# A GENERAL-PURPOSE PARALLEL SORTING ALGORITHM

ANDREW TRIDGELL and RICHARD P. BRENT
*Computer Sciences Laboratory*
*Australian National University*
*Canberra, ACT 0200, Australia*

ABSTRACT

A parallel sorting algorithm is presented for general purpose internal sorting on MIMD machines. The algorithm initially sorts the elements within each node using a serial sorting algorithm, then proceeds with a two-phase parallel merge. The algorithm is comparison-based and requires additional storage of order the square root of the number of elements in each node. Performance of the algorithm on the Fujitsu AP1000 MIMD supercomputer is discussed.

*Keywords:* Batcher's merge-exchange sort, distributed memory, Fujitsu AP1000, nCUBE2, parallel sorting, sorting, Thinking Machines CM5

**1. Introduction.** Many papers have discussed the task of sorting on parallel computers. See, for example, [1, 2, 11]. Most of these papers have dealt with the problem from a theoretical point of view, neglecting many issues which are important in a practical implementation of a parallel sorting algorithm [3, 9, 13]. This paper describes a practical parallel sorting algorithm which is suitable for efficient general-purpose internal sorting.

A description of the algorithm is given in Sections 2 and 3. In Section 4 the algorithm is compared with some other parallel sorting algorithms. Finally, the performance of our implementation on the Fujitsu AP1000 is discussed in Section 5.

The algorithm presented in this paper was first described in [15], which gives a more detailed account of the implementation of the algorithm.

**1.1. Nomenclature.** $P$ is the number of nodes (also called cells or processors) available on the parallel machine, and $N$ is the total number of elements to be sorted. $N_p$ is the number of elements in a particular node

```
procedure hypercube_balance(integer base, integer num)
if num = 1 return
for all i in [0..num/2)
  pair_balance (base+i, base+i+(num+1)/2)
hypercube_balance (base+num/2, (num+1)/2)
hypercube_balance (base, num - (num+1)/2)
end
```

FIG. 1. *Pseudo-code for load balancing*

$p$ ($0 \leq p < P$). To avoid double subscripts we abbreviate $N_{p_j}$ to $N_j$ if no confusion will arise.

Elements within each node of the machine are referred to as $E_{p,i}$, for $0 \leq i < N_p$ and $0 \leq p < P$. We write $E_{j,i}$ instead of $E_{p_j,i}$ if no confusion will arise.

When giving "big $O$" time bounds we usually assume that $P$ is fixed. Thus, we do not usually distinguish between $O(N)$ and $O(N/P)$.

The only operation assumed for elements is binary comparison, written with the usual comparison symbols. For example, $A < B$ means that element $A$ precedes element $B$. The elements are considered sorted when they are in non-decreasing order in each node, and non-decreasing order between nodes. More precisely, this means that $E_{p,i} \leq E_{p,j}$ for all relevant $i < j$ and $p$, and that $E_{p,i} \leq E_{q,j}$ for $0 \leq p < q < P$ and all relevant $i, j$.

Logarithms are always to base 2.

**2. The Sorting Algorithm.** The algorithm has several phases. In Sections 2.1 to 2.4 below, we describe the purpose and implementation of each phase. In Section 3 we describe the merge-exchange operation which is used in both the primary merging and cleanup phases.

The pre-balancing and primary merging phases are logically unnecessary, and could in principle be omitted. They are included to improve the performance. Without them, the algorithm would still sort, but much more slowly.

**2.1. Pre-Balancing.** The pre-balancing phase moves the elements between nodes so as to achieve as close to an even distribution as possible. This phase is desirable to minimise the load imbalance between nodes in later phases of the algorithm. The balancing is achieved by exchanging elements between pairs of nodes. The communication pattern corresponds to the edges of a hypercube in the case that the number of nodes is a power of 2. This method produces approximately $N/P$ elements in each node, with an error of order $\log P$ for each node if the number of nodes is a power of 2.

Pseudo-code for the pre-balancing algorithm is shown in Figure 1. When the algorithm is called, *base* is initially set to the index of the smallest node in the system, and *num* is set to $P$, the number of nodes. The algorithm operates recursively and takes $\log P$ steps to complete. When the number of nodes is not a power of 2, the effect is to have some of the nodes idle in some phases of the algorithm. Because the nodes which remain idle change with each step, all nodes take part in a pair-balance with another node. With some distributions of elements and values for $P$ it may be beneficial to iterate the pre-balancing algorithm until some balancing criterion is met.

As can be seen from the code for the algorithm, the actual work of the balance is performed by another routine called *pair_balance*. This routine is designed to exchange elements between a pair of nodes so that both nodes end up with as close as possible to the same number of elements.

**2.2. Serial Sorting.** In the serial sorting phase there is no communication between nodes, but a fast comparison-based serial sorting algorithm is applied to the elements in each of the nodes. At the end of this phase the data is in the form of $P$ sorted lists of elements, with approximately $N/P$ elements in each list.

After some experimentation, a combination of quick-sort and insertion sort was chosen for serial sorting. This was found to give the best performance over a wide range of conditions on the AP1000.

As will be seen in Section 5, the time taken by the serial sorting phase of the algorithm dominates the total time taken by the sorting algorithm when the number of elements being sorted is large. It is thus very important that the serial sorting algorithm be optimised as much as possible. This optimisation might include the use of an alternate serial sorting algorithm that can take advantage of particular forms of the comparison function between data elements. An example would be the use of a radix sorting algorithm for the sorting of ordinal data types such as integers.

**2.3. Primary Merging.** The aim of the primary merge phase of the algorithm is to "almost sort" the data in minimum time. For this purpose an algorithm with a very high parallel efficiency was chosen to control merge-exchange operations between the nodes. This led to significant performance improvements over the use of an algorithm with lower parallel efficiency that is guaranteed to completely sort the data (for example, Batcher's algorithm as used in the cleanup phase).

The pattern of merge-exchange operations in the primary merge is identical to that used in the pre-balancing phase of the algorithm. The pseudo-code for the algorithm is given in Figure 2. When the algorithm is called, *base* is initially set to the index of the smallest node in the system, and *num*

```
procedure primary_merge(integer base, integer num)
if num = 1 return
for all i in [0..num/2)
  merge_exchange (base+i, base+i+(num+1)/2)
primary_merge (base+num/2, (num+1)/2)
primary_merge (base, num - (num+1)/2)
end
```

FIG. 2. *Pseudo-code for primary merge*

is set to $P$, the number of nodes.

This algorithm completes in $\log P$ steps per node, with each step consisting of a merge-exchange operation. As with the pre-balancing algorithm, if $P$ is not a power of 2 then some nodes may be left idle at each step of the algorithm.

If $P$ is a power of 2 and the initial distribution of the elements is random, then at each step of the algorithm each node has about the same amount of work to perform as the other nodes. In other words, the load balance between the nodes is very good. The symmetry is only broken by an unusual distribution of the original data, or if $P$ is not a power of 2. In both these cases load imbalances may occur.

**2.4. Cleanup.** The cleanup phase of the algorithm is similar to the primary merge phase, but it must be guaranteed to complete the sorting process. The method chosen to achieve this is Batcher's merge-exchange algorithm. This algorithm has some useful properties which make it ideal for a cleanup operation.

The pseudo-code for Batcher's merge-exchange algorithm is given in [7, Algorithm M, page 112]. The algorithm defines a pattern of comparison-exchange operations which will sort a list of elements of any length.

The way the algorithm is normally described, the comparison-exchange operation operates on two elements and exchanges the elements if the first element is greater than the second. In the application of the algorithm to the cleanup operation we generalise the notion of an element to include all elements in a node. This means that the comparison-exchange operation must make all elements in the first node greater than all elements in the second. This is identical to the operation of the merge-exchange algorithm, described in Section 3. A proof that it is possible to make this generalisation while maintaining the correctness of the algorithm is given in [7, solution to problem 5.3.4.38]. This proof assumes that each of the nodes contains the same number of elements. A method for extending this result to the case of an uneven distribution of elements is given in Section 2.5.

Batcher's merge-exchange algorithm is ideal for the cleanup phase because it is very fast for almost sorted data. This is a consequence of a unidirectional merging property: the merge operations always operate in a direction so that the lower numbered node receives the smaller elements. This is not the case for some other fixed sorting networks, such as the bitonic algorithm [4]. Algorithms that do not have the unidirectional merging property are a poor choice for the cleanup phase as they tend to un-sort the data (undoing the work done by the primary merge phase), before sorting it. In practice the cleanup time is of the order of 1 or 2 percent of the total sort time if Batcher's merge-exchange algorithm is used and the merge-exchange operation is implemented efficiently.

**2.5. $N$ not a multiple of $P$.** A proof given in [7, solution to problem 5.3.4.38] shows that comparison-exchange based sorting algorithms can be extended to sort lists of elements by replacing the comparison-exchange operation with a merge-exchange operation. Simple examples demonstrate that the condition that each of the nodes have the same number of elements is necessary.

In [15] we describe a method called "infinity padding" which allows the cleanup phase of our algorithm to operate correctly when each node does not have the same number of elements. The method works by adding implicit $\infty$-elements to each of the nodes and calculating the number of real elements that must remain in each node after each merge-exchange operation.

In this section we describe an alternative to infinity padding which operates not in the merge-exchange operation but instead as an additional balancing phase coming directly after the pre-balancing phase of the algorithm. The effect of this change is to simplify the merge-exchange operation.

Let the maximum number of elements in any one node be $M$. We transfer $M - N_i$ elements from the highest numbered node to each of the lower numbered nodes in turn, starting with node 0. If the highest numbered node runs out of elements (which can occur if $N < P$) then we continue the process with the next highest numbered node, and so on.

Now, when a merge-exchange operation is performed between nodes $P_i$ and $P_j$ with $i < j$, we can be sure that either $N_i = M$ and $N_j <= N_i$, or that $N_j = 0$. This means that if we pad the nodes that have $N_i < M$ elements with $M - N_i$ $\infty$-elements, then we can guarantee that these $\infty$-elements will not move, and thus do not have to be explicitly represented.

This method relies on the fact that the sorting algorithm used for the cleanup phase is unidirectional in order to guarantee that the implicit $\infty$-elements always remain in the higher numbered nodes.

**3. Merge-Exchange.** The aim of the merge-exchange algorithm is to exchange elements between two nodes so that we end up with one node containing elements which are all smaller than all the elements in the other node, while maintaining the order of the elements in the nodes. In our implementation of parallel sorting we always require the node with the smaller node number to receive the smaller elements. This would not be possible if we used Batcher's bitonic algorithm instead of his merge-exchange algorithm for the cleanup phase.

Secondary aims of the merge-exchange operation are that it should be very fast for data that is almost sorted already, and that the memory overhead should be minimised.

Suppose that a merge operation is needed between two nodes, $p_1$ and $p_2$, which initially contain $N_1$ and $N_2$ elements respectively. We assume that the smaller elements are required in node $p_1$ after the merge.

In principle, merging two already-sorted lists of elements to obtain a new sorted list is a very simple process. The most obvious implementation takes $N_1 + N_2$ steps, with each step requiring one copy and one comparison operation [4]. The problem with this algorithm is the storage requirements implied by the presence of a destination array. The use of this algorithm as a component of a parallel sorting algorithm would restrict the number of elements that can be sorted to the number that can fit in half the available memory of the machine. The question arises as to whether an algorithm can be developed that does not require this destination array. In order to achieve this, it is clear that the algorithm must re-use the space that is freed by moving elements from the two source lists. We now describe how this can be done. The algorithm has several parts, each of which is described separately.

Our merge algorithm has similarities to known algorithms for in-place merging [5, 8], but is considerably simpler due to its use of a temporary storage area of size $O(\sqrt{N/P})$. In practice the size of this temporary storage area is insignificant.

**3.1. Find-Exact Algorithm.** When a node takes part in a merge-exchange with another node, it will need to be able to access the other node's elements as well as its own. The simplest method for doing this is for each node to receive a copy of all of the other node's elements before the merge begins.

A much better approach is to first determine exactly how many elements from each node will be required to complete the merge, and to transfer only those elements. This reduces the communications cost by minimising the number of elements transferred, and at the same time reduces the memory overhead of the merge.

The find-exact algorithm allows each node to determine exactly how many elements are required from another node in order to produce the correct number of elements in a merged list.

When a comparison is made between element $E_{1,A-1}$ and $E_{2,N_1-A}$ then the result of the comparison determines whether node $p_1$ will require more or less than $A$ of its own elements in the merge. If $E_{1,A-1}$ is greater than $E_{2,N_1-A}$, then the maximum number of elements that could be required to be kept by node $p_1$ is $A-1$, otherwise the minimum number of elements that could be required to be kept by node $p_1$ is $A$.

The proof of correctness relies on counting the number of elements that could be less than $E_{1,A-1}$. If $E_{1,A-1}$ is greater than $E_{2,N_1-A}$, then there are at least $N_1 - A + 1$ elements in node $p_2$ that are less than $E_{1,A-1}$. If these are combined with the $A-1$ elements in node $p_1$ that are less than $E_{1,A-1}$, then we have at least $N_1$ elements less than $E_{1,A-1}$. This implies that the number of elements that must be kept by node $p_1$ is at most $A-1$.

A similar argument can be used to show that, if $E_{1,A-1} \leq E_{2,N_1-A}$, then the number of elements to be kept by node $p_1$ must be at least $A$. Combining these two results leads to a "bisection" algorithm that can find the exact number of elements required in at most $\log N_1$ steps by successively halving the range of possible values for the number of elements required to be kept by node $p_1$. Once this result is determined, it is a simple matter to derive the number of elements that must be sent from node $p_1$ to node $p_2$, and from node $p_2$ to node $p_1$.

On a machine with a high message latency, this algorithm could be costly, as a relatively large number of small messages are transferred. The number of messages could be reduced, but with a penalty of increased message size and algorithm complexity. One method for doing this is described in [17]. On the AP1000 the cost of the find-exact algorithm was found to be very small, due to the low message latency of the machine, so the method of [17] was not implemented [6, 12].

We assume for the remainder of the discussion on the merge-exchange algorithm that the find-exact algorithm has determined that node $p_1$ must retain $L_1$ elements and must receive $L_2$ elements from node $p_2$.

**3.2. Transferring Elements.** After the exact number of elements to be transferred has been determined, the transfer takes the form of an exchange of elements between the two nodes. The elements that are sent from node $p_1$ leave behind them spaces which must be filled with the incoming elements from node $p_2$. The reverse happens on node $p_2$, so the transfer process must be careful not to overwrite elements that have not yet been sent.

After the transfer is complete, the elements on node $p_1$ are stored in

two contiguous sorted lists, of lengths $L_1$ and $N_1 - L_1$. In the remaining steps of the merge-exchange algorithm we merge these two lists so that all the elements are in order.

**3.3. Unbalanced Merging.** Before considering the algorithm for minimum memory merging, it is worth considering a special case where the result of the find-exact algorithm determines that the number of elements to be kept on node $p_1$ is much larger than the number of elements to be received from node $p_2$. In this case the task which node $p_1$ must undertake is to merge two lists of very different sizes. There is a very efficient algorithm for this special case, which may occur if the data is almost sorted (for example, near the end of the cleanup phase).

Suppose that $L_1$ is much greater than $L_2$. First we determine where each of the $L_2$ elements that have been transferred from $p_2$ belongs in the list of length $L_1$. This can be done with at most $L_2 \log L_1$ comparisons using a method similar to the find-exact algorithm. As $L_2$ is small, this number of comparisons is small, and the results take only $O(L_2)$ storage.

Once this is done we can copy all the elements in list 2 to a temporary storage area and begin the process of moving elements from list 1 and list 2 to their proper destinations. This takes at most $L_1 + L_2$ element copies, but in practice it often takes only about $2L_2$ copies. This is because when only a small number of elements are transferred between nodes there is often only a small overlap between the ranges of elements in the two nodes, and only the elements in the overlap region have to be moved. Thus the unbalanced merge performs very quickly in practice, and the overall performance of the sorting procedure is significantly better than it would be if we did not take advantage of this special case.

**3.4. Blockwise Merging.** The blockwise merge is a solution to the problem of merging two sorted lists of elements into one, while using only a small amount of additional storage. The first phase in the operation is to break the two lists into blocks of an equal size $B$. The exact value of $B$ is unimportant for the correctness of the algorithm, but influences its efficiency and memory usage. We assume that $B$ is $O(\sqrt{L_1 + L_2})$, which is small relative to the memory available on each node. To simplify the exposition we also assume, for the time being, that $L_1$ and $L_2$ are multiples of $B$.

The merge takes place by merging from the two blocked lists of elements into a destination list of blocks. The destination list is initially primed with two empty blocks which comprise a temporary storage area. As each block in the destination list becomes full the algorithm selects a new, empty block, choosing the next one in the destination list. As each block in either of the

two source lists becomes empty it is added to the destination list.

As the merge proceeds there are always exactly $2B$ free spaces in the three lists. This implies that there must always be at least one free block on the destination list, whenever a new destination block is required.

The algorithm actually takes no more steps than the simple merge mentioned earlier. Each element moves only once. The drawback, however, is that the elements end up in a blocked list structure rather than in a simple linear array.

The simplest method for resolving this problem, at the expense of additional memory accesses, is to go through a re-arrangement phase to put the blocks back in the standard form. This is what has been done in our current implementation. It would be possible to modify the algorithm so that all references to elements are performed with the elements in a block list format. This would improve performance at the cost of some additional coding complexity.

So far, we have assumed that $L_1$ and $L_2$ are both multiples of $B$. In general this is not the case. If $L_1$ is not a multiple of $B$ then this introduces the problem that the initial breakdown of list 2 into blocks of size $B$ will not produce blocks that are aligned on multiples of $B$ relative to the first element in list 1. To solve this problem we copy $L_1 \bmod B$ elements from the tail of list 1 and use this copy as a final source block. Then we offset the blocks when transferring them from source list 2 to the destination list so that they end up aligned on the proper boundaries. Finally we increase the amount of temporary storage to $3B$ and prime the destination list with three blocks to account for the fact that we cannot use the partial block from the tail of list 1 as a destination block.

**4. Comparison with other algorithms.** When designing the parallel sorting algorithm we wished to produce a general-purpose algorithm, similar in functionality to the many general-purpose serial sorting algorithms that already exist. Specifically, we aimed for the following properties:

1. Speed. The algorithm should be competitive with the fastest known algorithms.
2. Good memory utilisation. The number of elements that can be sorted should be close to the number that can be stored in the memory of the machine.
3. Flexibility. No restrictions should be placed on the number of records to sort or the number of processors.
4. Determinism. The algorithm should not use a random number generator.
5. Comparison-based. The only operation used on keys is binary comparison.

A parallel algorithm based on radix sorting has been shown to perform very well [3], but it relies on the ordinal properties of the data type being sorted. This limits its application as it can not be used to sort with an arbitrary comparison function.

Radix sorting, in its usual implementation, requires the use of additional workspace at least as large as the data being sorted. This limits the number of elements that can be sorted to half the number that can be fitted in the machine's memory.

The parallel sample sort algorithm has a similar problem of memory utilisation [3]. It also fails to satisfy property 4, as it relies on taking a pseudo-random sample of the input data.

Several methods which rely on the number of elements being a multiple of the number of nodes or the number of nodes being a power of 2 are ruled out by property 3 (see [2]).

**5. Performance.** Several runs were made on the AP1000 to examine the performance of our implementation of the sorting algorithm under a variety of conditions.

An evaluation of the algorithm has also been performed on the Thinking Machines CM5 [16]. The results for the CM5 may be found in [15]. Our code has also recently been ported to the nCUBE2 [10]. The results for this machine show a greater parallel speedup than for either the AP1000 or CM5. This is due to the hypercube connections of the nCUBE2, which suit our algorithm, along with a high ratio of network speed to processing speed.

When we refer to the parallel speedup of the algorithm we are referring to the ratio of the time the sorting task takes to complete with the best available serial algorithm on a single node of the parallel machine to the time for our parallel algorithm. The fact that our algorithm naturally reduces to the best available serial algorithm when run on one node of the parallel machine makes the comparison simpler.

**5.1. Timing Results.** The first results are shown in Figure 3. This figure shows the performance of the algorithm on the 128-node AP1000 as $N$ spans a wide range of values, from values which would be easily dealt with on a workstation, to those at the limit of the AP1000's memory capacity (2 Gbyte). The elements are 16-byte random strings. The comparison function is the C library function strcmp().

The results give the number of elements that can be sorted per second of real time. This time includes all phases of the algorithm, and gives an overall indication of performance.

Shown on the same graph is the performance of a hypothetical serial computer that operates $P$ times as fast as the $P$ individual nodes of the

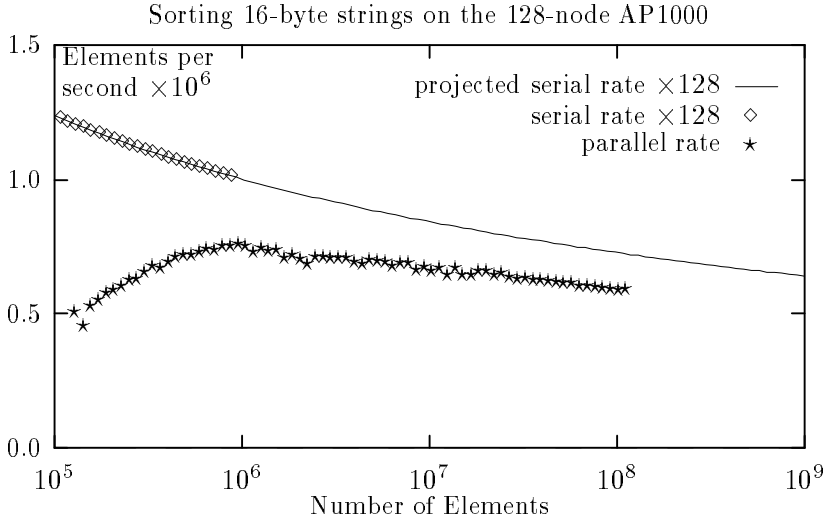


FIG. 3. *Sorting 16-byte strings on the AP1000*

parallel computer. This performance is calculated by sorting the elements on a single node and multiplying the resulting elements per second result by $P$. An extrapolation of this result to larger values of $N$ is also shown using a least squares fit to the function:

$$\text{time}(N) = a + b \log N + cN + dN \log N.$$

The extrapolation is necessary due to memory restrictions in individual nodes of the machine.

The graph shows that the performance of the sorting algorithm increases quickly as the number of elements approaches 1 million, after which a slow fall-off occurs which closely follows the profile of the ideal parallel speedup. The roll-off point of 1 million elements corresponds to the number of elements that can be held in the 128KB cache of each node. This indicates the importance of caching to the performance of the algorithm.

It is encouraging to note how close the algorithm comes to the ideal speedup of $P$ for a $P$-node machine. The algorithm achieves 85% of the ideal performance for a 128-node machine.

**5.2. Scalability.** An important aspect of a parallel algorithm is its scalability, which depends on the ability of the algorithm to utilise additional nodes. Shown in Figure 4 is the result of sorting 100,000 16-byte strings per node on the AP1000 as the number of nodes is varied. The percentages refer

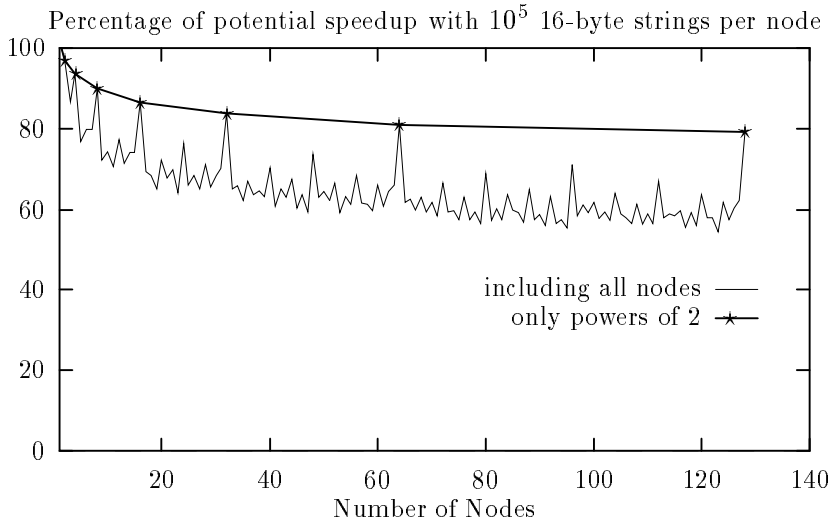


FIG. 4. *Scalability of sorting on the AP1000*

to the proportion of the ideal speedup $P$ that is achieved. The number of elements per node is kept constant to ensure that caching factors do not influence the result.

The left-most data point shows the speedup for a single node. This is equal to 1 as the algorithm reduces to our optimised quick-sort when only a single node is used. As the number of nodes increases, the proportion of the ideal speedup decreases, as communication costs and load imbalances begin to appear. The graph flattens out for larger numbers of nodes, which indicates that the algorithm should have a good efficiency when the number of nodes is large.

The two curves in the graph show the trend when all configurations are included and when only configurations with $P$ a power of 2 are included. The difference between these two curves clearly shows the preference for powers of two in the algorithm. It is also clear that certain values for $P$ are preferred to others. In particular even numbers of nodes perform better than odd numbers. Sums of adjacent powers of two also seem to be preferred, so that when $P$ takes on values of 24, 48 and 96 the efficiency is quite high.

Parallel versions of radix sort [3] and sample sort [3, 13] require "all-to-all" communications, so they may not scale as well as our algorithm to machines with a large number of processors.

**5.3. Where The Time Goes.** It is interesting to look at the proportion of the total time spent in each of the phases of the sort. In Figure 5

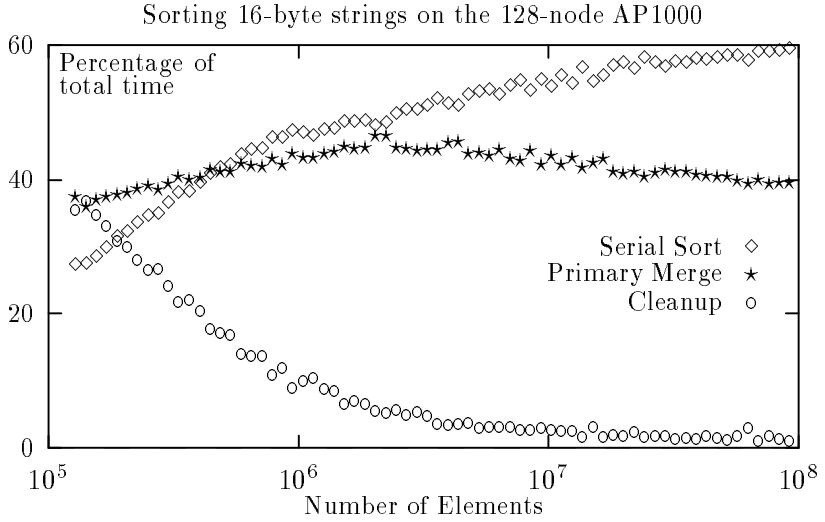


FIG. 5. *Timing breakdown by phase*

this is done over a wide range of values of $N$ for sorting 16-byte strings on the AP1000. The three phases that are examined are the initial serial sort, the primary merging and the cleanup phase.

This graph shows that as $N$ increases to a significant proportion of the memory of the machine the greatest time is taken by the initial serial sort of the elements in each cell. This is the case because this phase of the algorithm is $O(N \log N)$ whereas all other phases of the algorithm are $O(N)$ or lower. It is the fact that this component of the algorithm is able to dominate the time while $N$ is still a relatively small proportion of the capacity of the machine which leads to the practical efficiency of the algorithm. Many sorting algorithms are asymptotically optimal in the sense that their speedup approaches $P$ for large $N$, but few can get close to this speedup for values of $N$ which are of interest in practice [9].

It is interesting to note the small impact that the cleanup phase has for larger values of $N$. This demonstrates the fact that the primary merge does produce an almost sorted data set, and that the cleanup algorithm can take advantage of this.

A second way of splitting the time taken for the parallel sort to complete is by task. In this case we look at what kind of operation each of the nodes is performing, which provides a finer division of the time.

Figure 6 shows the result of this kind of split for the sorting of 16-byte strings on the 128-node AP1000, over a wide range of values of $N$. Again

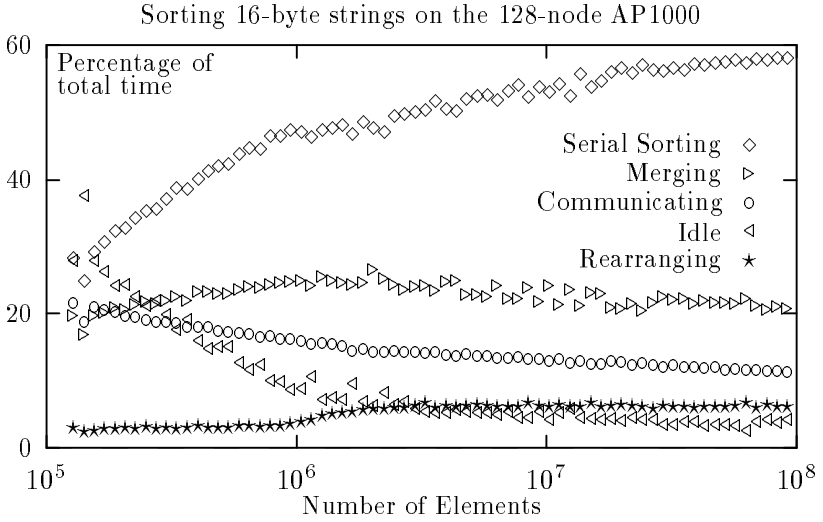


FIG. 6. *Timing breakdown by task*

it is clear that the serial sorting dominates for large values of $N$, for the same reasons as before. What is more interesting is that the proportion of time spent idling while waiting for messages and in actually communicating decreases steadily as $N$ increases. From the point of view of the parallel speedup of the algorithm these tasks are wasted time and need to be kept to a minimum.

**6. Conclusions.** We have presented a practical general-purpose parallel internal sorting algorithm that comes close to achieving the best possible speedup over an optimised serial algorithm. An implementation of the algorithm on a real machine has been discussed.

The algorithm derives its generality from the fact that it is comparison-based, and allows for a user-supplied comparison function. This corresponds to the commonly available serial sorting procedures that are the mainstay of internal sorting on serial computers.

The algorithm is frugal in its memory requirements, which allows data to be sorted almost to the limit of a parallel machine's memory. This is important because it is unreasonable to expect data sets being sorted on a parallel machine to occupy only a small fraction of the machine's memory.

**7. Availability.** Source code for the algorithm is available via anonymous ftp from `andosl.anu.edu.au` (Internet number 150.203.15.95) in the directory `pub/tridge/sorting/par_sort`. Also included in this directory

is a technical report [15]. The authors may be contacted by e-mail at `tridge@cslab.anu.edu.au` and `rpb@cslab.anu.edu.au`

**Acknowledgements.** Support by Fujitsu Laboratories, Fujitsu Limited, and Fujitsu Australia Limited via the Fujitsu-ANU CAP Project is gratefully acknowledged. Support of Andrew Tridgell via an ATERB postgraduate scholarship is also gratefully acknowledged.

## REFERENCES

[1] M. Ajtai, J. Kolmos and E. Szermeredi, "Sorting in $c \log n$ parallel steps", *Combinatorica* 3, 1983, 1-19.

[2] S. G. Akl, *Parallel Sorting Algorithms*, Academic Press, Toronto, 1985.

[3] G. E. Blelloch, C. E. Leiserson, B. M. Maggs, C. G. Plaxton, S. J. Smith and M. Zagha, "A comparison of sorting algorithms for the Connection Machine CM-2", *Proc. Symposium on Parallel Algorithms and Architectures*, Hilton Head, South Carolina, July 1991.

[4] G. C. Fox, M. A. Johnson, G. A. Lyzenga, S. W. Otto, J. K. Salmon and D. W. Walker, *Solving Problems on Concurrent Processors, Volume 1*, Prentice-Hall, Englewood Cliffs, New Jersey, 1988.

[5] B.-C. Huang and M. A. Langston, "Practical in-place merging", *Communications of the ACM* 31, 1988, 348-352.

[6] H. Ishihata, T. Horie and T. Shimizu, "Architecture for the AP1000 highly parallel computer", *Fujitsu Sci. Tech. J.* 29, 1993, 6-14.

[7] D. E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching* (second edition), Addison-Wesley, Menlo Park, 1981.

[8] M. A. Kronrod, "An optimal ordering algorithm without a field of operation", *Dokl. Akad. Nauk SSSR* 186, 1969, 1256-1258 (in Russian).

[9] L. Natvig, "Logarithmic Time Cost Optimal Parallel Sorting is Not Yet Fast in Practice!", *Proc Supercomputing 90*, IEEE Press, 1990, 486-494.

[10] T. Rashid, personal communication, March 1993.

[11] J. H. Reif and L. G. Valiant, "A logarithmic time sort for linear size networks", *J. ACM* 34, 1987, 60-76.

[12] T. Shimizu, T. Horie and H. Ishihata, "Performance evaluation of the AP1000", *Fujitsu Sci. Tech. J.* 29, 1993, 15-24.

[13] H. Shu and J. Schaeffer, "Parallel sorting by regular sampling", *J. Parallel and Distributed Computing* 14, 1992, 361-372.

[14] K. Thearling and S. Smith, "An Improved Supercomputing Sorting Benchmark", *Proc Supercomputing 92*, IEEE Press, 1992, 14-19.

[15] A. Tridgell and R. P. Brent, *An Implementation of a General-Purpose Parallel Sorting Algorithm*, Report TR-CS-93-01, Computer Sciences Laboratory, Australian National University, February 1993, 24 pp. Available by anonymous ftp (see Section 7).

[16] *CM-5 Technical Summary*, Thinking Machines Corporation, October 1991.

[17] B. B. Zhou, R. P. Brent and A. Tridgell, *Efficient Implementation of Sorting Algorithms on Asynchronous Distributed-Memory Machines*, Technical Report, CS Lab, ANU, to appear.