

The Performance of Configurable Collective Communication for LAM-MPI in Clusters and Multi-Clusters

John Markus Bjørndalen¹, Otto J. Anshus¹
Brian Vinter², Tore Larsen¹

¹) Department of Computer Science, University of Tromsø

²) Department of Mathematics and Computer Science, University of Southern Denmark

Abstract

Using a cluster of eight four-way computers, *PastSet*, an experimental tuple space based shared memory system, has been measured to be 1.83 times faster on global reduction than using the *Allreduce* operation of LAM-MPI. Our hypothesis is that this is due to *PastSet* using a better mapping of processes to computers resulting in less messages and more use of local processor cycles to compute partial sums. *PastSet* achieved this by utilizing the *PATHS* system for configuring multi-cluster applications. This paper reports on an experiment to verify the hypothesis by using *PATHS* on LAM-MPI to see if we can get better performance, and to identify the contributing factors.

By adding configurability to *Allreduce* by using *PATHS*, we achieved a performance gain of 1.52, 1.79, and 1.98 on respectively two, four and eight-way clusters. We conclude that the LAM-MPI algorithm for mapping processes to computers when using *Allreduce* was the reason for its poor performance relative to the implementation using *PastSet* and *PATHS*.

We then did a set of experiments to examine whether we could improve the performance of the *Allreduce* operation when using two clusters interconnected by

a WAN link with 30-50ms roundtrip latency. The experiments showed that even a bad mapping of *Allreduce* which resulted in multiple messages being sent across the WAN did not add significant performance penalty to the *Allreduce* operation for packet sizes up to 4KB. We believe this is due to multiple small messages concurrently in transit on the WAN.

1 Introduction

For efficient support of synchronization and communication in parallel systems, these systems require fast collective communication support from the underlying communication subsystem as, for example, is defined by the Message Passing Interface (MPI) Standard [1]. Among the set of collective communication operations broadcast is fundamental and is used in several other operations such as barrier synchronization and reduction [12]. Thus, it is advantageous to reduce the latency of broadcast operations on these systems.

In our work with the *PATHS*[5] configuration and orchestration system, we have experimented with micro-benchmarks and applications to study the effects of configurable communication.

In one of the experiments[18], we used

the configuration mechanisms to reduce the execution times of collective communication operations in PastSet. To get a baseline, we compared our reduction operation with the equivalent operation in MPI (Allreduce).

By trying a few configurations, we found that we could improve our Tuple Space system to be 1.83 times faster than LAM-MPI (Local Area for Multicomputer MPI) [11][14]. Our hypothesis was that this advantage came from a better usage of resources in the cluster rather than a more efficient implementation.

If anything, LAM-MPI should be faster than PastSet since PastSet stores the results of each global sum computation in a tuple space inducing more overhead than simply computing and distributing the sum.

This paper reports on an experiment where we have added configurable communication to the Broadcast and Reduce operations in LAM-MPI (both of which are used by Allreduce) to validate or falsify our hypothesis. In [4] we report on complimentary experiments where we also varied the number of computers per experiment.

The paper is organized as follows: Section 2 summarizes the main features of the PastSet and PATHS system. Section 3 describes the Allreduce, Reduce and Broadcast operations in LAM-MPI. Section 4 describes the configuration mechanism that was added to LAM-MPI for the experiments reported on in this paper. Section 5 describes the experiments and results, section 7 presents related work, section 6 presents our multicluster experiments, and section 8 concludes the paper.

2 PATHS: Configurable Orchestration and Mapping

Our research platform is PastSet[2][17], a structured distributed shared memory system in the tradition of Linda[6]. Past-

Set is a set of Elements, where each Element is an ordered collection of tuples. All tuples in an Element follow the same template.

The PATHS[5] system is an extension of PastSet that allows for mappings of processes to hosts at load time, selection of physical communication paths to each element, and distribution of communications along the path. PATHS also implements the X-functions[18], which are PastSet operation modifiers.

A path specification for a single thread needing access to a given element is represented by a list of stages. Each stage is implemented using a wrapper object hiding the rest of the path after that stage. The stage specification includes parameters used for initialisation of the wrapper.

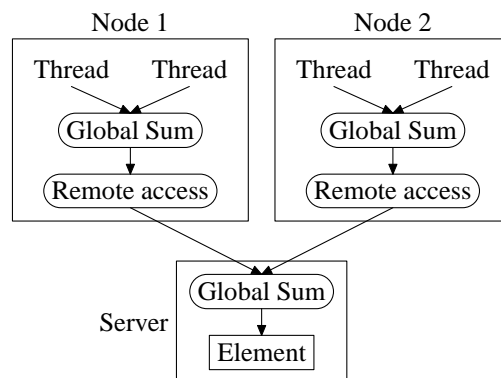


Figure 1. Four threads on two hosts accessing shared element on a separate server. Each host computes a partial sum that is forwarded to the global-sum wrapper on the server. The final result is stored in the element.

Paths can be shared whenever path descriptions match and point to the same element (see figure 1). This can be used to implement functionality such as, for instance, caches, reductions and broadcasts.

The collection of all paths in a system pointing to a given element forms a tree.

The leaf nodes in the tree are the application threads, while the root is the element.

Figure 1 shows a global reduction tree. By modifying the tree and the parameters to the wrappers in the tree, we can specify and experiment directly with factors such as which processes participate in a given partial sum, how many partial sum wrappers to use, where each sum wrapper is located, protocols and service requirements for remote operations and where the root element is located. Thus, we can control and experiment with tradeoffs between placement of computation, communication, and data location.

Applications tend to use multiple trees, either because the application uses multiple elements, or because each thread might use multiple paths to the same element.

To get access to an element, the application programmer can either choose to use lower-level functions to specify paths before handing it over to a path builder, or use a higher level function which retrieves a path specification to a named element and then builds the specified path. The application programmer then gets a reference to the topmost wrapper in the path.

The path specification can either be retrieved from a combined path specification and name server, or be created with a high-level language library loaded at application load-time¹

Since the application program invokes all operations through its reference to the topmost wrapper, the application can be mapped to different cluster topologies simply by doing one of the following:

- Updating a map description used by the high-level library.
- Specifying a different high-level library that generates path-specifications. This library may be written by the user.

¹Currently, a Python module is loaded for this purpose.

- Update the path mappings in the name server.

Profiling is provided via trace wrappers that log the start and completion time of operations that are invoked through it. Any number of trace wrappers can be inserted anywhere in the path.

Specialized tools to examine the performance aspects of the application can later read trace data stored with the path specifications from a system. We are currently experimenting with different visualizations and analyses of this data to support optimization of a given application.

The combination of trace data, a specification of communication paths, and computations along the path has been useful in understanding performance aspects and tuning benchmarks and applications that we have run in cluster and multi-cluster environments.

3 LAM-MPI implementation of Allreduce

LAM-MPI is an open source implementation of MPI available from [11]. It was chosen over MPICH [7] for our work in [18] since it had lower latency with less variance than MPICH for the benchmarks we used in our clusters.

The MPI Allreduce operation combines values from all processes and distribute the result back to all processes. LAM-MPI implements Allreduce by first calling Reduce, collecting the result in the root process, then calling Broadcast, distributing the result from the root process. For all our experiments, the root process is the process with rank 0 (hereafter called process 0).

The Reduce and Broadcast algorithms use a linear scheme (every process communicates directly with process 0) up to and including 4 processes. From there on they use a scheme that organizes the processes into a logarithmic spanning tree.

The shape of this tree is fixed, and doesn't change to reflect the topology of the computing system or cluster. Figure 2 shows the reduction trees used in LAM-MPI for 32 processes in a cluster. We observe that broadcast and reduction trees are different.

By default, LAM-MPI evenly distributes processes onto nodes. When we combine this mapping for 32 processes with the reduction tree, we can see in Figure 3 that a lot of messages are sent across nodes in the system. The broadcast operation has a better mapping for this cluster though.

4 Adding configuration to LAM-MPI

To minimize the necessary changes to LAM-MPI for this experiment, we didn't add a full PATHS system at this point. Instead, a mechanism was added that allowed for scripting the way LAM-MPI communicates during the broadcast and reduce operations.

There were two main reasons for this. Firstly, our hypothesis was that PastSet with PATHS allowed us to map the communication and computation better to the resources and cluster topology. For global reduction and broadcast, LAM-MPI already computes partial sums at internal nodes in the trees. This means that experimenting with different reduction and broadcast trees should give us much of the effect that we observed with PATHS in [18] and [5].

Secondly, we wanted to limit the influence that our system would have on the performance aspects of LAM-MPI such that any observable changes in performance would come from modifying the reduce and broadcast trees.

Apart from this, the amount of code changed and added was minimal, which reduced the chances of introducing errors into the experiments.

When reading the LAM-MPI source code, we noticed that the reduce operation was, for any process in the reduction tree, essentially a sequence of N receives from the N children directly below it in the tree, and one send to the process above it. For broadcast, the reverse was true; one receive followed by N sends.

Using a different reduction or broadcast tree would then simply be a matter of examining, for each process, which processes are directly above and below it in the tree and construct a new sequence of send and receive commands.

To implement this, we added new reduce and broadcast functions which used the rank of the process and size of the system to look up the sequence of sends and receives to be executed (including which processes to send and receive from). This is implemented by retrieving and executing a script with send and receive commands.

As an example, when using a scripted reduce operation with a mapping identical to the original LAM-MPI reduction tree, the process with rank 12 (see figure 2) would look up and execute a script with the following commands:

- Receive (and combine result) from rank 13
- Receive (and combine result) from rank 14
- Send result to 8

The new scripted functions are used instead of the original logarithmic Reduce and Broadcast operations in LAM-MPI. No change was necessary to the Allreduce function since it is implemented using the Reduce and Broadcast operations.

The changes to the LAM-MPI code was thus limited to 3 code lines, replacing the calls to the logarithmic reduction and broadcast functions as well as adding and a call in `MPI_Init` to load the scripts.

Remapping the application to another cluster configuration, or simply trying new mappings for optimization purposes, now consists of specifying new communication trees and generating the scripts. A Python program generates these scripts as Lisp symbolic expressions.

5 Experiments

Figure 4 shows the code run in the experiments. The code measures the average execution time of 1000 Allreduce operations. The average of 5 runs is then plotted. To make sure that the correct sum is computed, the code also checks the result on each iteration.

For each experiment, the number of processes was varied from 1 to 32. LAM-MPI used the default placement of processes on nodes, which evenly spread the processes over the nodes.

The hardware platforms consists of three clusters, each with 32 processors:

- 2W: 16*2-Way Pentium III 450 MHz, 256MB RAM
- 4W: 8*4-Way Pentium Pro 166 MHz, 128MB RAM
- 8W: 4*8-Way Pentium Pro 200 MHz, 2GB RAM

All clusters used 100MBit Ethernet for intra-cluster communication.

5.1 4-way cluster

Figure 5 shows experiments run on the 4-way cluster. The following experiments were run:

Original LAM-MPI The experiment was run using LAM-MPI 6.5.6.

Modified LAM-MPI, original scheme

The experiment was run using a modified LAM-MPI, but using the same broadcast and reduce trees

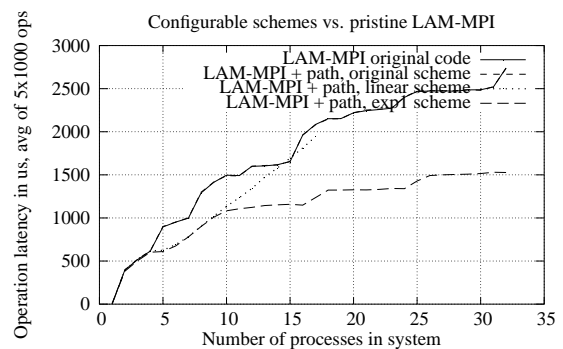


Figure 5. Allreduce, 4-way cluster

for communication as the original version used.

This graph completely overlaps the graph from the original LAM-MPI code, showing us that the scripting functionality is able to replicate the performance aspects of the original LAM-MPI for this experiment.

Modified LAM-MPI, linear scheme

This experiment was run using a linear scheme where all processes report directly to the process with rank 0, and rank 0 sends the results directly back to each client.

This experiment crashed at 18 processes due to TCP connection failures in LAM-MPI. We haven't examined what caused this, but it is likely to be a limitation in the implementation.

Modified LAM-MPI, hierarchal scheme 1

Rank 0 is always placed on node 0, so all other processes on node 0 contribute directly to rank 0. On the other nodes, a local reduction is done within each node before the nodes partial sum is sent to process 0.

The modified LAM-MPI with the hierarchal scheme is 1.79 times faster than the original LAM-MPI. This is close to the factor of 1.83 reported in [18], which was based on measurements on the same cluster. It is possible that further experiments




Figure 6. Allreduce, 8-way cluster

with configurations would have brought us closer to the original speed improvement.

5.2 8-way cluster

For the 8-way cluster, we expected the best configuration to involve using an internal root process on each node and limit the inter-node communications to one send (of the local sum) and one receive (of the global result) resembling the best configuration for the 4-way cluster. This turned out not to be true.

Due to the 8 processes on each node, the local root process on each node would get 7 children in the sum tree (with the exception of the root processes in the root node). The extra latency added per child of the local root processes was enough to increase the total latency for 32 processes to 1444 microseconds.

Splitting the processes to allow further subgrouping within each node introduced another layer in the sum tree. This extra layer brought the total latency for 32 processes to 1534 microseconds.

The fastest solution found so far involves partitioning the processes on each node into two groups of 4 processes. Each group computes a partial sum for that group, and forwards it to directly to the root process. This doubled the communication over the network, but the total execution time was shorter (1322 microseconds).

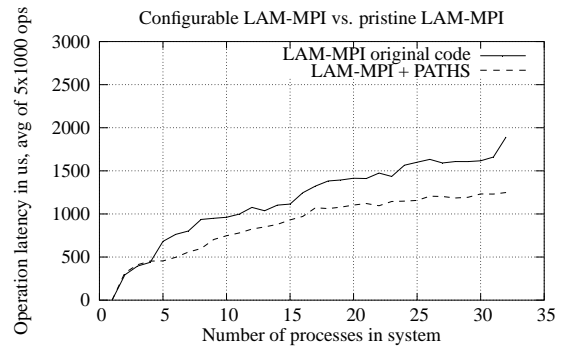


Figure 7. Allreduce, 2-way cluster.

The fastest configuration is 1.98 times faster than the original LAM-MPI Allreduce operation. In figure 6, the fastest configuration and the original LAM-MPI Allreduce are plotted for 1-32 processes.

5.3 2-way cluster

On the 2-way cluster, after trying a few configurations, we ended up with a configuration that was 1.52 times faster than the original LAM-MPI code. This was not as good as for the other clusters. We expected this since a reduction in this cluster would need more external messages.

Based on the experiments, we observed three main factors that influenced the scaling of the application:

- The number of children for a process in the operation tree. More children adds communication load for the node hosting this process.
- The depth of a branch in the operation tree. That is, the number of nodes a value has to be added through before reaching the root node.
- Whether a message was sent over the network or internally on a node.

The depth factor seemed to introduce a higher latency than adding a few more children to a process. On the other hand, adding another layer in the tree would also allow us to use partial sums. We have

not arrived at any model which allows us predict which tree organization would in practice lead to the shortest execution time of the Allreduce operation.

One difference between the PastSet/PATHS implementation and LAM-MPI is that PastSet/PATHS process incoming partial sums by order of arrival, while LAM-MPI process them in the order they appear in the sequence specified either by the original LAM-MPI code, or the script. This might influence the overhead when receiving messages from a larger number of children. We are investigating this further.

6 Multi-cluster experiments

In [5] we also ran multi-cluster global reductions using PastSet and PATHS, and [10] shows an approach where they reduce the number of messages over WAN connections to reduce the latency of collective operations over wide-area links.

To study the effect of choosing various configurations on multi-clusters, we added an experiment where we installed an IP tunnel between the 4W cluster in Tromsø and the 2W cluster in Denmark and ran the Allreduce benchmark using those clusters as one larger cluster.

Unfortunately, during the experiments two things happened: two of the nodes in Denmark went down, and the 2W cluster was allocated to other purposes after we finished the measurements with the unmodified LAM-MPI Allreduce, so we were unable to run the benchmark using configurable Allreduce on the multi-cluster system.

However, the performance measurements using the unmodified Allreduce did produce interesting results. An experiment with 32 processes both in Denmark and Tromsø documented that the reduce phase of the Allreduce operation had a very good mapping for minimizing messages over the WAN link: only one mes-

sage was sent from the Tromsø cluster to the root process in the Denmark cluster. However, in the broadcast phase each process in Tromsø was sent a message from a process in Denmark.

We expected this configuration of Allreduce to perform gradually worse when we scaled up the number of processes in Tromsø from 1 to 32. This did not turn out to be the case. With 32 processes in Denmark, the difference between the Allreduce latency of using 1 and 32 processes in Tromsø was small enough to be masked by the noise in latency.

We suspected that the reason for this was that the numerous routers and links down to Denmark allowed multiple small messages to propagate towards Tromsø concurrently.

This suggests that for the cluster sizes that we have available, a configurable mechanism that reduces the number of *small* messages on a long WAN link does not significantly reduce the latency. However, for larger packet sizes, the bandwidth of the link is important, and a reduction in the number of messages should matter more.

To study this, the benchmark was extended to include not only scaling from 1 to 64 processes, but also to calling Allreduce with value vectors from 4 byte (1 integer) to 4MB.

Since two computers in Denmark had gone down, we ran the benchmark with 28 processors (14 nodes) in Denmark, and 32 processors in Tromsø reducing the total cluster size to 60 processors. This produced a similar mapping to the original: when running with 60 processes, all Tromsø nodes would receive a message directly from Denmark during broadcast, and only two messages would be sent down to Denmark on reduce.

We also ran the benchmark with 61 processes. This resulted in a configuration where the process with rank 60 (the last process added when scaling the system to

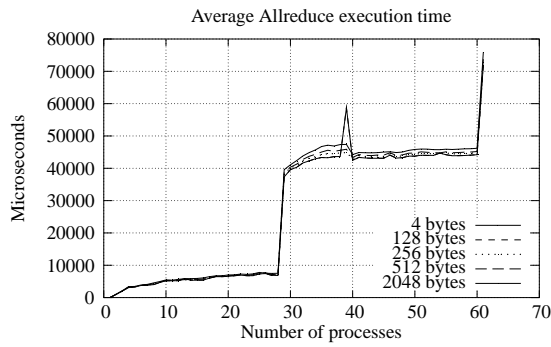


Figure 8. Allreduce performance using two clusters.




Figure 10. Roundtrip TCP/IP latency for one of N threads. 4W node to 8W node (LAN)

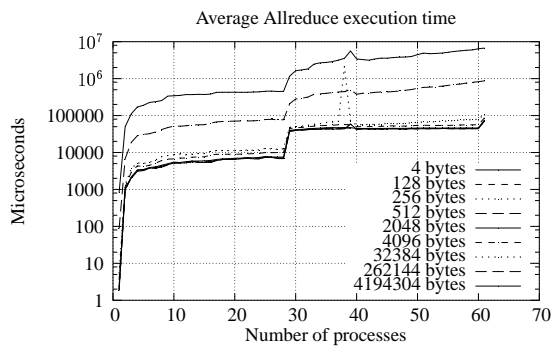


Figure 9. Allreduce performance using two clusters. Note logarithmic scale of y-axis.

61 processes) is located in Denmark but only communicates directly with nodes in Tromsø for Reduce and Broadcast. This forces the nodes in Tromsø to wait for process 60 before one of the two messages can be sent down to Denmark again, and the corresponding broadcast message for 60 has to propagate through Tromsø before reaching it. Figure 8 shows that this configuration nearly doubles the latency of Allreduce compared to running Allreduce with 60 processes. Clearly, configurations can be encountered where it pays off to control who communicates with who even for small messages.

Figure 8 also shows us that sending many messages (there is one message per process in the figure) for vectors up to 2KB does not add much performance

penalty to Allreduce compared to sending only a few messages. The figure shows a spike around 38 processes, which comes from a temporary network routing problem somewhere between Denmark and Tromsø when the experiment was performed.

Figure 9 shows that the performance of Allreduce rapidly deteriorates when the size and number of packets increase. The graph uses a logarithmic scale for the y-axis.

To examine this effect further, we devised another experiment. We had one thread measuring the latency of roundtrip messages over TCP/IP. To simulate the bad mapping of the Broadcast operation, a number of threads were added sending roundtrip messages concurrently with the first thread.

The hypotheses were that adding the extra threads would impact the first thread the least using the WAN to Denmark, and that the background threads would influence the first thread more for small messages when using the LAN.

Figure 10 shows that communicating between two cluster nodes, using only a single router between them, the background communication would influence the roundtrip latency of the benchmark thread even at the smallest packet sizes. As we add background communication,

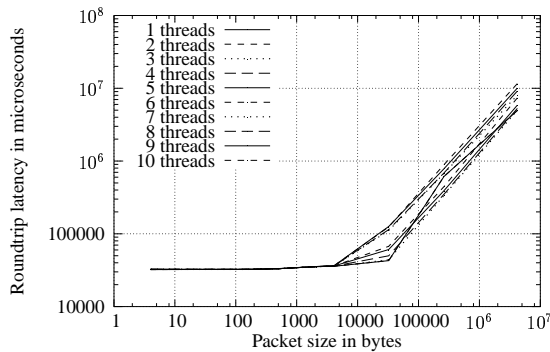


Figure 11. Roundtrip TCP/IP latency for one of N threads. Gateway node in Denmark to 4W node (WAN)

the latency increases for all packet sizes.

The cluster in Denmark was busy at the time of the experiment, so we had to use the clusters gateway node when running the WAN experiment. The gateway was where we originally installed the IP tunnel for the LAM-MPI Allreduce benchmark, so we still used the tunnel for the experiment.

Figure 11 shows that for messages up to 4KB, we can hardly see any difference in the performance of Allreduce as we add background communication. Only with larger packet sizes, where we assume that the bandwidth plays a larger role, is the performance influenced by the background communication.

7 Related work

Jacunski et al.[9] shows that selection of the best performing algorithm for all-to-all broadcast on clusters of workstations based on commodity switch-based networks is a function of both network characteristics as well as the message length and number of participating nodes in the all-to-all broadcast operation. In contrast our work used clusters with multiprocessors, we did performance measurements on the reduce operation as well as on the broadcast, and we documented the effect of the actual system at run time includ-

ing the workload, communication performance of the hosts.

Bernashci et al.[3] study the performance of the MPI broadcast operation on large shared memory systems using a-trees.

Kielmann et al.[10] show how the performance for collective operations like broadcast depend upon cluster topology, latency, and bandwidth. They develop a library of collective communication operations optimized for wide area systems where a majority of the speedup comes from limiting the number messages passing over wide-area links.

Husbands et al.[8] optimizes the performance of MPI_Bcast in clusters of SMP nodes by using a two-phase scheme; first messages are sent to each SMP node, then messages are distributed within the SMP nodes. Sistare et al.[13] uses a similar scheme, but focus on improving the performance of collective operations on large-scale SMP's.

Tang et al.[15] also use a two-phase scheme for a multithreaded implementation for MPI. They separate the implementation of the point-to-point and collective operations to allow for optimizations of the collective operations, which would otherwise be difficult.

In contrast to the above works, this paper is not focused on a particular optimization of the spanning trees. Instead, we focus on making the shape of the spanning trees configurable to allow easy experimentation on various cluster topologies and applications.

Vadhiyar et al.[16] shows an automatic approach to selecting buffer sizes and algorithms to optimize the collective operation trees by conducting a series of experiments.

8 Conclusions

We have observed that the broadcast and reduction trees in LAM-MPI are dif-

ferent, and do not necessarily take into account the actual topology of the cluster.

By introducing configurable broadcast and reduction trees, we have shown a simple way of mapping the reduction and broadcast trees to the actual clusters in use. This gave us a performance improvement up to a factor of 1.98.

For the cluster where we observed the performance difference of a factor 1.83 between PastSet/PATHS and LAM-MPI, we arrived at a reduction and broadcast tree that gave us an improvement of 1.79 for Allreduce over the original LAM-MPI implementation. This supports our hypothesis that the majority of the performance-difference between LAM-MPI and Past-Set/PATHS was a better mapping to the resources and topology in the cluster.

We have also observed that the assumption that doing a reduction internally in each node before sending a message on the network did not lead to the best performance on our cluster of 8-way nodes. Instead, increasing the number of messages on the network to reduce the depth of the reduction and broadcast as well as the number of direct children for each internal node proved to be a better strategy.

The reason for this may be found by studying three different factors that add to the cost of the global reduction and broadcast trees:

- The number of children directly below an internal node in the spanning tree.
- The depth of the spanning tree.
- Whether an arc between a child and a parent in the spanning tree is a message on the network or internally in the node.

We suspect that these factors are not increasing linearly, and that the cost of, for instance, adding another child to an internal node in the spanning tree depends on

factors such as the order of the sequence of receive commands as well as contention on the network layer in the host computer.

For the multi-cluster experiments, we observed that configurations which sent more messages than necessary over the WAN link did not perform as bad as we had expected. For message sizes up to 4KB, the extra messages did not add a noticeable operation time to the Allreduce operation. We believe that multiple messages concurrently in transit through the routers along the WAN link masks some of the overhead of the extra messages sent by LAM-MPI.

For larger messages, we see that the bandwidth of the WAN link is starting to penalize the bad configurations.

9 Acknowledgements

Thanks to Ole Martin Bjørndalen for reading the paper and suggesting improvements.

References

- [1] Mpi: A message-passing interface standard. *Message Passing Interface Forum* (Mar 1994).
- [2] ANSHUS, O. J., AND LARSEN, T. Macroscopic: The abstractions of a distributed operating system. *Norsk Informatikk Konferanse* (October 1992).
- [3] BERNASCHI, M., AND RICHELLI, G. Mpi collective communication operations on large shared memory systems. *Proceedings of the Ninth Euromicro Workshop on Parallel and Distributed Processing (EUROPDP.01)* (2001).
- [4] BJØRNDALLEN, J. M., ANSHUS, O., VINTER, B., AND LARSEN, T. Configurable collective communication

- in lam-mpi. *Proceedings of Communicating Process Architectures 2002, Reading, UK* (September 2002).
- [5] BJØRNDALEN, J. M., ANSHUS, O., LARSEN, T., AND VINTER, B. Paths - integrating the principles of method-combination and remote procedure calls for runtime configuration and tuning of high-performance distributed application. *Norsk Informatikk Konferanse* (November 2001), 164–175.
- [6] CARRIERO, N., AND GELERNTER, D. Linda in context. *Commun. ACM* 32, 4 (April 1989), 444–458.
- [7] GROPP, W., LUSK, E., DOSS, N., AND SKJELLUM, A. A high-performance, portable implementation of the mpi message passing interface standard. *Parallel Computing, Volume 22, Issue 6* (September 1996).
- [8] HUSBANDS, P., AND HOE, J. C. Mpi-start: delivering network performance to numerical applications. *Proceedings of the 1998 ACM/IEEE conference on Supercomputing* (1998). San Jose, CA.
- [9] JACUNSKI, M., SADAYAPPAN, P., AND PANDA, D. All-to-all broadcast on switch-based clusters of workstations. *13th International Parallel Processing Symposium and 10th Symposium on Parallel and Distributed Processing* (12 - 16 April 1999). San Juan, Puerto Rico.
- [10] KIELMANN, T., HOFMAN, R. F. H., BAL, H. E., PLAAT, A., AND BHOEDJANG, R. A. F. Magpie: Mpi's collective communication operations for clustered wide area systems. *Proceedings of the seventh ACM SIGPLAN symposium on Principles and practice of parallel programming* (1999). Atlanta, Georgia, United States.
- [11] <http://www.lam-mpi.org/>.
- [12] PANDA, D. Issues in designing efficient and practical algorithms for collective communication in wormhole-routed systems. *Proc. ICPP Workshop Challenges for Parallel processing* (1995), 8–15.
- [13] SISTARE, S., VANDEVAART, R., AND LOH, E. Optimization of mpi collectives on clusters of large-scale smp's. *Proceedings of the 1999 conference on Supercomputing* (1999). Portland, Oregon, United States.
- [14] SQUYRES, J. M., LUMSDAINE, A., GEORGE, W. L., HAGEDORN, J. G., AND DEVANEY, J. E. The interoperable message passing interface (IMPI) extensions to LAM/MPI. In *Proceedings, MPIDC'2000* (March 2000).
- [15] TANG, H., AND YANG, T. Optimizing threaded mpi execution on smp clusters. *Proceedings of the 15th international conference on Supercomputing* (2001). Sorrento, Italy.
- [16] VADHIYAR, S. S., FAGG, G. E., AND DONGARRA, J. Automatically tuned collective communications. *Proceedings of the 2000 conference on Supercomputing* (2000). Dallas, Texas, United States.
- [17] VINTER, B. *PastSet a Structured Distributed Shared Memory System*. PhD thesis, Tromsø University, 1999.
- [18] VINTER, B., ANSHUS, O. J., LARSEN, T., AND BJØRNDALEN, J. M. Extending the applicability of software dsm by adding user redefinable memory semantics. *Parallel Computing (ParCo) 2001, Naples, Italy* (September 2001).

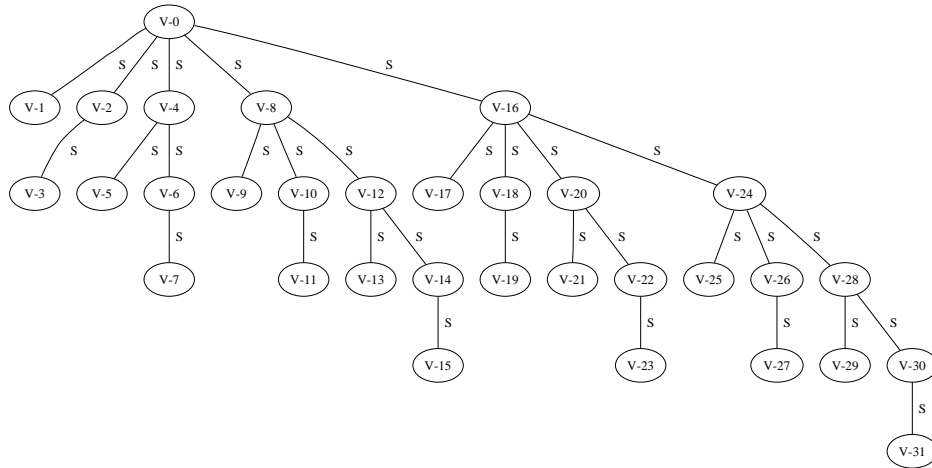


Figure 2. Log-reduce tree for 32 processes. The arcs represent communication between two nodes. Partial sums are computed at a node in the tree before passing the result further up in the tree.

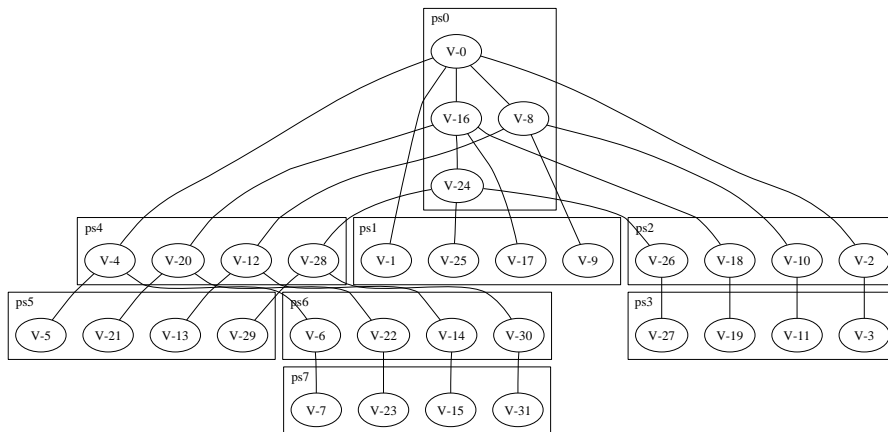


Figure 3. Log-reduce tree for 32 processes mapped onto 8 nodes.

```

t1 = get_usecs();
MPI_Allreduce(&hit, &ghit, 1, MPI_INT, MPI_SUM, MPI_COMM_WORLD);
t1 = get_usecs();
for (i = 0; i < ITERS; i++) {
    MPI_Allreduce(&i, &ghit, 1, MPI_INT, MPI_SUM, MPI_COMM_WORLD);
    if (ghit != (i * size))
        printf("oops at %d. %d != %d\n", i, ghit, (i * size));
}
t2 = get_usecs();
MPI_Allreduce(&hit, &ghit, 1, MPI_INT, MPI_SUM, MPI_COMM_WORLD);

```

Figure 4. Global reduction benchmark