

Optimization of Collective Communication in Intra-Cell MPI

M. K. Velamati¹, A. Kumar¹, N. Jayam¹, G. Senthilkumar¹, P.K. Baruah¹,
R. Sharma¹, S. Kapoor², and A. Srinivasan³

¹ Dept. of Mathematics and Computer Science, Sri Sathya Sai University

² IBM, Austin

³ Dept. of Computer Science, Florida State University
asriniva@cs.fsu.edu

Abstract. The Cell is a heterogeneous multi-core processor, which has eight co-processors, called SPEs. The SPEs can access a common shared main memory through DMA, and each SPE can directly operate on a small distinct local store. An MPI implementation can use each SPE as if it were a node for an MPI process. In this paper, we discuss the efficient implementation of collective communication operations for intra-Cell MPI, both for cores on a single chip, and for a Cell blade. While we have implemented all the collective operations, we describe in detail the following: barrier, broadcast, and reduce. The main contributions of this work are (i) describing our implementation, which achieves low latencies and high bandwidths using the unique features of the Cell, and (ii) comparing different algorithms, and evaluating the influence of the architectural features of the Cell processor on their effectiveness.

Keywords: Cell Processor, MPI, heterogeneous multicore processor.

1 Introduction

The Cell is a heterogeneous multi-core processor from Sony, Toshiba and IBM. There has been much interest in using it in High Performance Computing, due to the high flop rates it provides. However, applications need significant changes to fully exploit the novel architecture. A few different models of the use of MPI on the Cell have been proposed to deal with the programming difficulty, as explained later. In all these, it is necessary to implement collective communication operations efficiently within each Cell processor or blade.

In this paper, we describe the efficient implementation of a variety of algorithms for a few important collective communication operations, and evaluate their performance. The outline of the rest of the paper is as follows. In §2, we describe the architectural features of the Cell that are relevant to the MPI implementation, and MPI based programming models for the Cell. We explain common features of our implementations in §3.1. We then describe the implementations and evaluate the performance of MPI_Barrier, MPI_Broadcast, and MPI_Reduce in §3.2, §3.3, and §3.4 respectively. We summarize our conclusions in §4. *Further details on this work are available in a technical report [4].*

2 Cell Architecture and MPI Based Programming Models

Architecture. Figure 1 shows an overview of the Cell processor. It consists of a cache coherent PowerPC core (PPE), which controls eight SIMD cores called Synergistic Processing Elements (SPEs). All cores run at 3.2 GHz and execute instructions in-order. The Cell has a 512 MB to 2 GB external main memory, and an XDR memory controller provides access to it at a rate of 25.6 GB/s. The PPE, SPE, DRAM and I/O controllers are all connected via four data rings, collectively known as the EIB. Up to 128 outstanding DMA requests between main storage and SPEs can be in process concurrently on the EIB. The EIB's maximum bandwidth is 204.8 GB/s. The Cell Broadband Engine Interface (BEI) manages data transfers between the EIB and I/O devices. One of its channels can be used to connect to another Cell processor at 25.6 GB/s, creating a Cell blade with a logical global shared memory.

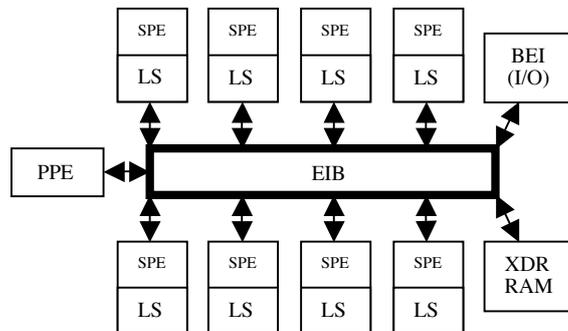


Fig. 1. Overview of the Cell processor.

Each SPE has its own 256 KB local store from which it fetches code and reads and writes data, with access latency of 6 cycles. All loads and stores issued from the SPE can only access the SPE's local store. Any main memory data needed by the SPE must be moved into the local store explicitly through a DMA. An SPE can have up to sixteen pending requests in its DMA queue. The maximum DMA size is 16 KB.

The DMAs may execute out-of-order. Partial ordering can be ensured by fenced or barriered DMAs. The former executes only after all previous DMAs with the same tag on the same SPE have completed. The latter has the same guarantee, but also ensures that all subsequent DMAs issued on the same SPE with the same tag execute after it has completed. A DMA list can be used to scatter data to or gather data from multiple locations. It occupies only one slot in the DMA queue.

We observe the following regarding the performance of DMAs [4]: (i) SPE-SPE DMAs are much faster than SPE-main memory DMAs on the same chip, (ii) sending multiple small DMAs is slower than sending fewer long ones from the same SPE, (iii) latency between SPEs on different chips are significantly higher than those on the same chip, (iv) maximum bandwidth between SPE and main memory is around 7 GB/s, while between SPE and SPE it is around 25 GB/s, (v) latency is higher in the

presence of congestion, when multiple SPEs are transferring data, and (vi) the variance of the latency is higher with congestion.

MPI Based Cell Programming Models. The heterogeneous architecture and the small local stores of the SPEs make programming the Cell difficult. Some of the programming models to deal with this challenge are based on MPI. In the MPI microtask model [6], the application is divided into several smaller tasks with small code and data size. A scheduler schedules the tasks on the SPEs. In another model [3], an existing application is ported to the Cell by treating each SPE as if it were a node for an MPI process, using the main memory to store the application data, and the local store as software controlled cache. Large code size can be dealt with by bringing in the code as needed through code overlaying. *This is the model for which we target our MPI implementation, assuming that application data is in main memory, and that the MPI calls are provided the effective addresses of these locations.* If the application data is in local store, then more efficient implementations can be developed. *We also discuss only the case of contiguous data.* Use of non-contiguous data will lead to higher latencies. Note that this implementation can also be helpful with clusters of Cell processors or blades – this implementation can be used for the portion of communication that happens within a chip or blade, which is combined with MPI communication connecting different Cell processors or blades. Similar strategies have been developed for SMP clusters on other hardware [5, 7, 9].

3 Algorithms for Collectives

3.1 Common Features of Intra-Cell Collectives

Let P be the desired number of MPI processes. In our MPI implementation, a PPE process spawns P SPE threads, which perform the actual computation. We sometimes refer to an SPE thread as a process, and an SPE as a processor, for the sake of consistency with usual usage in MPI. Each SPE runs one thread at most, and so P SPEs are involved in the computation. Each SPE maintains a metadata array of P elements in its local store. (This memory usage can be decreased for some algorithms.) Each entry is 16 Bytes; smaller space would suffice, but this size is forced by DMA alignment requirements. With the maximum of 16 SPEs on a blade, this requires 256 B, which is small. The barrier call has a separate metadata array to avoid interference with other calls. The implementation also allocates two buffers of 16 KB each on the local store to use as software controlled cache. Timing results [4] indicate that buffers of 4 KB each would yield comparable performance. The implementation tries to minimize data transfers involving the main memory, because of the larger latencies involved in such transfers, compared with that to local store on-chip. The bandwidth to main memory is also the bottleneck to most algorithms, and thus access to it should be minimized. The two buffers above are used instead; the use of multiple buffers helps reduce latency by enabling double buffering – when one buffer has data being transferred out of it, another buffer is used to transfer data into the SPE.

SPE i typically transfers data to SPE j by DMAing data to metadata array location i on SPE j . SPE j polls this entry, and then DMA's data from a local store buffer on SPE i to one on SPE j . It then typically acknowledges receipt to SPE i by DMAing to metadata entry j on SPE i . Serial numbers are often used in the metadata entries to deal correctly with multiple transfers. Writing to a metadata entry is atomic, because DMA's of size up to 128 B are atomic. However, the DMA's may be executed out of order, and the data sent may also differ from the data at the time the DMA was queued, if that location was updated in the meantime. We don't discuss the implementation details to ensure correctness in the presence of these issues, in order to present a clearer high level view of the algorithms. In order to simplify some of the implementation, we made the collective calls synchronize at the end of each call, using a barrier. We later show that the barrier implementation on the Cell is very efficient.

The experimental platform was a Cell IBM QS20 revision 5.1 blade at Georgia Tech, running Linux. The *xlc* compiler for the Cell, with optimization flag `-O5`, was used. The timings were performed using the decremter register on the Cell. This has a resolution of around 70 nano-seconds. The variances of the timing results for collective calls, other than the barrier, were fairly small. The variance for the barrier, however, was somewhat higher.

3.2 Barrier

This call blocks the calling process until all the other members of the group have also called it. It can return at any process only after all the group members have entered the call.

Algorithms. We have implemented three classes of algorithms, with a few variants in one of them.

Gather/Broadcast. In this class of algorithms, one special process, which we call the root, waits to be informed that all the processes have entered the barrier. It then broadcasts this information to all processes. On receiving the information broadcast, a process can exit the barrier. We have implemented the following three algorithms based on this idea. Along with an algorithm's name, we also give an abbreviation which will be used to refer to the algorithm later.

(OTA) One-To-All. Here, an SPE informs the root about its arrival by setting a flag on a metadata entry in the root. The root waits for all its entries to have their flag set, unsets these flags, and then sets a flag on a metadata entry of each SPE. These SPEs poll for this flag to be set, then unset it and exit. Note that polling is quite fast because the metadata entry is in the local store for each SPE performing the polling; the bottlenecks are (i) DMA latency and (ii) processes arriving late. The broadcast phase of this algorithm, where the root sets flags, has two variants. In the first one, the root uses a DMA *put* to transfer data to each SPE. An SPE can have sixteen entries in its own DMA queue, and so the root can post the DMA commands without blocking. In the second variant, the root issues a single *putl DMA List* command.

(*SIG*) *Use a Signal Register.* The signal registers on each SPE support one-to-many semantics, whereby data DMAed by an SPE is *ORed* with the current value. The broadcast phase of this algorithm is as in OTA, but the gather phase differs; each SPE sets a different bit of a signal register in the root, and the root waits for all signals to be received.

(*TREE*) *Tree.* This gathers and broadcasts data using the usual tree based algorithm [10]. In the broadcast phase of a binomial tree algorithm, the root starts by setting a metadata flag on another SPE. In each subsequent phase, each process that has its flag set in turn sets the flag of one other SPE. Thus, after i phases, 2^i processes have their flags set. Therefore $\lceil \log_2 P \rceil$ phases are executed for P SPEs. In a tree of *degree* k [1, 10], in each phase SPEs, which have their flag set, set the flags of $k - 1$ other distinct SPEs, leading to $\lceil \log_k P \rceil$ phases. The gather step is similar to the broadcast phase, but has the directions reversed.

Pairwise-Exchange (PE). This is a commonly used algorithm for barriers [10]. If P is a power of 2, then we can conceptually think of the SPEs as organized as a hypercube. In each phase, an SPE exchanges messages with its neighbor along a specific dimension. The barrier is complete in $\log_2 P$ phases. If P is not a power of two, then a slight modification to this algorithm [10] takes $2 + \lceil \log_2 P \rceil$ steps.

Dissemination (DIS). This is another commonly used algorithm for barriers [10]. In the i th phase here, SPE j sets a flag on SPE $j+2^i \pmod{P}$ and waits for its flag to be set by SPE $P+j-2^i \pmod{P}$. This algorithm takes $\lceil \log_2 P \rceil$ steps, even if P is not a power of two.

Performance Evaluation. We next evaluate the performance of the above algorithms. We found that the use of DMA lists does not improve the performance of the barrier [4] – in fact, the performance is worse when P is greater than four. We also found that the use of the signal register in the gather phase does not improve performance compared with the use of plain DMAs, which are used in OTA.

Figure 2 (left) evaluates the influence of tree degree in the TREE algorithm. We optimized the implementation when the tree degree is a power of 2, replacing modulo operations with bit-wise operations. This difference is not sufficient to explain the large difference in times seen for degree 2 and 4, compared with other degrees. We believe that the compiler is able to optimize a for loop involved in the computation better with power of two degrees. However, increasing the degree to eight lowers the performance. This can be explained as follows. As the tree degree increases, the number of phases decreases. However, the number of DMA issued by the root increases. Even though it can queue up to sixteen DMAs, and the message sizes are small enough that the bandwidth is not a limiting factor, each DMA in the queue has to wait for its turn. Consequently, having multiple DMAs in the queue can lower the performance. This trend is also shown by DMA timing results not presented here.

The PE and DIS algorithms perform substantially better than the gather/broadcast type of algorithms, with PE being clearly better than DIS when P is greater than eight. Before explaining this, we first discuss a factor that sometimes influences the per-

formance of DIS. In contrast to PE, where pairs of processes exchange information, in DIS, each process sends and receives messages to different processes. On some networks, exchange between processes is faster than sending and receiving between different processes, which can cause DIS to be slower than PE. This is not the case here. DMA tests show that exchanging data is no faster than communication between different SPEs. The reason for the difference in performance is that when the number of processes is greater than eight, some of the processes are on a different chip. The DMA latency between these is higher. In PE, all the inter-chip DMAs occur in the same phase. In DIS, this occurs in each phase. Thus each phase gets slower, whereas in PE, only one of the phases is slowed down due to this fact. This slower phase also explains the sudden jump in latency from eight to ten processes.

Further details on alternate algorithms and related work are given in [4].

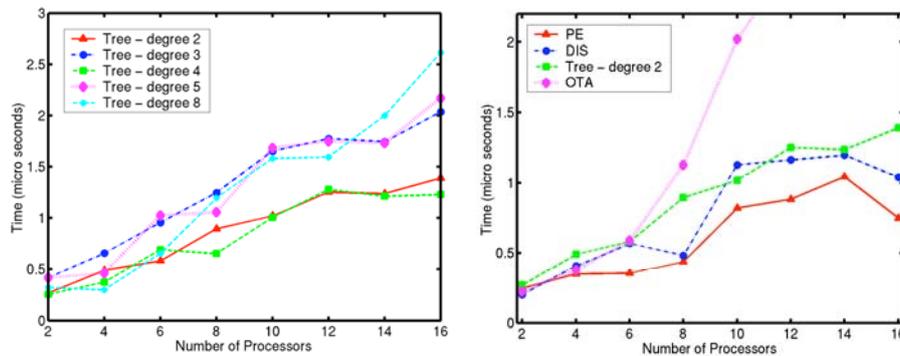


Fig. 2. Barrier latencies. *Left:* Comparison of TREE with different degrees. *Right:* Comparison of four barrier algorithms.

3.3 Broadcast

Algorithms. We discuss below five algorithms for broadcast.

(TREETM) Send/Receive. This algorithm is the usual tree based Send/Receive algorithm [2, 8], with modifications given below. The tree structure is as in the broadcast phase of TREE for the barrier. However, instead of just setting a flag, a process that sends data also passes the main memory location of its application data. A receiving process copies this data to its own main memory location. This cannot be performed directly, because DMA is possible only between a local store address and an effective address. So, an SPE first copies a chunk of data from the source location to its local store, and then copies this back from the local store to the destination location in main memory. While this seems wasteful, a similar process occurs in regular cache-based processors, where copying a line can involve two or three cache misses. We ameliorate the DMA latency by double buffering. Performance tests on memory to memory

copy shows [4] that double buffering yields a significant improvement in performance over single buffering. TREEMM's communication structure is similar to an implementation built on top of MPI_Send and MPI_Recv. However, it avoids the extra overheads of MPI calls by directly implementing the DMA calls in this routine. Furthermore, it avoids extra copy overheads, and uses double buffering to reduce the memory access latency (the latter is, in some sense, like prefetching to cache).

(OTA) Each SPU Copies its Data. In this implementation, the root broadcasts its metadata, as in the broadcast phase of barrier TREE. It sends the main memory location of the source data, in addition to setting the flag. Once an SPU receives this information, it copies data from the root's locations to its own location in main memory, using double buffering. On some systems, simultaneous access to the same memory location can degrade performance by making this a hotspot. We include a *shift S* to avoid this. That is, SPE *i* first copies with an offset of $i \times S$, and then copies the initial portion. If $i \times S$ is greater than the data size, then this index wraps around.

(G) Root Copies All Data. In this implementation, the root gathers metadata from all processes in a tree structured manner. The metadata contains the destination addresses, in addition to the flag. The root then copies its data to each of the destination addresses. This is, again, done through double buffering. Each time data is brought in to local store, it is DMAed to all destination locations. With the maximum of sixteen SPEs possible, we need at most fifteen *puts* and one *get* pending, and so the DMA requests can be placed in the queue without blocking.

(AG) Each Process Transfers a Piece of Data. In this implementation, all processes perform an *allgather* on their metadata to get the destination addresses of each process, and the source address of the root. Each process is then responsible for getting a different piece of data from the source and transferring it to the corresponding location in each destination. This is done in a double buffered manner, as with broadcast G. We also specify a *minimum size* for the piece of data any process can handle, because it may be preferable for a few processes to send large DMAs, than for many processes to send small DMAs, when the total data size is not very large. Increasing the minimum size decreases parallelism in the data transfer, with the potential benefit of fewer DMAs, for small messages.

(TREE) Local Store Based Tree. In this implementation, the root gets a piece of data from main memory to its local store, and broadcasts this piece in a tree structured manner to the local stores of all processes. Each piece can be assigned an index, and the broadcast is done by having an SPE with data sending its children (in the tree) metadata containing the index of the latest piece that is available. A child issues a DMA to actually get this data. After receiving the data, the child acknowledges to the parent that the data has been received. Once all children have acknowledged receiving a particular piece, the parent is free to reuse that local store buffer to get another piece of data. A child also DMAed received data to its main memory location, and sends metadata to its children in the tree. In this implementation too, we use double buffer-

ing, so that a process can receive a piece into one buffer, while another piece is waiting to be transferred to its children. In this implementation, we denote *pipelined* communication between the local stores by a tree of *degree 1*.

Performance Evaluation. We first determined the optimal parameter for each algorithm, such as the tree degree, shift size, or minimum piece size. We have also evaluated the effect of other implementation choices, such as use of fenced DMAs and DMA lists, but do not discuss these.

We found that on four processors, including some shift to avoid hotspots improves performance of OTA [4], though it is not very sensitive to the actual shift used. On larger numbers of processors, all shifts (including no shift) perform equally well. The likely reason for this is that with more processors, the time at which the DMA requests are executed varies more, and so we don't have a large number of requests arriving at the same time. We also found that on sixteen processors, using a minimum piece size of 2K or 4K in AG yields better performance than larger minimum sizes. (Lower sizes – 1KB and 128 B – perform worse for intermediate data sizes.) At small data sizes, there is only one SPE performing one DMA for all these sizes, and so performances are identical. For large data, all pieces are larger than the minimum, and so this size makes no difference. At intermediate sizes, the smaller number of DMAs does not compensate for the decrease in parallelism for minimum sizes larger than 4 KB. The same trend is observed with smaller numbers of processes.

We next compared the effect of different tree degrees on the performance of the tree-based algorithms [4]. In TREEMM, a tree degree of 3 yields either the best performance, or close to the best, for all process counts. A tree of degree two yields the worst performance, or close to that. However, the relative performances do not differ as much as they do in TREE. Furthermore, the differences show primarily for small messages. Note that for small messages, a higher tree degree lowers the height of the tree, but increases the number of metadata messages certain nodes send (unlike in a Send/Receive implementation, a parent sends only metadata to its children, and not the actual data). It appears that the larger number of messages affects the time more than the benefits gained in decrease of tree heights, beyond tree degree 3. A similar trend is demonstrated in TREE too, though the differences there are greater. For large messages, performances of the different algorithms are similar, though the pipelined implementation is slightly better for very large data sizes. The latter observation is not surprising, because the time taken for the pipeline to fill is then negligible related to the total time, and the number of DMAs issued by any SPE subsequently is lowest for pipelining.

Figure 3 compares the performance of the different algorithms. The trend for the different algorithms on eight processes (not shown here) is similar to that on sixteen processes. We can see that AG has the best, or close to the best, performance for large messages. TREE degree 3 is best for small messages with more than four processes. Up to four processes, broadcast G is best for small messages. Pipelining is also good at large message lengths, and a few other algorithms perform well under specific parameters. As a *good choice of algorithms*, we use broadcast AG for data of size 8192 B or more, broadcast TR degree 3 for small data on more than four processes, and broadcast G from small data on four or fewer processes. The maximum bandwidth that can be served by the main memory controller is 25.6 GB/s. We can see that

with this choice of algorithms, we reach close to the peak total bandwidth (for $P-1$ writes and one read) for all process counts of four or more, with data size 16 KB or more. The bandwidth per process can be obtained by dividing the total bandwidth by the number of processes.

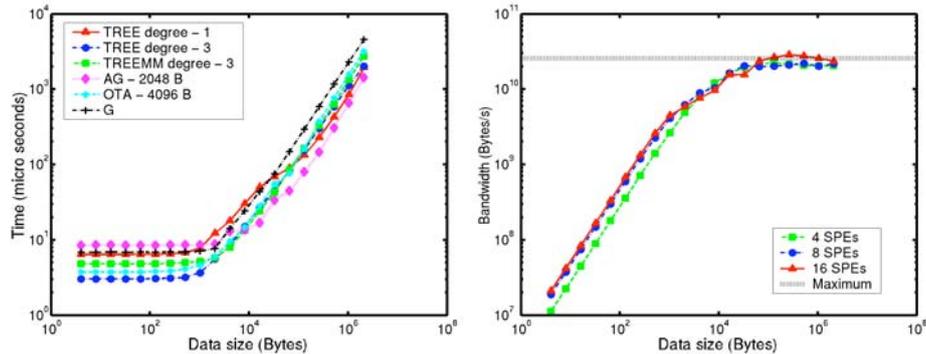


Fig. 3. Broadcast performance. *Left:* Timing on sixteen processes. *Right:* Main memory bandwidth with a "good" choice of algorithms.

3.4 Reduce

In this call, data from all the processes are combined using an associative operator, such as `MPI_SUM` for addition, and the result placed in the root. Two of the algorithms also assume that the operation is commutative, which is true for all the built-in operators.

Algorithms. The communication structure of this operation is similar to that of the broadcast, but with the directions reversed. In addition, each time a processor gets data, it also applies the operator to a current value and the new data. Since the communication structure is similar to the broadcast, we considered only the types of algorithms that worked well for the broadcast, namely, `TREE` and `AG`. In both these algorithms, the computation can also be parallelized efficiently, unlike with `OTA`.

(TREE) Local Store Based Tree. In this implementation, the communication direction of the broadcast `TREE` is reversed. A process gets a piece of data from its memory location to local store, gets data from a child's local store to its own local store when that data is available, and combines the two using the specified operator. It repeats this process for each child, except that it does not need to get its own data from main memory for subsequent children. Once it has dealt with all the children, it informs the parent about the availability of the data by DMAing a metadata entry, as in the broadcast. It repeats this for each piece of data in main memory. Double buffering is used to reduce the latency overhead by bringing data from main memory or the next child into a new buffer. Unlike with the broadcast, we need four buffers, two for each oper-

and of a reduce operation, and two more because of double buffering. Due to space constraints on the local store, we used the same buffers as in the broadcast, but conceptually treated them as having half the size (four buffers of 8KB each instead of two buffer of 16K each with broadcast).

(AG) Each Process Handles a Piece of Data. In this implementation, each process is responsible for reducing a different piece of the data, and then writing this to the destination location of the root. An initial all gather is used to get addresses of all SPEs, as in broadcast.

(TREEMM) Send/Receive. We implemented the usual tree based reduction algorithm on top of our implementation of MPI_Send and MPI_Recv [3] for the Cell processor, just for comparison purposes. The MPI_Send and MPI_Recv operations themselves make effective use of the features of the Cell. However, the extra copying to memory makes the performance worse.

Performance Evaluation. We give performance results in figure 4, for different numbers of processes. We can see that, for small data sizes, TREE degree 3 is either the best, or close to it, on greater than four processes. This is consistent with the behavior expected from the broadcast timings. On four processes, a tree of degree four, which has height 1, performs best. But, degrees 2, 3, and 4 are very close to each other in most cases. Reduce AG is worse for small messages, because of the overhead of all-gathering the metadata initially. TREE degree 1 is best for very large data, except on four processes. Reduce AG is the best at intermediate data sizes, except on four processes, where it is the best even for large messages. This can be explained as follows. Reduce AG parallelizes the computations perfectly, but issues $P+1$ DMAs for each piece of data (P gets and one put). As mentioned earlier, sending a large number of DMAs is less efficient. On four processes, this number is not very large, and so AG still performs well. The better performance of AG at intermediate data sizes can be explained as follows. The pipelined algorithm (TREE degree 1) requires communication of the order of the number of processes before results start coming out. If the data size is not very large, then this time plays a relatively large role. For large data, the initial time taken does not play as important a role, and so the pipelined algorithm is better. The reduction takes more time than the broadcast because of the extra overhead of the reduction operation. We have also evaluated the performance of the Send/Receive based reduce TREEMM. Its performance varies between $5 \mu s$ for $128 B$ to around $2000 \mu s$ for $1 MB$ on eight processors, and is worse than the best algorithm for each data size and processor count considered.

4 Conclusions and Future Work

We have described in detail implementation of three MPI collective operations: *barrier*, *bcast*, and *reduce*. We have also implemented the following: *gather*, *allreduce*, *scan*, *allgather*, *alltoall*, and vector versions of these calls. Our results show good per-

formance, both within a chip and in blade consisting of two Cell processors. We have implemented a variety of algorithms. While we use the availability of the common shared memory and high bandwidths available, the main benefit are obtained through effective use of the local store, and hiding the cost of access to main memory through double buffering.

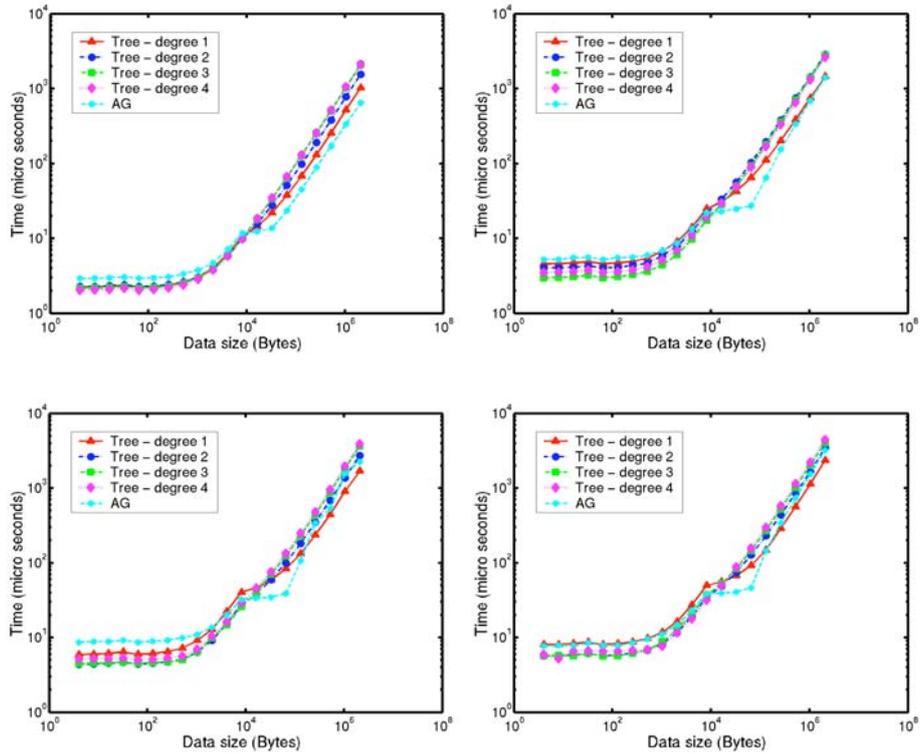


Fig. 4. Reduce timings for MPI_SUM on MPI_INT (four Bytes per int). *Top left:* Four processes. *Top right:* Eight processes. *Bottom left:* Twelve processes. *Bottom right:* Sixteen processes.

Some of the future work is as follows. (i) Since the metadata size is sixteen bytes, we can consider using some of the extra space in it for small messages, thereby reducing the number of communication operations for small messages. (ii) We have performed barrier synchronization at the end of each collective call, in order to prevent successive calls from interfering with each other. This can be avoided by using counters, as some implementations have done, and may be beneficial when applications reach the collective call at much different times. (iii) It will be useful to consider the integration of this intra-Cell implementation with implementations that connect Cell blades using networks, such as Infiniband. (iv) If the application data is in local store,

then our implementation can be made faster. This can be useful, for example, in the MPI microtask model for programming the Cell.

Acknowledgements. We thank the Sony-Toshiba-IBM Cell Center of Competence at Georgia Tech for providing use of their IBM QS20 Cell blades, and IBM for providing access to their Cell blades under the VLP program. Most of all, we express our gratitude to Sri Sathya Sai Baba for bringing us all together to perform this work, and for inspiring and helping us toward our goal.

References

1. R. Gupta, P. Balaji, D.K. Panda, and J. Nieplocha, Efficient Collective Operations Using Remote Memory Operations on VIA-Based Clusters, Proceedings of IPDPS, (2003)
2. S.P. Kini, J. Liu, J. Wu, P. Wyckoff, and D.K. Panda, Fast and Scalable Barrier Using RDMA and Multicast Mechanisms for Infiniband-Based Clusters, Proceedings of Euro PVM/MPI Conference, (2003)
3. M. Krishna, A. Kumar, N. Jayam, G. Senthilkumar, P.K. Baruah, S. Kapoor, R. Sharma, and A. Srinivasan, A Buffered Mode MPI Implementation for the Cell BE Processor, Proceedings of the International Conference on Computational Science (ICCS), Lecture Notes in Computer Science 4487 (2007) 603–610
4. M. Krishna, A. Kumar, N. Jayam, G. Senthilkumar, P.K. Baruah, S. Kapoor, R. Sharma, and A. Srinivasan, Optimization of Collective Communication in Intra-Cell MPI, Technical Report TR-070724 (www.cs.fsu.edu/research/reports/TR-070724.pdf), Dept. of Computer Science, Florida State University, (2007)
5. A.R. Mamidala, L. Chai, H-W. Jin, and D.K. Panda, Efficient SMP-Aware MPI-Level Broadcast over Infiniband's Hardware Multicast, Communication Architecture for Clusters Workshop, in Proceedings of IPDPS, (2006)
6. M. Ohara, H. Inoue, Y. Sohda, H. Komatsu, and T. Nakatani, MPI Microtask for Programming the Cell Broadband Engine™ Processor, IBM Systems Journal, 45 (2006) 85–102
7. S. Sistare, R. vande Vaart, and E. Loh, Optimization of MPI Collectives on Clusters of Large-Scale SMP's, Proceedings of SC'99, (1999)
8. R. Thakur, R. Rabenseifner, and W. Gropp, Optimization of Collective Communication Operations in MPICH, International Journal of High Performance Computing Applications, 19 (2005) 49-66
9. V. Tipparaju, J. Nieplocha, and D.K. Panda, Fast Collective Operations Using Shared and Remote Memory Access Protocols on Clusters, Proceedings of IPDPS, (2003)
10. W. Yu, D. Buntinas, R.L. Graham, and D.K. Panda, Efficient and Scalable Barrier over Quadrics and Myrinet with a New NIC-Based Collective Message Passing Protocol, Workshop on Communication Architecture for Clusters, in Proceedings of IPDPS, (2004)