of these machines.

In this paper we present two algorithms for sorting $N = n^2$ elements on an $n \times n$ mesh-type processor array that require $O(n)$ unit-distance routing steps and $O(n)$ comparison steps (n is assumed to be a power of 2). The best previous algorithm takes time $O(n \log n)$ (Orcutt [1974]). One of our algorithms, the s^2 -way merge sort, is shown optimal within a factor of 2 in time for sufficiently large n , if one comparison step takes no more than twice the time of a routing step. Our other $O(n)$ algorithm, an adaptation of Batcher's bitonic merge sort, is much less complex but optimal under the same assumption to within a factor of 4.5 for all n , and is more efficient for moderate n .

We believe that the algorithms of this paper will give the most efficient sorting algorithms for ILLIAC IV-type parallel computers.

Our algorithms can be generalized to higher-dimensional array interconnection patterns. For example, our second algorithm can be modified to sort N elements on a j -dimensionally mesh-connected N -processor computer in $O(N^{1/j})$ time, which is optimal within a small constant factor.


Efficient sorting algorithms have been developed for interconnection patterns other than the "mesh" considered in this paper. Stone [1971] maps Batcher's bitonic merge sort onto the "perfect shuffle" interconnection scheme, obtaining an N -element sort time of $O(\log^2 N)$ on N processors. The odd-even transposition sort (see Appendix) requires an optimal $O(N)$ time on a linearly connected N -processor computer. Sorting time is thus seen to be strongly dependent on the interconnection pattern. Exploration of this dependence for a given problem is of interest from both an architectural and an algorithmic point of view.

In Section 2 we give the model of computation. The sorting problem is defined precisely in Section 3. A lower bound on the sorting time is given in Section 4. Batcher's 2-way odd-even merge is mapped on our 2-dimensional mesh-connected processor array in the next section. Generalizing the 2-way odd-even merge, we introduce a $2s$ -way merge algorithm in Section 6. This is further generalized to an s^2 -way merge in Section 7, from which our most efficient sorting algorithm for large n is developed. Section 8 shows that Batcher's bitonic sort can be performed efficiently on our model by choosing an appropriate processor indexing scheme. Some extensions and implications of our results are discussed in Section 9. The Appendix contains a description of the odd-even transposition sort.

2. Model of Computation

We assume a parallel computer with $N = nxn$ identical processors. The architecture of the machine is similar to that of the ILLIAC IV (Barnes, et. al. [1968]). The major assumptions are as follows:

- (i) The interconnections between the processors are a subset of those on the ILLIAC IV, and are defined by the following two dimensional array:



where the p's denote the processors. That is, each processor is connected to all its neighbors. Processors at the perimeter have two or three rather than four neighbors; there are no "wrap-around" connections as found on the ILLIAC IV.

The bounds obtained in this paper would be affected at most by a factor of 2 if "wrap-around" connections were included, but we feel that this addition would obscure the ideas of this paper without substantially strengthening the results.

- (ii) It is a SIMD (Single Instruction stream Multiple Data stream) machine (Flynn[1966]). During each time unit, a single instruction is broadcast to all processors, but only executed by the set of processors specified in the instruction. For the purpose of the paper, only two instruction types are needed: the routing instruction for interprocessor data moves, and the comparison instruction on two data elements in each processor. The comparison instruction is a conditional interchange on the contents of two registers in each processor. Actually, we need both "types" of such comparison instructions to allow either register to receive the minimum; normally both types will be issued during "one comparison step".

- (iii) Define

t_R = time required for one unit-distance routing step, i.e., moving one item from a processor to one of its neighbors,

t_C = time required for one comparison step.


For example, a comparison-interchange step between two items in adjacent processors can be done in time $2t_P + t_C$. Of course, concurrent data movement is allowed, so long as it is all in the same direction; and also any number (up to N) of concurrent comparisons may be performed simultaneously.

3. The Sorting Problem

The processors may be indexed by any function that is a one-to-one mapping from $\{1, 2, \dots, n\} \times \{1, 2, \dots, n\}$ onto $\{0, 1, \dots, N-1\}$. Assume that N elements from a linearly ordered set are initially loaded in the N processors, each receiving exactly one element. With respect to any index function, the sorting problem is defined to be the problem of moving the j th smallest element to the processor indexed by j for all $j = 0, 1, \dots, N-1$.


Example 3.1.

Suppose that $n = 4$ (hence $N = 16$) and that we want to sort 16 elements initially loaded as follows:




Three ways of indexing the processors will be considered in this paper.

(i) Row-Major Indexing: After sorting we have




(ii) Shuffled Row-Major Indexing: After sorting we have



Note that this indexing is obtained by shuffling the binary representation of the row-major index. For example, the row-major index 5 has the binary representation 0101. Shuffling the bits gives 0011 which is 3. (In general, the shuffled binary number, say, "abcdefgh" is "aebfcgdh".)

(iii) Snake-Like Row-Major Indexing: After sorting we have



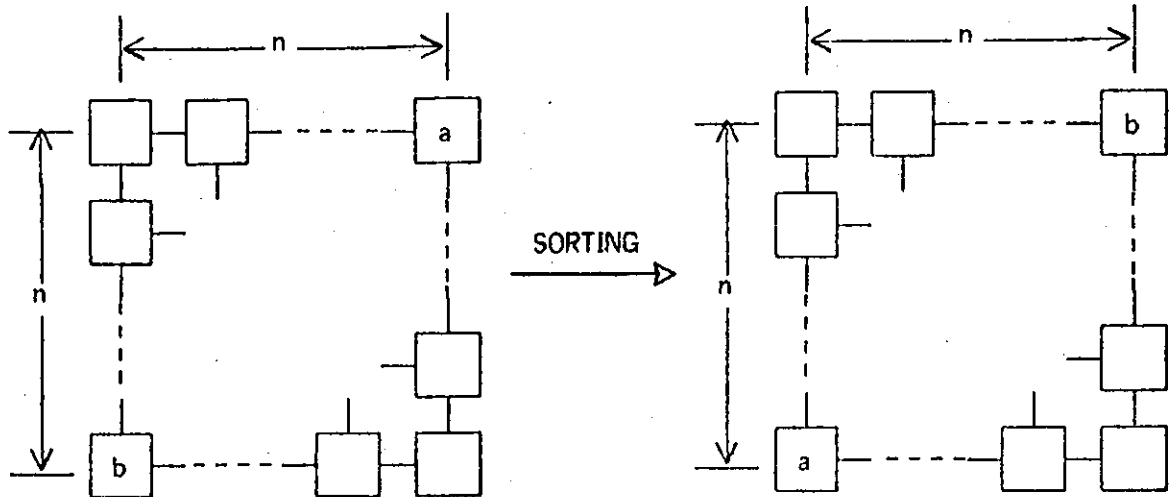
This indexing is obtained from the row-major indexing by reversing the ordering in even rows.

The choice of a particular indexing scheme depends upon how the sorted elements will be used (or accessed), and upon which sorting algorithm is to be used. For example, we found that the row-major indexing is poor for merge-sorting.

It is clear that the sorting problem with respect to any index scheme can be solved by using the routing and comparison steps. We are interested in designing algorithms which minimize the time spent in routing and comparing.

4. A Lower Bound

Observe that for any index scheme there are situations where the two elements initially loaded at the opposite corner processors have to be transposed during the sorting.




It is easy to argue that even for this simple transposition we need at least $4(n-1)$ unit-distance routing steps. This implies that no algorithm can sort n^2 elements in time less than $O(n)$. In this paper, we shall show two algorithms which can sort n^2 elements in time $O(n)$. One will be developed in Sections 5 through 7, the other in Section 8.


5. The 2-Way Odd-Even Merge

Batcher's odd-even merge (Batcher[1968], Knuth [1973, pp. 224-226]) of two sorted sequences $\{u(i)\}$ and $\{v(i)\}$ is performed in two stages. First, the "odd sequences" $\{u(1), u(3), u(5), \dots, u(2j+1), \dots\}$ and $\{v(1), v(3), \dots, v(2j+1), \dots\}$ are merged concurrently with the merging of the "even sequences" $\{u(2), u(4), \dots, u(2j), \dots\}$ and $\{v(2), v(4), \dots, v(2j), \dots\}$. Then the two merged sequences are interleaved, and a single parallel comparison-interchange step produces the sorted result. The merges in the first stage are done in the same manner (i.e., recursively).


We first illustrate how the odd-even method can be performed efficiently on linearly connected processors, then the idea is generalized to 2-dimensionally connected arrays. If two sorted sequences $\{1, 3, 4, 6\}$ and $\{0, 2, 5, 7\}$ are initially loaded in 8 linearly-connected processors, then Batcher's odd-even merge can be diagrammed as:




L1. Unshuffle: Odd-indexed elements to left, evens to right.




L2. Merge the "odd sequences" and the "even sequences".




L3. Shuffle.



L4. Comparison-interchange (the \vec{C} 's indicate comparison-interchanges).



Step L3 above is the "perfect shuffle" (Stone[1971]) and step L1 is its inverse, the "unshuffle". Note that the perfect shuffle can be achieved by using the triangular interchange pattern below:




where the \leftrightarrow 's indicate interchanges. Similarly, an inverted triangular interchange pattern will do the unshuffle. Therefore, both the perfect shuffle and unshuffle can be done in $k-1$ interchanges (i.e., $2k-2$ routing steps) when performed on a row of length $2k$ in our model.


We now give an implementation of the odd-even merge on a rectangular region of our model. Let $M(j,k)$ denote our algorithm of merging two j by $k/2$ sorted adjacent subarrays to form a sorted j by k array, where j,k are powers of 2, $k>1$, and all the arrays are arranged in the snake-like row major ordering. We first consider the case where $k=2$. If $j=1$, a single comparison-interchange step suffices to sort the two unit "sub-arrays". Given two sorted columns of length $j>1$, $M(j,2)$ consists of the following steps:

- J1. Move all odds to the left column and all evens to the right. Time: $2t_R$
- J2. Use the "odd-even transposition sort" (see Appendix) to sort each column. Time: $2jt_R + jt_C$
- J3. Interchange on even rows. Time: $2t_R$
- J4. One step of comparison-interchange (every "even" with the next "odd"). Time: $2t_R + t_C$


The following diagram illustrates the algorithm $M(j,2)$ for $j=4$.




For $k > 2$, $M(j,k)$ is defined recursively in the following way. Steps $M1$ and $M2$ unshuffle the elements, step $M3$ recursively merges the "odd sequences" and the "even sequences", steps $M4$ and $M5$ shuffle the "odds" and "evens" together, and step $M5$ performs the final comparison-interchange. The accompanying diagrams illustrate the algorithm $M(4,4)$, where the two given sorted 4 by 2 subarrays are initially stored in 16 processors as follows:




M1. Single Interchange step on even rows if $j > 2$, so that columns contain either all evens or all odds. If $j = 2$, do nothing: the columns are already segregated. Time: $2t_R$




M2. Unshuffle each row. Time: $(k-2)t_p$




M3. Merge by calling $M(j, k/2)$ on each half. Time: $T(j, k/2)$




M4. Shuffle each row. Time: $(k-2)t_p$



M5. Interchange on even rows. Time: $2t_R$



M6. Comparison-interchange of adjacent elements (every "even" with the next "odd").
Time: $4t_R + t_C$



Let $T(j,k)$ be the time needed by $M(j,k)$. Then we have

$$T(j,2) = (2j + 6)t_R + (j + 1)t_C$$

and for $k > 2$,

$$T(j,k) = (2k + 4)t_R + t_C + T(j,k/2).$$

These imply that

$$T(j,k) \leq (2j + 4k + 4\log k)t_R + (j + \log k)t_C.$$

(All logarithms in this paper are taken to base 2.)

An $n \times n$ sort may be composed of $M(j,k)$ by sorting all columns in $O(n)$ routes and compares by, say, the odd-even transposition sort, then using $M(n,2)$, $M(n,4)$, $M(n,8)$, ..., $M(n,n)$, for a total of $O(n \log n)$ routes and compares. This poor performance may be assigned to two inefficiencies in the algorithm. First, the recursive sub-problems ($M(n,n/2)$, $M(n,n/4)$, ..., $M(n,1)$) generated by $M(n,n)$ are not decreasing in size along both dimensions: they are all $O(n)$ in complexity. Second, the method is extremely "local" in the sense that no comparisons are made between elements initially in different halves of the array until the last possible moment, when each half has already been independently sorted.

The first inefficiency can be attacked by designing an "upwards" merge to complement the "sideways" merge just described. Even more powerful is the idea of combining many upwards merges with a sideways one. This idea is used in the next section.

6. The 2s-Way Merge

In this section we give an algorithm $M'(j,k,s)$ for merging $2s$ arrays of size j/s by $k/2$ in a j by k region of our processors, where j,k,s are powers of 2, $s \geq 1$, and the arrays are in the snake-like row-major ordering. The algorithm $M'(j,k,s)$ is almost the same as the algorithm $M(j,k)$ described in the previous section, except that $M'(j,k,s)$ requires a few more comparison-interchanges during step M6. These steps are exactly those performed in the initial portion of the odd-even transposition sort mapped onto our "snake" (see Appendix). More precisely, for $k > 2$, M1 and M6 are replaced by

M1'. Single interchange step on even rows if $j \geq s$, so that columns contain either all evens or all odds. If $j < s$, do nothing; the columns are already segregated. Time: $2t_R$

M6'. Perform the first $2s-1$ parallel comparison-interchange steps of the odd-even transposition sort on the "snake". It is not difficult to see that the time needed is at most

$$s(4t_R + t_C) + (s-1)(2t_R + t_C) = (6s-2)t_R + (2s-1)t_C.$$

Note that the original step M6 is just the first step of an odd-even transposition sort. Thus the 2-way merge is seen to be a special case of the $2s$ -way merge.

Similarly, for $M'(j,2,s)$, $j > s$, J4 is replaced by M6', which takes time

$$(2s-1)(2t_R + t_C).$$

$M'(s,2,s)$ is a special case analogous to $M(1,2)$, and may be performed by the odd-even transposition sort (see Appendix) in time $4st_R + 2st_C$.

The validity of this algorithm may be demonstrated by use of the 0-1 principle (Knuth [1973], p.224): if a network sorts all sequences of 0's and 1's, then it will sort any arbitrary sequence of elements chosen from a linearly ordered set. Thus, we may assume that the inputs are 0's and 1's. It is easy to check that there may be as many as $2s$ more zeros on the left as on the right after unshuffling (i.e., after step J1 or step M2). After the shuffling, the first $2s-1$ steps of an odd-even transposition sort suffice to sort the resulting array.

Let $T'(j,k,s)$ be the time required by the algorithm $M'(j,k,s)$. Then we have that

$$T'(j,2,s) \leq (2j + 4s + 2)t_R + (j + 2s - 1)t_C$$

and that for $k > 2$,

$$T'(j,k,s) \leq (2k + 6s - 2)t_R + (2s - 1)t_C + T'(j,k/2,s).$$

These imply that

$$T'(j,k,s) = (2j + 4k + (6s)\log k + O(s+\log k))t_P \\ + (j + (2s)\log k + O(s+\log k))t_C.$$

For $s = 2$, a merge sort may be derived that has the following time behavior

$$S'(n,n) = S'(n/2,n/2) + T'(n,n,2).$$

Thus,

$$S'(n,n) = (12n + O(\log^2 n))t_P + (2n + O(\log^2 n))t_C.$$

Suddenly, we have an algorithm that sorts in linear time. In the following section, the constants will be reduced by a factor of 2 by the use of a more complicated multi-way merge algorithm.

7. The s^2 -Way Merge

The s^2 -way merge $M''(j,k,s)$ to be introduced in this section is a generalization of the 2-way merge $M(j,k)$. Input to $M''(j,k)$ is s^2 sorted j/s by k/s arrays in a j by k region of our processors, where j,k,s are powers of 2 and $s > 1$. Steps M1 and M2 still suffice to move odd-indexed elements to the left and evens to the right so long as $j > s$ and $k > s$; $M''(j,s,s)$ is a special case analogous to $M(j,2)$ of the 2-way merge. Steps M1 and M6 are now replaced by

M1'': Single interchange step on even rows if $j \geq s$, so that columns contain either all evens or all odds. If $j = s$, do nothing: the columns are already segregated. Time: $2t_R$

M6'': Perform the first $s^2 - 1$ parallel comparison-interchange steps of the odd-even transposition sort on the "snake" (see Appendix). The time required for this is

$$(s^2/2)(4t_R + t_C) + (s^2/2 - 1)(2t_R + t_C) = (3s^2 - 2)t_R + (s^2 - 1)t_C.$$

The motivation for this new step comes from the realization that when the inputs are 0's and 1's, there may be as many as s^2 more zeros on the left half as the right after unshuffling.

$M''(j,s,s)$, $j \geq s$, can be performed in the following way:

- N1. $(\log s/2)$ 2-way merges: $M(j/s,2), M(j/s,4), \dots, M(j/s,s/2)$.
- N2. A single $2s$ -way merge: $M'(j,s,s)$.

If $T''(j,k,s)$ is the time taken by $M''(j,k,s)$, we have for $k = s$

$$T''(j,s,s) = (2j + O((s + j/s)\log s))t_R + (j + O((s + j/s)\log s))t_C$$

and for $k > s$,

$$T''(j,k,s) = (2k + 3s^2 + O(1))t_R + (s^2 + O(1))t_C + T''(j,k/2,s).$$

Therefore,

$$\begin{aligned} T''(j,k,s) &= (4k + 2j + 3s^2 \log(k/s) + O((s + j/s)\log s + \log k))t_R \\ &\quad + (j + s^2 \log(k/s) + O((s + j/s)\log s + \log k))t_C. \end{aligned}$$

A sorting algorithm may be developed from the s^2 -way merge; a good value for s is approximately $n^{1/3}$ (remember that s must be a power of 2). Then the time of sorting $n \times n$ elements satisfies

$$S''(n,n) = S''(n^{2/3}, n^{2/3}) + T''(n,n, n^{1/3}).$$

This leads immediately to the following result.

Theorem 7.1

If the snake-like row-major indexing is used, the sorting problem can be done in time

$$(6n + O(n^{2/3} \log n))t_R + (n + O(n^{2/3} \log n))t_C.$$

If $t_C \leq 2t_R$, Theorem 7.1 implies that $(6n + 2n + O(n^{2/3} \log n))t_R$ is sufficient time for sorting. In section 4, we showed that $4(n-1)t_P$ time is necessary. Thus, for large enough n , the s^2 -way algorithm is optimal to within a factor of 2. Preliminary investigation indicates that a careful implementation of the s^2 -way merge sort is optimal within a factor of 7 for all n , under the assumption that $t_C \leq 2t_R$.

8. The Bitonic Merge

In this section we shall show that Batcher's bitonic merge algorithm (Batcher [1968], Knuth [1973, pp. 232-237]) lends itself well to sorting on a mesh-connected parallel computer, once the proper indexing scheme has been selected. Two indexing schemes will be considered, the "row-major" and the "shuffled row-major" indexing schemes defined in Section 3.

The bitonic merge of two sections of a bitonic array of $j/2$ elements each takes $\log j$ passes, where pass i consists of a comparison-interchange between processors with indices differing only in the i^{th} bit of their binary representations. (This operation will be termed "comparison-interchange on the i^{th} bit".) Sorting an entire array of 2^k elements by the bitonic method requires k comparison-interchanges on the 0^{th} bit (the least significant bit), $k-1$ comparison-interchanges on the first bit, ..., $(k-1)$ comparison-interchanges on the i^{th} bit, ..., and 1 comparison-interchange on the most significant bit. For any fixed indexing scheme, in general a comparison-interchange on the j^{th} bit will take a different amount of time than when done on the j^{th} bit: an optimal processor indexing scheme for the bitonic algorithm minimizes the time spent on comparison-interchange steps. A necessary condition for optimality is that a comparison-interchange on the j^{th} bit be no more expensive than the $(j+1)^{\text{st}}$ bit for all j . If this were not the case for some j , then a better indexing scheme could immediately be derived from the supposedly optimal one by interchanging the j^{th} and the $(j+1)^{\text{st}}$ bits of all processor indices (since more comparison-interchanges will be done on the original j^{th} bit than on the $(j+1)^{\text{st}}$ bit).

The bitonic algorithm has been analyzed for the row-major indexing scheme: It takes


$$O(n \log n)t_R + O(\log^2 n)t_C$$

time to sort n^2 elements on n^2 processors (Orcutt [1974]). However, the row-major indexing scheme is decidedly non-optimal. For the case $n^2 = 64$, processor indices have six bits. A comparison-interchange on bit 0 takes just $2t_R + t_C$, for the processors are horizontally adjacent. A comparison-interchange on bit 1 takes $4t_R + t_C$, since the processors are two units apart. Similarly, a comparison-interchange on bit 2 takes $8t_R + t_C$, but a comparison-interchange on bit 3 takes only $2t_R + t_C$ because the processors are vertically adjacent. This phenomenon may be analyzed by considering the row-major index as the concatenation of a 'Y' and an 'X' binary vector: in the case $n^2 = 64$, the index is $Y_2Y_1Y_0X_2X_1X_0$. A comparison-interchange on X_i takes more time than one on X_j when $i > j$; however, a comparison-interchange on Y_i takes exactly the same time as on X_i . Thus a better indexing scheme may be derived by "shuffling" the 'X' and 'Y' vectors, obtaining (in the case $n^2 = 64$) $Y_2X_2Y_1X_1Y_0X_0$; this "shuffled row-major" indexing scheme satisfies our condition on optimality.


Geometrically, "shuffling" the 'X' and 'Y' vectors ensures that all arrays encountered in the merging process are nearly square, so that routing time will not be

excessive in either direction. The standard row-major indexing causes the bitonic sort to contend with sub-arrays that are always at least as wide as they are tall; the aspect ratio can be as high as n on an $n \times n$ processor array.


Programming the bitonic sort would be a little tricky, as the "direction" of a comparison-interchange step depends on the processor index. Orcutt [1974] covers these gory details for row-major indexing; his algorithm may easily be modified to handle the shuffled row-major indexing scheme. Here is an example of the bitonic merge sort on a 4×4 processor array for the shuffled row-major indexing; the comparison "directions" were derived from the following diagram (Knuth [1973], p. 237):



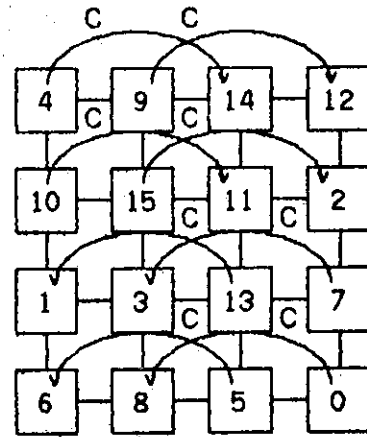

Initial data configuration:





Stage 1. Merge pairs of adjacent 1×1 matrices by the comparison-interchange indicated. Time: $2t_R + t_C$.




Stage 2. Merge pairs of 1x2 matrices; note that one member of a pair is sorted in ascending order, the other in descending order. This will always be the case in any bitonic merge. Time: $4t_R + 2t_C$.




Stage 3. Merge pairs of 2x2 matrices. Time: $8t_R + 3t_C$.





Stage 4. Merge the two 2x4 matrices. Time: $12t_R + 4t_C$.



Let $T'''(2^j)$ be the time to merge the bitonically sorted elements in processors #0 through $\#2^j-1$, where the shuffled row-major indexing is used. Then after one pass of comparison-interchange, which takes time $2^{\lceil i/2 \rceil} t_R + t_C$, the problem is reduced to the bitonic merge of elements in processors #0 through $\#2^{j-1}-1$, and that of elements in processors $\#2^{j-1}$ to $\#2^j-1$. It may be observed that the latter two merges can be done concurrently. Thus we have

$$T'''(1) = 0,$$

$$T'''(2^j) = T'''(2^{j-1}) + 2^{\lceil i/2 \rceil} t_R + t_C.$$

Hence

$$T^m(2^i) = \begin{cases} (3 \cdot 2^{(i+1)/2} - 4)t_R + it_C, & \text{if } i \text{ is odd, and} \\ (4 \cdot 2^{i/2} - 4)t_R + it_C, & \text{if } i \text{ is even.} \end{cases}$$

Let $S^m(2^{2j})$ be the time taken by the corresponding sorting algorithm (for a square array). Then

$$\begin{aligned} S^m(1) &= 0, \\ S^m(2^{2j}) &= S^m(2^{2j-1}) + T^m(2^{2j}) \\ &= S^m(2^{2(j-1)}) + T(2^{2j}) + T(2^{2j-1}). \end{aligned}$$

Hence

$$S^m(2^{2j}) = (14(2^j - 1) - 8j)t_R + (2j^2 + j)t_C.$$

In our model, we have $2^{2j} = N = n^2$ processors, leading to the following theorem.

Theorem 8.1

If the shuffled row-major indexing is used, the bitonic sort can be done in time

$$(14(n-1) - 8 \log n)t_R + (2 \log^2 n + \log n)t_C.$$

If $t_C \leq 2t_R$, it may be seen that the bitonic merge sort algorithm is optimal to within a factor of 4.5 for all n (since $4(n-1)t_R$ time is necessary, as shown in Section 4). Preliminary investigation indicates that the bitonic merge sort is faster than the s^2 -way odd-even merge sort for $n \leq 512$, under the assumption that $t_C \leq 2t_R$.

9. Extensions and Implications


- (i) By Theorem 7.1 or 8.1, the elements may be sorted into snake-like row-major ordering or in the shuffled row-major ordering in $O(n)$ time. By the following lemma we know that they can be rearranged to obey any other index function with relatively insignificant extra costs, provided each processor has sufficient memory size.

Lemma 9.1

If $N = nxn$ elements have already been sorted with respect to some index function and if each processor can store n elements, then the N elements can be sorted with respect to any other index function by using an additional $4(n-1)t_R$ units of time.

The proof follows from the fact that all elements can be moved to their destinations by four sweeps of $n-1$ routing steps in all four directions.

- (ii) If the processors are connected in a $k \times m$ rectangular array,



instead of a square array, similar results can still be obtained. For example, corresponding to Theorem 7.1, we have

Theorem 9.1

If the snake-like row-major indexing is used, the sorting problem for a $k \times m$ processor array (k, m powers of 2) can be done in time

$$(4m + 2k + O(h^{2/3} \log h))t_R + (k + O(h^{2/3} \log h))t_C$$

where $h = \min(k,m)$, by using the s^2 -way merge sort with $s=O(n^{1/3})$.

- (iii) The number of elements to be sorted could be larger than N , the number of processors. An efficient means of handling this situation is to distribute an approximately equal number of elements to each processor initially and to use a merge-splitting operation for each comparison-interchange operation. This idea is discussed by Knuth [1973, Exercise 5.3.4-38], and used by Baudet and Stevenson [1975]. Baudet and Stevenson's results will be immediately improved if the algorithms of this paper are used, since they used Orcutt's $O(n \log n)$ algorithm.
- (iv) Higher-dimensional array interconnection patterns, i.e., $N=n^j$ processors each connected to its $2j$ nearest neighbors, may be sorted by algorithms generalized from those presented in this paper. For example, $N = n^j$ elements may be sorted in time

$$((3j^2 + j)(n-1) - 2j \log N)t_R + (1/2)(\log^2 N + \log N)t_C$$

by adapting Batcher's bitonic merge sort algorithm to the "j-way shuffled row-major ordering". This new ordering is derived from the binary representation of the row-major indexing by a j-way bit shuffle. If $n=2^3$, $j=3$, and $Z_2Z_1Z_0Y_2Y_1Y_0X_2X_1X_0$ is a row-major index, then the j-way shuffled index is $Z_2Y_2X_2Z_1Y_1X_1Z_0Y_0X_0$.

This formula may be derived in the following way. The t_C term is not dimension-dependent: the same number of comparison is performed in any mapping of the bitonic sort onto an N processor array. The t_R term is the solution of

$$\sum_{1 \leq i \leq \log n} 2^i \sum_{1 \leq k \leq j} ((\log N) - ij + k),$$

where the 2^i term is the cost of a comparison-interchange on the $(i-1)^{st}$ bit of any of the " k^{th} -dimension indices" (i.e., Z_{i-1}, Y_{i-1} , and X_{i-1} when $j=3$ as in the example above). The $((\log N) - ij + k)$ term is the number of times a comparison-interchange is performed on the $(ij-k)^{th}$ bit of the j-way shuffled row-major index during the bitonic sort. Therefore we have the following theorem.

Theorem 9.2

If N processors are j -dimensionally mesh-connected, then the bitonic sort can be performed in time $O(N^{1/j})$, using the j-way shuffled row-major index scheme.

By using the argument of Section 4, one can easily check that the bound in the theorem is asymptotically optimal for large N .

Appendix. The Odd-Even Transposition Sort

The odd-even transposition sort (Knuth [1973], p. 241) may be mapped onto our 2-dimensional arrays with snake-like row-major ordering following way. Given N processors initially loaded with a data value, repeat $N/2$ times:

- O1. "Expensive comparison-interchange" of processors $\#(2i+1)$ with processors $\#(2i+2)$, $0 \leq i < N/2-1$. Time: $4t_R + t_C$ if processor array has more than two columns and more than one row; 0 if $N=2$; and $2t_R + t_C$ otherwise.
- O2. "Cheap comparison-interchange" of processors $\#(2i)$ with processors $\#(2i+1)$, $0 \leq i < N/2-1$. Time: $2t_R + t_C$.

If $T_{oe}(j,k)$ is the time required to sort jk elements in a $j \times k$ region of our processor by the odd-even transposition sort into the snake-like row-major ordering, then

$$T_{oe}(j,k) = \begin{array}{l} 0 \\ 2t_R + t_C \\ jk(2t_R + t_C) \\ jk(3t_R + t_C) \end{array} \begin{array}{l} , \text{ IF } jk=1 \text{ ELSE} \\ , \text{ IF } jk=2 \text{ ELSE} \\ , \text{ IF } j=1 \text{ OR } k=2 \text{ ELSE} \\ \end{array}$$

Step J2 of the 2-way odd-even merge (Section 4) cannot be performed by the version of the odd-even transposition sort indicated above. Since N is even here ($N = 2j$), step O2 may be placed before step O1 in the algorithm description above (see Knuth [1973]). Now step O2 may be performed in the normal time of $2t_R + t_C$, even starting from the non-standard initial configuration depicted in Section 4 as the result of step J1.

References

- Barnes, G. H., et. al [1968]. "The ILLIAC IV Computer," IEEE Trans. on Computers, C-17, pp. 746-757.
- Batcher, K. E. [1968]. "Sorting Networks and Their Applications," Proc. AFIPS Spring Joint Computer Conference, Vol. 32, pp. 307-314.
- Baudet, G., and Stevenson, D. [1975]. "Optimal Sorting Algorithms for Parallel Computers", Computer Science Department report, Carnegie-Mellon University, May 1975.
- Flynn, M. J. [1966]. "Very High-Speed Computing Systems," Proc. IEEE, Vol.54, pp. 1901-1909.
- Knuth, D. E. [1973]. The Art of Computer Programming: Vol. 3, Sorting and Searching, Reading, Mass.: Addison-Wesley.
- Orcutt, S. E. [1974]. Computer Organization and Algorithms for Very-High Speed Computations, Ph.D. Thesis, Stanford University, September 1974, Chap. 2, pp. 20-23.
- Stone, H. S. [1971]. "Parallel Processing with the Perfect Shuffle," IEEE Trans. on Computers, Vol. C-20, pp. 153-161.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) SORTING ON A MESH-CONNECTED PARALLEL COMPUTER		5. TYPE OF REPORT & PERIOD COVERED Interim
7. AUTHOR(s) C. D. Thompson & H. T. Kung		6. PERFORMING ORG REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Computer Science Dept. Pittsburgh, PA 15213		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0370, NR 044-422
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE March 1976
		13. NUMBER OF PAGES 29
		15. SECURITY CLASS (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Two algorithms are presented for sorting n^2 elements on an $n \times n$ mesh-connected processor array that require $O(n)$ routing and comparison steps. The best previous algorithm takes time $O(n \log n)$. The algorithms of this paper are shown to be optimal in time within small constant factors. Extensions to higher-dimensional mesh-connected processor arrays are also given.		