Designing Practical Efficient Algorithms for Symmetric Multiprocessors* (Extended Abstract)

David R. Helman   Joseph JáJá

Institute for Advanced Computer Studies &
Department of Electrical Engineering,
University of Maryland, College Park, MD 20742
{helman, joseph}@umiacs.umd.edu

October 8, 1998

Abstract

Symmetric multiprocessors (SMPs) dominate the high-end server market and are currently the primary candidate for constructing large scale multiprocessor systems. Yet, the design of efficient parallel algorithms for this platform currently poses several challenges. In this paper, we present a computational model for designing efficient algorithms for symmetric multiprocessors. We then use this model to create efficient solutions to two widely different types of problems - linked list prefix computations and generalized sorting. Our novel algorithm for prefix computations builds upon the sparse ruling set approach of Reid-Miller and Blelloch. Besides being somewhat simpler and requiring nearly half the number of memory accesses, we can bound our complexity *with high probability* instead of merely *on average*. Our algorithm for generalized sorting is a modification of our algorithm for sorting by regular sampling on distributed memory architectures. The algorithm is a stable sort which appears to be asymptotically faster than any of the published algorithms for SMPs. Both of our algorithms were implemented in C using POSIX threads and run on three symmetric multiprocessors - the DEC AlphaServer, the Silicon Graphics Power Challenge, and the HP-Convex Exemplar. We ran our code for each algorithm using a variety of benchmarks which we identified to examine the dependence of our algorithm on memory access patterns. In spite of the fact that the processors must compete for access to main memory, both algorithms still yielded scalable performance up to 16 processors, which was the largest platform available to us. For some problems, our prefix computation algorithm actually matched or exceeded the performance of the best sequential solution using only a single thread. Similarly, our generalized sorting algorithm always beat the performance of sequential merge sort by at least an order of magnitude, even with a single thread.

*Supported in part by NSF grant No. CCR-9627210 and NSF HPCC/GCAG grant No. BIR-9318183. It also utilized the NCSA HP-Convex Exemplar SPP-2000 and the NCSA SGI/CRAY POWER CHALLENGEarray with the support of the National Computational Science Alliance under grant No. ASC970038N.

1 Introduction

Symmetric multiprocessors (SMPs) dominate the high-end server market and are currently the primary candidate for constructing large scale multiprocessor systems. Yet, the design of efficient parallel algorithms for this platform currently poses several challenges. The reason for this is that the rapid progress in microprocessor speed has left main memory access as the primary limitation to SMP performance. Since memory is the bottleneck, simply increasing the number of processors will not necessarily yield better performance. Indeed, memory bus limitations typically limit the size of SMPs to 16 processors. This has at least two implications for the algorithm designer. First, since there are relatively few processors available on an SMP, any parallel algorithm must be competitive with its sequential counterpart with as little as one processor in order to be relevant. Second, for the parallel algorithm to scale with the number of processors, it must be designed with careful attention to minimizing the number and type of main memory accesses.

In this paper, we present a computational model for designing efficient algorithms for symmetric multiprocessors. We then use this model to create efficient solutions to two widely different types of problems - linked list prefix computations and generalized sorting. Both problems are memory intensive, but in different ways. Whereas generalized sorting algorithms typically require a large number of memory accesses, they are usually to contiguous memory locations. By contrast, prefix computation algorithms typically require a more modest quantity of memory accesses, but they are are usually to non-contiguous memory locations.

Our novel algorithm for prefix computations builds upon the sparse ruling set approach of Reid-Miller and Blelloch [10]. Unlike the original algorithm, we choose the ruling set in such a way as to avoid the need for conflict resolution. Besides making the algorithm simpler, this change allows us to achieve a stronger bound on the complexity. Whereas Reid-Miller and Blelloch claim an *expected* complexity of $O\left(\frac{n}{p}\right)$ for $n >> p$, we claim a complexity *with high probability* of $O\left(\frac{n}{p}\right)$ for $n > p^2 \ln n$. Additionally, our algorithm incurs approximately half the memory costs of their algorithm, which we believe to be the smallest of any parallel algorithm of which we are aware. Our algorithm for generalized sorting is a modification of our algorithm for sorting by regular sampling on distributed memory architectures. The algorithm is a stable sort which appears to be asymptotically faster than any of the published algorithms we are aware of.

Both of our algorithms were implemented in C using POSIX threads and run on three symmetric multiprocessors - the DEC AlphaServer, the Silicon Graphics Power Challenge, and the HP-Convex Exemplar. We ran our code for each algorithm using a variety of benchmarks which we identified to examine the dependence of our algorithm on memory access patterns. In spite of the fact that the processors must compete for access to main memory, both algorithms still yielded scalable performance up to 16 processors, which was the largest platform available to us. For some problems, our prefix computation algorithm actually matched or exceeded the performance of the best sequential solution using only a single thread. Similarly, our generalized sorting algorithm always beat the performance of sequential merge sort by at least an order of magnitude, which from our experience is the best sequential sorting algorithm on these platforms.

The organization of our paper is as follows. **Section 2** presents our computational model for analyzing algorithms on symmetric multiprocessors. **Section 3** describes our prefix computation algorithm for this platform and its experimental performance. Similarly, **Section 4** describes our generalized sorting algorithm for this platform and its experimental performance.

2 A Computational Model for Symmetric Multiprocessors

For our purposes, the cost of an algorithm needs to include a measure that reflects the number and type of memory accesses. A number of models have already been proposed which focus on the cost of accessing different levels of memory, including the D-disk model of Vitter and Shriver [13], the Hierarchical Memory with Block Transfer Model of Aggarwal et al. [1], and the Uniform Memory Hierarchy Model of Alpern et al. [2]. However, we believe that these models are unnecessarily complicated to describe the behavior of existing symmetric multiprocessors. Other models have been proposed which focus instead on the contention caused by multiple processors competing to access the same location in main memory, including the (d,x)-BSP model of Blelloch et al. [4] and the Queuing Shared Memory (QSM) of Gibbons et al. [5]. The difficulty with these models is that while they address an issue which has an important impact on performance, the contention they describes depends on specific implementation details such as the memory map which may be entirely beyond the control of the algorithm designer.

In our SMP model, we acknowledge the dominant expense of memory access. Indeed, it has been widely observed that the rapid progress in microprocessor speed has left main memory access as the primary limitation to SMP performance. The problem can be minimized by insisting where possible on a pattern of contiguous data access. This exploits the contents of each cache line and takes full advantage of the pre-fetching of subsequent cache lines. However, since it does not always seem possible to direct the pattern of memory access, our complexity model needs to include an explicit measure of the number of non-contiguous main memory accesses required by an algorithm. Additionally, we recognize that efficient algorithm design requires the efficient decomposition of the problem amongst the available processors, and, hence, we also include the cost of computation in our complexity.

More precisely, we measure the overall complexity of an algorithm by the triplet of values $\langle M_A, M_E, T_C \rangle$, where M_A is the maximum number of accesses made by any processor to main memory, M_E is the maximum amount of data exchanged by any processor with main memory, and T_C is an upper bound on the local computational complexity of any of the processors. Note that M_A is simply a measure of the number of non-contiguous main memory accesses, where each such access may involve an arbitrary sized contiguous block of data. Notice also that while we report T_C using the customary asymptotic notation, we report M_A and M_E as *approximations* of the actual values. By approximations, we mean that if M_A or M_E is described by a quadratic expression, then we report the highest order term *and* its coefficient. We distinguish between memory access cost and computational cost in this fashion because of the dominant expense of memory access on this architecture. With so few processors available, this coefficient is usually crucial in determining whether or not a parallel algorithm can be a viable replacement to the sequential alternative. On the other hand, despite the importance of these memory costs, we report only the highest order term, since otherwise the expression can easily become unwieldy.

In practice, it is often possible to focus on either M_A or M_E when examining the cost of algorithmic alternatives. For example, we observed when comparing prefix computation algorithms that the number of contiguous and non-contiguous memory accesses were always of the same asymptotic order, and therefore we only report M_A, which describes only the much more expensive non-contiguous accesses. Subsequent experimental analysis of the step-by-step costs has validated this simplification. On the other hand, algorithms for generalized sorting are usually all based on the idea of repeatedly merging sorted sequences, which are accessed in a contiguous fashion. Moreover, since our model is concerned only with the cost of main memory access, once values are stored in cache they may be accessed in any pattern at no cost. As a consequence, the number of non-contiguous memory accesses are always much less than the number of contiguous memory accesses, and in this situation we only report M_E, which includes the much more numerous contiguous memory accesses. Again, subsequent experimental analysis of the step-by-step costs has validated this simplification.

3 Prefix Computations

Consider the problem of performing a prefix computation on a linked list of n elements stored in arbitrary order in an array X. For each element X_i, we are given $X_i.succ$, the array index of its successor, and $X_i.data$, its input value for the prefix computation. Then, for any binary associative operator $\otimes$, the prefix computation is defined as:

$$X_i.prefix = \begin{cases} X_i.data & \text{if } X_i \text{ is the head of the list.} \\ X_i.data \otimes X_{(pre)}.prefix & \text{otherwise.} \end{cases}, \tag{1}$$

where pre is the index of the predecessor of X_i. The last element in the list is distinguished by a negative index in its successor field, and nothing is known about the location of the first element.

Any of the known parallel prefix algorithms in the literature can be considered for implementation on an SMP. However, to be competitive, a parallel algorithm must contend with the extreme simplicity of the obvious sequential solution. A prefix computation can be performed by a single processor with two passes through the list, the first to identify the head of the list and the second to compute the prefix values. The pseudocode for this obvious sequential algorithm is as follows:

- **(1):** Visit each list element X_i in order of ascending array index. If X_i is not the terminal element, then label its successor with index $X_i.succ$ as having a predecessor.
- **(2):** Find the one element not labeled as having a predecessor by visiting each list element X_i in order of ascending array index - this unlabeled element is the head of the list.
- **(3):** Beginning at the head, traverse the elements in the list by following the successor pointers. For each element traversed with index i and predecessor pre, set *List[i].prefix_data* = *List[i].prefix_data* $\otimes$ *List[pre].prefix_data*.

To compute the complexity, note that Step (1) requires at most n non-contiguous accesses to label the successors. Step (2) involves a single non-contiguous memory access to a block of n contiguous elements. Step (3) requires at most n non-contiguous memory accesses to update the successor of each element. Hence, this algorithm requires approximately $2n$ non-contiguous memory accesses and runs in in $O(n)$ computation time.

According to our model, however, the obvious algorithm is not necessarily the best sequential algorithm. The non-contiguous memory accesses of Step (1) can be replaced by a single contiguous memory access by observing that the index of the successor of each element is a unique value between 0 and $n-1$ (with the exception of the tail, which by convention has been set to a negative value). Since only the head of the list does not have a predecessor, it follows that together the successor indices comprise the set $\{0, 1, .., h-1, h+1, h+2, .., n-1\}$, where h is the index of the head. Since the sum of the complete set $\{0, 1, .., n-1\}$ is given by $\frac{1}{2}n(n-1)$, it easy to see that the identity of the head can be found by simply subtracting the sum of the successor indices from $\frac{1}{2}n(n-1)$. The importance of this lies in the fact that the sum of the successor indices can be found by visiting the list elements in order of ascending array index, which according to our model requires only a single non-contiguous memory access. The pseudocode for this improved sequential algorithm is as follows:

- **(1):** Compute the sum Z of the successor indices by visiting each list element X_i in order of ascending array index. The index of head of the list is $h = \left(\frac{1}{2}n(n-1) - Z\right)$.
- **(2):** Beginning at the head, traverse the elements in the list by following the successor pointers. For each element traversed with index i and predecessor *pre*, set *List[i]._prefix_data* = *List[i].prefix_data* $\otimes$ *List[pre].prefix_data*.

Since this modified algorithm requires no more than approximately n non-contiguous memory accesses while running in $O(n)$ computation time, it is optimal according to our model.

The first fast parallel algorithm for prefix computations was probably the list ranking algorithm of Wyllie [14], which requires at least $n \log n$ non-contiguous accesses. Other parallel algorithms which improved upon this result include those of of Miller and Reif [9] ($5n$ non-contiguous accesses), Anderson and Miller [3] ($4n$ non-contiguous accesses), and Reid-Miller and Blelloch [10] ($2n$ non-contiguous accesses - see [6] for details of this analysis). Clearly, however, none of these approach the memory requirement of our optimal sequential algorithm, which seems necessary to be practically significant on the relatively small number of processors available on the SMP.

3.1 A New Algorithm for Prefix Computations

A high-level description of our algorithm proceeds as follows. We first identify the head of the list using the same procedure as in our optimal sequential algorithm. We then partition the input list into s sublists by randomly choosing exactly one *splitter* from each memory block of $\frac{n}{(s-1)}$ elements, where s is $\Omega(p \log n)$ (the list head is also designated as a splitter). Corresponding to each of these sublists is a record in an array called *Sublists*. We then traverse each of these sublists, making a note at each list element of the index of its sublist and the prefix value of that element within the sublist. The results of these sublist traversals are also used to create a linked list of the records in *Sublists*, where the input value of each node is simply the sublist prefix value of the last element in the previous sublist. We then determine the prefix values of the records in the *Sublists* array by sequentially traversing this list from its head. Finally, for each element in the input list, we apply the prefix operation between its current prefix input value (which is its sublist prefix value) and the prefix value of the corresponding *Sublists* record to obtain the desired result.

The pseudo-code of our algorithm is as follows, in which the input consists of an array of n records called *List*. Each record consists of two fields, *successor* and *prefix_data*, where *successor* gives the integer index of the successor of that element and *prefix_data* initially holds the input value for the prefix operation. The output of the algorithm is simply the *List* array with the properly computed prefix value in the *prefix_data* field. Note that as mentioned above we also make use of an intermediate array of records called *Sublists*. Each *Sublists* record consists of the four fields *head*, *scratch*, *prefix_data*, and *successor*, whose purpose is detailed in the pseudo-code.

- **(1):** Processor P_i $(0 \leq i \leq p-1)$ visits the list elements with array indices $\frac{in}{p}$ through $\left(\frac{(i+1)n}{p} - 1\right)$ in order of increasing index and computes the sum of the successor indices. Note that in doing this a negative valued successor index is ignored since by convention it denotes the terminal list element - this negative successor index is however replaced by the value $(-s)$ for future convenience. Additionally, as each element of *List* is read, the value in the successor field is preserved by copying it to an identically indexed location in the array *Succ*. The resulting sum of the successor indices is stored in location i of the array Z.

- **(2):** Processor P_0 computes the sum T of the p values in the array Z. The index of the head of the list is then $h = \left(\frac{1}{2}n(n-1) - T\right)$.

- **(3):** For $j = \frac{is}{p}$ up to $\left(\frac{(i+1)s}{p} - 1\right)$, processor P_i randomly chooses a location x from the block of list elements with indices $\left((j-1)\frac{n}{(s-1)}\right)$ through $\left(j\frac{n}{(s-1)} - 1\right)$ as a splitter which defines the head of a sublist in *List* (processor P_0 chooses the head of the list as its first splitter). This is recorded by setting *Sublists[j].head* to x. Additionally, the value of *List[x].successor* is copied to *Sublists[j].scratch*, after which *List[x].successor* is replaced with the value $(-j)$ to denote both the beginning of a new sublist and the index of the record in *Sublists* which corresponds to its sublist.

- **(4):** For $j = \frac{is}{p}$ up to $\left(\frac{(i+1)s}{p} - 1\right)$, processor P_i traverses the elements in the sublist which begins with *Sublists[j].head* and ends at the next element which has been chosen as a splitter (as evidenced by a negative value in the *successor* field). For each element traversed with index x and predecessor *pre* (excluding the first element in the sublist), we set *List[x].successor* = -j to record the index of the record in *Sublists* which corresponds to that sublist. Additionally, we record the prefix value of that element within its sublist by setting *List[x].prefix_data* = *List[x].prefix_data* $\otimes$ *List[pre].prefix_data*. Finally, if x is also the last element in the sublist (but not the last element in the list) and k is the index of the record in *Sublists* which corresponds to the successor of x, then we also set *Sublists[j].successor* $= k$ and *Sublists[k].prefix_data* = *List[x].prefix_data*. Finally, the *prefix_data* field of *Sublists[0]*, which corresponds to the sublist at the head of the list is set to the prefix operator identity.

- **(5):** Beginning at the head, processor P_0 traverses the records in the array *Sublists* by following the successor pointers from the head at *Sublists[0]*. For each record traversed with index j and predecessor *pre*, we compute the prefix value by setting *Sublists[j].prefix_data* = *Sublists[j].prefix_data* $\otimes$ *Sublists[pre].prefix_data*.
- **(6):** Processor P_i visits the list elements with array indices $\frac{in}{p}$ through $\left(\frac{(i+1)n}{p} - 1\right)$ in order of increasing index and completes the prefix computation for each list element x by setting *List[x].prefix_data* = *List[x].prefix_data* $\otimes$ *Sublists[-(List[x].successor)].prefix_data*. Additionally, as each element of *List* is read, the value in the successor field is replaced with the identically indexed element in the array *Succ*. Note that is reasonable to assume that the entire array of s records which comprise *Sublists* can fit into cache.

We can establish the complexity of this algorithm with high probability - that is with probability $\geq (1 - n^{-\epsilon})$ for some positive constant ϵ. But before doing this, we need the results of the following Lemma, whose proof has been omitted for brevity [6].

Lemma 1: The number of list elements traversed by any processor in Step (4) is at most $\alpha\frac{n}{p}$ *with high probability*, for any $\alpha(s) \geq 2.62$ (read $\alpha(s)$ as "the function α of s"), $s \geq (p \ln n + 1)$, and $n > p^2 \ln n$.

With this result, the analysis of our algorithm is as follows. In Step (1), each processor moves through a contiguous portion of the list array to compute the sum of the indices in the *successor* field and to preserve these indices by copying them to the array *Succ*. When this task is completed, the sum is written to the array Z. Since this is done in order of increasing array index, it requires only three non-contiguous memory accesses to exchange approximately $\frac{2n}{p}$ elements with main memory and $O\left(\frac{n}{p}\right)$ computation time. In Step (2), processor P_0 computes the sum of the p entries in the array Z. Since this is done in order of increasing array index, this step requires only a single non-contiguous memory accesses to exchange p elements with main memory and $O(p)$ computation time. In Step (3), each processor randomly chooses $\frac{s}{p}$ splitters to be the heads of sublists. For each of these sublists, it copies the index of the corresponding record in the *Sublists* array into the successor field of the splitter. While the *Sublists* array is traversed in order of increasing array index, the corresponding splitters may lie in mutually non-contiguous locations and so the whole process may require $\frac{s}{p}$ non-contiguous memory accesses to exchange $\frac{2s}{p}$ elements with main memory and $\frac{s}{p}$ computation time. In Step (4), each processor traverses the sublist associated with each of its $\frac{s}{p}$ splitters, which together contain at most $\alpha(s)\frac{n}{p}$ elements *with high probability*. As each sublist is completed, the prefix value of the last element in the subarray is written to the record in the *Sublists* array which corresponds to the succeeding sublist. Since we can reasonably assume that $(s << n)$ and can therefore ignore the cost of writing to the *Sublists* array, this step requires approximately $\alpha(s)\frac{n}{p}$ non-contiguous memory accesses to exchange approximately $\alpha(s)\frac{n}{p}$ elements with main memory and $O\left(\frac{n}{p}\right)$ computation time *with high probability* . However, it is important to note that an $\frac{s}{n}$-biased binomial process requires *on average* $\frac{n}{s}$ events before encountering the first success and so *on average* each processor traverses about $\frac{n}{p}$ list elements (which is what we observe experimentally in the next section). In Step (5), processor P_0 traverses the the linked list of s records in the *Sublists* array established in Step (4) to compute their prefix values, which requires s non-contiguous memory accesses to exchange s elements with main memory and $O(s)$ computation time. Finally, in Step (6), each processor completes the prefix values for a contiguous chunk of the input list by first looking up the prefix value of the record in *Sublists* which maps to the head of its sublist. Since we make the reasonable assumption that the entire array of s records which comprise *Sublists* will fit into the cache, which is the case for all three platforms considered in this paper and the choices for n, accessing the prefix values in the *Sublists* array will only require s non-contiguous memory accesses (non-contiguous because we are assuming they are accessed in the order of request). As the computation of the prefix value for an element is completed, the correct value is restored to its *successor* field from the array *Succ*. Hence, this step will require approximately $(s+1)$ non-contiguous memory accesses to exchange approximately $\frac{2n}{p}$ elements with main memory and $O\left(\frac{n}{p}\right)$ computation time. Thus, *with high probability*, the overall complexity of our prefix computation algorithm is given by

$$T(n,p) = \langle M_A(n,p); M_E(n,p); T_C(n,p)\rangle \quad (2)$$

$$= \langle \alpha(s)\frac{n}{p}; \left((\alpha(s)+4)\frac{n}{p}\right); O\left(\frac{n}{p}\right)\rangle \quad (3)$$

for $\alpha(s) \geq 2.62$, $s \geq (p \ln n + 1)$, $n >> s$, and $n > p^2 \ln n$. Noting that the relatively expensive M_A non-contiguous memory accesses comprise a substantial proportion of the M_E total elements exchanges with memory, and recalling that *on average* each processor traverses only about $\frac{n}{p}$ elements in Step (4), we would expect that in practice the complexity of our algorithm could be characterized as

$$T(n,p) = \langle M_A(n,p); T_C(n,p)\rangle \quad (4)$$

$$= \langle \frac{n}{p}; O\left(\frac{n}{p}\right)\rangle, \quad (5)$$

Notice that our algorithm's requirement of approximately n non-contiguous memory accesses is nearly half the cost of Reid-Miller and Blelloch and compares very closely with the requirements of the optimal sequential algorithm.

3.2 Performance Evaluation

Both our parallel algorithm and the optimal sequential algorithm were implemented in C using POSIX threads and run on a DEC AlphaServer 2100A system, an SGI Power Challenge, and an HP-Convex Exemplar. To evaluate these algorithms, we examined the prefix operation of of floating point addition on three different benchmarks, which were selected to compare the impact of various memory access patterns. These benchmarks are the **Random [R]**, in which each successor is randomly chosen, the **Stride [S]**, in which each successor is (wherever possible) some stride S away, and the **Ordered [O]**, in which which element is tb]

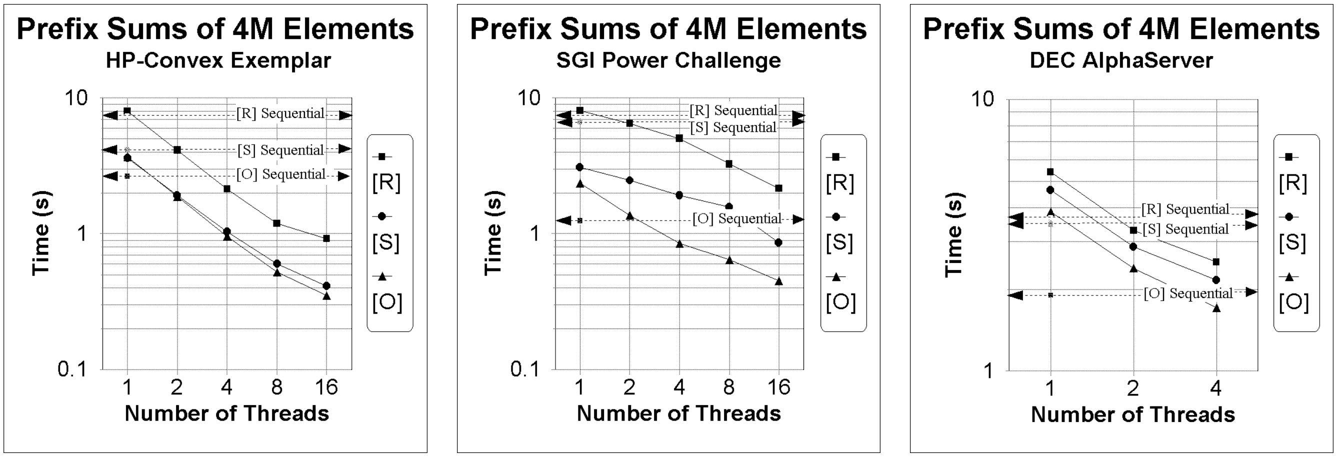

Figure 1: Comparison between the performance of our parallel algorithm and our optimal sequential algorithm on three different platforms using three different benchmarks.

paced in the array according to its rank. See [6] for a more detailed description and justification of these benchmarks.

The graphs in **Figure 1** compare the performance of our optimal parallel prefix computation algorithm with that of our optimal sequential algorithm. Notice first that our parallel algorithm almost always outperforms the optimal sequential algorithm with only one or two threads. The only exception is the [O] benchmark, where the successor of an element is always the next location in memory. Notice also that for a given algorithm, the [O] benchmark is almost always solved more quickly than the [S] benchmark, which in turn is always solved more quickly than the [R] benchmark. A step by step breakdown of the execution time in **Table I** verifies that these differences are entirely due to the time required for the sublist traversal in Step (4). This agrees well with our theoretical expectations, since in the [R] benchmark, the location of the successor is randomly chosen, so almost every step in the traversal involves accessing a non-contiguous location in memory. By contrast, in the [O] benchmark, the memory location of the successor is always the successive location in memory, which in all likelihood is already present in cache. Finally, the [S] benchmark is designed so that where possible the successor is always a constant stride away. Even though we chose the stride to be 1001, so that each step of the sublist traversal should involve accessing a non-contiguous location

Step:	**Number of Threads & Benchmark**														
	[1]			[2]			[4]			[8]			[16]		
	[R]	[S]	[O]	[R]	[S]	[O]	[R]	[S]	[O]	[R]	[S]	[O]	[R]	[S]	[O]
(1)-(3):	0.59	0.87	0.66	0.34	0.40	0.34	0.18	0.21	0.18	0.10	0.12	0.10	0.08	0.08	0.08
(4):	6.69	1.86	2.33	3.40	1.08	1.17	1.75	0.57	0.59	0.96	0.31	0.30	0.74	0.22	0.18
(5):	0.01	0.12	0.01	0.01	0.04	0.01	0.01	0.05	0.01	0.01	0.06	0.01	0.01	0.02	0.01
(6):	0.69	0.75	0.69	0.37	0.38	0.35	0.21	0.20	0.19	0.11	0.12	0.11	0.09	0.12	0.08
Total:	7.97	3.60	3.68	4.12	1.91	1.87	2.14	1.03	0.97	1.19	0.60	0.52	0.92	0.41	0.35

Table I: Comparison of the time (in seconds) required as a function of the benchmark for each step of computing the prefix sums of 4M list elements on an HP-Convex Exemplar, for a variety of threads.

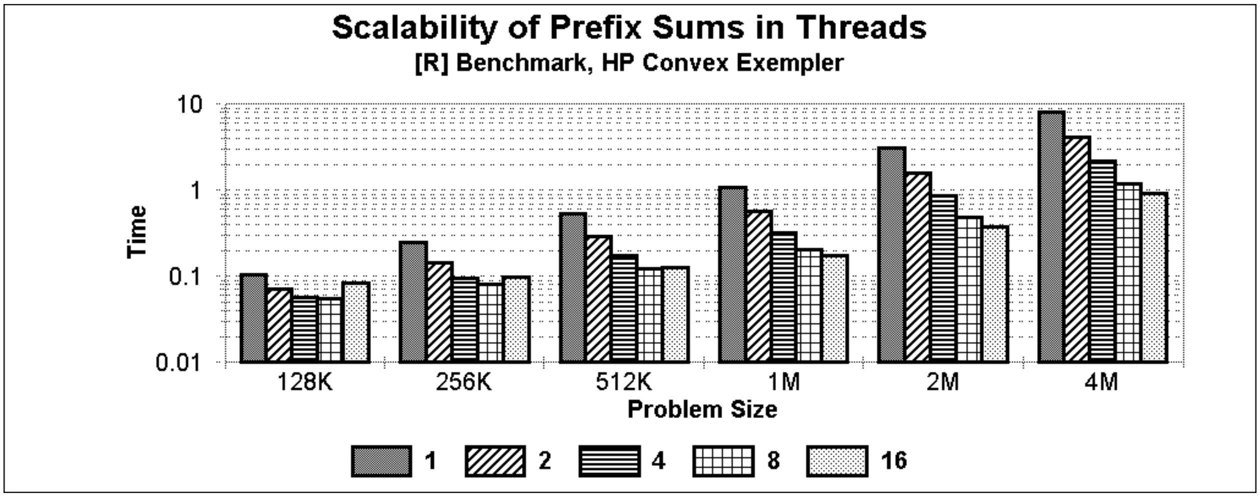

Figure 2: Scalability of our prefix computation algorithm on the HP-Convex Exemplar with respect to the number of threads, for differing problem sizes.

in memory, the constant stride still allows it to take advantage of cache pre-fetching. Lastly, notice that, in Table I, the n noncontiguous memory required by the [R] benchmark in Step (4) consume on average almost five time as much time as the $4n$ contiguous memory accesses of Steps (1) and (6). Taken as a whole, these results strongly support the emphasis we place on minimizing the number of non-contiguous memory accesses in this problem.

The graph in **Figure 2** examines the scalability of our prefix computation algorithm as a function of the number of threads. Bearing in mind that these graphs are log-log plots, they show that for large enough inputs, the execution time decreases as we increase the number of threads p, which is the expectation of our model. For smaller inputs, this inverse relationship between the execution time and the number of threads deteriorates. In this case, such performance is quite reasonable if we consider the fact that for small problem sizes the size of the cache approaches that of the problem. This introduces a number of issues which are beyond the intended scope of our algorithm.

4 Generalized Sorting

Consider the problem of sorting n elements equally distributed amongst p processors, where we assume without loss of generality that p divides n evenly. Any of the algorithms that have been proposed in the literature for sorting on hierarchical memory models can be considered for possible implementation on an SMP. However, without modifications, most are unnecessarily complex or inefficient for a relatively simple platform such as ours. A notable exception is the algorithm of Varman et al. [12]. Yet another approach is an adaptation of our sorting by regular sampling algorithm [8], which we originally developed for distributed memory machines.

4.1 A New Algorithm for Generalized Sorting

The idea behind sorting by regular sampling is to find a set of $p-1$ *splitters* to partition the n input elements into p groups indexed from 0 up to $p-1$ such that every element in the i^{th} group is less than or equal to each of the elements in the $(i+1)^{th}$ group, for $(0 \leq i \leq p-1)$. Then the task of sorting each of the p groups can be turned over to the correspondingly indexed processor, after which the n elements will be arranged in sorted order. One way to chose the splitters is by regularly sampling the input elements - hence the name Sorting by Regular Sampling. As modified for an SMP, this algorithm is similar to the parallel sorting by regular sampling (PSRS) algorithm of Shi and Schaeffer [11], except that our algorithm can be easily implemented as a stable sort.

The pseudocode for our algorithm is as follows, where C is the size of the cache and L is the cache line size:

- **(1)** Each processor P_i $(0 \leq i \leq p-1)$ sorts the subsequence of the n input elements with indices $\left(\frac{in}{p}\right)$ through $\left(\frac{(i+1)n}{p} - 1\right)$ as follows:
  - **(A)** Sort each block of m input elements using sequential merge sort, where $m \leq \frac{C}{2}$.
  - **(B)** For $j = 0$ up to $\left(\frac{\log(n/pm)}{\log(z)} - 1\right)$, merge the sorted blocks of size $\left(mz^j\right)$ using z-way merge, where $z \leq \frac{C}{L}$.
- **(2)** Each processor P_i selects each $\left(\frac{in}{p} + (k+1)\frac{n}{ps}\right)^{th}$ element as a sample, for $(0 \leq k \leq s-1)$ and a given value of s $\left(p \leq s \leq \frac{n}{p^2}\right)$.
- **(3)** Processor $P_{(p-1)}$ merges the p sorted subsequences of samples and then selects each $((k+1)s)^{th}$ sample as Splitter$[k]$, for $(0 \leq k \leq p-2)$. By default, the p^{th} splitter is the largest value allowed by the data type used. Additionally, binary search is used to compute for the set of samples with indices 0 through $((k+1)s-1)$ the number of samples Est$[k]$ which share the same value as Splitter$[k]$.
- **Step (4):** Each processor P_i uses binary search to define an index $b_{(i,j)}$ for each of the p sorted input sequences created in Step (1). If we define $T_{(i,j)}$ as a subsequence containing the first $b_{(i,j)}$ elements in the j^{th} sorted input sequence, then the set of p subsequences $\{T_{(i,1)}, T_{(i,2)}, ..., T_{(i,p)}\}$ will contain all those values in the input set which are strictly less than Splitter$[i]$ and *at most* $\left(\text{Est}[i] \times \frac{n}{ps}\right)$ elements with the same value as Splitter$[i]$. The term *at most* is used because there may not actually be this number of elements with the same value as Splitter$[i]$.
- **Step (5):** Each processor P_i merges those subsequences of the sorted input sequences which lie between indices $b_{((i-1),j)}$ and $b_{(i,j)}$ using p-way merge. It can be shown [8] that no processor will merge more than $\left(\frac{n}{p} + \frac{n}{s} - p\right)$ elements.

Before establishing the complexity of this algorithm, we need the results of the following lemma, whose proof has been omitted for brevity [8]:

Lemma 2: At the completion of the partitioning in Step (4), no more than $\left(\frac{n}{p}+\frac{n}{s}-p\right)$ elements will be associated with any splitter, for $n \geq p^3$ and $\left(p \leq s \leq \frac{n}{p^2}\right)$.

With this result, the analysis of our algorithm is as follows. In Step (1A), each processor moves through a contiguous portion of the input array to sort it in blocks of size m using sequential merge sort. If we assume that $\left(m \leq \frac{C}{2}\right)$, this will require only a single non-contiguous memory accesses to exchange $\frac{2n}{p}$ elements with main memory and $O\left(\frac{n}{p}\log m\right)$ computation time. Step (1B) involves $\frac{\log(n/pm)}{\log(z)}$ rounds of z-way merge. Since round j will begin with $\frac{n}{pmz^j}$ blocks of size mz^j, this will require at most $\frac{2nz}{pm(z-1)}$ non-contiguous memory accesses to exchange $\frac{2n\log(n/pm)}{p\log(z)}$ elements with main memory memory and $O\left(\frac{n}{p}log\left(\frac{n}{pm}\right)\right)$ computation time. The selection of s noncontiguous samples by each processor in Step (2) requires s non-contiguous memory accesses to exchange $2s$ elements with main memory and $O(s)$ computation time. Step (3) involves a p-way merge of blocks of size s followed by p binary searches on segments of size s. Hence, it requires approximately $p\log(s)$ non-contiguous memory accesses to exchange approximately $2sp$ elements with main memory and $O(sp\log p)$ computation time. Step (4) involves p binary searches by each processor on segments of size $\frac{n}{p}$ and hence requires approximately $p\log\left(\frac{n}{p}\right)$ non-contiguous memory accesses to exchange approximately $p\log\left(\frac{n}{p}\right)$ elements with main memory and $O\left(p\log\left(\frac{n}{p}\right)\right)$ computation time. Step (5) involves a p-way merge of p blocks of total maximum size $\left(\frac{n}{p}+\frac{n}{s}-p\right)$, requiring approximately p non-contiguous memory accesses to exchange approximately $2\left(\frac{n}{p}+\frac{n}{s}\right)$ elements with main memory and $O\left(\frac{n}{p}\right)$ computation time. Hence, the overall complexity of our sorting algorithm is given by

$$\begin{aligned} T(n,p) &= \langle M_A(n,p); M_E(n,p); T_C(n,p)\rangle && (6)\\ &= \langle\left(\frac{nz}{pm(z-1)}+s+p\log\left(\frac{n}{p}\right)\right); \left(\left(2\frac{\log(n/pm)}{\log(z)}+2\right)\frac{n}{p}+2\frac{n}{s}\right); O\left(\frac{n}{p}\log n\right)\rangle && (7)\end{aligned}$$

for $n \geq p^2\log\left(\frac{n}{p}\right)$, $\left(p \leq s \leq \frac{n}{p^2}\right)$, $m \leq \frac{C}{2}$, and $z \leq \frac{C}{L}$. Since the analysis suggests that the parameters m and z should be as large as possible subject to the stated constraints while selecting s so that $\left(p << s << \frac{n}{p}\right)$, we would expect that in practice the complexity of our algorithm could be characterized as

$$\begin{aligned} T(n,p) &= \langle M_E(n,p); T_C(n,p)\rangle && (8)\\ &= \langle\left(4+2\frac{\log(n/pm)}{\log(z)}\right)\frac{n}{p}; O\left(\frac{n}{p}\log n\right)\rangle. && (9)\end{aligned}$$

4.2 Performance Evaluation

Both the sequential algorithm and our parallel algorithm were implemented in C using POSIX threads and run on a DEC AlphaServer 2100A system, an SGI Power Challenge, and an HP-Convex Exemplar. We ran our code using six widely different benchmarks which were selected to test the dependence of our algorithm on the input distribution. A detailed description and justification of these benchmarks is presented in [7]. The results in **Table II** verify that as expected performance does not significantly depend on the input distribution. Because of this independence, the remainder of this section will only discuss performance on the single benchmark [U], in which the input data forms a uniform random distribution.

Input Size	Benchmark					
	[U]	**[G]**	**[Z]**	**[WR]**	**[DD]**	**[RD]**
512K	0.397	0.394	0.320	0.421	0.337	0.348
1M	0.868	0.856	0.741	0.844	0.724	0.710
2M	1.64	1.72	1.39	1.73	1.40	1.51
4M	3.50	3.47	3.00	3.52	3.01	2.98

Table II: Sorting doubles (in seconds) using 4 threads on a DEC AlphaServer 2100A.

Block Size	Denomination of z-Way Merge											
	2	**4**	**8**	**16**	**32**	**64**	**128**	**256**	**512**	**1024**	**2048**	**4096**
1K	30.43	21.49	19.22	17.66	17.87	16.20	16.74	16.91	16.87	16.82	16.73	16.29
2K	29.14	21.65	19.15	18.02	17.95	16.63	16.91	16.98	16.79	18.29	15.86	
4K	27.96	20.55	19.62	18.29	16.65	17.00	17.14	17.06	16.86	15.91		
8K	27.59	21.55	19.18	19.27	17.90	18.07	18.04	17.84	17.08			
16K	26.69	20.73	19.84	18.21	18.50	18.53	18.43	17.83				
32K	27.77	23.14	22.14	20.81	20.88	20.90	20.41					
64K	29.98	25.46	24.26	24.51	24.51	23.06						
128K	37.54	34.19	33.26	33.36	31.84							
256K	39.85	36.51	36.74	35.37								
512K	39.78	37.81	36.54									
1M	39.53	37.62										
2M	39.25											
4M	38.86 - (No z-way merge is necessary for this block size)											

Table III: Time (in seconds) required on the HP-Convex Exemplar to sort 4M doubles using a single thread as a function of M and z.

Table III displays the times required to sort 4M *doubles* (i.e. double precision floating point values) on the HP-Convex Exemplar using a single thread as a function of m and z. Notice first that performance suffers dramatically when the block size reaches 1MB (128K eight byte double precision numbers), which is the limit of the cache on the Exemplar. This is expected, since sorting a block in Step (1A) now requires that data be repeatedly swapped to main memory. Consider also the data for a given block size - say 1K. The execution time drops as we move from $z = 2$ to $z = 16$. This is reasonable since we require 12 rounds of 2-way merge, 6 rounds of 4-way merge, 4 rounds of 8-way merge, and only 3 rounds of 16-way merge, and each round of z-way merge is obviously another round where all the input elements must be brought in from main memory. Moving from $z = 16$ to $z = 32$ has little effect on the execution time since it does nothing to reduce the memory requirements, but moving to $z = 64$ saves a round of memory access and, hence, the execution time is further reduced. However, the most dramatic illustration of the importance of minimizing secondary memory access can be found by comparing the optimal sorting time of 15.86 seconds for $m = 2K$ and $z = 2048$ with the time of 39.25 seconds required to sort using only binary merge sort. Reducing memory access by a combination of block sorting and z-way merging improved the performance by 60%. Clearly, such results strongly support the attention that we place in this algorithm on the number of contiguous memory accesses.

The graph in **Figure 3** examines the scalability of our sorting algorithm as a function of the number

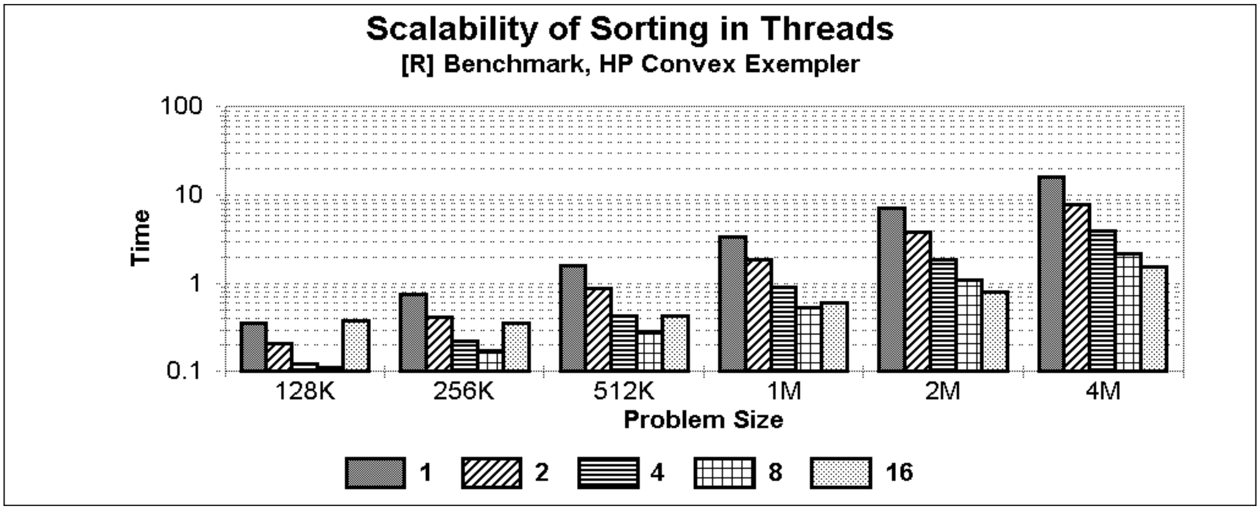

Figure 3: Scalability of our generalized sorting algorithm on the HP-Convex Exemplar with respect to the number of threads, for differing problem sizes.

of threads, for a variety of problem sizes. Bearing in mind that these graphs are log-log plots, they show that for large enough inputs, the execution time decreases as we increase the number of threads p, which is the expectation of our model. For smaller inputs, this inverse relationship between the execution time and the number of threads deteriorates when we move to 16 threads. This explanation for this problem may lie in the fact that when we moved to 16 threads on this platform, the data suddenly became very erratic, perhaps because some threads now had to compete with operating system processes for access to one of the 16 processors.

References

[1] A. Aggarwal, A. Chandra, and M. Snir. Heirarchical Memory with Block Transfer. In *Proceedings of the 28th Annual IEEE Symposium on Foundations of Computer Science*, pages 204–216, October 1987.

[2] B. Alpern, L. Carter, E. Feig, and T. Selker. The Uniform Memory Hierarchy Model of Compuatation. *Algorithmica*, 12:72–109, 1994.

[3] R. Anderson and G. Miller. Deterministic Parallel List Ranking. In *Proceedings Third Aegean Workshop on Computing, AWOC 88*, pages 81–90, Corfu, Greece, June/July 1988. Springer-Verlag.

[4] G.E. Blelloch, P.B. Gibbons, Y. Matias, and M. Zagha. Accounting for Memory Bank Contention and Delay in High-Bandwidth Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 8(9):943–958, 1997.

[5] P.B. Gibbons, Y. Matias, and V. Ramachandran. Can a Shared-Memory Model Serve as a Bridging-Model for Parallel Compatation? In *Proceedings of the 9th ACM Symposium on Parallel Algorithms and Architectures*, pages 72–83, June 1997.

[6] D.R. Helman and J. JáJá. Prefix Computations on Symmetric Multiprocessors. Technical Report CS-TR-3915 and UMIACS-TR-98-38, UMIACS and Electrical Engineering, University of Maryland, College Park, MD, July 1998.

[7] D.R. Helman and J. JáJá. Sorting on Clusters of SMPs. In *Proceedings of the 12th International Parallel Processing Symposium*, Orlando, Florida, April 1998.

[8] D.R. Helman, J. JáJá, and D.A. Bader. A New Deterministic Parallel Sorting Algorithm With an Experimental Evaluation. Technical Report CS-TR-3670 and UMIACS-TR-96-54, UMIACS and Electrical Engineering, University of Maryland, College Park, MD, August 1996. To appear in the *Journal of Experimental Algorithmics.*

[9] G. L. Miller and J. H. Reif. Parallel Tree Contraction and its Application. In *Proceedings Twenty-Sixth Annual IEEE Symposium on Foundations of Computer Science*, pages 478–489, Portland. OR, October 1985.

[10] M. Reid-Miller. List Ranking and List Scan on the Cray C90. *Journal of the Computer and System Sciences*, 53:344–356, 1996.

[11] H. Shi and J. Schaeffer. Parallel Sorting by Regular Sampling. *Journal of Parallel and Distributed Computing*, 14:361–372, 1992.

[12] P. Varman, B. Iyer, and S. scheufler. A Multiprocessor Algorithm for Merging Multiple Sorted Lists. In *Proceedings of the 1990 International Conference on Parallel Processing*, pages 22–26.

[13] J. Vitter and E. Shriver. Algorithms for Parallel Memory I: Two-Level Memories. *Algorithmica*, 12:110–147, 1994.

[14] J.C. Wyllie. *The Complexity of Parallel Computations.* PhD thesis, Department of Computer Science, Cornell University, Ithica, NY, 1979.