

Accurately Measuring Collective Operations at Massive Scale

Torsten Hoefler,^{1,2} Timo Schneider² and Andrew Lumsdaine¹

¹Open Systems Laboratory
Indiana University
501 N. Morton Street
Bloomington, IN 47404 USA
{htor,lums}@cs.indiana.edu

²Dept. of Computer Science
Technical University of Chemnitz
Strasse der Nationen 62
Chemnitz, 09107 GERMANY
{htor,timos}@cs.tu-chemnitz.de

Abstract

Accurate, reproducible and comparable measurement of collective operations is a complicated task. Although different measurement schemes are implemented in well-known benchmarks, many of these schemes introduce different systematic errors in their measurements. We characterize these errors and select a window-based approach as the most accurate method. However, this approach complicates measurements significantly and introduces a clock synchronization as a new source of systematic errors. We analyze approaches to avoid or correct those errors and develop a scalable synchronization scheme to conduct benchmarks on massively parallel systems. Our results are compared to the window-based scheme implemented in the SKaMPI benchmarks and show a reduction of the synchronization overhead by a factor of 16 on 128 processes.

Keywords: benchmarking, collective operations, MPI, time synchronization, scalable synchronization

1 Introduction

Collective operations, i.e., operations that are defined on a group of processes rather than a single process, are an important part of parallel scientific computing. Their advantages comprise the separation of communication and computation which enables machine specific communication optimization, performance portability among different parallel machines, better programmability and the reduction of implementation errors. All those benefits have been recognized by the scientific community and collective operations play an important role in many parallel applications (cf. [21]). Thus, the performance, i.e., latency, of the collective operations is crucial for the performance and running time of numerous applications.

However, the complex interaction between communication and computation in real-world applications requires models to understand and optimize parallel codes. Serial computation models, such as Modeling Assertions [2], have to be combined with network models such as the LogP [4]. Unfortunately, today's systems are too complex to be described entirely by an execution or communication model. It is necessary to assess the model parameters for every real-world system, in the serial execution case in [2] as well as in the communication case in [5, 8, 12]. Pjesivac-Grbovic et al. showed in [20] that the latency of collective operations that are implemented on top of point-to-point messages can be modeled with the LogP model family and we showed in an earlier work [7] that it is also possible to predict application performance by modeling the communication and computation separately. Thus, it is crucial to the modeling of parallel applications to have accurate models for the latency of collective operations. Several benchmarking studies have been done for different systems [3, 19, 22].

In this paper, we discuss and analyze different established measurement methods for collective operations on parallel computing systems and point out common systematic errors. Our measurement method is derived from the SKaMPI benchmarks [25] and is universal and not limited to point-to-point based methods or specific algorithms. The following section discusses established benchmark methods.

1.1 Related Work

Different benchmark schemes have been proposed. Currently known methods can be divided into three groups. The first group synchronizes the processes explicitly with the use of synchronization routines (i.e., `MPI_Barrier`). The second scheme, presented in [25], establishes the notion of a global time and the processes start the operation synchronously. The third scheme assesses the quality of a col-

```

MPI_Gather(...); /* warmup */
MPI_Barrier(...); /* synchronization */
t0 = MPI_Wtime(); /* take time */
for (i=0; i<reps; i++) {
    MPI_Gather(...); /* execute benchmark */
}
t1 = MPI_Wtime(); /* take time */
MPI_Barrier(...);
time = t1 - t0;

```

Listing 1. MPPTTEST Benchmark Scheme

```

t0 = TIMER(); /* take time */
for (i=0; i<reps; i++) {
    MPI_Alltoall(...); /* execute benchmark */
}
t1 = TIMER(); /* take time */
MPI_Barrier(...);
time = t1 - t0;

```

Listing 2. MPBench Benchmark Scheme

lective implementation by comparison to point-to-point operations [23] and is thus limited to algorithms using point-to-point messages. We investigate several publicly available benchmarks in the following and characterize them in the three groups.

MPPTTEST implements the discussions on reproducible MPI performance measurements [6]. As described in the article, the operation to measure is executed in a warm-up round before the actual benchmark is run. The nodes are synchronized with a single `MPI_Barrier` operation before the operation is run N times in a loop. A pseudo-code is shown in Listing 1. Only the time measurements at rank 0 are reported to the user.

MPBench was developed by Mucci et al. [16]. `MPBench` does not synchronize at all before the benchmarks. Rank 0 takes the start time, runs N times the collective operation to benchmark and takes the end time. A pseudo-code is shown in Listing 2. The timer can use the `RDTSC` CPU instruction of `gettimeofday()`. Time measurement is only performed and printed on rank 0.

Intel MPI Benchmarks (formerly Pallas MPI benchmarks [18]), measure a wide variety of MPI calls including many collective functions. The code issues a definable number of `MPI_Barrier` operations before every benchmark and measures the collective operation in a loop afterwards. The time needed to execute the loop is taken as a measurement point. The scheme is shown in Listing 3. The benchmark prints minimum, maximum and average time over all processes.

SKaMPI The `SKaMPI` benchmark uses a time-window based approach, described in [25], that ensures that all processes start the operation at the same time. No explicit syn-

```

for(i=0; i<numbarr; i++) MPI_Barrier(...);
t0 = MPI_Wtime(); /* take time */
for (i=0; i<reps; i++) {
    MPI_Alltoall(...); /* execute benchmark */
}
t1 = MPI_Wtime(); /* take time */
time = (t1 - t0)/reps;

```

Listing 3. Intel MPI Benchmark Scheme

chronization is used and the times are either reported per process or cumulative.

1.2 Systematic Errors in Common Measurement Methods

Benchmarking collective operations is a controversial field. It is impossible to find a single correct scheme to measure collective operations because the variety of real-world applications is tremendously high. Thus, every benchmark may have its justification and is not erroneous in this case. However, microbenchmarks are often used to compare implementations and to model the influence of the communication several applications. This is why a benchmark should represent the “average” or at least the majority of applications. Our model application for this work is a well balanced application that issues at least two different collective operations in a computational loop (cf. [21]). This model application would benefit from well balanced collective operations that do not introduce process skew. The following paragraphs describe common systematic errors done in the measurement of collective operations. This section is concluded with the selection of a benchmark method.

Implementation Assumptions The implementation of collective communication operations is usually not standardized to provide as much optimization space as possible to the implementer. Thus, any point-to-point algorithm or hardware supported operation that offers the functionality defined in the interface is a valid implementation. Some research groups used elaborate techniques (e.g., hardware optimization and/or specialized algorithms) to optimize collective communication on different systems (cf. [11, 24, 26]). A portable benchmark that uses the collective interface, for example MPI, can not make any assumptions about the internal implementation.

Results on multiple processes A second problem is that benchmarks are usually providing a single number to the user, while all processes benchmark their own execution time. Some benchmarks just return the time measured on a single node (e.g., the rank 0), some use the average of all times and some the maximum time of all processes. The decision which time to use for the evaluation is not trivial it is even desirable to include the times of all ranks in the evaluation of the implementation. Worsch et al. define three schemes to reduce the times to a single number, (1) the time

needed at a designated process, (2) maximum time on all processes and (3) the time between the start of the first process and the finish of the last. This list can be extended further, e.g., (4) the average time of all processes or (5) the minimum time might play a role and is returned by certain benchmarks.

A simplified LogP model derived from [4] is used to model the network transmissions and effects in collective communication. The LogP model uses four parameters to describe a parallel system. The parameter L models the network latency and g is the time that has to be waited between two packets. The CPU overhead o does not influence the network transmission and is thus omitted in our simplified model. The number of participating processes P is constantly four in our examples.

Pipelined Measurements A second source for systematic errors are pipelining effects that occur when many operations are executed in a row. A common scheme is to execute N operations in a loop, measure the time and divide this time by N . This scheme was introduced to avoid the relative high inaccuracy of timers when short intervals are measured. We show in Section 2 that this is not necessary for high precision timers. An example LogP modeling for MPI_Bcast with root 0, implemented with a linear scheme, is shown in Figure 1. A single execution is much

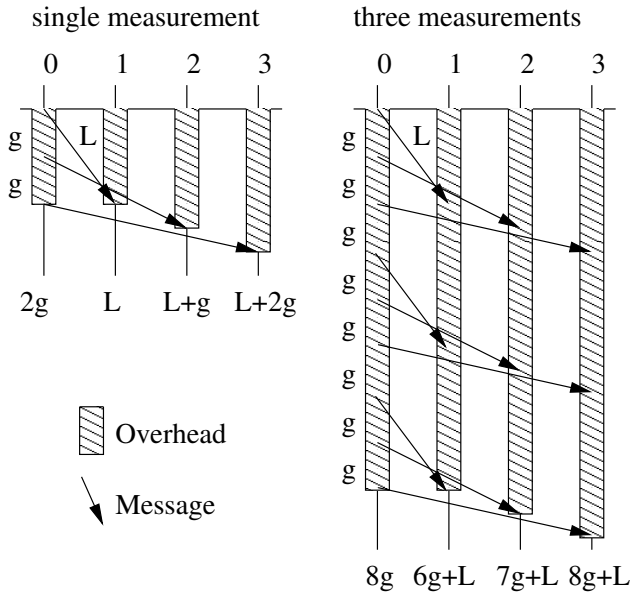


Figure 1. Pipeline Effects for a linear Broadcast Implementation

more likely to model the behavior of real applications (multiple successive collective operations should be merged into a single one with more data). Both schemes result in differ-

ent execution times, e.g., the worst-case (maximum among all nodes) returned latency for a single execution is $L + 2g$ for a single operation and $(8g + L)/3$ for three successive operations. The pipelined measurement tends to underestimate the latency in this example.

Process Skew The LogP models in the previous paragraph assumed that the first operation started at exactly the same (global) time. This is hardly possible in real parallel systems. The processes arrive at the benchmark in random order and at undefined times. Process skew is influenced by operating system noise [1, 14] or other collective operations (cf. Figure 1 where rank 0 leaves the operation after $2g$ and rank 3 finishes only after $L + 2g$). A LogP example for the influence of process skew is shown in Figure 2. The left

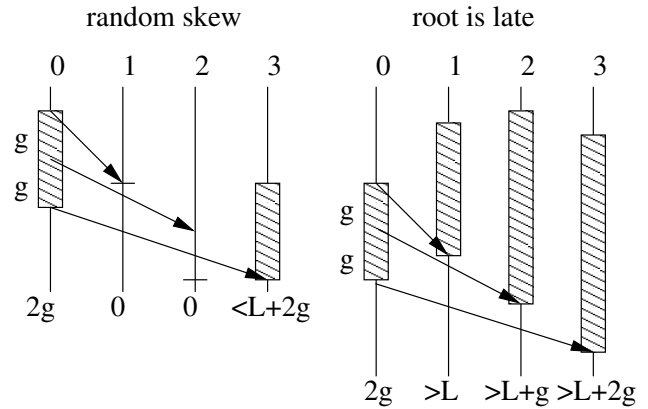


Figure 2. Influence of Process Skew on a Broadcast Benchmark

side shows a random skew pattern where rank 1 and 2 arrive relatively late and do not have to wait for their message while the right side shows a situation where the root (rank 0 in this example) arrives late and all ranks have to wait much longer than usual. Process skew can not be avoided and is usually introduced during the runtime of the parallel program. This effect is well known and several benchmarks use an MPI_Barrier before the measurement to correct skew.

Synchronization Perturbation and Congestion The MPI_Barrier operation has two problems, the first one is that this operation may be implemented with any algorithm because it only guarantees that all processes arrived before the first leaves the call but not that the processes leave at the same time (i.e., the barrier operation may introduce process skew). The second one is that it may use the same network as other collective operations which may influence the messages of the investigated collective operation. An example with a linear barrier is shown in Figure 3.

Network Congestion Network congestion can occur if multiple operations are started successively or synchroniza-

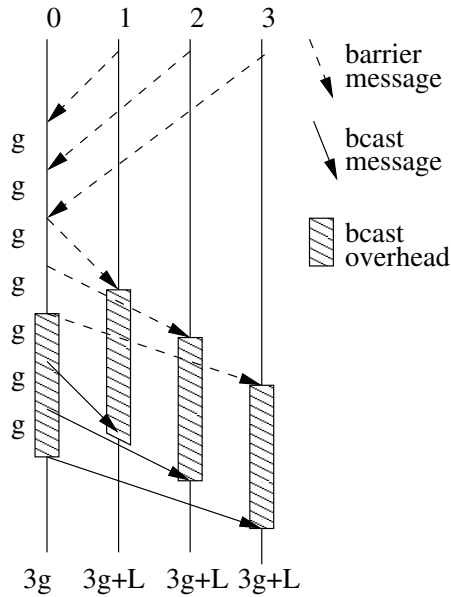


Figure 3. Possible Effect of Barrier Synchronization on Measurements

tion messages interfere with the measurement. This influences the measured latencies.

1.3 Selecting a Benchmark Scheme

The SKaMPI benchmarks avoid most of the systematic errors with the window-based mechanism [25]. This mechanism relies heavily on the assumption, that the difference between the local clocks does not change (drift) or changes in a predictable way (can be corrected). We analyze the clock drift in the following sections. Another problem might be the variation in the latency of point-to-point messages. This variation is also analyzed and a new fast point-to-point synchronization scheme is presented. Furthermore, we propose and implement a new scalable group synchronization algorithm which scales logarithmically instead of the linear SKaMPI approach.

2 Measurements in Parallel Systems

Several restrictions apply to measurements in parallel systems. One of the biggest problems is the missing time synchronization. Especially in cluster systems, where every node is a complete and independent system with its own local time source, one has to assume that potentially all processes of a parallel job run at slightly different clock speeds. However, to perform the necessary measurements, we need to synchronize all clocks or have at least the time offsets of all processes to a global time. It can also not be assumed that the processes start in any synchronous state. To ensure portability, MPI mechanisms have to be used to synchronize. However, collective operations semantics do not guarantee any timing, thus, we need to synchronize the

processes with point-to-point operations. Those operations do not guarantee timing either but are less complex than collective operations (no communication patterns). We analyze local time sources and their accuracy in the following. This is followed by an analysis of the clock skew in parallel systems and the distribution of latencies. This analyses are used to derive a novel and precise synchronization scheme in the next section.

2.1 Local Time Measurement

All time sources in computing systems work in a similar way: They use a crystal that oscillates with a fixed frequency and a register that counts the number of oscillations since a certain point in time (for example the system startup). However, the way to access this information can be different. Some timing devices can be configured to issue an interrupt when the register reaches a certain value or overflows and others just enable the programmer to read the register.

A very important timer in a modern PC is the Real Time Clock (RTC) which is powered by a battery or capacitor so that it continues to tick even when the PC is turned off. It is used to get the initial time of the day at system startup, but since it is often very inaccurate (it is optimized for low power consumption, not for accuracy), it should not be used to measure short time differences, like they might occur in benchmarking scenarios.

Another time-source is the Programmable Interval Timer (PIT). It can be configured to issue interrupts at a certain frequency (Linux 2.6 uses 1000.15 Hz). These interrupts are used to update the system time and perform several operating system functions. When using the system time (for example via `gettimeofday()`) one must be aware that the returned value might be influenced by `ntp` or other time-synchronization mechanisms and that there is a whole software stack behind this simple system call which might add additional perturbation (i.e., interrupts, scheduling, etc.).

The resolution of the discussed time sources is not very high and not accurate enough for the benchmarking of fast events like single message transmissions. Thus, many instruction set architectures (ISA) offer calls to read the CPU's clock register which is usually incremented at every tick. For example the x86 and x86-64 ISAs offer the atomic instruction `RDTSC` call [13] to read a 64 bit CPU tick register. In fact most modern ISAs support similar features. The resolution of those timers is usually very high (e.g., $0.5ns$ on a 2GHz system). It has to be noted that this mechanism introduces several problems on modern CPUs. The first issue is caused by techniques that dynamically change the CPU frequency (e.g., to save energy) such as "Intel-SpeedStep" or "AMD-PowerNow". Changing the CPU clock results in invalid time measurements. Thus, we recommend to disable those mechanisms in cluster systems. A second problem, called "process hopping", may occur on

multi-processor systems where the process is re-scheduled between multiple CPUs. The counters on the CPUs are not necessarily identical. This might also influence the measurement. This problem is also minor because most modern operating systems (e.g., Linux 2.6) offer interfaces to bind a process to a specific CPU (e.g., CPU affinity), and in fact, most operating systems avoid “process hopping” by default.

2.2 Clock Skew and Network Latencies

As you can imagine the crystals used for hardware timers are not ideal with respect to their frequency. They may be a little bit slower or faster than their nominal rate. This drift is also temperature dependent and has been analyzed in [15] and [17]. It was shown that this effect is significant enough to distinguish / identify single computers and sometimes even the timezone which they are located in. Due to other effects, such as the NTP daemon that synchronizes every 11 minutes (local time!) by default, the clock difference between two nodes may behave very unpredictable, which will lead to erroneous results. Therefore the usage of a software independent clock like the TSC is generally a viable alternative. Of course those effects could have a negative influence on collective benchmarks which rely on time synchronization, especially if the synchronization is done only once before a long series of benchmarks. In our experiments it turns out that two clocks (even on identical hardware) always run at slightly different speeds. Therefore we analyzed the clock skew between various nodes in a cluster system over a longer period of time.

Similar to the clocks, that do not run totally synchronously, we do also expect a variance in the network transmission parameters for different messages. The most important parameter for benchmarks and synchronization is the network latency (or round-trip-time (RTT) in ping-pong benchmarks). Thus, we have to analyze the variance of RTT for different networks.

We used a simple ping pong scheme to determine the RTT: Rank 1 sends its local time t_1 as an eight byte message with a blocking send to rank 2. As soon as rank 2 has completed the corresponding recv, it sends its local time t_2 back to rank 1. After rank 1 is finished receiving that timestamp it checks his local time t_3 . To use a portable high-precision timing interface and to support many network interconnects, the benchmark scheme was implemented in the Netgauge performance measurement framework [10]. A pseudo-code is shown in Listing 4. The difference between the first and the second timestamp obtained by node 1 is the *roundtrip time* ($t_{rtt} \leftarrow t_3 - t_1$). On our x86 systems, we used `RDTSC` in the `take_time()` macro because this gives us a very high resolution and accuracy. However, we double-checked our findings with `MPI_Wtime` (Which uses `gettimeofday()` or similar functions) to avoid common pitfalls described in Section 2.1.

This measurement was repeated 50,000 times, once ev-

```

if (rank == 0) {
    t1 = take_time();
    module→ send(1, &t1, 8);
    module→ recv(1, &t2, 8);
    t3 = take_time();
} else {
    module→ recv(0, &t1, 8);
    t2 = take_time();
    module→ send(0, &t2, 8);
}

```

Listing 4. The RTT Benchmark

ery second over a period of 14 hours. We gather data to get information about the RTT distribution and also the clock skew between two nodes. Thus, we define the difference between t_1 and t_2 as *clock-difference* ($t_{diff} \leftarrow |t_1 - t_2|$) and collect statistical data for the clock differences too.

The benchmark results, a histogram of 50,000 RTT measurements in 200 uniform bins, for the latency (RTT/2) distributions of InfiniBand, Myrinet and Gigabit Ethernet are shown in Figure 4.




Figure 4. Distribution of Roundtrip Times

The clock skew results of 6 different node pairs on our test systems are shown in Figure 5. We see that the clock difference behaves relatively linear and can thus be corrected by a linear error function. That way synchronization error below 1 ppm (parts per million) can be achieved. Without linear error correction the error can be up to 300 ppm. The detailed methodology that was used to determine the clock difference is described in the next section.

3 Time Synchronization

We describe a new time synchronization schemes that base on our analysis of clock skew and latency variation. We begin by defining a scheme to synchronize two processes and derive a scalable scheme to synchronize large process groups. We used this schemes to synchronize the processes in our collective benchmark NBCBench [9] that




Figure 5. Clock Drift for Different Pairs

uses a window-based benchmark scheme.

3.1 Synchronizing two Processes

A clock synchronization between two peers is often accomplished with a ping pong scheme similar to the one described above: Two nodes calculate their clock difference so that the client node knows his clock offset relative to the server node. This offset can be subtracted from the clients local time when clock synchronization is required. However this procedure has certain pitfalls one has to be aware of.

Many implementations of the scheme described above, for example the one found in the SKaMPI code, use `MPI_Wtime` to acquire timestamps. This is of course the most portable solution and works for homogeneous nodes as well heterogeneous ones. But you can not be sure which timing source is used by `MPI_Wtime`, for example the usage of `gettimeofday()` is, due to its portability, quite likely. But the clock tick rate of this clock can vary (cf. Section 2.1). We showed in Section 2 that a linear correction can be used for nodes running at different clock speeds if software errors (e.g., NTP) are avoided.

We also showed in Section 2 that measured network latencies are varying with an unpredictable distribution. The effect of this pseudo-random variation of the latency to the clock synchronization has to be minimized. Using the average or median of the clock difference, like many codes do today, is not a viable option because, as you can see in the histograms, they vary a lot. A better approach is the measurement of roundtrip time and clock difference at the same time and only use the clock differences obtained in measurements which showed a latency below a certain threshold.

The distribution of roundtrip times can not be known in advance. That implies that the threshold can not be selected easily. We chose a different approach to ensure accurate measurements. For the measurements in Section 2, we used

only the 25% of the results that had the smallest roundtrip time. However, this requires a fixed number of measurements to be conducted every time which makes this scheme unusable for online measurements. For this purpose, we developed another approach. We conduct as many measurements as we needed so that the minimal observed roundtrip time does not become smaller for N consecutive measurements.

While this scheme is guaranteed to converge, it is not possible to predict how many measurements have to be conducted for a certain N . It is also not possible to select an N mathematically that ensures a certain quality of the observed “minimal” roundtrip time. Thus, we performed a simulation to find suitable values for N for different inter-connection networks.

The simulation takes a random roundtrip time from our list¹ and checks if it was bigger than the smallest one observed in this run. If this condition is met N consecutive times the run is completed and we compute the difference between the smallest RTT in our whole dataset and the one observed in the current run. The averaged values for several networks are graphed in Figure 6. It shows that the quality of the measurement and the measurement costs (i.e., the number of measurements) grow relatively independent of the network with N (the costs for the different networks look nearly identical on the double-logarithmic plot, thus we plotted it only once). However, the quality of the results grows pretty fast for small N and seems to “saturate” around 5%. To support every network, we chose $N = 100$, which has an error of less than 10% for our tested networks and converges after approximately 180 measurements.




Figure 6. Influence of N on Quality

3.2 Scalable Group Time Synchronization

Collective time synchronization means that all nodes know the time difference to a single node so that every node

¹we used the 50,000 RTTs gathered as described in Section 2

can compute a global time locally. Rank 0 is conveniently chosen as global time source, i.e., after the time synchronization, every rank knows the difference between its own clock and rank 0's clock. This enables globally synchronous events initiated by rank 0. For example, rank 0 can broadcast a message that some function is to be executed at its local time x . Every rank can now calculate its local time when this operation has to be executed.

A common scheme to synchronize all ranks is start the point-to-point synchronization procedure between rank 0 and every other rank. This disadvantage of this scheme is that it takes $P - 1$ synchronization steps to synchronize P processes. We propose a new and scalable time synchronization algorithm for our scalable benchmark. Our algorithm uses $\lceil \log_2 P \rceil$ communication steps to synchronize P processes.

The algorithm divides all ranks in two groups. The first group consists of the maximum power of two ranks, $t = 2^k \forall k \in N, t < P$ beginning from rank 0 to rank $t - 1$. The second group includes the remaining ranks t to $P - 1$.

The algorithm works in two steps, the first group synchronizes with a tree-based scheme in $\log_2 t$ synchronization rounds. The point-to-point scheme, described in Section 3.1 is used to synchronize client and server. Every rank r in round r acts as a client if $r \bmod 2^r = 0$ and as a server if $r \bmod 2^r = 2^{(r-1)}$. All clients r use rank $r + 2^{(r-1)}$ as server and all servers rank $r - 2^{(r-1)}$ as client. All client-server groups do the point-to-point synchronization scheme in parallel. Every server gathers some time difference data in every round. This gathered data has to be communicated at the end of every round. To be precise, a server communicates $2^{(r-1)} - 1$ time differences to its client at the end of every round. The clients receive the data and update their local differences to pass them on in the next step. After $\log_2 t$ rounds, rank 0 knows the time differences to all processes in group 1.

In the second step, all processes in group 2 choose peer $r - t$ in group 0 to synchronize with. All nodes in group 2 synchronize in a single step with their peers and send the result to rank 0 which in turn calculates all the time offsets for all nodes and scatters them accordingly. After this step, all nodes are time synchronized, i.e., know their time difference to the global clock of rank 0. The whole algorithm for an example with $P = 7$ is shown in Figure 7.

Figure 8 shows the difference in synchronization time between a linear scheme and the proposed tree-like algorithm. The benchmark, which measures the synchronization time at rank 0, was run on the odin cluster at Indiana University, the cluster consists of 128 Dual Opteron dual-core nodes. To simulate a real application, we used all available cores of the machine. The synchronization time is greatly reduced (up to a factor of more than 16 for 128 processes) with the new scheme. The gain of the new method is even

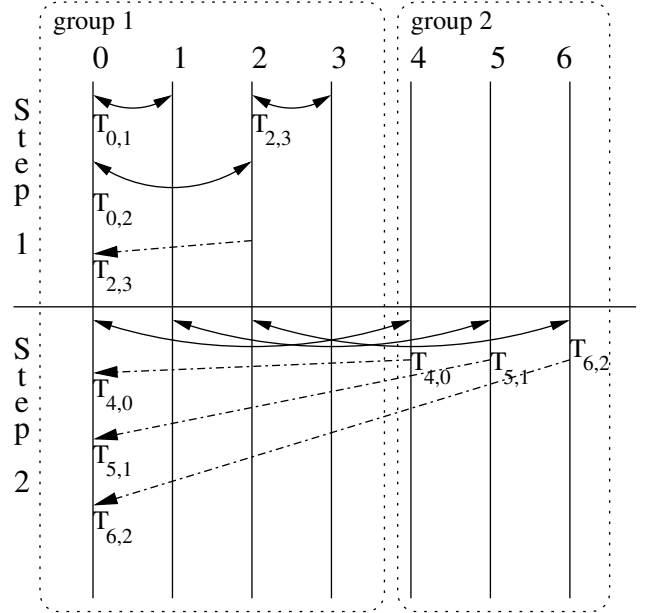


Figure 7. Synchronization Method

higher for Gigabit Ethernet but graphs have been omitted due to space restrictions.

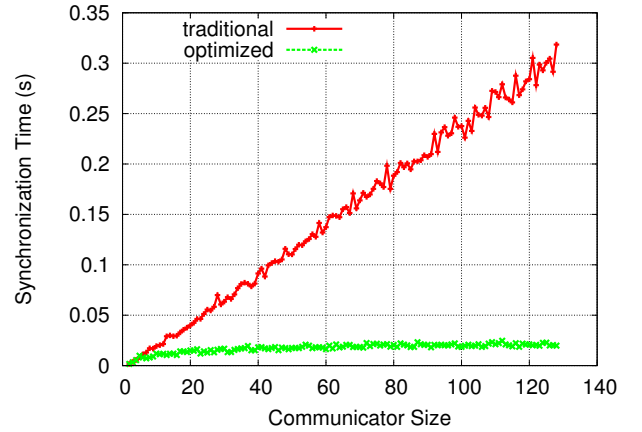


Figure 8. Synchronization Time over IB

4 Conclusions

In this work, we analyzed different schemes to benchmark the latency of collective operations. We defined certain systematic errors in common methods and proposed a new and scalable scheme that bases on the window mechanism used in SKaMPI. Therefore, we propose and analyze a new scalable group synchronization method for collective benchmarks. Our method is more than 16 times faster on 128 processes and promises to be more accurate than currently used schemes. We implemented the network analysis benchmark in the open source performance analysis tool Netgauge. The described scheme to benchmark collective

operations has been implemented in the open source tool NBCBench.

Future work on this area includes the application of this new scheme to measure latency and communication/computation overlap of non-blocking collective operations.

Acknowledgements

Pro Siobhan. This research was funded by a gift from the Silicon Valley Community Foundation, on behalf of the Cisco Collaborative Research Initiative of Cisco Systems.

References

- [1] S. Agarwal, R. Garg, and N. Vishnoi. The impact of noise on the scaling of collectives: A theoretical approach. In *12th Annual IEEE International Conference on High Performance Computing*, Goa, India, Dec. 2005.
- [2] S. R. Alam, N. Bhatia, and J. S. Vetter. An exploration of performance attributes for symbolic modeling of emerging processing devices. In R. H. Perrott, B. M. Chapman, J. Subhlok, R. F. de Mello, and L. T. Yang, editors, *HPCC*, volume 4782 of *Lecture Notes in Computer Science*, pages 683–694. Springer, 2007.
- [3] T. Bönisch, M. M. Resch, and H. Berger. Benchmark evaluation of the message-passing overhead on modern parallel architectures. In *Parallel Computing: Fundamentals, Applications and New Directions, Proceedings of the Conference ParCo97*, pages 411–418, 1997.
- [4] D. Culler, R. Karp, D. Patterson, A. Sahay, K. E. Schauer, E. Santos, R. Subramonian, and T. von Eicken. LogP: towards a realistic model of parallel computation. In *Principles Practice of Parallel Programming*, pages 1–12, 1993.
- [5] D. Culler, L. T. Liu, R. P. Martin, and C. Yoshikawa. LogP Performance Assessment of Fast Network Interfaces. *IEEE Micro*, February 1996.
- [6] W. Gropp and E. L. Lusk. Reproducible measurements of mpi performance characteristics. In *Proceedings of the 6th European PVM/MPI Users' Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 11–18, London, UK, 1999. Springer-Verlag.
- [7] T. Hoefler, R. Janisch, and W. Rehm. Parallel scaling of Teter's minimization for Ab Initio calculations. 11 2006. Presented at the workshop HPC Nano in conjunction with SC'06.
- [8] T. Hoefler, A. Lichei, and W. Rehm. Low-Overhead LogGP Parameter Assessment for Modern Interconnection Networks. 03 2007.
- [9] T. Hoefler, A. Lumsdaine, and W. Rehm. Implementation and performance analysis of non-blocking collective operations for mpi. In *Accepted for publication in the proceedings of Supercomputing '07*, 2007.
- [10] T. Hoefler, T. Mehlan, A. Lumsdaine, and W. Rehm. Netgauge: A Network Performance Measurement Framework. In *High Performance Computing and Communications, Third International Conference, HPCC 2007, Houston, USA, September 26-28, 2007, Proceedings*, volume 4782, pages 659–671. Springer, 9 2007.
- [11] T. Hoefler, T. Mehlan, F. Mietke, and W. Rehm. Fast Barrier Synchronization for InfiniBand. In *Proceedings, 20th International Parallel and Distributed Processing Symposium IPDPS 2006 (CAC 06)*, April 2006.
- [12] T. Hoefler, T. Mehlan, F. Mietke, and W. Rehm. LogfP - A Model for small Messages in InfiniBand. In *Proceedings, 20th International Parallel and Distributed Processing Symposium IPDPS 2006 (PMEO-PDS 06)*, April 2006.
- [13] Intel Corporation. Intel Application Notes - Using the RDTSC Instruction for Performance Monitoring. Technical report, Intel, 1997.
- [14] K. Iskra, P. Beckman, K. Yoshii, and S. Coghlan. The influence of operating systems on the performance of collective operations at extreme scale. In *Proceedings of Cluster Computing, 2006 IEEE International Conference*, 2006.
- [15] T. Kohno, A. Broido, and K. Claffy. Remote physical device fingerprinting. *Dependable and Secure Computing, IEEE Transactions on*, 2(2):93–108, 2005.
- [16] P. J. Mucci, K. London, and J. Thurman. The MPIBench Report. Technical report, CEWES/ERDC MSRC/PET, 1998.
- [17] S. Murdoch. Hot or not: revealing hidden services by their clock skew. *Proceedings of the 13th ACM conference on Computer and communications security*, pages 27–36, 2006.
- [18] Pallas GmbH. Pallas MPI Benchmarks - PMB, Part MPI-1. Technical report, 2000.
- [19] J. Pjesivac-Grbovic. Open MPI Collective Operation Performance on Thunderbird. Technical report, The University of Tennessee, Computer Science Department, Knoxville, Technical Report, UT-CS-07-594, 2007.
- [20] J. Pjesivac-Grbovic, T. Angskun, G. Bosilca, G. E. Fagg, E. Gabriel, and J. J. Dongarra. Performance Analysis of MPI Collective Operations. In *Proceedings of the 19th International Parallel and Distributed Processing Symposium*, Denver, CO, April 2005.
- [21] R. Rabenseifner. Automatic MPI Counter Profiling. In *42nd CUG Conference*, 2000.
- [22] S. Saini, R. Ciotti, B. T. N. Gunney, T. E. Spelce, A. E. Koniges, D. Dossa, P. A. Adamidis, R. Rabenseifner, S. R. Tiyyagura, M. Miller, and R. Fatoohi. Performance evaluation of supercomputers using hpcc and imb benchmarks. In *20th International Parallel and Distributed Processing Symposium (IPDPS 2006), Proceedings, 25-29 April 2006, Rhodes Island, Greece*. IEEE, 2006.
- [23] M. Shro and R. Geijn. Collmark mpi collective communication benchmark, 2001.
- [24] S. S. Vadhiyar, G. E. Fagg, and J. Dongarra. Automatically tuned collective communications. In *Supercomputing '00: Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)*, page 3, Washington, DC, USA, 2000. IEEE Computer Society.
- [25] T. Worsch, R. Reussner, and W. Augustin. On benchmarking collective mpi operations. In *Proceedings of the 9th European PVM/MPI Users' Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 271–279, London, UK, 2002. Springer-Verlag.
- [26] 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. In *18th International Parallel and Distributed Processing Symposium (IPDPS 2004), CD-ROM / Abstracts Proceedings, 26-30 April 2004, Santa Fe, New Mexico, USA*, 2004.