

Optimization of Collective Communication Operations in MPICH

Rajeev Thakur*

Rolf Rabenseifner†

William Gropp*

Abstract

We describe our work on optimizing the collective communication operations in MPICH for clusters connected by switched networks. For each collective operation, we use multiple algorithms depending on the message size, with the goal of minimizing latency for short messages and minimizing bandwidth use for long messages. Although we have implemented new algorithms for all MPI collective operations, because of limited space we describe only the algorithms for allgather, broadcast, all-to-all, reduce-scatter, reduce, and allreduce. Performance results on a Myrinet-connected Linux cluster and an IBM SP indicate that, in all cases, the new algorithms significantly outperform the old algorithms used in MPICH on the Myrinet cluster, and, in many cases, they outperform the algorithms used in IBM's MPI on the SP. We also explore in further detail the optimization of two of the most commonly used collective operations, allreduce and reduce, particularly for long messages and non-power-of-two numbers of processes. The optimized algorithms for these operations perform several times better than the native algorithms on a Myrinet cluster, IBM SP, and Cray T3E. This work demonstrates that to achieve the best performance for a collective communication operation, we need to use a number of different algorithms and select the right algorithm for a particular message size and number of processes.

Keywords: Message Passing, MPI, Collective Communication, Reduction

1 Introduction

Collective communication is an important and frequently used component of MPI and offers implementations considerable room for optimization.

*Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439, USA. {thakur, gropp}@mcs.anl.gov

†High Performance Computing Center (HLRS), University of Stuttgart, Allmandring 30, D-70550 Stuttgart, Germany. rabenseifner@hlrs.de, www.hlrs.de/people/rabenseifner/

MPICH [12], although widely used as an MPI implementation, has until now had fairly rudimentary implementations of the collective operations. We have recently focused on improving the performance of all the collective operations in MPICH. Our initial target architecture is the one that is the most popular among our users, namely, clusters of machines connected by a switch, such as Myrinet or the IBM SP switch. For each collective operation, we use multiple algorithms based on message size: The short-message algorithms aim to minimize latency, and the long-message algorithms aim to minimize bandwidth use. Our approach has been to identify the best algorithms known in the literature, improve on them or develop new algorithms where necessary, and implement them efficiently. We have done extensive experimental analysis to determine the cutoff points for switching between different algorithms depending on the message size and number of processes. We have implemented new algorithms for all the MPI collective operations, namely, scatter, gather, allgather, broadcast, reduce, allreduce, reduce-scatter, scan, and barrier. Because of limited space, however, we describe and present performance results for only the new algorithms for allgather, broadcast, all-to-all, reduce-scatter, reduce, and allreduce.

Based on the results of a five-year profiling study of applications running in production mode on a Cray T3E at the University of Stuttgart [15], we chose to optimize further the two most commonly used collective operations: allreduce and reduce. We present a detailed study of different ways of implementing allreduce and reduce, particularly for long messages and non-power-of-two numbers of processes, which occur frequently according to the profiling study.

The rest of this paper is organized as follows. In Section 2, we describe related work in the area of collective communication. In Section 3, we describe the cost model used to guide the selection of algorithms. In Section 4, we describe the new algorithms in MPICH and their performance. In Section 5, we study in detail the further optimization of reduce and allreduce. In Section 6, we conclude with a brief discussion of future work.

2 Related Work

Early work on collective communication focused on developing optimized algorithms for particular architectures, such as hypercube, mesh, or fat tree, with an emphasis on minimizing link contention, node contention, or the distance between communicating nodes [2, 4, 5, 18]. More recently, Vadhiyar et al. have developed automatically tuned collective communication algorithms [24]. Their approach consists of running tests to measure system parameters and then tuning their algorithms for those parameters. Researchers in Holland and at Argonne have optimized MPI collective communication for wide-area distributed environments [9, 10]. In such environments, the goal is to minimize communication over slow wide-area links at the expense of more communication over faster local-area connections. Researchers have also developed collective communication algorithms for clusters of SMPs [17, 20, 22, 23], where communication within an SMP is done differently from communication across a cluster. Some efforts have focused on using different algorithms for different message sizes, such as the work by Van de Geijn et al. for the Intel Paragon [1, 11, 19], by Rabenseifner on reduce and allreduce [14], and by Kale et al. on all-to-all communication [8]. Benson et al. studied the performance of the allgather operation in MPICH on Myrinet and TCP networks and developed a dissemination allgather based on the dissemination barrier algorithm [3]. Bruck et al. proposed algorithms for allgather and all-to-all that are particularly efficient for short messages [6].

3 Cost Model

We use a simple model to estimate the cost of the collective communication algorithms in terms of latency and bandwidth use and to guide the selection of algorithms for a particular collective communication operation. We assume that the time taken to send a message between any two nodes can be modeled as $\alpha + n\beta$, where α is the latency (or startup time) per message, independent of message size, β is the transfer time per byte, and n is the number of bytes transferred. We assume further that the time taken is independent of how many pairs of processes are communicating with each other, independent of the distance between the communicating nodes, and that the communication links are bidirectional (that is, a message can be transferred in both directions on the link in the same time as in one direction).

The node's network interface is assumed to be single ported; that is, at most one message can be sent and one message can be received simultaneously. In the case of reduction operations, we assume that γ is the computation cost per byte for performing the reduction operation locally on any process.

4 Algorithms

In this section we describe the new algorithms and their performance. We measured performance by using the SKaMPI benchmark [25] on two platforms: a Linux cluster at Argonne connected with Myrinet 2000 and the IBM SP at the San Diego Supercomputer Center. On the Myrinet cluster we used MPICH-GM and compared the performance of the new algorithms with the old algorithms in MPICH-GM. On the IBM SP, we used IBM's MPI and compared the performance of the new algorithms with the algorithms used in IBM's MPI. On both systems, we ran one MPI process per node. We implemented the new algorithms as functions on top of MPI point-to-point operations, so that we could compare performance simply by linking or not linking the new functions.

4.1 Allgather

`MPI_Allgather` is a gather operation in which the data contributed by each process is gathered on all processes, instead of just the root process as in `MPI_Gather`. The old algorithm for allgather in MPICH uses a ring method in which the data from each process is sent around a virtual ring of processes. In the first step, each process i sends its contribution to process $i + 1$ and receives the contribution from process $i - 1$ (with wrap-around). From the second step onward each process i forwards to process $i + 1$ the data it received from process $i - 1$ in the previous step. If p is the number of processes, the entire algorithm takes $p - 1$ steps. If n is the total amount of data to be gathered on each process, then at every step each process sends and receives $\frac{n}{p}$ amount of data. Therefore, the time taken by this algorithm is given by $T_{ring} = (p - 1)\alpha + \frac{p-1}{p}n\beta$. Note that the bandwidth term cannot be reduced further because each process must receive $\frac{n}{p}$ data from $p - 1$ other processes. The latency term, however, can be reduced by using an algorithm that takes $\lg p$ steps. We consider two such algorithms: recursive doubling and the Bruck algorithm [6].

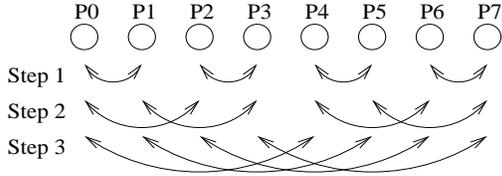


Figure 1: Recursive doubling for allgather

4.1.1 Recursive Doubling

Figure 1 illustrates how recursive doubling works. In the first step, processes that are a distance 1 apart exchange their data. In the second step, processes that are a distance 2 apart exchange their own data as well as the data they received in the previous step. In the third step, processes that are a distance 4 apart exchange their own data as well the data they received in the previous two steps. In this way, for a power-of-two number of processes, all processes get all the data in $\lg p$ steps. The amount of data exchanged by each process is $\frac{n}{p}$ in the first step, $\frac{2n}{p}$ in the second step, and so forth, up to $\frac{2^{\lg p - 1}n}{p}$ in the last step. Therefore, the total time taken by this algorithm is $T_{rec_dbl} = \lg p \alpha + \frac{p-1}{p}n\beta$.

Recursive doubling works very well for a power-of-two number of processes but is tricky to get right for a non-power-of-two number of processes. We have implemented the non-power-of-two case as follows. At each step of recursive doubling, if the current subtree is not a power of two, we do additional communication (in a logarithmic fashion) to ensure that all processes get the data they would have gotten had the subtree been a power of two. This extra communication is necessary for the subsequent steps of recursive doubling to work correctly. The total number of steps for the non-power-of-two case is bounded by $2 \lceil \lg p \rceil$.

4.1.2 Bruck Algorithm

The Bruck algorithm for allgather [6] (referred to as concatenation) is a variant of the dissemination algorithm for barrier, described in [7]. Both algorithms take $\lceil \lg p \rceil$ steps in all cases, even for non-power-of-two numbers of processes. In the dissemination algorithm for barrier, in each step k ($0 \leq k < \lceil \lg p \rceil$), process i sends a (zero-byte) message to process $(i + 2^k)$ and receives a (zero-byte) message from process $(i - 2^k)$ (with wrap-around). If the same order were used to perform an allgather, it would require communicating noncontiguous data in each step in order to get the right data to the right process (see [3] for details). The Bruck algorithm avoids this problem nicely by a

simple modification to the dissemination algorithm in which, in each step k , process i sends data to process $(i - 2^k)$ and receives data from process $(i + 2^k)$, instead of the other way around. The result is that all communication is contiguous, except that at the end, the blocks in the output buffer must be shifted locally to place them in the right order, which is a local memory-copy operation.

Figure 2 illustrates the Bruck algorithm for an example with six processes. The algorithm begins by copying the input data on each process to the top of the output buffer. In each step k , process i sends to the destination $(i - 2^k)$ all the data it has so far and stores the data it receives (from rank $(i + 2^k)$) at the end of the data it currently has. This procedure continues for $\lceil \lg p \rceil$ steps. If the number of processes is not a power of two, an additional step is needed in which each process sends the first $(p - 2^{\lceil \lg p \rceil})$ blocks from the top of its output buffer to the destination and appends the data it receives to the data it already has. Each process now has all the data it needs, but the data is not in the right order in the output buffer: The data on process i is shifted “up” by i blocks. Therefore, a simple local shift of the blocks downwards by i blocks brings the data into the desired order. The total time taken by this algorithm is $T_{bruck} = \lceil \lg p \rceil \alpha + \frac{p-1}{p}n\beta$.

4.1.3 Performance

From the cost model for the three algorithms, we observe that they all have the same bandwidth cost; recursive doubling and the Bruck algorithm have a lower latency cost than the ring algorithm; and the Bruck algorithm has lower latency than recursive doubling for non-power-of-two numbers of processes. Therefore, it appears that the Bruck algorithm would work well in all cases. In practice, however, we find that the Bruck algorithm is best for short messages and non-power-of-two numbers of processes; recursive doubling is best for power-of-two numbers of processes and short or medium-sized messages; and the ring algorithm is best for long messages and any number of processes and also for medium-sized messages and non-power-of-two numbers of processes.

Figure 3 shows the advantage of the Bruck algorithm over recursive doubling for short messages and non-power-of-two numbers of processes because it takes fewer additional steps. For power-of-two numbers of processes, however, recursive doubling performs better because of the pairwise nature of its communication pattern. As the message size increases, the Bruck algorithm suffers because of the memory

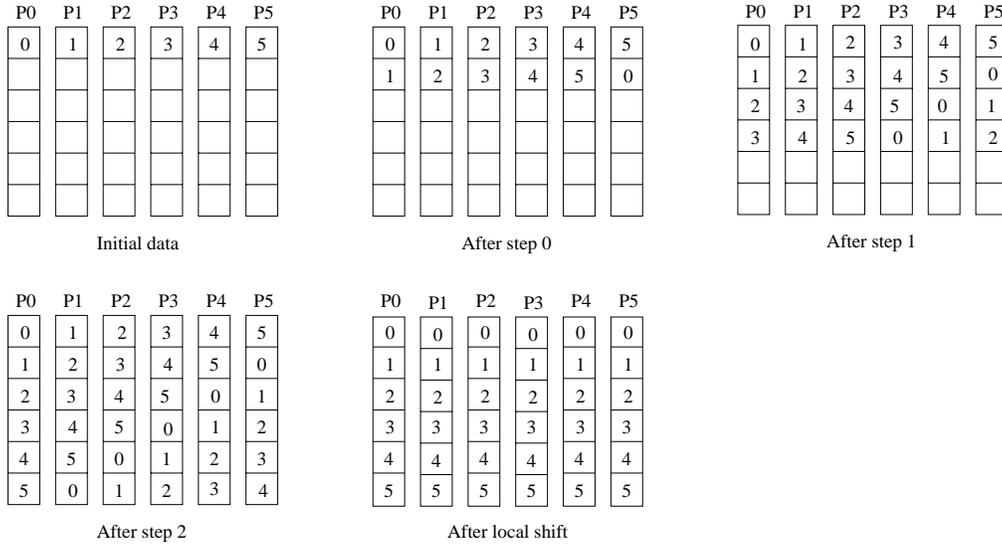


Figure 2: Bruck allgather

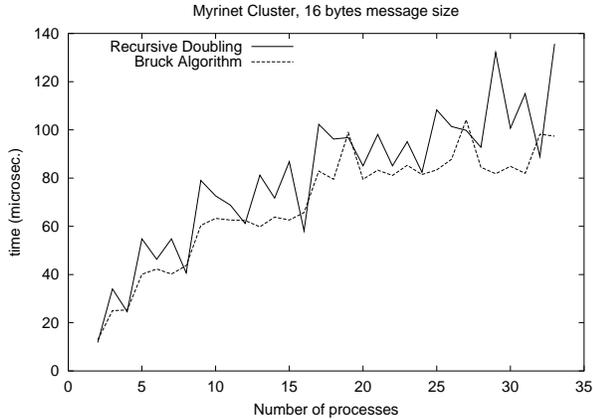


Figure 3: Performance of recursive doubling versus Bruck allgather for power-of-two and non-power-of-two numbers of processes (message size 16 bytes per process).

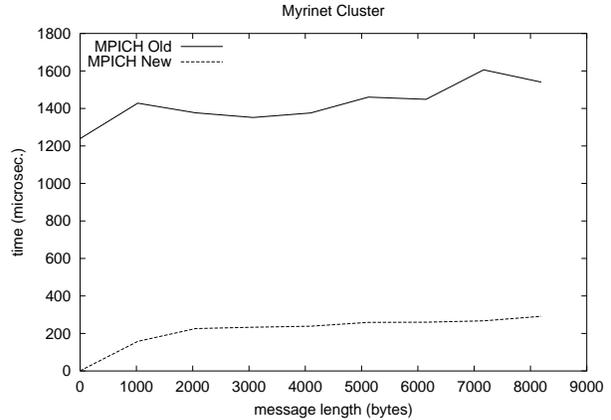


Figure 4: Performance of allgather for short messages (64 nodes). The size on the x-axis is the total amount of data gathered on each process.

copies needed for reordering data in the last step. In MPICH, therefore, we use the Bruck algorithm for short messages (< 80 KB total data gathered) and non-power-of-two numbers of processes, and recursive doubling for power-of-two numbers of processes and short or medium-sized messages (< 512 KB total data gathered). For short messages, the new allgather performs significantly better than the old allgather in MPICH, as shown in Figure 4.

For long messages, the ring algorithm performs better than recursive doubling (see Figure 5). We believe this is because it uses a nearest-neighbor communi-

cation pattern, whereas in recursive doubling, processes that are much farther apart communicate. To confirm this hypothesis, we used the *b_eff* MPI benchmark [13], which measures the performance of about 48 different communication patterns, and found that, for long messages on both the Myrinet cluster and the IBM SP, some communication patterns (particularly nearest neighbor) achieve more than twice the bandwidth of other communication patterns. Therefore, for long messages (≥ 512 KB total data gathered) and any number of processes and also for medium-sized messages and non-power-of-two numbers of processes, we use the ring algorithm.

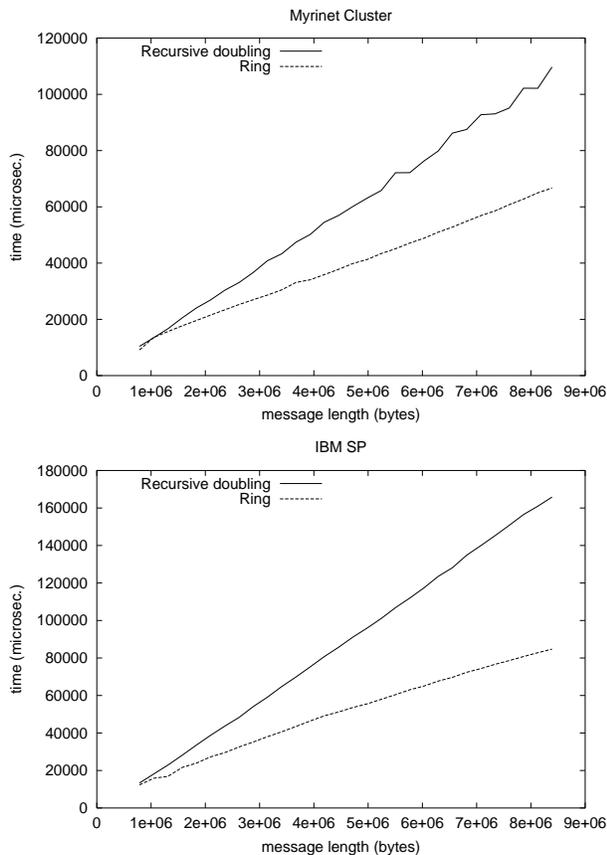


Figure 5: Ring algorithm versus recursive doubling for long-message allgather (64 nodes). The size on the x-axis is the total amount of data gathered on each process.

4.2 Broadcast

The old algorithm for broadcast in MPICH is the commonly used binomial tree algorithm. In the first step, the root sends data to process $(root + \frac{p}{2})$. This process and the root then act as new roots within their own subtrees and recursively continue this algorithm. This communication takes a total of $\lceil \lg p \rceil$ steps. The amount of data communicated by a process at any step is n . Therefore, the time taken by this algorithm is $T_{tree} = \lceil \lg p \rceil (\alpha + n\beta)$.

This algorithm is good for short messages because it has a logarithmic latency term. For long messages, however, a better algorithm has been proposed by Van de Geijn et al. that has a lower bandwidth term [1, 19]. In this algorithm, the message to be broadcast is first divided up and scattered among the processes, similar to an `MPI_Scatter`. The scattered data is then collected back to all processes, similar to

an `MPI_Allgather`. The time taken by this algorithm is the sum of the times taken by the scatter, which is $(\lg p \alpha + \frac{p-1}{p} n\beta)$ for a binomial tree algorithm, and the allgather for which we use either recursive doubling or the ring algorithm depending on the message size. Therefore, for very long messages where we use the ring allgather, the time taken by the broadcast is $T_{vandegeijn} = (\lg p + p - 1)\alpha + 2\frac{p-1}{p}n\beta$.

Comparing this time with that for the binomial tree algorithm, we see that for long messages (where the latency term can be ignored) and when $\lg p > 2$ (or $p > 4$), the Van de Geijn algorithm is better than binomial tree. The maximum improvement in performance that can be expected is $(\lg p)/2$. In other words, the larger the number of processes, the greater the expected improvement in performance. Figure 6 shows the performance for long messages of the new algorithm versus the old binomial tree algorithm in MPICH as well as the algorithm used by IBM's MPI on the SP. In both cases, the new algorithm performs significantly better. Therefore, we use the binomial tree algorithm for short messages (< 12 KB) and the Van de Geijn algorithm for long messages (≥ 12 KB).

4.3 All-to-All

All-to-all communication is a collective operation in which each process has unique data to be sent to every other process. The old algorithm for all-to-all in MPICH does not attempt to schedule communication. Instead, each process posts all the `MPI_Irecv`s in a loop, then all the `MPI_Isend`s in a loop, followed by an `MPI_Waitall`. Instead of using the loop index i as the source or destination process for the `irecv` or `isend`, each process calculates the source or destination as $(rank + i) \% p$, which results in a scattering of the sources and destinations among the processes. If the loop index were directly used as the source or target rank, all processes would try to communicate with rank 0 first, then with rank 1, and so on, resulting in a bottleneck.

The new all-to-all in MPICH uses four different algorithms depending on the message size. For short messages (≤ 256 bytes per message), we use the index algorithm by Bruck et al. [6]. It is a store-and-forward algorithm that takes $\lceil \lg p \rceil$ steps at the expense of some extra data communication ($\frac{p}{2} \lg p \beta$ instead of $n\beta$, where n is the total amount of data to be sent or received by any process). Therefore, it is a good algorithm for very short messages where latency is an issue.

Figure 7 illustrates the Bruck algorithm for an example with six processes. The algorithm begins by do-

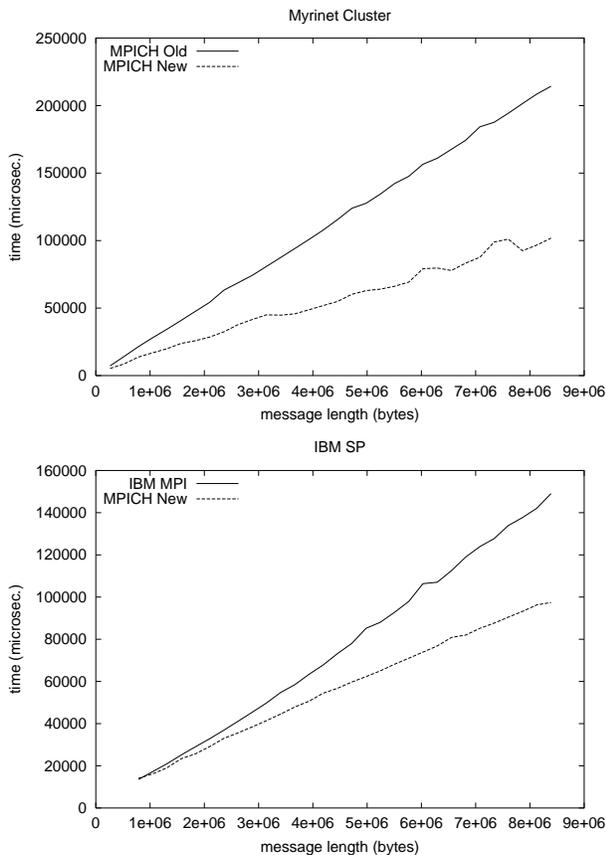


Figure 6: Performance of long-message broadcast (64 nodes)

ing a local copy and “upward” shift of the data blocks from the input buffer to the output buffer such that the data block to be sent by each process to itself is at the top of the output buffer. To achieve this, process i must rotate its data up by i blocks. In each communication step k ($0 \leq k < \lceil \lg p \rceil$), process i sends to rank $(i + 2^k)$ (with wrap-around) all those data blocks whose k th bit is 1, receives data from rank $(i - 2^k)$, and stores the incoming data into blocks whose k th bit is 1 (that is, overwriting the data that was just sent). In other words, in step 0, all the data blocks whose least significant bit is 1 are sent and received (blocks 1, 3, and 5 in our example). In step 1, all the data blocks whose second bit is 1 are sent and received, namely, blocks 2 and 3. After a total of $\lceil \lg p \rceil$ steps, all the data gets routed to the right destination process, but the data blocks are not in the right order in the output buffer. A final step in which each process does a local inverse shift of the blocks (memory copies) places the data in the right order.

The beauty of the Bruck algorithm is that it is

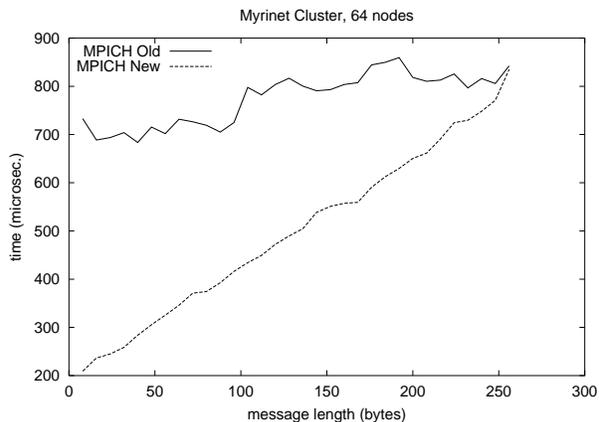


Figure 8: Performance of Bruck all-to-all versus the old algorithm in MPICH (isend-irecv) for short messages. The size on the x-axis is the amount of data sent by each process to every other process.

a logarithmic algorithm for short-message all-to-all that does not need any extra bookkeeping or control information for routing the right data to the right process—that is taken care of by the mathematics of the algorithm.

If n is the total amount of data a process needs to send to or receive from all other processes, the time taken by the Bruck algorithm can be calculated as follows. If the number of processes is a power of two, each process sends and receives $\frac{n}{2}$ amount of data in each step, for a total of $\lg p$ steps. Therefore, the time taken by the algorithm is $T_{bruck} = \lg p \alpha + \frac{n}{2} \lg p \beta$. If the number of processes is not a power of two, in the final step, each process must communicate $\frac{n}{p}(p - 2^{\lfloor \lg p \rfloor})$ data. Therefore, the time taken in the non-power-of-two case is $T_{bruck} = \lceil \lg p \rceil \alpha + (\frac{n}{2} \lg p + \frac{n}{p}(p - 2^{\lfloor \lg p \rfloor})) \beta$.

Figure 8 shows the performance of the Bruck algorithm versus the old algorithm in MPICH (isend-irecv) for short messages. The Bruck algorithm performs significantly better because of its logarithmic latency term. As the message size is increased, however, latency becomes less of an issue, and the extra bandwidth cost of the Bruck algorithm begins to show. Beyond a per process message size of about 256 bytes, the isend-irecv algorithm performs better. Therefore, for medium-sized messages (256 bytes to 32 KB per message), we use the irecv-isend algorithm, which works well in this range.

For long messages and power-of-two number of processes, we use a pairwise-exchange algorithm, which takes $p - 1$ steps. In each step k , $1 \leq k < p$, each

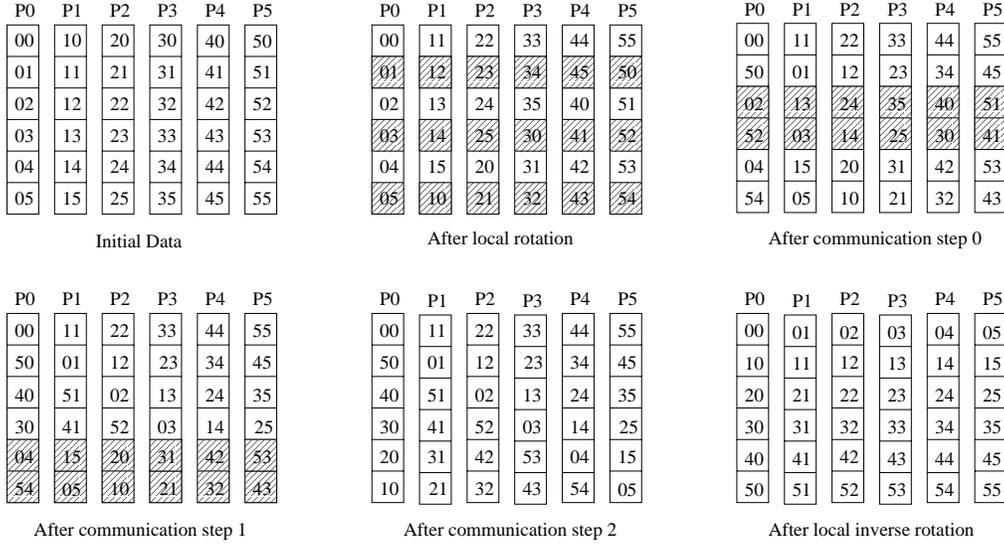


Figure 7: Bruck algorithm for all-to-all. The number ij in each box represents the data to be sent from process i to process j . The shaded boxes indicate the data to be communicated in the next step.

process calculates its target process as $(rank \hat{=} k)$ (exclusive-or operation) and exchanges data directly with that process. This algorithm, however, does not work if the number of processes is not a power of two. For the non-power-of-two case, we use an algorithm in which, in step k , each process receives data from $rank - k$ and sends data to $rank + k$. In both these algorithms, data is directly communicated from source to destination, with no intermediate steps. The time taken by these algorithms is given by $T_{long} = (p - 1)\alpha + n\beta$.

4.4 Reduce-Scatter

Reduce-scatter is a variant of reduce in which the result, instead of being stored at the root, is scattered among all processes. The old algorithm in MPICH implements reduce-scatter by doing a binomial tree reduce to rank 0 followed by a linear scatter. This algorithm takes $\lg p + p - 1$ steps, and the bandwidth term is $(\lg p + \frac{p-1}{p})n\beta$. Therefore, the time taken by this algorithm is $T_{old} = (\lg p + p - 1)\alpha + (\lg p + \frac{p-1}{p})n\beta + n \lg p \gamma$.

In our new implementation of reduce-scatter, for short messages, we use different algorithms depending on whether the reduction operation is commutative or noncommutative. The commutative case occurs most commonly because all the predefined reduction operations in MPI (such as `MPI_SUM`, `MPI_MAX`) are commutative.

For commutative operations, we use a recursive-

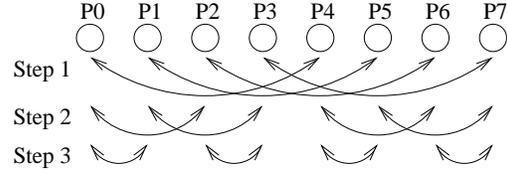


Figure 9: Recursive halving for commutative reduce-scatter

halving algorithm, which is analogous to the recursive-doubling algorithm used for allgather (see Figure 9). In the first step, each process exchanges data with a process that is a distance $\frac{p}{2}$ away: Each process sends the data needed by all processes in the other half, and receives the data needed by all processes in its own half, and performs the reduction operation on the received data. The reduction can be done because the operation is commutative. In the second step, each process exchanges data with a process that is a distance $\frac{p}{4}$ away. This procedure continues recursively, halving the data communicated at each step, for a total of $\lg p$ steps. Therefore, if p is a power of two, the time taken by this algorithm is $T_{rec_half} = \lg p \alpha + \frac{p-1}{p}n\beta + \frac{p-1}{p}n\gamma$. We use this algorithm for messages up to 512 KB.

If p is not a power of two, we first reduce the number of processes to the nearest lower power of two by having the first few even-numbered processes send their data to the neighboring odd-numbered process ($rank + 1$). These odd-numbered processes do a reduce on the received data, compute the result for themselves and their left neighbor during the recur-

recursive halving algorithm, and, at the end, send the result back to the left neighbor. Therefore, if p is not a power of two, the time taken by the algorithm is $T_{rec_half} = (\lceil \lg p \rceil + 2)\alpha + 2n\beta + n(1 + \frac{p-1}{p})\gamma$. This cost is approximate because some imbalance exists in the amount of work each process does, since some processes do the work of their neighbors as well.

If the reduction operation is not commutative, recursive halving will not work. Instead, we use a recursive-doubling algorithm similar to the one in allgather. In the first step, pairs of neighboring processes exchange data; in the second step, pairs of processes at distance 2 apart exchange data; in the third step, processes at distance 4 apart exchange data; and so forth. However, more data is communicated than in allgather. In step 1, processes exchange all the data except the data needed for their own result ($n - \frac{n}{p}$); in step 2, processes exchange all data except the data needed by themselves and by the processes they communicated with in the previous step ($n - \frac{2n}{p}$); in step 3, it is ($n - \frac{4n}{p}$); and so forth. Therefore, the time taken by this algorithm is $T_{short} = \lg p \alpha + n(\lg p - \frac{p-1}{p})\beta + n(\lg p - \frac{p-1}{p})\gamma$. We use this algorithm for very short messages (< 512 bytes).

For long messages (≥ 512 KB in the case of commutative operations and ≥ 512 bytes in the case of noncommutative operations), we use a pairwise exchange algorithm that takes $p-1$ steps. In step i , each process sends data to ($rank + i$), receives data from ($rank - i$), and performs the local reduction. The data exchanged is only the data needed for the scattered result on the process ($\frac{n}{p}$). The time taken by this algorithm is $T_{long} = (p-1)\alpha + \frac{p-1}{p}n\beta + \frac{p-1}{p}n\gamma$. Note that this algorithm has the same bandwidth requirement as the recursive halving algorithm. Nonetheless, we use this algorithm for long messages because it performs much better than recursive halving (similar to the results for recursive doubling versus ring algorithm for long-message allgather).

The SKaMPI benchmark, by default, uses a non-commutative user-defined reduction operation. Since commutative operations are more commonly used, we modified the benchmark to use a commutative operation, namely, `MPI_SUM`. Figure 10 shows the performance of the new algorithm for short messages on the IBM SP and on the Myrinet cluster. The performance is significantly better than that of the algorithm used in IBM’s MPI on the SP and several times better than the old algorithm (reduce + scatterv) used in MPICH on the Myrinet cluster.

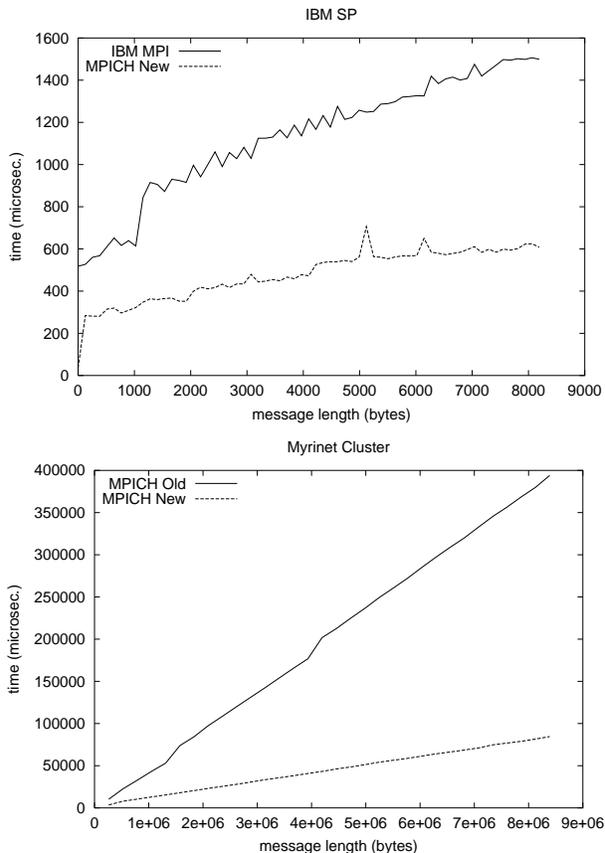


Figure 10: Performance of reduce-scatter for short messages on the IBM SP (64 nodes) and for long messages on the Myrinet cluster (32 nodes)

4.5 Reduce and Allreduce

`MPI_Reduce` performs a global reduction operation and returns the result to the specified root, whereas `MPI_Allreduce` returns the result on all processes. The old algorithm for reduce in MPICH uses a binomial tree, which takes $\lg p$ steps, and the data communicated at each step is n . Therefore, the time taken by this algorithm is $T_{tree} = \lceil \lg p \rceil (\alpha + n\beta + n\gamma)$. The old algorithm for allreduce simply does a reduce to rank 0 followed by a broadcast.

The binomial tree algorithm for reduce is a good algorithm for short messages because of the $\lg p$ number of steps. For long messages, however, a better algorithm exists, proposed by Rabenseifner [14]. The principle behind Rabenseifner’s algorithm is similar to that behind Van de Geijn’s algorithm for long-message broadcast. Van de Geijn implements the broadcast as a scatter followed by an allgather, which reduces the $n \lg p \beta$ bandwidth term in the binomial tree algorithm to a $2n\beta$ term. Rabenseifner’s algorithm implements

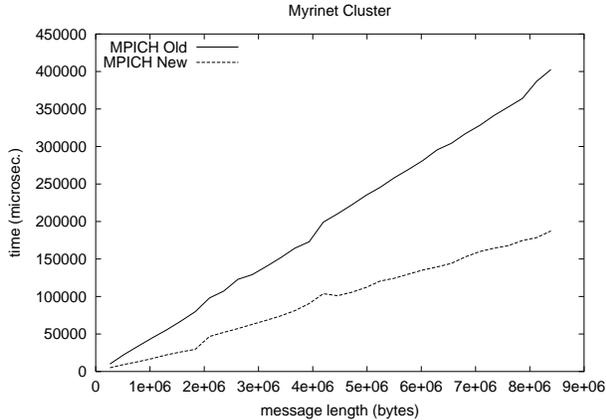


Figure 11: Performance of reduce (64 nodes)

a long-message reduce effectively as a reduce-scatter followed by a gather to the root, which has the same effect of reducing the bandwidth term from $n \lg p \beta$ to $2n\beta$. The time taken by Rabenseifner’s algorithm is the sum of the times taken by reduce-scatter (recursive halving) and gather (binomial tree), which is $T_{rabenseifner} = 2 \lg p \alpha + 2 \frac{p-1}{p} n \beta + \frac{p-1}{p} n \gamma$.

For reduce, in the case of predefined reduction operations, we use Rabenseifner’s algorithm for long messages (> 2 KB) and the binomial tree algorithm for short messages (≤ 2 KB). In the case of user-defined reduction operations, we use the binomial tree algorithm for all message sizes because, unlike with predefined reduction operations, the user may pass derived datatypes, and breaking up derived datatypes to do the reduce-scatter is tricky. Figure 11 shows the performance of reduce for long messages on the Myrinet cluster. The new algorithm is more than twice as fast as the old algorithm in some cases.

For allreduce, we use a recursive doubling algorithm for short messages and for long messages with user-defined reduction operations. This algorithm is similar to the recursive doubling algorithm used in allgather, except that each communication step also involves a local reduction. The time taken by this algorithm is $T_{rec-dbl} = \lg p \alpha + n \lg p \beta + n \lg p \gamma$.

For long messages and predefined reduction operations, we use Rabenseifner’s algorithm for allreduce [14], which does a reduce-scatter followed by an allgather. If the number of processes is a power of two, the cost for the reduce-scatter is $\lg p \alpha + \frac{p-1}{p} n \beta + \frac{p-1}{p} n \gamma$. The cost for the allgather is $\lg p \alpha + \frac{p-1}{p} n \beta$. Therefore, the total cost is $T_{rabenseifner} = 2 \lg p \alpha + 2 \frac{p-1}{p} n \beta + \frac{p-1}{p} n \gamma$.

5 Further Optimization of Allreduce and Reduce

A five-year profiling study of applications running in production mode on the Cray T3E 900 at the University of Stuttgart revealed that more than 40% of the time spent in MPI functions was spent in the two functions `MPI_Allreduce` and `MPI_Reduce` and that 25% of all execution time was spent on program runs that involved a non-power-of-two number of processes [15]. We therefore investigated in further detail how to optimize allreduce and reduce for both power-of-two and non-power-of-two numbers of processes. We experimented with several variations of the basic algorithm for long-message allreduce and reduce described above. We ran the best-performing algorithms for the usage scenarios indicated by the profiling study and found that the new algorithms reduce the communication time by 20% for allreduce and 54% for reduce compared with the vendor’s implementation on the T3E, as shown in Figure 12. We describe these algorithms and the results below.

The cost model described in Section 3 assumes that all processes can send and receive one message at the same time, regardless of the source and destination. Although this is a good approximation, many networks are faster if pairs of processes exchange data with each other, rather than if a process sends to and receives from different processes [3]. Therefore, we refine the cost model further by defining two costs: $\alpha + n\beta$ is the time taken for bidirectional communication between a pair of processes, and $\alpha_{uni} + n\beta_{uni}$ is the time taken for unidirectional communication from one process to another. We also define the ratios $f_\alpha = \alpha_{uni}/\alpha$ and $f_\beta = \beta_{uni}/\beta$. These ratios are normally in the range 0.5 (simplex network) to 1.0 (full-duplex network).

We consider five different algorithms for implementing allreduce and reduce. The first two algorithms are binomial tree and recursive doubling, which were explained above. Binomial tree for reduce is well known. For allreduce, it involves doing a binomial-tree reduce to rank 0 followed by a binomial-tree broadcast. Recursive doubling is used for allreduce only. The other three algorithms are recursive halving and doubling, binary blocks, and ring. For describing these algorithms, we define the following terms:

- *Recursive vector halving*: The vector to be reduced is recursively halved in each step.
- *Recursive vector doubling*: Small pieces of the vector scattered across processes are recursively gathered or combined to form the large vector

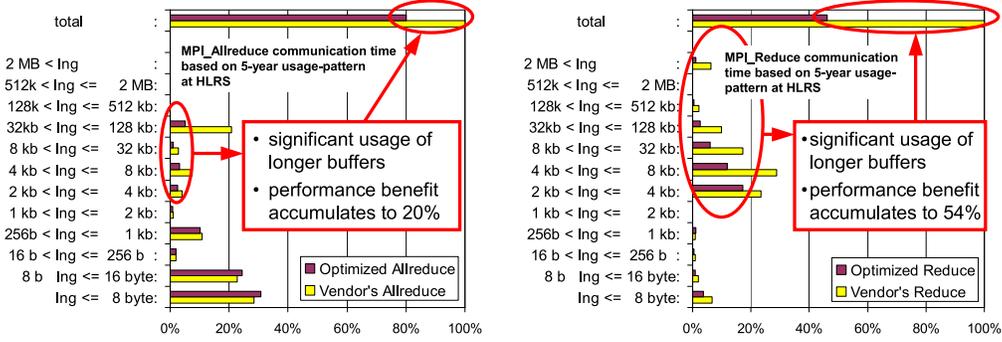


Figure 12: Benefit of new allreduce and reduce algorithms optimized for long vectors on the Cray T3E

- *Recursive distance halving*: The distance over which processes communicate is recursively halved at each step $(\frac{p}{2}, \frac{p}{4}, \dots, 1)$.
- *Recursive distance doubling*: The distance over which processes communicate is recursively doubled at each step $(1, 2, 4, \dots, \frac{p}{2})$.

5.1 Recursive Halving and Doubling

This algorithm is a combination of a reduce-scatter implemented with recursive vector halving and distance doubling followed either by an allgather implemented by a recursive vector doubling combined with recursive distance halving (for allreduce) or by a gather implemented with a binomial tree (for reduce).

Since these recursive algorithms require a power-of-two number of processes, if the number of processes is not a power of two, we first reduce it to the nearest lower power of two ($p' = 2^{\lfloor \lg p \rfloor}$) by removing $r = p - p'$ extra processes as follows. In the first $2r$ processes (ranks 0 to $2r - 1$), all the even ranks send the second half of the input vector to their right neighbor ($rank + 1$), and all the odd ranks send the first half of the input vector to their left neighbor ($rank - 1$), as illustrated in Figure 13. The even ranks compute the reduction on the first half of the vector and the odd ranks compute the reduction on the second half. The odd ranks then send the result to their left neighbors (the even ranks). As a result, the even ranks among the first $2r$ processes now contain the reduction with the input vector on their right neighbors (the odd ranks). These odd ranks are removed from participating in the rest of the algorithm, which leaves behind a power-of-two number of processes. The first r even-ranked processes and the last $p - 2r$ processes are now renumbered from 0 to $p' - 1$, p' being a power of two.

Figure 13 illustrates the algorithm for an example

on 13 processes. The input vectors and all reduction results are divided into 13 parts (A, B, ..., H) and denoted as A–H_{ranks}. After the first reduction, process P0 has computed A–D₀₋₁, denoting the reduction result of the first half (A–D) of the input vector from processes 0 and 1. Similarly, P1 has computed E–H₀₋₁, P2 has computed A–D₂₋₃, and so forth. The odd ranks then send their half to the even ranks on their left: P1 sends E–H₀₋₁ to P0, P3 sends E–H₂₋₃ to P0, and so forth. This completes the first step, which takes $(1 + f_\alpha)\alpha + \frac{n}{2}(1 + f_\beta)\beta + \frac{n}{2}\gamma$ time. This step is not necessary if the number of processes is a power of two.

Now we begin with the first step of recursive vector halving and distance doubling, in which the even-ranked processes send the second half of their buffer to $rank' + 1$ and the odd-ranked processes send the first half of their buffer to $rank' - 1$. All processes then compute the reduction between the local buffer and the received buffer. In the next $\lg p' - 1$ steps, the buffers are recursively halved, and the distance is doubled. At the end, each of the p' processes has $\frac{1}{p'}$ of the total reduction result vector. All these recursive steps take $\lg p' \alpha + (\frac{p'-1}{p'})n\beta + n\gamma$ time. The second part implements either an allgather or gather to complete the allreduce or reduce operation.

Allreduce: To implement allreduce, we do an allgather using recursive vector doubling and distance halving. In the first step, process pairs exchange $\frac{1}{p'}$ of the buffer to achieve $\frac{2}{p'}$ of the result vector, in the next step $\frac{2}{p'}$ of the buffer is exchanged to get $\frac{4}{p'}$ of the result, and so forth. After $\lg p'$ steps, the p' processes receive the total reduction result. This allgather part costs $\lg p' \alpha + (\frac{p'-1}{p'})n\beta$. If the number of processes is not a power of two, the total result vector must be sent to the r processes that were removed in the first step, which results in additional overhead of $\alpha_{uni} + n\beta_{uni}$.

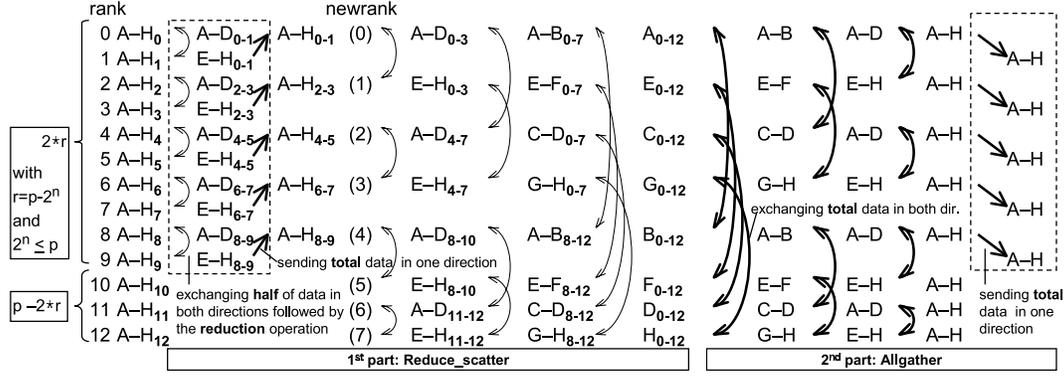


Figure 13: Recursive halving and doubling for allreduce. The figure shows the intermediate results after each communication step (including the reduction operation in the reduce-scatter phase). The dotted frames show the additional overhead caused by a non-power-of-two number of processes.

The total allreduce operation therefore takes the following amount of time:

- If p is a power of two: $T_{all,h\&d,p=2^{\epsilon}p} = 2 \lg p \alpha + 2n\beta + n\gamma - \frac{1}{p}(2n\beta + n\gamma) \simeq 2 \lg p \alpha + 2n\beta + n\gamma$
- If p is not a power of two: $T_{all,h\&d,p \neq 2^{\epsilon}p} = (2 \lg p' + 1 + 2f_{\alpha})\alpha + (2 + \frac{1+3f_{\beta}}{2})n\beta + \frac{3}{2}n\gamma - \frac{1}{p'}(2n\beta + n\gamma) \simeq (3 + 2 \lfloor \lg p \rfloor)\alpha + 4n\beta + \frac{3}{2}n\gamma$

This algorithm is good for long vectors and power-of-two numbers of processes. For non-power-of-two numbers of processes, the data transfer overhead is doubled, and the computation overhead is increased by $\frac{3}{2}$. The binary blocks algorithm described below can reduce this overhead in many cases.

Reduce: A binomial tree algorithm is used to gather the data to the root using recursive vector doubling and distance halving. The time needed for this gather operation is $\lg p' \alpha_{uni} + \frac{p'-1}{p'} n\beta_{uni}$. In the non-power-of-two case, if the root happens to be one of those odd-ranked processes that would normally be removed in the first step, then the role of this process and its partner in the first step are interchanged after the first reduction in the reduce-scatter phase, which causes no additional overhead. The total reduce operation therefore takes the following amount of time:

- If p is a power of two: $T_{red,h\&d,p=2^{\epsilon}p} = \lg p(1 + f_{\alpha})\alpha + (1 + f_{\beta})n\beta + n\gamma - \frac{1}{p}((1 + f_{\beta})n\beta + n\gamma) \simeq 2 \lg p \alpha + 2n\beta + n\gamma$
- If p is not a power of two: $T_{red,h\&d,p \neq 2^{\epsilon}p} = \lg p'(1 + f_{\alpha})\alpha + (1 + f_{\alpha})\alpha + (1 + \frac{1+f_{\beta}t_{\alpha}}{2} + f_{\beta})n\beta + \frac{3}{2}n\gamma - \frac{1}{p'}((1 + f_{\beta})n\beta + n\gamma) \simeq (2 + 2 \lfloor \lg p \rfloor)\alpha + 3n\beta + \frac{3}{2}n\gamma$

5.2 Binary Blocks Algorithm

This algorithm reduces some of the load imbalance in the recursive halving and doubling algorithm when the number of processes is a non-power-of-two. The algorithm starts with a binary-block decomposition of all processes in blocks with power-of-two numbers of processes (see the example in Figure 14). Each block executes its own reduce-scatter with the recursive vector halving and distance doubling algorithm described in the previous section. Then, starting with the smallest block, the intermediate result (or the input vector in case of a 2^0 block) is split into the segments of the intermediate result in the next higher block and sent to the processes in that block, and those processes compute the reduction on the segment. This does cause a load imbalance in computation and communication compared with the execution in the larger blocks. For example, in the third exchange step in the 2^3 block, each process sends one segment, receives one segment, and computes the reduction of one segment (P0 sends B, receives A, and computes the reduction on A). The load imbalance is introduced by the smaller blocks 2^2 and 2^0 : In the 2^2 block, each process receives and reduces two segments (for example, A-B on P8), whereas in the 2^0 block (P12), each process has to send as many messages as the ratio of the two block sizes (here $2^2/2^0$). At the end of the first part, the highest block must be recombined with the next smaller block, and the ratio of the block sizes again determines the overhead.

We see that the maximum difference between the ratio of two successive blocks, especially in the low range of exponents, determines the imbalance. Let

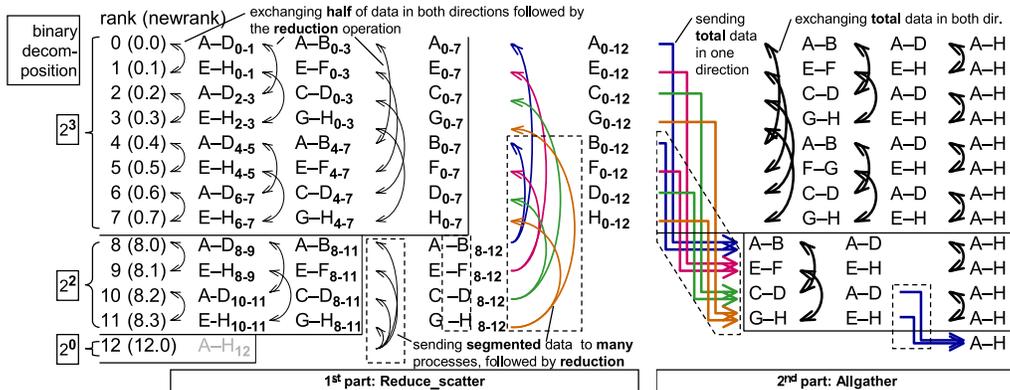


Figure 14: Allreduce using the binary blocks algorithm

us define $\delta_{\text{expo,max}}$ as the maximal difference of two consecutive exponents in the binary representation of the number of processes. For example, $100 = 2^6 + 2^5 + 2^2$, $\delta_{\text{expo,max}} = \max(6 - 5, 5 - 2) = 3$. If $\delta_{\text{expo,max}}$ is small, the binary blocks algorithm can perform well. The profiling study on the T3E revealed that 12 % of all applications ran with 96 processes and 8 % ran with 36, 60, 61, 77, 80, 81, 100, 235, or 251 processes [15]. For these cases, $\delta_{\text{expo,max}}$ is 1 (96, 60), 2 (61, 80, 235, 251), or 3 (36, 77, 100). Therefore, this algorithm is likely to perform well in the common cases used in practice.

Allreduce: For allreduce, the second part is an allgather implemented with recursive vector doubling and distance halving in each block. For this purpose, data must be provided to the processes in the smaller blocks with a pair of messages from processes of the next larger block, as shown in Figure 14.

Reduce: For reduce, if the root is outside of the largest block, then the intermediate result segment of rank 0 is sent to the root, and the root plays the role of rank 0. A binomial tree is used to gather the result segments to the root process.

We note that if the number of processes is a power of two, the binary blocks algorithm is identical to the recursive halving and doubling algorithm.

5.3 Ring Algorithm

While the algorithms in the preceding two sections are good for power-of-two numbers of processes and long vectors, for medium non-power-of-two numbers of processes and long vectors there exists another good algorithm. It uses a pairwise-exchange algorithm for the reduce-scatter phase (see Section 4.4). For allreduce, it uses a ring algorithm to do the allgather, and, for reduce, all processes directly send their result seg-

ment to the root. This algorithm is good in bandwidth use for non-power-of-two numbers of processes, but the latency scales with the number of processes. Therefore this algorithm should be used only for a small or medium number of processes. The time taken is $T_{\text{all,ring}} = 2(p-1)\alpha + 2n\beta + n\gamma - \frac{1}{p}(2n\beta + n\gamma)$ for allreduce and $T_{\text{red,ring}} = (p-1)(\alpha + \alpha_{\text{uni}}) + n(\beta + \beta_{\text{uni}}) + n\gamma - \frac{1}{p}(n(\beta + \beta_{\text{uni}}) + n\gamma)$ for reduce.

5.4 Choosing the Fastest Algorithm

Based on the number of processes and the vector length, the reduction routine must decide which algorithm to use. The decision is not an easy one and depends on a number of factors. We experimentally determined which algorithm works the best for different buffer sizes and number of processes on a Cray T3E 900. The results for allreduce are shown in Figure 15. For buffer sizes less than or equal to 32 bytes, recursive doubling is the best; for buffer sizes less than or equal to 1 KB, the vendor's algorithm (for power-of-two) and binomial tree (for non-power-of-two) are the best, but not much better than recursive doubling; for longer buffer sizes, the ring algorithm is good for some buffer sizes and some number of processes less than 32. In general, on a Cray T3E 900, the binary blocks algorithm is faster if $\delta_{\text{expo,max}} < \lg(\text{vector length in bytes})/2.0 - 2.5$ and $\text{vector size} \geq 16$ KB and more than 32 processes are used. In a few cases, for example, 33 processes and less than 32 KB, recursive halving and doubling is the fastest.

Figure 16 shows the bandwidths obtained by the various algorithms for a 32 KB buffer size on the T3E. For this buffer size, the new protocols are clearly better than the vendor's protocol (Cray MPT.1.4.0.4) as well as the binomial tree algorithm for all numbers

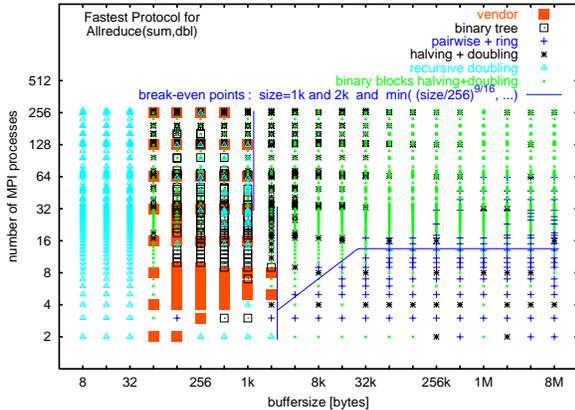


Figure 15: The fastest algorithm for allreduce (MPI_DOUBLE, MPI_SUM) on a Cray T3E 900

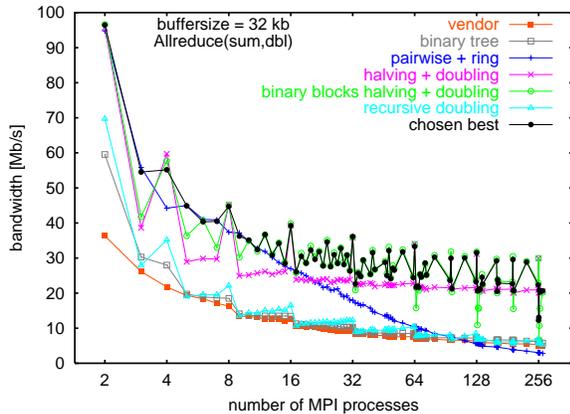


Figure 16: Bandwidth comparison for allreduce (MPI_DOUBLE, MPI_SUM) with 32 KB vectors on a Cray T3E 900.

of processes. For up to 32 processes, we measured the performance for all numbers of processes. For more than 32 processes, we measured only selected values with small and large $\delta_{\text{exp}, \text{max}}$. We observe that the bandwidth of the binary blocks algorithm depends strongly on $\delta_{\text{exp}, \text{max}}$ and that recursive halving and doubling is faster on 33, 65, 66, 97, 128–131 processes. The ring algorithm is faster on 3, 5, 7, 9–11, and 17 processes.

5.5 Comparison with Vendor's MPI

We also ran some experiments to compare the performance of the best of the new algorithms with the algorithm in the native MPI implementations on the IBM SP at San Diego Supercomputer Center, a Myrinet cluster at the University of Heidelberg, and the Cray T3E. Figure 17 shows that with a pure MPI program-

ming model (1 MPI process per CPU) on the IBM SP, the fastest algorithm performs about 1.5 times better than the vendor's algorithm for buffer sizes of 8–64 KB and 2–5 times better for larger buffers. With a hybrid programming model comprising 1 MPI process per SMP node, where each MPI process is itself SMP-parallelized (with OpenMP, for example) and only the master thread calls MPI functions (the master-only style in [16]), the performance is about 1.5–3 times better for buffer sizes 4–128 KB and more than 4 processes. Figure 18 compares the new algorithm with the old MPICH-1 algorithm on the Heidelberg Myrinet cluster. The new algorithms show a performance benefit of 3–7 times with pure MPI and 2–5 times with the hybrid model. Figure 19 shows that on the T3E, the new algorithms are 3–5 times faster than the vendor's algorithm for the operation MPI_SUM and, because of the very slow implementation of structured derived datatypes in Cray's MPI, up to 100 times faster for MPI_MAXLOC.

6 Conclusions and Future Work

Our results demonstrate that optimized algorithms for collective communication can provide substantial performance benefits and, to achieve the best performance, we need to use a number of different algorithms and select the right algorithm for a particular message size and number of processes. Determining the cutoff points for switching between the different algorithms is tricky, however. At present, we use experimentally determined cutoff points, but we plan to develop a model to calculate the cutoff points automatically based on system parameters. We also plan to extend this work to incorporate topology awareness, particularly algorithms that are optimized for architectures comprising clusters of SMPs and clusters distributed over a wide area, such as the Tera-Grid [21]. Furthermore, we plan to explore the use of one-sided communication to improve the performance of collective operations.

The source code for the new algorithms will be available in the next release of MPICH (1.2.6) and in version 1.0 of our new MPI implementation, MPICH2. Several of the algorithms are already implemented in MPICH 1.2.5 and early releases of beta versions of MPICH2. Both MPICH-1 and MPICH-2 can be downloaded from www.mcs.anl.gov/mpi/mpich.

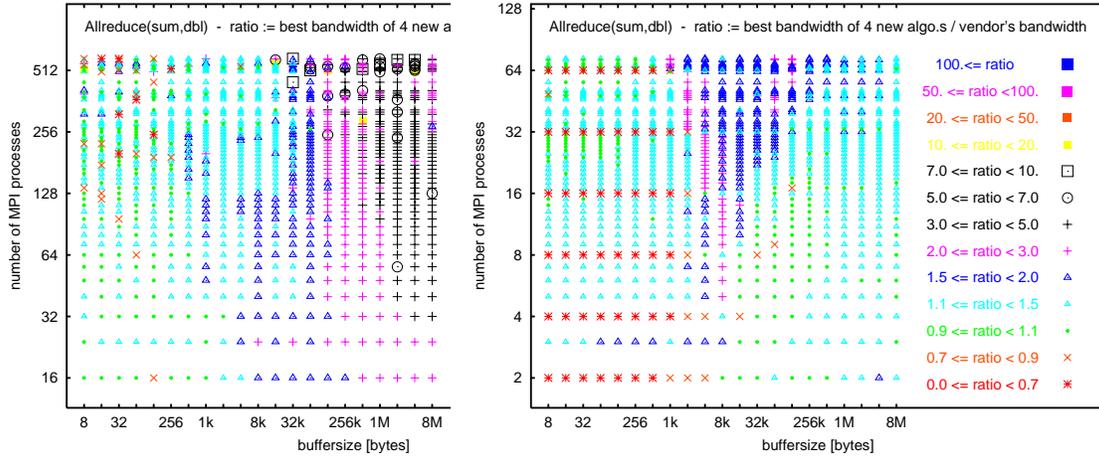


Figure 17: Ratio of the bandwidth of the fastest of the new algorithms (not including recursive doubling) and the vendor's allreduce on the IBM SP at SDSC with 1 MPI process per CPU (left) and per SMP node (right)

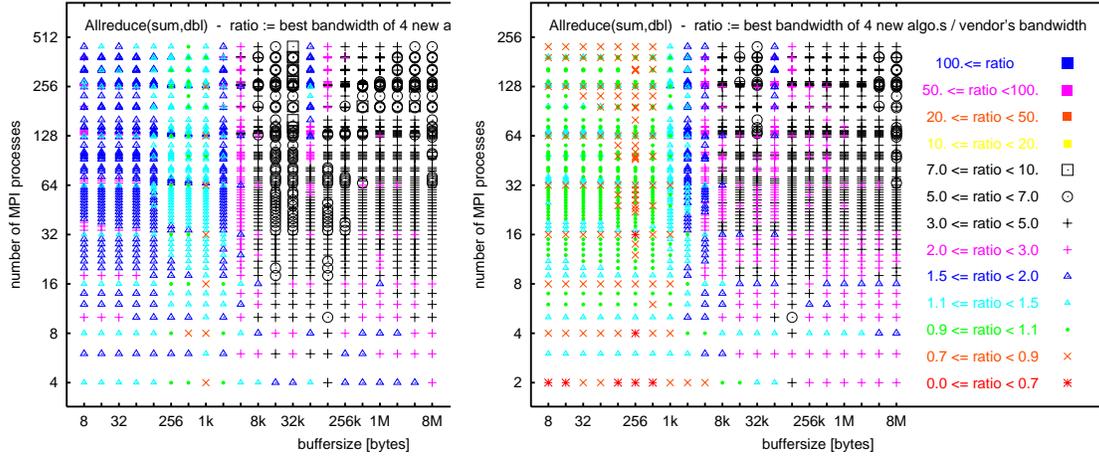


Figure 18: Ratio of the bandwidth of the fastest of the new algorithms (not including recursive doubling) and the old MPICH-1 algorithm on a Myrinet cluster with dual-CPU PCs (HELICS cluster, University of Heidelberg) and 1 MPI process per CPU (left) and per SMP node (right)

Acknowledgments

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

The authors would like to acknowledge their colleagues and others who provided suggestions and help-

ful comments. They would especially like to thank Jesper Larsson Träff for helpful discussion on optimized reduction algorithms and Gerhard Wellein, Thomas Ludwig, and Ana Kovatcheva for their benchmarking support.

References

- [1] M. Barnett, S. Gupta, D. Payne, L. Shuler, R. van de Geijn, and J. Watts. Interprocessor collective communication library (InterCom). In *Proceedings of Supercomputing '94*, November 1994.

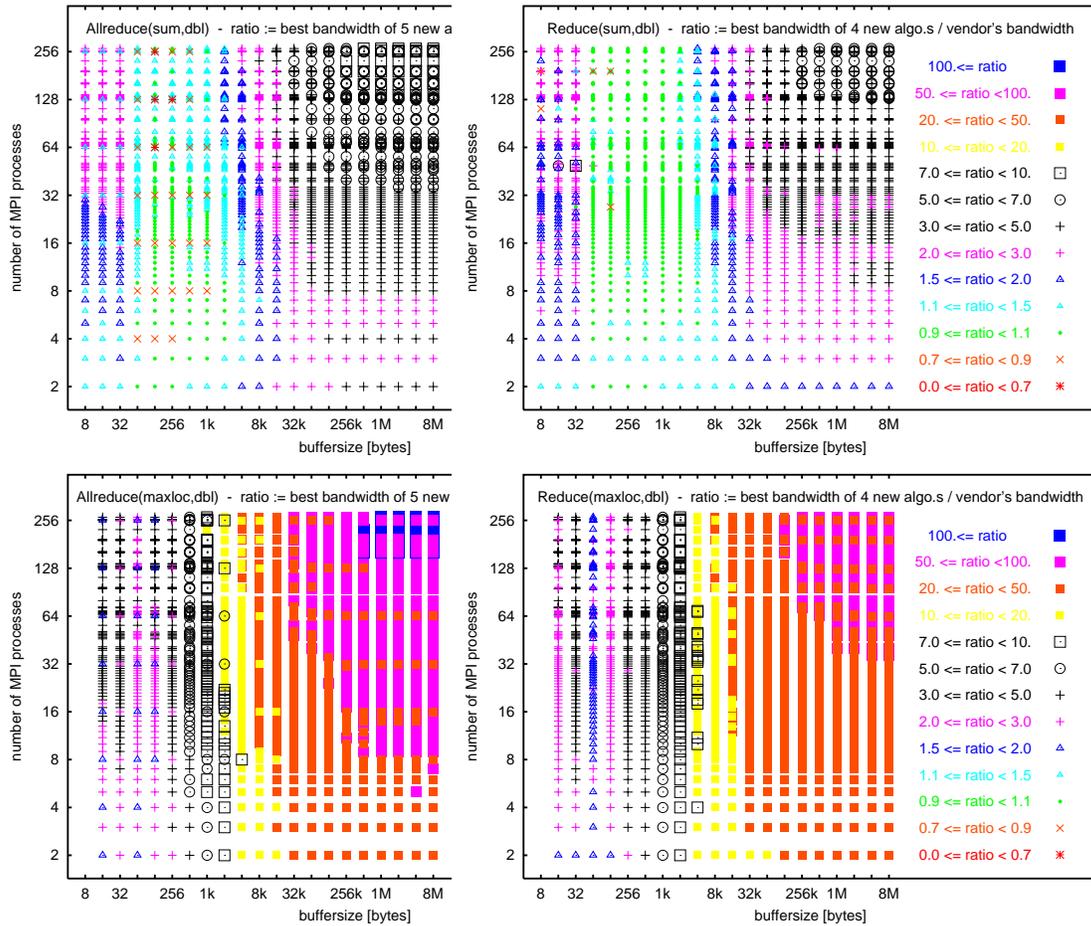


Figure 19: Ratio of the bandwidth of the fastest of the new algorithms and the vendor's algorithm for allreduce (left) and reduce (right) with operation MPI_SUM (first row) and MPI_MAXLOC (second row) on a Cray T3E 900

- [2] M. Barnett, R. Littlefield, D. Payne, and R. van de Geijn. Global combine on mesh architectures with wormhole routing. In *Proceedings of the 7th International Parallel Processing Symposium*, April 1993.
- [3] Gregory D. Benson, Cho-Wai Chu, Qing Huang, and Sadik G. Caglar. A comparison of MPICH allgather algorithms on switched networks. In Jack Dongarra, Domenico Laforenza, and Salvatore Orlando, editors, *Recent advances in Parallel Virtual Machine and Message Passing Interface, 10th European PVM/MPI Users' Group Meeting*, pages 335–343. Lecture Notes in Computer Science 2840, Springer, September 2003.
- [4] S. Bokhari. Complete exchange on the iPSC/860. Technical Report 91–4, ICASE, NASA Langley Research Center, 1991.
- [5] S. Bokhari and H. Berryman. Complete exchange on a circuit switched mesh. In *Proceedings of the Scalable High Performance Computing Conference*, pages 300–306, 1992.
- [6] Jehoshua Bruck, Ching-Tien Ho, Schlomo Kipnis, Eli Upfal, and Derrick Weathersby. Efficient algorithms for all-to-all communications in multiport message-passing systems. *IEEE Transactions on Parallel and Distributed Systems*, 8(11):1143–1156, November 1997.
- [7] Debra Hensgen, Raphael Finkel, and Udi Manbet. Two algorithms for barrier synchronization. *International Journal of Parallel Programming*, 17(1):1–17, 1988.
- [8] L. V. Kale, Sameer Kumar, and Krishnan Vardarajan. A framework for collective personalized communication. In *Proceedings of the 17th In-*

- ternational Parallel and Distributed Processing Symposium (IPDPS '03)*, 2003.
- [9] N. Karonis, B. de Supinski, I. Foster, W. Gropp, E. Lusk, and J. Bresnahan. Exploiting hierarchy in parallel computer networks to optimize collective operation performance. In *Proceedings of the Fourteenth International Parallel and Distributed Processing Symposium (IPDPS '00)*, pages 377–384, 2000.
- [10] T. Kielmann, R. F. H. Hofman, H. E. Bal, A. Plaat, and R. A. F. Bhoedjang. MagPIe: MPI's collective communication operations for clustered wide area systems. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'99)*, pages 131–140. ACM, May 1999.
- [11] P. Mitra, D. Payne, L. Shuler, R. van de Geijn, and J. Watts. Fast collective communication libraries, please. In *Proceedings of the Intel Supercomputing Users' Group Meeting*, June 1995.
- [12] MPICH – A portable implementation of MPI. <http://www.mcs.anl.gov/mpi/mpich>.
- [13] Rolf Rabenseifner. Effective bandwidth (b_eff) benchmark. http://www.hlrs.de/mpi/b_eff.
- [14] Rolf Rabenseifner. New optimized MPI reduce algorithm. <http://www.hlrs.de/organization/par/services/models/mpi/myreduce.html>.
- [15] Rolf Rabenseifner. Automatic MPI counter profiling of all users: First results on a CRAY T3E 900-512. In *Proceedings of the Message Passing Interface Developer's and User's Conference 1999 (MPIDC '99)*, pages 77–85, March 1999.
- [16] Rolf Rabenseifner and Gerhard Wellein. Communication and optimization aspects of parallel programming models on hybrid architectures. *International Journal of High Performance Computing Applications*, 17(1):49–62, 2003.
- [17] Peter Sanders and Jesper Larsson Träff. The hierarchical factor algorithm for all-to-all communication. In B. Monien and R. Feldman, editors, *Euro-Par 2002 Parallel Processing*, pages 799–803. Lecture Notes in Computer Science 2400, Springer, August 2002.
- [18] D. Scott. Efficient all-to-all communication patterns in hypercube and mesh topologies. In *Proceedings of the 6th Distributed Memory Computing Conference*, pages 398–403, 1991.
- [19] Mohak Shroff and Robert A. van de Geijn. CollMark: MPI collective communication benchmark. Technical report, Dept. of Computer Sciences, University of Texas at Austin, December 1999.
- [20] Steve Sistare, Rolf vandeVaart, and Eugene Loh. Optimization of MPI collectives on clusters of large-scale SMPs. In *Proceedings of SC99: High Performance Networking and Computing*, November 1999.
- [21] Teragrid. <http://www.teragrid.org>.
- [22] V. Tipparaju, J. Nieplocha, and D. K. Panda. Fast collective operations using shared and remote memory access protocols on clusters. In *Proceedings of the 17th International Parallel and Distributed Processing Symposium (IPDPS '03)*, 2003.
- [23] Jesper Larsson Träff. Improved MPI all-to-all communication on a Giganet SMP cluster. In Dieter Kranzlmüller, Peter Kacsuk, Jack Dongarra, and Jens Volkert, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 9th European PVM/MPI Users' Group Meeting*, pages 392–400. Lecture Notes in Computer Science 2474, Springer, September 2002.
- [24] Sathish S. Vadhiyar, Graham E. Fagg, and Jack Dongarra. Automatically tuned collective communications. In *Proceedings of SC99: High Performance Networking and Computing*, November 1999.
- [25] Thomas Worsch, Ralf Reussner, and Werner Augustin. On benchmarking collective MPI operations. In Dieter Kranzlmüller, Peter Kacsuk, Jack Dongarra, and Jens Volkert, editors, *Recent advances in Parallel Virtual Machine and Message Passing Interface, 9th European PVM/MPI Users' Group Meeting*, pages 271–279. Lecture Notes in Computer Science 2474, Springer, September 2002.