Efficient MPI Collective Operations for Clusters in Long-and-Fast Networks

Motohiko Matsuda, Tomohiro Kudoh, Yuetsu Kodama, Ryousei Takano
Grid Technology Research Center
National Institute of Advanced Industrial Science and Technology

Yutaka Ishikawa
University of Tokyo, and
Grid Technology Research Center
National Institute of Advanced Industrial Science and Technology

Abstract

Several MPI systems for Grid environment, in which clusters are connected by wide-area networks, have been proposed. However, the algorithms of collective communication in such MPI systems assume relatively low bandwidth wide-area networks, and they are not designed for the fast wide-area networks that are becoming available. On the other hand, for cluster MPI systems, a bcast algorithm by van de Geijn, et al and an allreduce algorithm by Rabenseifner have been proposed, which are efficient in a high bi-section bandwidth environment. We modify those algorithms so as to effectively utilize fast wide-area inter-cluster networks and to control the number of nodes which can transfer data simultaneously through wide-area networks to avoid congestion. We confirmed the effectiveness of the modified algorithms by experiments using a 10 Gbps emulated WAN environment. The environment consists of two clusters, where each cluster consists of nodes with 1 Gbps Ethernet links and a switch with a 10 Gbps upper link. The two clusters are connected through a 10 Gbps WAN emulator which can insert latency. In a 10 millisecond latency environment, when the message size is 32 MB, the proposed bcast and allreduce are 1.6 and 3.2 times faster, respectively, than the algorithms used in existing MPI systems for Grid environment.

1 Introduction

There are several MPI systems for clusters in Grid environment, such as MPICH-G2, MagPIe, and PACX-MPI, and they have implemented a set of efficient algorithms of collective communication for high latency networks [5, 10, 11]. The set of algorithms assumes that the bandwidth of inter-cluster links is much lower than that of the network links of the cluster nodes.

However, recently, the bandwidth of wide-area networks has become much wider, and far exceeds the bandwidth of the network interfaces of typical cluster nodes [6, 18]. This situation will continue at least for a while, because the optical network technology for wide-area networks will continue to advance, while providing the fastest network interfaces to all the cluster nodes costs too much and is impractical. Thus, new algorithms of collective communication are needed, which match the fast inter-cluster networks.

Among the set of collective operations, *bcast* and *allreduce* are two important ones [15]. It is reported that the *allreduce* operation took 37% of MPI execution time in 5-year profiling on a Cray T3E [15, 19]. For these two operations, efficient algorithms have been invented targeting a high bi-section bandwidth environment: the *bcast* algorithm by van de Geijn, et al [1, 2], and the *allreduce* algorithm by Rabenseifner [16]. Both the *bcast* and *allreduce* algorithms are based on a similar idea of splitting a message, where a message is split and distributed to the nodes, and then gathering the split messages again in every node. The van de Geijn *bcast* can be implemented by *scatter* followed by *allgather*. The Rabenseifner *allreduce* can be implemented by *reduce-scatter* followed by *allgather*. Here, *scatter*, *reduce-scatter*, and *allgather* are all implemented by log(P)-step algorithms, and can effectively utilize the available bandwidth of each node.

We have been investigating algorithms of collective communication for systems where multiple clusters are connected by wide-area networks. In such systems, the network interface of each node has one bi-directional link, and intra-cluster communication is through a switched network, while inter-cluster communication is through a long-and-

---

Part of this research was supported by a grant from the Ministry of Education, Sports, Culture, Science and Technology (MEXT) of Japan through the NAREGI (National Research Grid Initiative) Project.

1-4244-0328-6/06/$20.00

fast network. Here, long-and-fast roughly means:

$$\text{long} : L(\text{inter-cluster}) \gg L(\text{inter-node}) * \log(P)$$

$$\text{fast} : B(\text{inter-cluster}) > B(\text{node-link})$$

where, L is the latency, B is the bandwidth, and P is the number of nodes in a cluster. The reason for adding $\log(P)$ factor to the latency is to assume that arbitrary $\log(P)$-step algorithms can be performed inside a cluster without considering inter-cluster communication issues. Also, we focus only on large messages, since the efficiency of communication for short messages depends almost only on the inter-cluster latency. Under these assumptions, the inter-node latency is marginal and can be omitted from consideration.

A low bandwidth environment, in which the inter-cluster bandwidth is less than the bandwidth of a network link of a node, is out of the scope of this paper. In such an environment, only one node will perform inter-cluster communication at a time, and regulating the amount of transmission from that node is the issue. Bandwidth limiting mechanisms such as PSPacer [17] and Linux's Token Bucket Filter will work for that purpose, but they are not discussed in this paper.

We have shown in a previous paper that the practical upper-bound of the latency of inter-cluster communication is about 10 milliseconds, when benchmark programs for clusters, such as the NAS Parallel Benchmarks, are not modified to tolerate latency [13]. Although the effect of latency naturally depends on the application, most benchmarks have shown good performance up to a 10 millisecond latency. Out of this range, however, most benchmarks run poorly and connecting two clusters is meaningless in terms of the computing performance, whereas there is still a benefit of using large amounts of resources such as memory and disks. A 10 millisecond latency roughly corresponds to 1000 miles in actual networks, and there may exist some large-scale clusters in this range. Thus, all experiments in this paper are performed with a 10 millisecond latency.

In the following, the designs and the implementations of the algorithms for long-and-fast networks are described in Section 2, then, the experimental results are shown in Section 3. We mention very briefly related work in Section 4, and conclude the paper in Section 5.

2 Design and Implementation

2.1 Design Overview

Our objective is to design algorithms of collective communication to utilize the available bandwidth of wide-area networks, which is a number of times larger than the bandwidth of the network link of each node, e.g., the inter-cluster bandwidth is 10 Gbps while the inter-node bandwidth is 1 Gbps. In such an environment, a number of nodes should send messages simultaneously to the inter-cluster network to fully utilize the bandwidth. However, contention among the messages should be avoided when nodes send messages simultaneously, especially when the TCP/IP protocol is used on long-and-fast networks [14]. The total transmission rate of the sending nodes should be limited to the bandwidth of the inter-cluster network.

For a cluster environment, good algorithms for *bcast* and *allreduce* operations have been proposed. van de Geijn, et al [1, 2] proposed a *bcast* algorithm. Rabenseifner [16] proposed an *allreduce* algorithm. Our algorithms are based on those algorithms, which are modified to efficiently utilize the bandwidth of wide-area networks which connect clusters.

The van de Geijn *bcast* is algorithmically equivalent to *scatter* followed by *allgather*. *scatter* splits a message and distributes the parts of that message to all nodes. *allgather* collects parts of the message from all nodes, and rebuilds the message from the collected parts on each node. We modified this algorithm to extend it for inter-cluster communication. The modified van de Geijn *bcast* inserts a copy operation between the *scatter* and *allgather* stages, which copies the message between the clusters through a wide-area network. Details are described following this section.

The Rabenseifner *allreduce* is algorithmically equivalent to *reduce-scatter* followed by *allgather*. *reduce-scatter* splits a message and distributes the parts of that message to all nodes. In addition, it performs a reduction on the part of the message. *allgather* collects parts of the reduced message from all nodes, and rebuilds the result in each node. We modified this algorithm to extend it for inter-cluster communication. The modified Rabenseifner *allreduce* inserts copy and reduction operations between the *reduce-scatter* and *allgather* stages in a way similar to the case of *bcast*. Details are described following this section, too.

Based on the modifications of the algorithms as stated above, all nodes in a cluster participate in inter-cluster communication. Now, we need to regulate the total transmission rate to avoid contention among the messages. We took a simple forwarding approach in the implementation, in which only selected nodes are allowed to send messages to the opposite cluster, while the other nodes send messages to the selected nodes inside the cluster.

2.2 Bcast

Figures 1 and 2 show an outline of the modified van de Geijn *bcast* and its data movement example. The algorithm is identical to the van de Geijn *bcast* except for the addition of `counter_copy` and working in the `hemisphere` com-

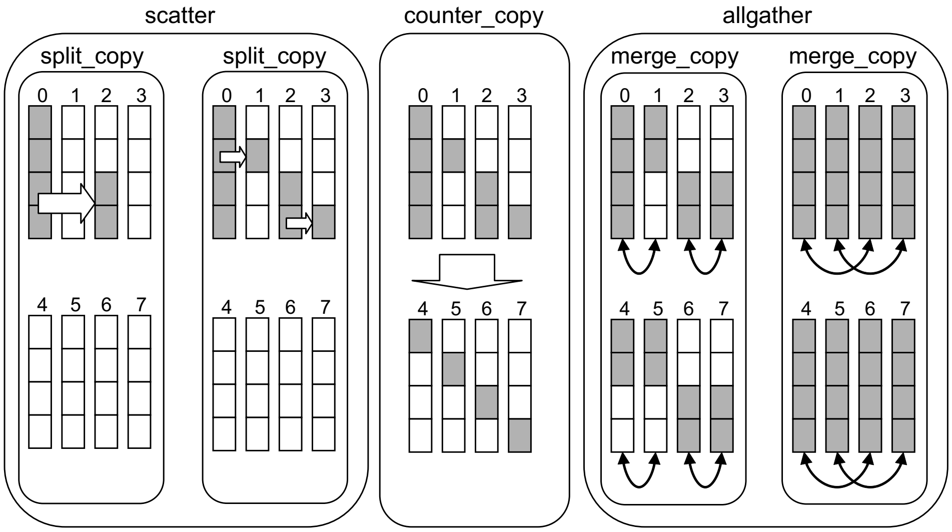

Note that in the beginning, data represented by the gray area is in node 0.

Figure 1. Data movement of the modified van de Geijn bcast.

```
void
vandegeijn(void *buf, int siz)
{
  if (rank < (nprocs/2)) {
    for (i = 0; i < (log2(nprocs)-1); i++) {
      split_copy(buf, siz, i, hemisphere);
    }
  }
  counter_copy(buf, siz);
  for (i = 0; i < (log2(nprocs)-1); i++) {
    merge_copy(buf, siz, i, hemisphere);
  }
}
```

Figure 2. Skeleton of the modified van de Geijn bcast.

municator. The `hemisphere` communicator represents a half of `MPI_COMM_WORLD` and corresponds to each cluster. The `nprocs` variable holds the number of processes in `MPI_COMM_WORLD`.

The steps of the modified van de Geijn *bcast* are as follows:

1. The `split_copy` function performs the steps of *scatter* in the original *bcast* algorithm. In the i-th step, it splits a message in half and copies the half to a $2^{(nprocs/2-i)}$-apart process. It follows a pattern called *recursive-halving*. Note that the `split_copy` function is only performed in a cluster which contains the root rank (the source of the *bcast* data).
2. The `counter_copy` function copies the scattered messages from the cluster containing the root rank to the other cluster. Basically, each node sends its part of the message to the opposite node in the other cluster as shown in Figure 1. That is, there is one-to-one correspondence between nodes in each cluster.
3. The `merge_copy` function performs the steps of *allgather* in the original *bcast* algorithm. In the i-th step, it merges a message from a 2^i-apart process to its holding message. It follows a pattern called *recursive-doubling*.

The `counter_copy` function can overflow the bottleneck link when all the nodes simultaneously send their messages. Thus, `counter_copy` takes a parameter for the number of nodes which are allowed to send messages simultaneously. Only selected nodes may send messages to the other cluster. The other nodes forward their messages to selected nodes inside the cluster. The number of selected nodes can be any value between 1 and $nprocs/2$ according to the available inter-cluster bandwidth.

2.3 Allreduce

Figures 3 and 4 show an outline of the modified Rabenseifner *allreduce* and its data movement example. The algorithm is identical to the Rabenseifner *allreduce* except for the addition of `counter_reduce` and working with `hemisphere`. The *allreduce* algorithm is very similar to *bcast*, where each step performs a reduction in place of a simple copy.

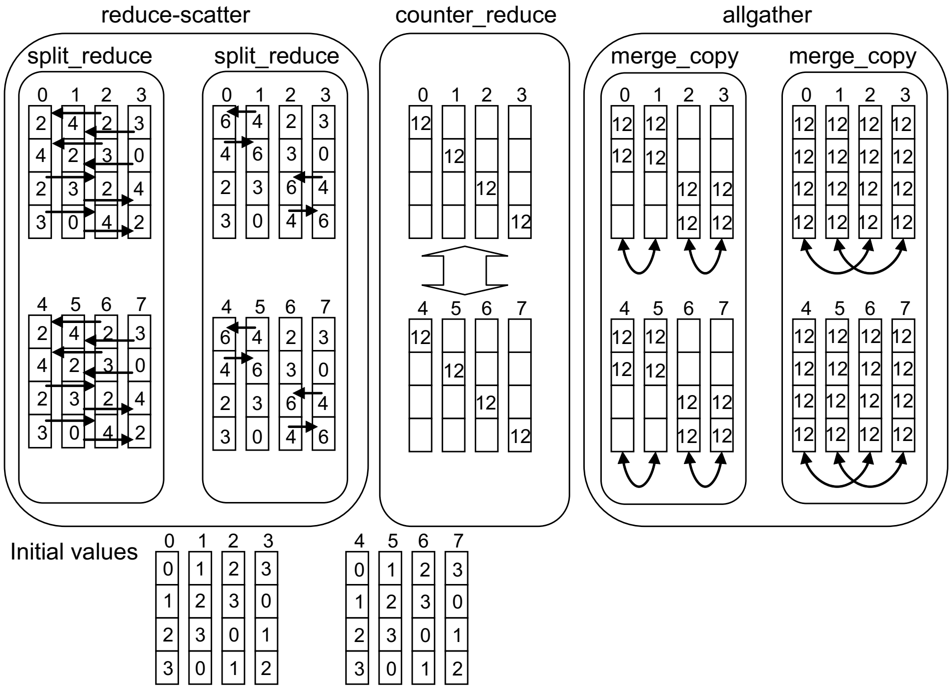

Figure 3. Data movement in the modified Rabenseifner allreduce with the addition operation.

```
void
rabenseifner(void *buf, void *tmp, int siz)
{
  for (i = 0; i < (log2(nprocs)-1); i++) {
    split_reduce(buf, tmp, siz, i, hemisphere);
  }
  counter_reduce(buf, tmp, siz);
  for (i = 0; i < (log2(nprocs)-1); i++) {
    merge_copy(buf, siz, i, hemisphere);
  }
}
```

Figure 4. Skeleton of the modified Rabenseifner allreduce.

The steps of the modified Rabenseifner *allreduce* are as follows:

1. The `split_reduce` function performs the steps of *reduce-scatter* in the original *allreduce* algorithm. In the i-th step, it splits a message in half and copies the half to a $2^{(nprocs/2-i)}$-apart process. It then performs reduction. It follows a pattern called *recursive-halving*.

2. The `counter_reduce` function performs a bi-directional copy and a reduction on the exchanged message. The `counter_reduce` function works in the same way as `counter_copy`, but it is followed by a reduction.

3. The `merge_copy` function performs the steps of *allgather* in the original *allreduce* algorithm. It is the same as the one in the modified van de Geijn *bcast* algorithm.

`counter_reduce` takes a parameter for the number of nodes which are allowed to send messages simultaneously, too. It can take any value between 1 and $nprocs/2$.

Note that `counter_reduce` is bi-directional and receives messages from the opposite cluster. Therefore, forwarding inside a cluster may cause conflicting use of the receiving link of the node. In the implementation, the order of sends was ad hocly skewed in order to avoid concentration, but we observed no significant effect on the performance in the experiments.

2.4 Design Choice: Avoiding Contention

The forwarding scheme described in the previous subsections is one choice to reduce the inter-cluster traffic, and there are other ways to reduce the number of nodes simultaneously communicating. We will discuss some design choices below.

One way is to algorithmically reduce the number of nodes. In *bcast*, the `split_copy` operation splits and distributes a message as a part of the *scatter* stage. If the `counter_copy` operation is performed just after the i-th iteration of `split_copy`, 2^i nodes hold the split message. Therefore, only 2^i nodes participate in `counter_copy`. After `counter_copy`, the remaining iterations of `split_copy` are performed. Although this scheme seems attractive, the number of communicating nodes is limited to powers of two, and it is not flexible.

Algorithmically reducing number of sending nodes cannot be done without some penalty in *allreduce*. The `split_reduce` steps in the *reduce-scatter* stage cannot be suspended before completion, because fully reduced values are needed to minimize the message exchanged between clusters. Therefore, to reduce the number of nodes allowed to communicate, a few extra `merge_copy` operations must be performed before the `counter_reduce`. Repeating `merge_copy` i times makes 2^i nodes hold identical copies of the reduced message. Thus, after the i-th `merge_copy`, the $1/2^i$ nodes may participate in the `counter_reduce` function. However, the receiving side also needs to perform some extra `split_copy` operations to get back to the state at which *allgather* can be performed. Note that `split_copy` undoes the effect of `merge_copy` performed in the sending cluster. An arbitrary number of pairs of `split_copy` and `merge_copy` can be used without affecting the correctness of the algorithm.

Another way to reduce the traffic is to select a set of nodes by chaining one set after another, and phasing the use of the bottle-neck link. However, there is no portable way to know the end of transmission in the standard socket API. Messages are buffered in the socket, and there is no way to know when the socket buffer becomes empty. Therefore, chaining the nodes is not considered in our implementation, because it needs support from the system software.

The forwarding mechanism increases the traffic inside a cluster, and also occupies the receiving link on the forwarding node. However it is the simplest way, and allows regulating the sending nodes to an arbitrary number. Thus, simple forwarding is used in our implementation.

2.5 Adaptation for More Clusters

Naturally, the two-cluster algorithm described above can be extended to more clusters. Inter-cluster communication can be simple one-to-all for *bcast* and all-to-all for *allreduce*. The algorithms can also be adapted to the available bandwidth between each pair of clusters. When the wide-area network is shared, in a case such as one where one city exists in the middle of two other cities, the number of sending nodes should be reduced.

In addition, there is a situation where imbalance exists in the number of nodes of clusters. The above mentioned method to algorithmically reduce the number of communicating nodes can be used at the larger cluster, in case the numbers of nodes in clusters are different by more than a factor of two. By using the method, the number of communicating nodes in the larger cluster can be matched to that of the opposing cluster.

Table 1. Cost of Algorithms.

Bcast	
van de Geijn	$(L + M/nB + M/B + M/B)$
farfirst	$(L + M/B + (M/B + M/B))$
Allreduce	
Rabenseifner	$(L + M/nB + M/B + M/B)$
two-tier	$(L + M/B + (M/B + M/B) + (M/B + M/B))$

3 Evaluation

3.1 Simple Cost Model of Bcast

The communication cost is modeled by the following parameters in this section: M is the message size, B is the bandwidth of the link of the node, L is the latency of inter-cluster communication. n is the number of connections used in inter-cluster communication.

To compare the performance of the proposed *bcast* algorithm, a simple *bcast* algorithm is implemented. The algorithm is called *farfirst bcast* in this paper. In the algorithm, the whole message is sent first using inter-cluster communication. It minimizes time by using long links first. This algorithm is a simplified one used in existing MPI systems for wide-area networks. It uses the van de Geijn *bcast* algorithm inside a cluster.

The cost of *farfirst* is $(L + M/B + (M/B + M/B))$, when simply ignoring the latency in the intra-cluster communication. The term $(M/B + M/B)$ is for the *bcast* inside a cluster, where the first M/B in it corresponds to the *scatter* stage, and the second M/B to the *allgather* stage. Note that the cost of the *scatter* stage is M/B, because the `split_copy` operation halves the message in each step, and the sum of the cost of repeating it accumulates to M/B asymptotically. It is similar for the *allgather* stage.

Similarly, the cost of the modified van de Geijn *bcast* is $(L + M/nB + M/B + M/B)$. The second term is changed to M/nB by the effect of using multiple connections. Thus, the the modified van de Geijn wins by multiple uses of connections.

Table 1 summarizes the costs of the algorithms.

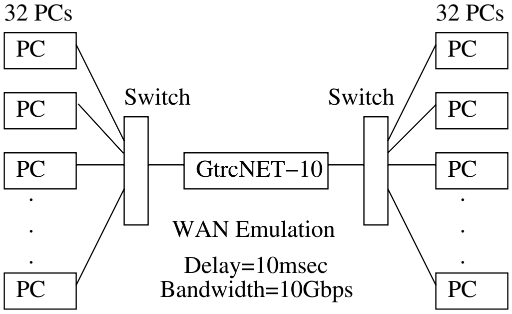

Figure 5. Experimental Setting.

Table 2. PC Cluster Specifications.

Node PC	
CPU	Opteron (2.0 GHz)
Memory	6GB DDR333
NIC	Broadcom BCM5704 (on-board)
OS	Fedora Core 5 (Linux-2.6.15)
Switch	
Huawei-3Com Quidway S5648 + optional 10 Gbps port	

3.2 Simple Cost Model of Allreduce

To compare the performance of the proposed *allreduce* algorithm, a simple *allreduce* algorithm is implemented. The algorithm is called *two-tier allreduce* in this paper. It first performs the reduction inside each cluster, exchanges the reduced messages between clusters, and then performs *bcast* with the messages inside each cluster. This algorithm is a simplified one used in existing MPI systems for wide-area networks. It performs the reduction inside a cluster by a variant of the Rabenseifner *allreduce*, in which *allgather* of the second stage is replaced with *gather*.

The cost of *two-tier allreduce* is $(L + M/B + (M/B + M/B) + (M/B + M/B))$. The third and fourth terms $(M/B + M/B)$ correspond to the reduction and *bcast* carried out inside a cluster.

Similarly, the cost of the modified Rabenseifner *allreduce* is $(L + M/nB + M/B + M/B)$. The third M/B term corresponds to the *reduce-scatter* stage and the fourth M/B term corresponds to the *allgather* stage.

Note that the modified Rabenseifner performs better algorithmically, without regard to the bandwidth of inter-cluster communication.

Table 1 summarizes the costs of the algorithms.

3.3 Experimental Setting

Figure 5 shows the experimental setting, and Table 2 shows the specification of the node PC and the Ethernet switches. Two clusters were connected via a WAN emulator. In the experiment, we used GtrcNET-10 [8] to emulate a WAN environment, which is a 10 Gbps successor of a well-established network testbed, GtrcNET-1 [12] for 1 Gbps Ethernet. GtrcNET-10 consists of a large-scale Field Programmable Gate Array (FPGA), three 10 Gbps Ethernet XENPAK ports, and three blocks of 1 GB DDR-SDRAM. The FPGA is a Xilinx XC2VP100, which includes three 10 Gbps Ethernet MAC and XAUI interfaces. GtrcNET-10 provides many functions such as traffic monitoring in millisecond resolution, traffic shaping, and WAN emulation at 10 Gbps wire speed. GtrcNET-10 was used to add latency between clusters and to observe precise network traffic. GtrcNET-10 added a 10 millisecond delay (one-way) in the experiment.

We used the MPI system, YAMPII [9], in the experiments, which almost fully implements the MPI-2.0 specification. YAMPII is the base of the MPI system, GridMPI [7], for Grid environment. Both YAMPII and GridMPI are fully functional, but YAMPII was used in the experiment, because GridMPI supports heterogeneity and has overheads in handling messages (e.g., byte-order conversion).

In the experiment, some TCP parameters were set as in the table below, because standard settings of the socket buffer sizes are not adequate for the experiment. tcp_no_metrics_save disables recording of the parameters of the previous connection to reuse them. These TCP parameters can be found in Linux in the directories /proc/sys/net/core and /proc/sys/net/ipv4.

tcp_no_metrics_save	1
wmem_max	3000000
rmem_max	3000000
tcp_rmem	3000000 3000000 3000000
tcp_wmem	3000000 3000000 3000000
tcp_mem	3000000 3000000 3000000

We ran each operation 10 times and took the maximum for stable results. TCP behaves disastrously at congestion, and the variance of performance sometimes reached near 50 percent in the experiment.

3.4 Bcast Performance

The left graph of Figure 6 shows the bandwidth achieved in *bcast* by varying the message size. The unit of the Y-axis is MB/s, but it just represents the value of the total user-level messages over the time $(M * nprocs/T)$. It does not count the actual messages sent by nodes because different algorithms send different amounts of messages.

The results labeled with *van de geijn* with n are for the modified van de Geijn algorithm, and n indicates the number of nodes simultaneously communicating. The

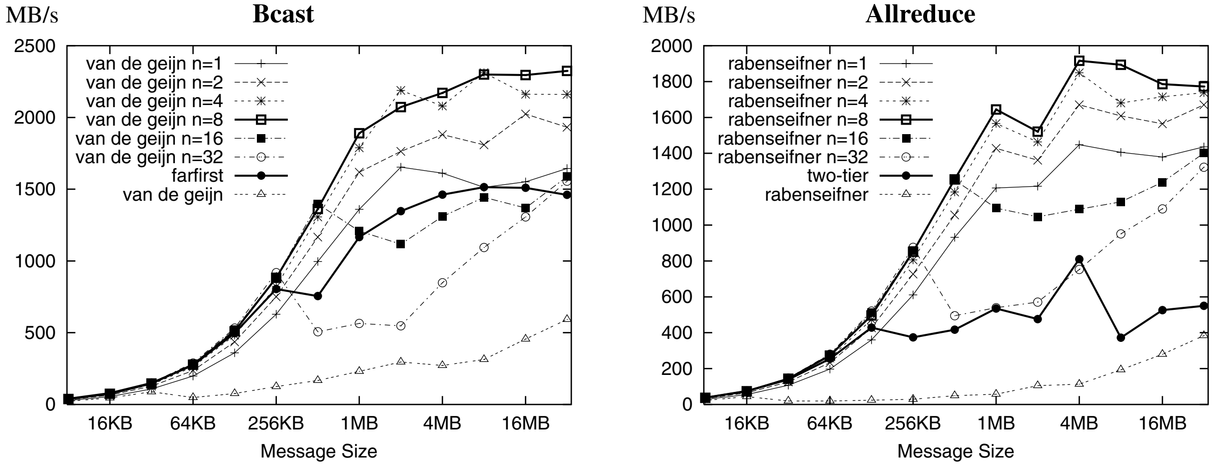

Figure 6. Throughput of bcast and allreduce (delay=10msec).

label *farfirst* indicates the *farfirst* algorithm. The label *van de geijn* alone indicates the original algorithm, which works in `MPI_COMM_WORLD` and does not distinguish a bottle-neck link. It is included for comparison.

As the cost model suggests, *farfirst bcast* behaves similarly to the modified van de Geijn algorithm for $n = 1$. However, the modified van de Geijn algorithm improves as n increases, up to $n = 8$. When the traffic overwhelms the bandwidth between clusters over $n = 8$, the performance drops but gradually improves as the message size increases.

No resend of TCP was observed in the modified van de Geijn algorithm for $n \leq 8$ in inter-cluster communication.

3.5 Allreduce Performance

The right graph of Figure 6 shows the bandwidth achieved in *allreduce* by varying the message size. The unit of the Y-axis is MB/s, but it just represents the value of the total message size over the time ($M * nprocs^2/T$).

The results labeled with *rabenseifner* with n are for the modified Rabenseifner algorithm, and n indicates the number of nodes simultaneously communicating. The label *two-tier* indicates the *two-tier* algorithm. The label *rabenseifner* alone indicates the original algorithm, which works in `MPI_COMM_WORLD` and does not distinguish a bottle-neck link.

As the cost model suggests, the modified Rabenseifner *allreduce* outperforms the *two-tier* algorithm, even at $n = 1$.

No resend of TCP was observed in the modified Rabenseifner algorithm for $n \leq 8$ in inter-cluster communication.

3.6 Clusters Environment

The proposed algorithms are expected to show good performance for clusters with limited bi-section bandwidth, such as clusters with multiple Ethernet switches or fat-trees with a reduced number of links at upper levels.

Figure 7 shows the results without delay. The two Ethernet switches were still connected via GtrcNET-10 and there was a bottle-neck at 10 Gbps bandwidth.

The results labeled with *van de geijn* and *rabenseifner* are for the original algorithms, which work in `MPI_COMM_WORLD` and do not distinguish a bottle-neck link. The performance of the original algorithms is not good. It is because they heavily use the bottle-neck link as if it had a full bi-section bandwidth. Although the modified algorithms generally perform well, the effect of reducing the number of nodes simultaneously communicating was almost not observed in a low latency environment.

4 Related Work

The van de Geijn *bcast* [1, 2] and the Rabenseifner *allreduce* [16] algorithms described in the design section are used in MPICH-1.2.6 and MPICH-2 [19].

PACX-MPI [5], MPICH-G2 [10], and MagPIe [11] are MPI systems designed for Grid environment. They are all designed to exploit the hierarchy of communication media ranging from memory systems to wide-area networks, and to adapt to the complex structure in latency and topology of wide-area networks. Although some optimality results have been presented for their algorithms, they are not designed to exploit the bandwidth of a single wide-area network. Thus, their *bcast* and *allreduce* algorithms are reduced to the simple *farfast* and *two-tier* algorithms shown in the evaluation section, when the setting is a simple two cluster configura-

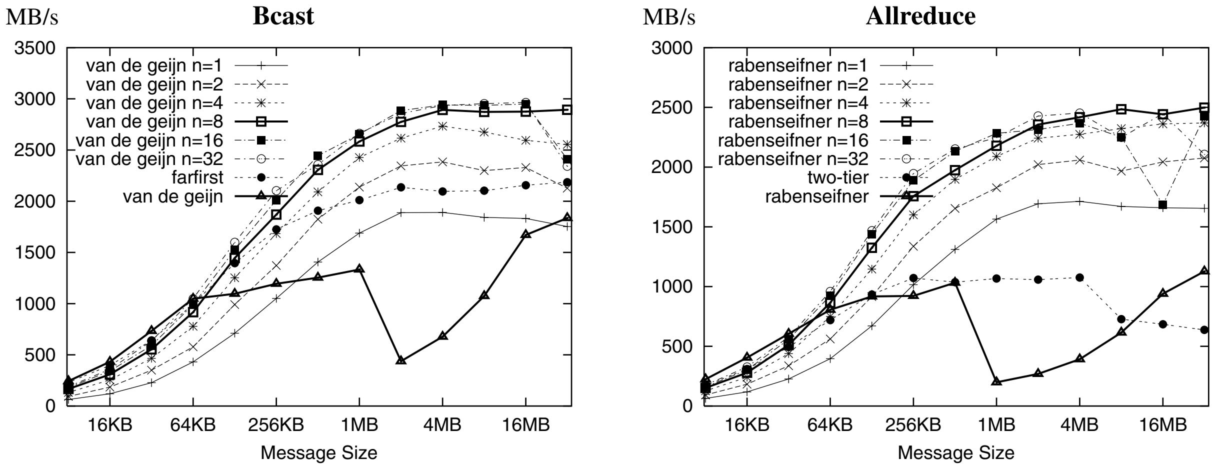

Figure 7. Throughput of bcast and allreduce (without delay).

tion connected by a fast network. We have already shown that the modified van de Geijn and modified Rabenseifner algorithms outperformed the *farfast* and *two-tier* algorithms in such a setting.

Chan, et al discuss algorithms of collective operations for machines capable of sending to or receiving from multiple links [4]. They have proposed the variations of algorithms depending on the capabilities and limitations of the torus network of BlueGene/L.

There is large amount of research in multicast overlay networks (including *bcast*) under the general setting where a network is represented by a weighted graph. Especially, den Burger, el al discuss the use of multiple multicast trees in clusters connected via wide-area networks [3]. However, it cannot directly be compared to the communication algorithms, because the work is general and proposes a method to find near-optimal multicast trees, and thus, the behavior of the algorithm is implied by the trees.

5 Conclusion

Since the assumption of low bandwidth wide-area networks is now false, the algorithms of collective communication need to be redesigned. In this paper, we have developed algorithms of inter-cluster collective communication for *bcast* and *allreduce*, which are based on the algorithms by van de Geijn, et al and by Rabenseifner. They are designed for fast wide-area networks, with bandwidth larger than that of a network link of a node. The algorithms utilize multiple node-to-node connections while regulating the number of nodes simultaneously communicating, and improve the performance of collective operations on large messages. Experiments using an emulated WAN environment with 10 Gbps bandwidth and a 10 millisecond latency have shown that our *bcast* and *allreduce* algorithms performed 1.6 and 3.2 times faster, respectively, than the algorithms used in existing MPI systems for Grid environment.

References

[1] M. Barnett, S. Gupta, D. G. Payne, L. Shuler, R. van de Geijn, and J. Watts. Interprocessor Collective Communication Library (InterCom). Proc. of the Scalable High Performance Computing Conference, 1994.

[2] M. Barnett, S. Gupta, D. G. Payne, L. Shuler, R. van de Geijn, and J. Watts. Building a High-Performance Collective Communication Library. Supercomputing 94 (SC94), 1994.

[3] M. den Burger, T. Kielmann, and H. E. Bal. Balanced Multicasting: High-throughput Communication for Grid Applications. Supercomputing 2005 (SC05), 2005.

[4] E. Chan, R. van de Geijn, W. Gropp, and R. Thakur. Collective Communication on Architectures that Support Simultaneous Communication over Multiple Links. Proc. of the ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming (PPoPP 2006), 2006.

[5] E. Gabriel, M. Resch, and R. Rühle. Implementing MPI with Optimized Algorithms for Metacomputing. Proc. of the Third MPI Developer's and User's Conference, 1999.

[6] GLIF: Global Lambda Integrated Facility. http://www.glif.is

[7] GridMPI Project. http://www.gridmpi.org

[8] GtrcNET-10. http://www.gtrc.aist.go.jp/gnet

[9] Y. Ishikawa. YAMPII Official Home Page. http://www.il.is.s.u-tokyo.ac.jp/yampii

[10] N. T. Karonis, B. R. de Supinski, I. T. Foster, W. Gropp, E. L. Lusk, and S. Lacour. A Multilevel Approach to Topology-Aware Collective Operations in Computational Grids. Tech. Rep. ANL/MCS-P948-0402, 2002.

[11] T. Kielmann, H. E. Bal, and S. Gorlatch. Bandwidth-efficient Collective Communication for Clustered Wide Area Systems. Proc. of the 14th Intl. Symp. on Parallel and Distributed Processing, 1999.

[12] Y. Kodama, T. Kudoh, R. Takano, H. Sato, O. Tatebe, and S. Sekiguchi. GNET-1: Gigabit Ethernet Network Testbed. IEEE Intl. Conf. on Cluster Computing (Cluster2004), 2004.

[13] M. Matsuda, Y. Ishikawa, and T. Kudoh. Evaluation of MPI Implementations on Grid-connected Clusters using an Emulated WAN Environment. CCGrid2003, 2003.

[14] M. Matsuda, Y. Ishikawa, T. Kudoh, Y. Kodama, and R. Takano. TCP Adaptation for MPI on Long-and-Fat Networks. IEEE Intl. Conf. on Cluster Computing (Cluster2005), 2005.

[15] R. Rabenseifner. Automatic MPI Counter Profiling of All Users: First Results on a CRAY T3E 900-512. Proc. of the Message Passing Interface Developers and Users Conference 1999 (MPIDC99), 1999.

[16] R. Rabenseifner. Optimization of Collective Reduction Operations. Intl. Conf. on Computational Science, LNCS, Springer-Verlag, 2004.

[17] R. Takano, T. Kudoh, Y. Kodama, M. Matsuda, H. Tezuka, and Y. Ishikawa. Design and Evaluation of Precise Software Pacing Mechanisms for Fast Long-Distance Networks. 3rd Intl. Workshop on Protocols for Fast Long-Distance Networks (PFLDnet05), 2005.

[18] TeraGrid. http://www.teragrid.org

[19] R. Thakur, R. Rabenseifner, and W. Gropp. Optimization of Collective Communication Operations in MPICH. Intl. Journal of High Performance Computing Applications, Vol.19, No.1, pp49-66, 2005.