A Study of Process Arrival Patterns for MPI Collective Operations

Ahmad Faraj
Blue Gene Software Development
IBM Corporation
Rochester, MN 55901
faraja@us.ibm.com

Pitch Patarasuk Xin Yuan
Department of Computer Science
Florida State University
Tallahassee, FL 32306
{patarasu, xyuan}@cs.fsu.edu

ABSTRACT

Process arrival pattern, which denotes the timing when different processes arrive at an MPI collective operation, can have a significant impact on the performance of the operation. In this work, we characterize the process arrival patterns in a set of MPI programs on two common cluster platforms, use a micro-benchmark to study the process arrival patterns in MPI programs with balanced loads, and investigate the impacts of the process arrival pattern on collective algorithms. Our results show that (1) the differences between the times when different processes arrive at a collective operation are usually sufficiently large to significantly affect the performance; (2) application developers in general cannot effectively control the process arrival patterns in their MPI programs in cluster environments: balancing loads at the application level does not balance the process arrival patterns; and (3) the performance of the collective communication algorithms is sensitive to process arrival patterns. These results indicate that the process arrival pattern is an important factor that must be taken into consideration in developing and optimizing MPI collective routines. We propose a scheme that achieves high performance with different process arrival patterns, and demonstrate that by explicitly considering process arrival pattern, more efficient MPI collective routines than the current ones can be obtained.

Categories and Subject Descriptors

C.4 [**Computer Systems Organization**]: Performance of Systems; D.0 [**Software**]: General

General Terms

Performance

Keywords

MPI, Collective Communication, Process arrival pattern

1. INTRODUCTION

MPI collective operations are used in most MPI applications and they account for a significant portion of the communication time in some applications [21]. Yet, compared to their point-to-point counterparts, MPI collective operations have received less attention, and some fundamental issues in collective operations are still not well understood [9].

The term *process arrival pattern* denotes the timing when different processes arrive at an MPI collective operation (the call site of the collective routine). A process arrival pattern is said to be *balanced* when all processes arrive at the call site roughly at the same time such that the arrival timing does not dramatically affect the performance of the operation, and *imbalanced* otherwise. The terms, balanced and imbalanced arrival patterns, are quantified in Section 3.

The process arrival pattern can have a profound impact on the performance because it decides the time when each process can start participating in an operation. Unfortunately, this important factor has been largely overlooked by the MPI developers community. We are not aware of any study that characterizes process arrival patterns in application programs. MPI developers routinely make the implicit assumption that all processes arrive at the same time (a balanced process arrival pattern) when developing and analyzing algorithms for MPI collective operations [9, 26]. However, as will be shown in this paper, the process arrival patterns in MPI programs, even well designed programs with balanced loads, are more likely to be sufficiently imbalanced to significantly affect the performance.

The imbalanced process arrival pattern problem is closely related to the application load balancing problem. MPI practitioners who have used a performance tool such as Jumpshot to visually see the process arrival times for their collectives should have noticed the imbalanced process arrival pattern problem. However, these two problems are significantly distinct in their time scales: the time differences that cause load imbalance at the application level are usually orders of magnitude larger than those causing imbalanced process arrival patterns. It is often possible to "balance" application loads by applying some load balancing techniques. However, as will be shown later, it is virtually impossible to balance the process arrival patterns in typical cluster environments: even programs with perfectly balanced loads tend to have imbalanced process arrival patterns.

This work is concerned about efficient implementations of MPI collective routines. Application load balancing, although important, requires techniques in the application level and is beyond the scope of this paper. For applications with balanced loads to achieve high performance, it is essential that the MPI library can deliver high performance with different (balanced and imbalanced) process arrival patterns. Hence, from the library implementer point of view, it is crucial to know (1) how application programs behave (the process arrival pattern characteristics); (2) whether the application behavior can cause performance problems in the library routines; and (3) how to deal with the problem and make the library most efficient for the given application behavior. These are the questions that we try to answer in this paper. Note that for an MPI library, there is no difference in applications with different load balancing characteristics. The library should try to deliver the best performance to applications with or without balanced loads.

We study the process arrival patterns of a set of MPI benchmarks on two commercial off-the-shelf (COTS) clusters: a high-end Alphaserver cluster and a low-end Beowulf cluster with Gigabit Ethernet connection. These two clusters are representative and our results can apply to a wide range of practical clusters. We characterize the process arrival patterns in MPI programs, use a micro-benchmark to examine the process arrival patterns in applications with balanced loads and to study the causes of the imbalanced process arrival patterns, and investigate the impacts of different process arrival patterns on some commonly used algorithms for MPI collective operations. The findings include:

- The process arrival patterns for MPI collective operations are usually imbalanced. Even in a micro-benchmark with a perfectly balanced load, the process arrival patterns are still imbalanced.
- In cluster environments, it is virtually impossible for application developers to control the process arrival patterns in their programs without explicitly invoking a global synchronized operation. Many factors that can cause imbalance in computation and communication are beyond the control of the developers. Balancing the loads at the application level is insufficient to balance the process arrival patterns.
- The performance of the MPI collective communication algorithms is sensitive to the process arrival pattern. In particular, the algorithms that perform better with a balanced process arrival pattern tend to perform worse when the process arrival pattern becomes more imbalanced.

These findings indicate that for an MPI collective routine to be efficient in practice, it must be able to achieve high performance with different (balanced and imbalanced) process arrival patterns. Hence, MPI library implementers must take process arrival pattern into consideration in developing and optimizing MPI collective routines. We propose a scheme that uses a dynamic adaptive mechanism to deal with the imbalanced process arrival pattern problem, and demonstrate that by explicitly considering process arrival pattern, more robust MPI collective routines than the current ones can be developed.

The rest of this paper is organized as follows. Section 2 discusses the related work. Section 3 formally describes the process arrival pattern and the parameters we use to characterize it. Section 4 presents the statistics of process arrival patterns in a set of benchmark programs. In Section 5, we study a micro-benchmark that has a perfectly balanced load and investigate the causes for such a program to have imbalanced process arrival patterns. In Section 6, we evaluate the impacts of process arrival patterns on some common algorithms for MPI collective operations. In Section 7, we propose and evaluate a potential solution to the imbalanced process arrival pattern problem. Finally, Section 8 concludes the paper.

2. RELATED WORK

Understanding the application/system behavior is critical for developing an efficient MPI library. Due to its importance, there are numerous research efforts focusing on analyzing MPI communication behavior. Examples include [5, 6, 11, 24, 25, 30]. In [30], the performance of parallel applications is analyzed using a technique that automatically classifies inefficiencies in point-to-point communications. The study analyzes the usage of MPI collective communication routines and their elapsed times. The studies in [6, 11] performed quantitative measures of the static and dynamic MPI routines in parallel applications. Work in [24] performed statistical analysis of all-to-all elapsed communication time on the IBM SP2 machine to understand the causes of performance drop as the number of processors increases. The researchers in [25, 5] examined the NAS parallel benchmarks [16] to quantitatively describe the MPI routines usage and distribution of message sizes. The analysis performed on the parallel applications in these studies (and other similar studies) often involves the investigation of communication attributes such as the type of MPI routines, message size, message volume, message interval, bandwidth requirement, and communication elapsed time. Our study focuses on a specific communication attribute for collective operations, the process arrival pattern, which to the best of our knowledge, has not been studied before. It must be noted that the process arrival pattern is affected not only by the application, but also by the operating system, the system hardware, and the communication library.

3. PROCESS ARRIVAL PATTERN

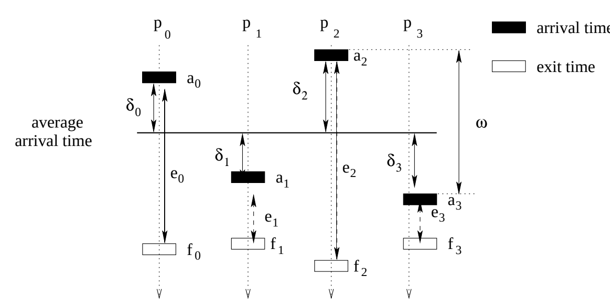

Figure 1: Process arrival pattern

Let n processes, $p_0, p_1, ..., p_{n-1}$, participate in a collective operation. Let a_i be the time when process p_i arrives at the collective operation. The *process arrival pattern* can be represented by the tuple $(a_0, a_1, ..., a_{n-1})$. The average process arrival time is $\bar{a} = \frac{a_0+a_1+...+a_{n-1}}{n}$. Let f_i be the time when process p_i finishes the operation. The *process exit pattern* can be represented by the tuple $(f_0, f_1, ..., f_{n-1})$. The elapsed time that process p_i spends in the operation is thus $e_i = f_i - a_i$, the total time is $e_0 + e_1 + ... + e_{n-1}$, and the average per node time is $\bar{e} = \frac{e_0+e_1+...+e_{n-1}}{n}$. In an application, the total time or the average per node time accurately reflects the time that the program spends on the operation. We will use the average per node time ($\bar{e}$) to denote the performance of an operation (or an algorithm).

We will use the term **imbalance** in the process arrival pattern to signify the differences in the process arrival times at a collective communication call site. Let δ_i be the time difference between p_i's arrival time a_i and the average arrival time $\bar{a}$, $\delta_i = |a_i - \bar{a}|$. The imbalance in the process arrival pattern can be characterized by the *average case imbalance time*, $\bar{\delta} = \frac{\delta_0+\delta_1+...+\delta_{n-1}}{n}$, and the *worst case imbalance time*, $\omega = max_i\{a_i\} - min_i\{a_i\}$. Figure 1 depicts the described parameters in a process arrival pattern.

An MPI collective operation typically requires each process to send multiple messages. A collective algorithm organizes the messages in the operation in a certain way. For example, in the *pair* algorithm for *MPI_Alltoall* [26], the messages in the all-to-all operation are organized in $n-1$ phases: in phase i, $0 \leq i \leq n-1$, process p_j exchanges a message with process $p_{j\ \oplus\ i}$ ($\oplus$ is the exclusive or operator). The impact of an imbalanced process arrival pattern is mainly caused by the early completions or late starts of some messages in the operation. In the *pair* algorithm, early arrivals of some processes will cause some processes to complete a phase and start the next phase while other processes are still in the previous phase, which may cause system contention and degrade the performance. Hence, the impacts of an imbalanced process arrival pattern can be better characterized by the number of messages that can be sent during the period when some processes arrive while others do not. To capture this notion, we normalize the worst case and average case imbalance times by the time to communicate one message. The normalized results are called the *average/worst case imbalance factor*. Let T be the time to communicate one message in the operation, the *average case imbalance factor* equals to $\frac{\bar{\delta}}{T}$ and the *worst case imbalance factor* equals to $\frac{\omega}{T}$. A worst case imbalance factor of 20 means that by the time the last process arrives at the operation, the first process may have sent twenty messages. In general, a process arrival pattern is **balanced** if the worst case imbalance factor is less than 1 (all processes arrive within a message time) and **imbalanced** otherwise.

4. PROCESS ARRIVAL PATTERNS IN MPI PROGRAMS

4.1 Platforms

The process arrival pattern statistics are collected on two representative platforms. The first is the Lemieux machine located in Pittsburgh Supercomputing Center (PSC) [20]. The machine consists of 750 Compaq Alphaserver ES45 nodes connected by Quadrics, each of the nodes includes four 1-GHz SMP processors with 4GB of memory. The system runs Tru64 Unix operating system. All benchmarks are compiled with the native *mpicc* and linked with the native MPI and ELAN library. ELAN is a low-level internode communication library for Quadrics. On Lemieux, the experiments are conducted with a batch partition of 32, 64, and 128 processors (4 processors per node). The second platform is a 16-node Beowulf cluster, whose nodes are Dell Dimension 2400, each with a 2.8GHz P4 processor and 128MB of memory. All machines run Linux (Fedora) with the 2.6.5-1.358 kernel. These machines are connected by a Dell Powerconnect 2624 1Gbps Ethernet switch. This system uses MPICH 2-1.0.1 for communication. All programs are compiled with the *mpicc* that comes with the MPICH package.

Some of the times and the corresponding bandwidths (BW) for one way point-to-point communications with different message sizes on the two platforms are summarized in Table 1. These numbers, which are obtained using a pingpong program, are used to compute imbalance factors.

Table 1: One way point-to-point communication time and bandwidth on Lemieux and Beowulf

message size	Lemieux		Beowulf	
	time (ms)	BW (MB/s)	time (ms)	BW (MB/s)
4B	0.008	0.50	0.056	0.07
256B	0.008	32.0	0.063	4.10
1KB	0.021	49.5	0.088	11.6
4KB	0.029	141	0.150	27.3
16KB	0.079	207	0.277	59.1
32KB	0.150	218	0.470	69.7
64KB	0.291	225	0.846	77.5
128KB	0.575	228	1.571	83.4

4.2 Benchmarks

Table 2 summarizes the seven benchmarks. For reference, we show the code size as well as the execution and collective communication elapsed times for running the programs on $n = 64$ processors on Lemieux. Table 3 shows the major collective communication routines in the benchmarks and their dynamic counts and message sizes (assuming $n = 64$). There are significant collective operations in all programs. Next, we briefly describe each benchmark and the related parameters/settings used in the experiments.

FT (Fast-Fourier Transform) is one of the parallel kernels included in NAS parallel benchmarks [16]. FT solves a partial differential equation using forward and inverse FFTs. The collective communication routines used in FT include *MPI_Alltoall*, *MPI_Barrier*, *MPI_Bcast*, and *MPI_Reduce* with most communications being carried out by *MPI_Alltoall*. We used class B problem size supplied by the NAS benchmark suite in the evaluation.

IS (Integer Sort) is a parallel kernel from NAS parallel benchmarks. It uses bucket sort to order a list of integers. The MPI collective routines in IS are *MPI_Alltoall*, *MPI_Alltoallv*, *MPI_Allreduce*, and *MPI_Barrier* with most communications carried out by the *MPI_Alltoallv* routine. We also used class B problem size.

Table 2: Summary of benchmarks (times are measured on Lemieux with 64 processors)

benchmark	#lines	total time	comm. time
FT	2234	13.4s	8.3s
IS	1091	2.2s	1.6s
LAMMPS	23510	286.7s	36.1s
PARADYN	6252	36.6s	33.1s
NBODY	256	59.5s	1.5s
NTUBE 1	4480	894.4s	32.3s
NTUBE 2	4570	852.9s	414.1s

LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) [13] is a classical parallel molecular dynam-

Table 3: The dynamic counts of major collective communication routines in the benchmarks ($n = 64$)

benchmark	routine	msg size	dyn. count
FT	alltoall	131076	22
	reduce	16	20
IS	alltoallv	33193*	11
	allreduce	4166	11
	alltoall	4	11
LAMMPS	allreduce	42392	2012
	bcast	4-704	48779
	barrier		4055
PARADYN	allgatherv	6-1290*	16188
	allreduce	4-48	13405
NBODY	allgather	5000	300
NTUBE 1	allgatherv	16000*	1000
NTUBE 2	allreduce	8	1000

* the average of all message sizes in the v-version routines.

ics code. It models the assembly of particles in a liquid, solid, or gaseous state. The code uses *MPI_Allreduce*, *MPI_Bcast*, and *MPI_Barrier*. We ran the program with 1720 copper atoms for 3000 iterations.

PARADYN (Parallel Dynamo) [18] is a molecular dynamics simulation. It utilizes the embedded atom method potentials to model metals and metal alloys. The program uses *MPI_Allgather*, *MPI_Allgatherv*, *MPI_Allreduce*, *MPI_Bcast*, and *MPI_Barrier*. In the experiments, we simulated 6750 atoms of liquid crystals in 1000 time steps.

NBODY [17] simulates over time steps the interaction, in terms of movements, positions and other attributes, among the bodies as a result of the net gravitational forces exerted on one another. The code is a naive implementation of the nbody method and uses *MPI_Allgather* and *MPI_Gather* collective communications. We ran the code with 8000 bodies and for 300 time steps.

NTUBE 1 performs molecular dynamics calculations of thermal properties of diamond [22]. This version of the code uses *MPI_Allgatherv* and *MPI_Reduce*. In the evaluation, the program ran for 1000 steps and each processor maintained 100 atoms.

NTUBE 2 is a different implementation of the Nanotube program. The functionality of NTUBE 2 is exactly the same as NTUBE 1. The collective communication routines used in this program are *MPI_Allreduce* and *MPI_Reduce*. In the evaluation, the program ran for 1000 steps with each processor maintaining 100 atoms.

4.3 Data collection

To investigate process arrival patterns and other statistics of MPI collective communications, we develop an MPI wrapper library. The wrapper records an event at each MPI process for each entrance and exit of an MPI collective communication routine. An event records information about the timing, the operation, the message size, etc. The times are measured using the *MPI_Wtime* routine. Events are stored in memory during program execution until *MPI_Finalize* is called, when all processors write the events to log files for post-mortem analysis. The functionality of our wrapper is similar to PMPI, we use our own wrapper for future extension. Accurately measuring the time on different machines requires a globally synchronized clock. On Lemieux, such a synchronized clock is available. On the Beowulf cluster, the time on different machines is not synchronized. We resolve the problem by calling an *MPI_Barrier* after *MPI_Init* and having all measured times normalized with respect to the exit time of the *MPI_Barrier*. Basically, we are assuming that all (16) machines exit a barrier operation at the same time. This introduces inaccuracy that is roughly equal to the time to transfer several small messages.

4.4 Process arrival pattern statistics

In this sub-section, we focus on presenting the process arrival pattern statistics. The causes for MPI applications to have such behavior will be investigated in the next section. Table 4 shows the average of the worst/average case imbalance factors among all collective routines in each benchmark on Lemieux and the Beowulf cluster. The table reveals several notable observations. First, the averages of the worst case imbalance factors for all programs on both clusters are quite large, even for FT, whose computation is fairly balanced. Second, the process arrival pattern depends heavily on the system architecture. For example, the imbalance factors for NTUBE 1 and NTUBE 2 are much larger on Lemieux than on the Beowulf cluster. This is because these two programs were designed for single CPU systems. When running them on Lemieux, an SMP cluster, the process arrival patterns become extremely imbalanced. Overall, the imbalance factors for all programs on both platforms are large: the best average worst case imbalance factor is 19 for Lemieux (LAMMPS) and 17 for Beowulf (NTUBE 1).

Table 4: The average of worst case ($\frac{\bar{\omega}}{T}$) and average case ($\frac{\bar{\delta}}{T}$) imbalance factors among all collective routines on two the platforms

benchmark	imbalance factor			
	Lemieux (n = 128)		Beowulf	
	average	worst	average	worst
FT	91.0	652	278	1.2K
IS	61.0	358	1.4K	11K
LAMMPS	4.00	19.0	273	630
PARADYN	9.10	46.0	12.0	79.0
NBODY	13.0	132	12.0	50.0
NTUBE 1	4.8K	38K	4.30	17.0
NTUNE 2	85K	347K	9.00	39.0

Operations that account for most of the communication times typically have large message sizes. In Figure 2, we distinguish operations with medium/large message sizes ($>$ 1000B) from those with small message sizes ($\leq$ 1000B). Part (a) of Figure 2 shows the distribution of the worst case imbalance factors for the results on Lemieux (128 processors) while part (b) shows the results on the Beowulf cluster. All benchmarks are equally weighted when computing the distribution. As expected, arrival patterns for operations with large messages are in general less imbalanced than those for operations with small messages. This is mainly due to the way the imbalance factors are computed: larger messages mean larger per message time (T). However, as can be seen from the figure, there is a significant portion of operations with both small sizes and medium/large sizes having large imbalance factors and only a small fraction of the operations are balanced. In particular, for operations with medium/large messages, only a small percentage (21% on Lemieux and 6% on Beowulf) have balanced process arrival patterns (a worst case imbalance factor less than 1). The percentage is smaller for operations with small messages. This indicates that imbalanced process arrival patterns are much more common than balanced process arrival patterns.

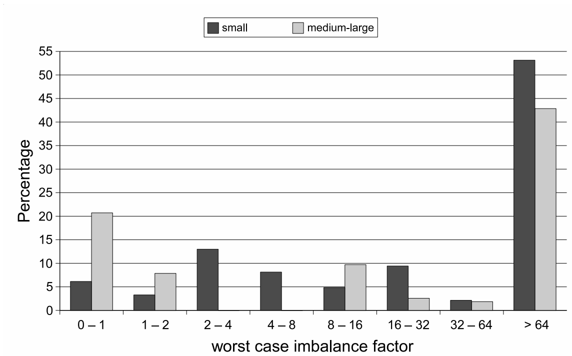

(a) Lemieux (128 processors)

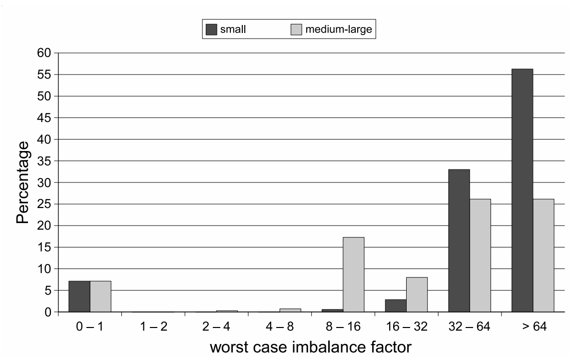

(b) Beowulf cluster

Figure 2: The distribution of worst case imbalance factors ($\frac{\bar{\omega}}{T}$)

In Table 5, we further narrow our focus on the imbalance factors for collective operations that are important in the benchmarks. These are the operations that appear in the main loop and account for a significant amount of application time. Compared with the imbalance factors shown in Table 4, we can see that the process arrival patterns for these important routines are generally less imbalanced than the average of all routines in the applications, which reflects the fact that programmers are more careful about the load balancing issue in the main loop. However, the process arrival patterns for these important routines are still quite imbalanced. On both platforms, only *MPI_Alltoallv* in IS can be classified as having balanced process arrival patterns. Examining the source code reveals that this routine is called right after another MPI collective routine.

Another interesting statistics is the characteristics of process arrival patterns for each individual call site. If the process arrival patterns for each call site in different invocations exhibit heavy fluctuation, the MPI routine for this call site must achieve high performance for all different types of process arrival patterns to be effective. On the other hand, if the process arrival patterns for the same call site is statistically similar, the MPI implementation will only need to optimize for the particular type of process arrival patterns. In the experiments, we observe that the process arrival patterns for different invocations of the same call site exhibit

Table 5: The imbalance factor for major routines

	major routine	imbalance factor			
		Lemieux (128)		Beowulf	
		ave.	worst	ave.	worst
FT	alltoall	2.90	24.0	26.0	124
IS	alltoallv	0.00	0.20	0.20	0.80
	allreduce	145	756	4.4K	34K
LAMMPS	bcast	0.20	3.40	299	671
	allreduce	16.3	91.3	24	132
	barrier	28.6	157.3	106	442
PARADYN	allgatherv	0.80	6.50	10.0	66.5
	allreduce	15.7	73.3	14.0	93.0
NBODY	allgather	13.0	132	12.0	50.0
NTUBE 1	allgatherv	78.8	345	3.50	14.0
NTUBE 2	allreduce	83K	323K	9.00	39.0

a *phased* behavior: the process arrival patterns are statistically similar for a period of time before they change. In some cases, the process arrival patterns for the same call site are statistically similar in the whole program. Figure 3 depicts a representative case: the imbalance factors for the *MPI_Allgather* routine in NBODY. As can be seen from the figure, the majority of the calls have similar worst case and average case imbalance factors despite some large spikes that occur once in a while. This indicates that it might be feasible to customize the routine for each MPI call site and get good performance.

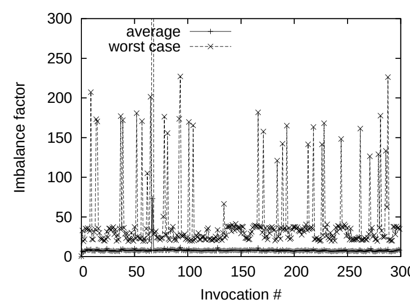

Figure 3: The imbalance factors for *MPI_Allgather* in NBODY on Lemieux ($n = 128$)

4.5 Summary

While we expect to see some imbalance process arrival patterns in MPI programs, it is surprising to see the very low percentage of balanced process arrival patterns. The low percentage applies to applications whose loads are fairly balanced, to collective operations in the main loops where load balancing is critical for the performance of the applications, and to operations with all different message sizes.

5. PROCESS ARRIVAL PATTERNS IN A MICRO-BENCHMARK

Since a well designed MPI program typically has a balanced computation load, understanding the process arrival patterns in this type of programs is particularly important. In this section, we study a simple micro-benchmark, shown in Figure 4, where all processes perform exactly the same amount of computation and communication (the load is perfectly balanced). The goal is (1) to determine whether application programmers can control the critical process arrival patterns in their MPI programs by balancing the load at the application level, and (2) to investigate the causes of the imbalanced process arrival patterns. In this micro-benchmark, a barrier is called before the main loop that is executed 1000 times. There are two components inside the loop: lines 4-6 simulating the computation and an MPI_Alltoall() operation in line 8 after the computation. The computation time can be adjusted by changing the parameter XTIME.

```
(1) MPI_Barrier(...);
(2) for (i=0; i<1000; i++) {
(3)    /* compute for roughly X milliseconds */
(4)    for (m=0; m< XTIME; m++)
(5)       for (k=1, k<1000; k++)
(6)          a[k] = b[k+1] - a[k-1] * 2;
(7)    arrive[i] = MPI_Wtime();
(8)    MPI_Alltoall(...);
(9)    leave[i] = MPI_Wtime()
(10)}
```

Figure 4: Code segment for a micro-benchmark

We measured the process arrival patterns for the all-to-all operation. Due to space limitation, we will only report results for message size 64KB. Smaller message sizes result in larger imbalance factors. The average computation time in each node is set to 200ms for both clusters. Figure 5 shows the worst and average case imbalance factors in each invocation in a typical execution on each of the two platforms. In both clusters, there is a substantial imbalance in the process arrival patterns even though all processors perform exactly the same operations. The imbalance factors on Lemieux are larger than those on the Beowulf cluster for several reasons. First, Lemieux has more processes and thus has a higher chance to be imbalanced. Second, on Lemieux, different jobs share the network in the system, the uncertainty in messaging can cause the imbalance. Third, Lemieux has a faster network, the same imbalance time results in a larger imbalanced factor.

We further investigate the causes of the imbalanced process arrival patterns in this simple benchmark. For the *MPI_Alltoall* routine to have the imbalanced process arrival patterns shown in Figure 5, there can be only two potential causes. First, it might take different processors different times to run the (same) computation. An earlier study [19] has shown that this is indeed happening in some clusters and has attributed this phenomenon to the asynchronous operating system events. Second, it might take different processors different times to perform the communication (*MPI_Alltoall*). This imbalance in the communication is reflected in the process exit patterns. In the following, we study the relationship among the imbalance in process arrival patterns, computation times, and process exit patterns in the micro-benchmark. The worst case imbalance factor for a process exit pattern is defined similarly to that of a process arrival pattern. The *computation imbalance time* is defined as the maximum time among all processes to execute the computation minus the minimum time among all processes. To be consistent, we use the imbalance factor in the comparison, which is equal to the imbalance time divided by the time to send one message (64KB).

We change the XTIME parameter such that the average computation time lasts for 50, 100, 200, 400, and 800ms. Due to the nondeterministic nature in the imbalance, we repeat each experiment 5 times, each on a different day. In each experiment, we collect data from the 1000 invocations of the all-to-all routine. We then use the data from the 5000 samples (5 experiments, 1000 samples per experiment) to compute the average values and the 95% confidence intervals of imbalance factors for process arrival patterns, process exit patterns, and computation.

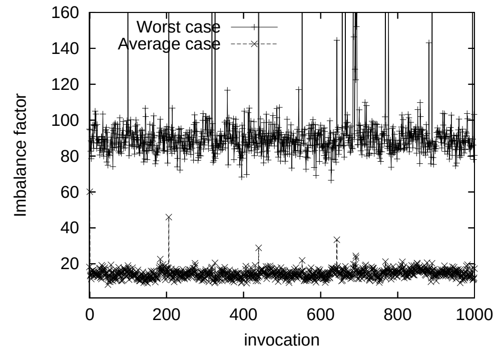

(a) Lemieux (128 processors)

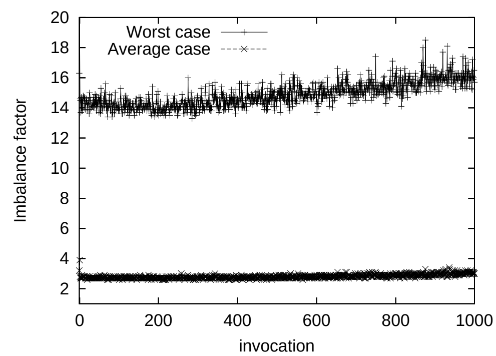

(b) Beowulf cluster

Figure 5: Process arrival patterns in the micro-benchmark (64KB message size, 200ms computation time) on the two platforms

Tables 6 and 7 show the worst case imbalance factors for exit patterns, computation, and arrival patterns in the micro-benchmark for different computation times on the two platforms. In the tables, for each worst case (exit, computation, arrival) imbalance factor, we show the average value along with the confidence interval in the format of $ave \pm \frac{interval}{2}$, which denotes that the 95% confidence interval is $[ave - \frac{interval}{2}, ave + \frac{interval}{2}]$. There are a number of observations in the tables. First, when changing the computation time from 50ms to 800ms, the computation imbalance in both clusters increases almost linearly. Such imbalance in computation is inherent to the system and is impossible for application developers to overcome. This explains why in our benchmark study of the previous section, we only observe balanced process arrival patterns in consecutive collective routine calls. Second, the worst case imbalance factors for process arrival patterns are consistently larger than the computation imbalance factors, which indicates that the imbalances in both computation and communication are contributing to the imbalance in the process arrival patterns. Third, on Lemieux, the imbalance factors for process exit patterns are almost the same with different process arrival patterns while on the Beowulf cluster, the imbalance factors for process exit patterns are quite different. This is because different algorithms are used to implement *MPI_Alltoall* on the two clusters. On the Beowulf cluster, since the imbalance factors for process exit patterns are somewhat related to those for process arrival patterns, the imbalance effect may be accumulated as the simple benchmark executes. This explains the slight upward trend in the worst case imbalance factor in Figure 5 (b). Nonetheless, the imbalance in communication, which is directly affected by the library implementation, is beyond the control of application developers.

Table 6: Effects of process exit patterns and computation imbalance on process arrival patterns on Lemieux (32 processors)

comp. time	worst case imbalance factor($\frac{\omega}{T}$)		
	exit	computation	arrival
50ms	15.2 ± 0.7	23.4 ± 0.3	32.5 ± 0.8
100ms	15.2 ± 0.6	46.8 ± 1.6	54.5 ± 1.9
200ms	15.0 ± 0.3	87.4 ± 1.8	92.7 ± 1.9
400ms	15.1 ± 0.8	160 ± 1.9	164 ± 2.0
800ms	15.0 ± 0.3	320 ± 3.6	322 ± 3.6

Table 7: Effects of process exit patterns and computation imbalance on process arrival patterns on the Beowulf cluster

comp. time	worst case imbalance factor ($\frac{\omega}{T}$)		
	exit	computation	arrival
50ms	5.07 ± 1.29	3.16 ± 0.02	7.02 ± 1.29
100ms	4.32 ± 1.00	7.52 ± 0.02	9.53 ± 0.99
200ms	3.71 ± 0.11	14.18 ± 0.02	15.17 ± 0.06
400ms	6.22 ± 0.23	31.41 ± 0.30	33.17 ± 0.35
800ms	11.62 ± 0.41	56.24 ± 0.05	56.29 ± 0.20

5.1 Summary

The way a program is coded is only one of many factors that can affect process arrival patterns. Other factors, such as system characteristics and library implementation schemes that can introduce the inherent imbalance in computation and communication, are beyond the control of application developers. Hence, it is unrealistic to assume that application programmers can balance the load at the application level to make the process arrival patterns balanced. The process arrival patterns in MPI programs are and will be imbalanced in most cases in a cluster environment.

6. IMPACTS OF IMBALANCED PROCESS ARRIVAL PATTERNS

We study the impact of the process arrival pattern on commonly used algorithms for *MPI_Alltoall* and *MPI_Bcast*. *MPI_Alltoall* and *MPI_Bcast* represent two types of MPI collective operations: *MPI_Alltoall* is an inherently synchronized operation, that is, a process can complete this operation only after all processes arrive; while *MPI_Bcast* is not an inherently synchronized operation. The impacts of imbalanced process arrival patterns on the algorithms are not clear. For example, some communication algorithms such as some *MPI_Bcast* algorithms, may be able to tolerate some degrees of imbalanced process arrivals while others may not. This section tries to systematically study the impacts of imbalanced process arrival patterns on different types of algorithms.

The evaluated *MPI_Alltoall* algorithms include the *simple*, *Bruck*, *pair*, and *ring* algorithms. The *simple* algorithm basically posts all receives and all sends, starts the communications, and waits for all communications to finish. The *Bruck* algorithm [4] is a $lg(n)$-step algorithm that is designed for achieving efficient all-to-all with small messages. The *pair* algorithm only works when the number of processes, n, is a power of two. It partitions the all-to-all communication into $n-1$ steps. In step i, process p_j exchanges a message with process $p_{j \oplus i}$. The *ring* algorithm also partitions the all-to-all communication into $n-1$ steps. In step i, process p_j sends a messages to process $p_{(j+i) \ mod \ n}$ and receives a message from process $p_{(j-i) \ mod \ n}$. More detailed description of these algorithms can be found in [26]. We also consider the native algorithm used in *MPI_Alltoall* on Lemieux, which is unknown to us.

The evaluated *MPI_Bcast* algorithms include the *flat tree*, *binomial tree*, *scatter-allgather*, and the native algorithm on Lemieux, which is unknown to us. In the flat tree algorithm, the root sequentially sends the broadcast message to each of the receivers. In the binomial tree algorithm [15], the broadcast follows a hypercube communication pattern and the total number of messages that the root sends is $lg(p)$. The scatter-allgather algorithm, used for broadcasting large messages in MPICH [15], first distributes the *msize*-byte message to all nodes by a scatter operation (each node gets $\frac{msize}{p}$ bytes), and then performs an all-gather operation to combine the scattered messages to all nodes.

```
(1) r = rand() % MAX_IF;
(2) for (i=0; i<ITER; i++) {
(3)    MPI_Barrier (...);
(4)    for (j=0; j<r; j++) {
(5)       ... /* computation time equal to one msg time */
(6)    }
(7)    t0 = MPI_Wtime();
(8)    MPI_Alltoall(...);
(9)    elapse += MPI_Wtime() - t0;
(10)}
```

Figure 6: Code segment for measuring the impacts of imbalanced process arrival patterns

Figure 6 outlines the code segment we use to measure the performance with a controlled imbalance factor in the random process arrival patterns. The worst case imbalance factor is controlled by a variable *MAX_IF* (maximum imbalance factor). Line 1 generates a random number r that is bounded by *MAX_IF*. Before the all–to–all routine (or broadcast) is measured (lines 7-9), the controlled imbalanced process arrival pattern is created by first calling a barrier (line 3) and then introducing some computation between the barrier and all-to-all routines. The time to complete the computation is controlled by r. The time spent in the loop body in line 5 is made roughly equal to the time for sending one message (see Table 1), and the total time for the computation is roughly equal to the time to send r messages. Hence, the larger the value of *MAX_IF* is, the more imbalanced the process arrival pattern becomes. Note that the actual worst case imbalance factor, especially for small message sizes, may not be bounded by *MAX_IF* since the process exit patterns of *MPI_Barrier* may not be balanced.

For each process arrival pattern, the routine is measured 100 times ($ITER = 100$) and the average elapsed time on each node is recorded. For each *MAX_IF* value, we perform 32 experiments (32 random process arrival patterns with the same value of *MAX_IF*). The communication time is reported by the confidence interval with a 95% confidence level, computed from the results of the 32 experiments.

Figure 7 (a) shows the results for 1B all-to-all communication on Lemieux (32 processors). When $MAX\_IF \leq 9$, the Bruck algorithm performs better than the ring and pair algorithms, and all three algorithms perform significantly better than the simple algorithm. However, when the imbalance factor is larger ($16 \leq MAX\_IF \leq 129$), the simple algorithm shows better results. The native algorithm performs much better than all algorithms in the case when $MAX\_IF \leq 129$. When $MAX\_IF = 257$, the native algorithm performs worse than the ring and simple algorithms. These results show that under different process arrival patterns with different worst case imbalance factors, the algorithms have different performance. When the imbalance factor increases, one would expect that the communication time should increase. While this applies to the Bruck, ring, pair and the native algorithms, it is not the case for the simple algorithm: the communication time actually decreases as *MAX_IF* increases when $MAX\_IF \leq 17$. The reason is that, in this cluster, 4 processors share the network interface card. With moderate imbalance in the process arrival pattern, different processors initiate their communications at different times, which reduces the resource contention and improves communication efficiency.

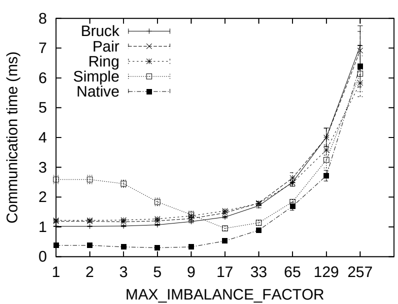

(a) Message size $= 1B$

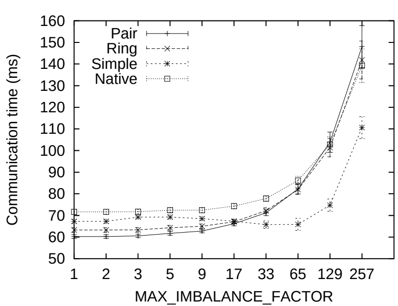

(b) Message size $= 64KB$

Figure 7: 1B and 64KB *MPI_Alltoall* on Lemieux (32 processors)

Figure 7 (b) shows the performance when the message size is 64KB. When $MAX\_IF \leq 9$, the pair algorithm is noticeably more efficient than the ring algorithm, which in turn is faster than the simple algorithm. However, the simple algorithm offers the best performance when $MAX\_IF \geq 33$. For this message size, the native algorithm performs worse than all three algorithms when $MAX\_IF \leq 65$. The figure also shows that each algorithm performs very differently under process arrival patterns with different imbalance factors. The trend observed in Lemieux is also seen in the Beowulf cluster, which is captured in Figure 8.

Since *MPI_Alltoall* is an inherent synchronized operation, when the imbalance factor is very large, all algorithms should have a similar performance. This is shown in all experiments except for the 64KB case on Lemieux where $MAX\_IF = 257$ is not sufficiently large. However, from the experiments, we can see that algorithms that perform better with a balanced process arrival pattern tend to perform worse when the process arrival pattern becomes more imbalanced.

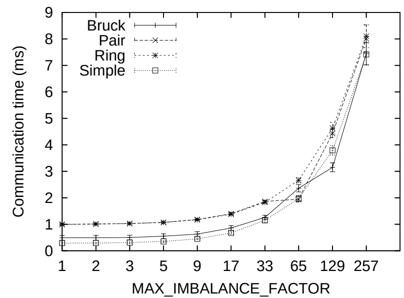

(a) Message size $= 1B$

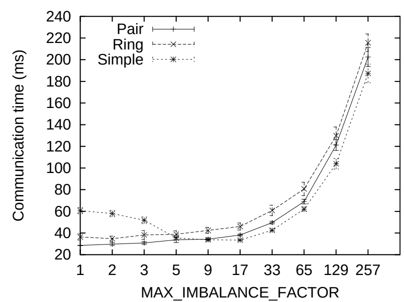

(b) Message size $= 64KB$

Figure 8: 1B and 64KB *MPI_Alltoall* on Beowulf cluster

Figure 9 (a) shows the results for 1B broadcast on Lemieux (32 processors). When $MAX\_IF \leq 8$, all algorithms perform similarly. When $MAX\_IF > 8$, the flat tree algorithm performs considerably better than the other algorithms. Part (b) of the figure shows the results for broadcasting 64KB messages. When $MAX\_IF < 8$, native, binomial, and scatter-allgather algorithms perform similarly and better than the flat tree algorithm. However, when $MAX\_IF > 16$, the flat tree algorithm performs better than all other algorithms. Moreover, the performance advantage of the flat tree algorithm increases as the imbalance factor increases. The results on the Beowulf cluster (not shown due to space limitation) have a similar trend.

The algorithms for *MPI_Bcast* that perform better under a balanced process arrival pattern also perform worse when the arrival pattern becomes imbalanced. In contrast to the results for *MPI_Alltoall*, the performance difference

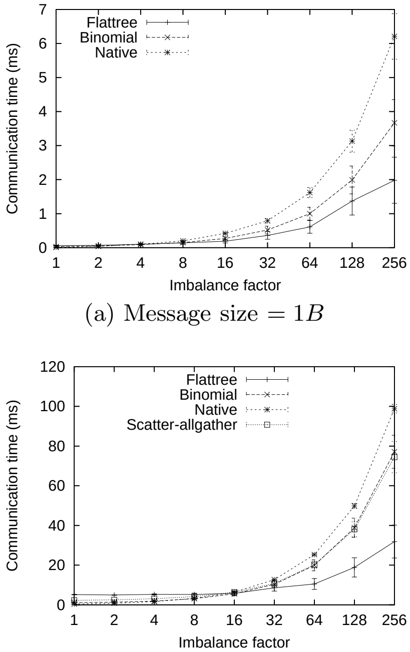

(a) Message size = $1B$

(b) Message size = $64KB$

Figure 9: 1B and 64KB *MPI_Bcast* on Lemieux (32 processors)

for different broadcast algorithms widens as the imbalance factor increases. Due to the implicit synchronization in *MPI_Alltoall*, there is a limit on the impacts of an imbalanced pattern (all algorithms will have a similar performance when the imbalance factor is very large). However, for the *MPI_Bcast* type of operations that are not inherently synchronized, the impacts can potentially be unlimited.

6.1 Summary

The common observation in the experiments in this section is that collective communication algorithms respond differently to different process arrival patterns. The algorithm that performs better with a balanced process arrival pattern tends to perform worse when the process arrival pattern becomes more imbalanced. Moreover, depending on the type of collective operations, the impact of imbalanced process arrival pattern can be large.

7. A POTENTIAL SOLUTION

From Section 6, we can see that different collective communication algorithms react to different process arrival patterns differently. Thus, ideally, to achieve the best performance, we should use a feedback mechanism to adapt the communication algorithms based on the process arrival patterns. Such a solution requires (1) the knowledge of the best algorithm for any given process arrival pattern, and (2) the knowledge of the process arrival pattern at the time when the algorithm is selected. Both requirements are difficult to be met. First, the number of different process arrival patterns is infinite. Even when the process arrival time of each process is quantified with a small number of classes, the number of patterns is still very large. Consider for example a system with 128 processors. Assuming that the process arrival time for each process is classified with two levels: early and late (an extremely coarse grain quantization), there are still 2^{128} process arrival patterns to be considered. Second, it should be clear from the results in Section 5 that the factors that cause imbalance process arrival patterns are mostly random and beyond the control of a programmer. Hence, the exact process arrival pattern for a given invocation cannot be determined unless measured: implementing the feedback mechanism would likely disturb the process arrival patterns for the operation.

Our proposed solution is based on two key observations. First, while the process arrival pattern for a collective operation is nondeterministic, the process arrival patterns for each individual call site tend to exhibit a phased behavior as discussed in Section 4, that is, the process arrival patterns are statistically similar for an extended period of time before they change (See Figure 3). Hence, if the library routine can find an algorithm that can provide good performance, it is likely that the algorithm will provide good performance for an extended period of time. Second, while different collective algorithms may perform the best for different process arrival patterns, the performance of a given algorithm changes slowly as the maximum imbalance factor changes, as shown by the small 95% confidence intervals and the smooth curves for all algorithms in Section 6. This indicates that when an algorithm gives the best performance for a particular process arrival pattern, it tends to give reasonable performance for other process arrival patterns that are not drastically different. Hence, to get reasonable performance, we do not need to find all best algorithms for different process arrival patterns. Instead, we just need to find *some* best algorithms for *some* representative points in the process arrival pattern space.

These two observations strongly suggest that it might be possible to develop a collective routine that performs well for different process arrival patterns by (1) identifying good algorithms for different process arrival patterns and (2) using a dynamic adaptive mechanism that selects the best performing algorithm at run-time. The STAR-MPI that we developed previously [8] provides such a dynamic adaptive mechanism. We apply the STAR-MPI idea to develop a robust *MPI_Alltoall* routine by incorporating process arrival pattern aware all-to-all algorithms. As shown in the performance evaluation, our robust routine consistently achieves higher performance for different platforms and applications (different process arrival patterns) than native MPI implementations. Next, we will describe the process arrival pattern aware all-to-all algorithms included in our robust all-to-all routine. Details about the dynamic adaptive mechanism can be found in [8].

7.1 Process arrival pattern aware all-to-all algorithms

To identify good algorithms for different process arrival patterns, we empirically test an extensive set of algorithms that we implemented [7] on different platforms. We will describe the selected algorithms and give rationale about why they provide good performance in different situations.

Pair/Ring algorithms. The *pair* and *ring* algorithms, described in Section 6, provide good performance when the process arrival pattern is balanced.

While the ring and pair algorithms are efficient when the process arrival pattern is balanced, they do not perform well when the imbalanced factor is larger. In particular, when the

worst case imbalanced factor is larger than 1, early arrivals of some processes in the pair/ring algorithms will cause some processes to complete a phase and start the next phase while other processes are still in the previous phase. This may destroy the phase structure, cause system contention, and degrade the performance. This problem can be resolved in two ways, each resulting a different type of efficient algorithms.
Ring/Pair + one MPI barrier. One solution is to prevent the imbalanced arrival patterns from happening. This can be achieved by adding a barrier operation before the ring/pair algorithm. This way, when the ring or pair algorithm is executed, it guarantees to have a balanced process arrival pattern. This approach forces processes that arrive at the operation early to idle. It provides good performance when the worst case imbalance factor is small (but not 0).
Ring/Pair + light barrier. The *ring/pair + one MPI barrier* algorithm forces processes that arrive at the operation early to idle. This may not be efficient when a large number of processes arrive at the operation significantly earlier than others since processes that arrive early could have used the idle time to perform some useful operations. The *ring/pair+light barrier* is another solution to the problem caused by the imbalanced process arrival patterns. The idea is (1) to allow the phases to proceed in an asynchronous manner and (2) to use a mechanism (light barrier) to minimize the impact of the imbalanced process arrival pattern. Basically, whenever there is possibility that two messages (in different phases) can be sent to the same processes at the same time and cause contention, a light barrier is added to sequentialize the two messages. Hence, the impact of the imbalanced process arrival pattern is minimized.
Simple. All the above algorithms are based on the concept of *phase*, which requires processes to coordinate. In the case when the imbalance factor is large, the coordination among processes may actually hinder the communication performance. The *simple* algorithm, described in Section 6, performs all communications in a single phase (step), eliminating the coordination among processes. As a result, this algorithm performs very well for sufficiently imbalanced process arrival patterns.

Besides these algorithms, our routine also includes the native *MPI_Alltoall*, which is selected in the native MPI library for a good reason. Hence, there are a total of 8 algorithms that are included in our robust *MPI_Alltoall* routine. As shown in the performance evaluation section, our routine performs better than the native *MPI_Alltoall* in most cases, which indicates that the native *MPI_Alltoall* implementation is not the best performing algorithm among the algorithms in many practical cases. Notice that some of these algorithms, such as *pair/ring* and *simple*, are included in MPICH, where they are used to realize the all-to-all operation with different message sizes. In our routine, all the algorithms can be selected to realize the operation with the same message size, but different process arrival patterns.

7.2 Performance results

We evaluate the performance of the robust *MPI_Alltoall* routine on the following high-end clusters: the Lemieux cluster at Pittsburgh Supercomputing Center [20], the UC/ANL Teragrid cluster at Argonne [28], the AURORA cluster at the University of Technology at Vienna [1], and the AVIDD-T cluster at Indiana University [2]. Table 8 summarizes the configurations of all the clusters besides Lemieux, whose configuration is described in Section 4. The benchmarks were compiled with the native *mpicc* or *mpif90* installed on the systems and linked with the native MPI library. We use the micro-benchmark in Figure 4 (Section 5) and a set of application benchmarks in the evaluation. In presenting the results, we will denote our robust routine as ROBUST and the native routine as NATIVE. The software used in this section, including our robust all-to-all routine and all benchmarks, are available to the public at http://www.cs.fsu.edu/∼xyuan/MPI/STAR-ALLTOALL.

cluster	UC-TG [28]	Aurora [1]	Avidd-T [2]
node	two 2.4GHz Xeon	two 3.6GHz Nocona	four 1.3 GHz Itanium II
memory	4GB	4GB	6GB
interconn.	Myrinet	Infiniband	Myrinet
MPI	MPICH-GM 1.2.7	MVAPICH 0.9.5	MPICH-GM 1.2.7

Table 8: Clusters used other then Lemieux

Micro-benchmark results

Table 9 shows the micro-benchmark (Figure 4 in Section 5) results. The table gives the all-to-all communication times for different average computation times (varying XTIME). The message size for the all-to-all operation is 64KB. Note again that the process arrival patterns are still imbalanced even in such cases when the load in the micro-benchmark is perfectly balanced. As shown in the table, as we increase the computation time from 50ms to 400ms, causing a relative increase in the process arrival imbalance, the communication time for both NATIVE and ROBUST increases as the imbalance increases. However, ROBUST is able to sustain its substantial speed ups over NATIVE across different clusters, different number of nodes, and different computation loads. We have also performed experiments when imbalanced computation loads are explicitly introduced, the trend in the results is similar. This demonstrates the robustness of ROBUST.

machine	implem-entation	computation time			
		50ms	100ms	200ms	400ms
Lemieux (128)	NATIVE	346	348	352	362
	ROBUST	263	266	273	283
	speed up	31.6%	30.8%	28.9%	27.9%
UC-TG (64)	NATIVE	117	108	147	185
	ROBUST	105	93.0	122	121
	speed up	11.4%	16.1%	20.5%	52.9%
Avidd-T (32)	NATIVE	76.0	77.5	77.5	80.5
	ROBUST	63.3	64.7	66.7	64.9
	speed up	20.1%	19.8%	16.2%	24.0%
Aurora (16)	NATIVE	8.90	9.20	9.20	9.60
	ROBUST	8.20	8.10	8.50	8.10
	speed up	8.50%	13.6%	8.30%	18.5%

Table 9: *MPI_Alltoall* (64KB) time (milli-seconds) in the micro-benchmark

Results for application benchmarks

We also evaluate the performance using four MPI all-to-all benchmarks: FT, VH-1, MT, and FFT-2D. FT is a parallel kernel from the NAS parallel benchmarks [16]. In the evaluation, we run the class C problem for 400 steps. The VH-1

program	implem.	LEMIEUX ($n = 128$)		UC-TG ($n = 64$)		AVIDD-T ($n = 32$)		AURORA ($n = 16$)	
		comm.	total	comm.	total	comm.	total	comm.	total
FT	NATIVE	265.0s	501.3s	6182s	10107s	1069s	1720s	616.0s	1690s
	ROBUST	221.0s	450.6s	5917s	9832s	865.0s	1583s	424.0s	1500s
	speed up	19.9%	11.3%	4.9%	2.8%	23.6%	8.7%	45.3%	12.7%
	algorithm	pair+one barrier		simple		pair+light barrier		pair+light barrier	
VH-1	NATIVE	2495s	3679s	836.0s	5489s	443.0s	1600s	45.50s	457.0s
	ROBUST	1602s	3429s	661.0s	5277s	337.0s	1506s	39.50s	451.0s
	speed up	55.7%	7.3%	26.5%	3.9%	31.5%	6.2%	15.2%	1.3%
	algorithm	simple		simple		simple		simple	
FFT-2D	NATIVE	178.3s	403.0s	78.60s	594.0s	91.20s	190.5s	47.30s	255.0s
	ROBUST	165.0s	399.0s	60.00s	576.0s	79.80s	180.3s	38.10s	244.0s
	speed up	8.1%	1.0%	31.0%	3.1%	14.3%	5.7%	24.2%	4.5%
	algorithm	simple		simple		ring+light barrier		simple	
MT	NATIVE	14.46s	15.97s	15.30s	16.50s	44.10s	47.40s	22.50s	27.74s
	ROBUST	12.62s	13.96s	14.70s	16.20s	37.80s	41.10s	21.93s	27.22s
	speed up	14.5%	14.3%	4.1%	1.9%	16.7%	14.8%	2.6%	1.9%
	algorithm	pair+one barrier		pair+light barrier		pair+light barrier		ring+light barrier	

Table 10: Performance of application benchmarks

(Virginia Hydrodynamics) [29] is a multidimensional ideal compressible hydrodynamics code based on the Lagrangian remap version of the Piecewise Parabolic Method. The code uses *MPI_Alltoall* to perform matrix transposes and runs for 500 steps. The FFT-2D [27] program performs a two-dimensional Fast Fourier Transform on a 4K×4K complex matrix. In the evaluation, the code executes for 300 steps. Finally, the MT (Matrix Transpose) [14] is a simple program that uses *MPI_Alltoall* to perform matrix transpositions on a $4K \times 4K$ matrix. The code executes for 300 steps. Note that the dynamic adaptation mechanism in ROBUST takes 80 invocations to determine the best performing algorithm. The process arrival patterns in these four benchmarks on different platforms are very different. Hence, the performance of these benchmarks on different platforms gives good indications about the performance of ROBUST in practical situations.

Table 10 shows the results for the benchmarks. Both total application time and total all-to-all communication time are shown. We can see clearly that ROBUST significantly improves the communication time over NATIVE across the different applications on different platforms. For example, ROBUST achieves a communication time speed up of 55.74% and 31.45% for the VH-1 benchmark on LEMIEUX and AVIDD-T clusters, respectively. For the overall application time, the speed up depends on several factors, including (1) the percentage of all-to-all time in the total application time, and (2) how the all-to-all operation interacts with computation and other (collective) communications. In particular, the interaction among the all-to-all operation, the computation, and other collective communications can either offset or enhance the performance improvement: the improvement in the communication time may or may not transfer into an improvement in the total application time. For example, for VH-1 on Lemieux, the communication time is improved by 895 seconds, but the improvement in the total application time is only 250 seconds. On the other hand, for VH-1 on UC-TG, the communication time is improved by 175 seconds while the total application time is improved by 212 seconds. Such interaction among the all-to-all operation, computation, and other (collective) communications is not considered by our current implementation and needs to be investigated further. Nonetheless, ROBUST achieves noticeable improvement over NATIVE for total benchmark times in all cases. The table also shows the best performing algorithm selected by ROBUST. In many cases, the algorithms selected by ROBUST are different across different programs on different platforms. This show the importance of having multiple process arrival pattern aware algorithms to deal with different applications of different arrival patterns.

7.3 Summary

Although the native implementations of the *MPI_Alltoall* routine across the different platforms exploit features of the underlying network architecture, these routines do not perform as good as ROBUST in many cases. This can mainly be attributed to the fact that the native routines were designed without taking process arrival pattern into consideration. As such, they do not provide high performance for many practical cases. By explicitly considering process arrival pattern and employing a dynamic adaptive technique, more robust collective routines than the current ones can be developed.

8. CONCLUSION

In this paper, we investigate the process arrival patterns in a set of MPI benchmarks on two representative cluster platforms. We show that in such environments, it is virtually impossible for application developers to control process arrival patterns in their applications without explicitly invoking global synchronization operations and that process arrival patterns are likely to be imbalanced. Since the process arrival pattern has a significant impact on the performance of a collective communication algorithm, we conclude that MPI developers must take the process arrival pattern characteristics into consideration when developing MPI collective communication routines that can provide high performance in practical clusters. This study advocates further investigation for understanding the impact of process arrival patterns on different MPI collective operations and different collective communication algorithms and for identifying efficient process arrival pattern aware algorithms. The current understanding of MPI collective algorithms, which assumes a balanced process arrival pattern, is insufficient for developing routines that are efficient in practice. We demonstrate that when process arrival pattern aware algorithms for an operation are identified, a dynamic adaptive scheme can be used to realize robust collective routines that provide high performance across different applications and platforms.

Acknowledgment

This research is supported in part by National Science Foundation (NSF) grants CCF-0342540, CCF-0541096, and CCF-0551555. The experiments on the Lemieux cluster at Pittsburgh Supercomputing Center are sponsored through an NSF Teragrid grant CCR-050010T.

9. REFERENCES

[1] The AURORA cluster of University of Technology at Vienna, http://aurora.tuwien.ac.at.

[2] The AVIDD-T cluster, http://rac.uits.iu.edu/rats/research/avidd-t/hardware.shtml.

[3] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert, "Pipelined Broadcasts on Heterogeneous Platforms." *IEEE Trans. on Parallel and Distributed Systems*, 16(4), 2005.

[4] J. Bruck, C. Ho, S. Kipnis, E. Upfal, and D. Weathersby, "Efficient Algorithms for All-to-all Communications in Multiport Message-Passing Systems." *IEEE TPDS*, 8(11):1143-1156, Nov. 1997.

[5] W. E. Cohen and B. A. Mahafzah, "Statistical Analysis of Message Passing Programs to Guide Computer Design," In *Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences*, volume 7, pages 544-553, 1998.

[6] A. Faraj and X. Yuan, "Communication Characteristics in the NAS Parallel Benchmarks," In *Fourteenth IASTED International Conference on Parallel andDistributed Computing and Systems (PDCS 2002)*, pages 724-729, November 2002.

[7] A. Faraj and X. Yuan, "Automatic Generation and Tuning of MPI Collective Communication Routines," The *19th ACM International Conference on Supercomputing* (ICS'05), pages 393-402, Cambridge, MA, June 20-22, 2005.

[8] A. Faraj, X. Yuan, and D. K. Lowenthal, "STAR-MPI: Self Tuned Adaptive Routines for MPI Collective Operations," *The 20th ACM International Conference on Supercomputing (ICS'06)*, Queensland, Australia, June 28-July 1, 2006.

[9] J.Pjesivac-Grbovic, T. Angskun, G. Bosilca, G. E. Fagg, E. Gabriel, and J. Dongarra, "Performance Analysis of MPI Collective Operations," *IEEE IPDPS*, 2005.

[10] W. Gropp, E. Lusk, N. Doss, and A. Skjellum, "A High-Performance, Portable Implementation of the MPI Message Passing Interface Standard," In *MPI Developers Conference*, 1995.

[11] D. Lahaut and C. Germain, "Static Communications in Parallel Scientific Propgrams," In *Proceedings of the 6th International PARLE Conference on Parallel Architectures and Languages*, pages 262-276, 1994.

[12] LAM/MPI Parallel Computing. http://www.lam-mpi.org.

[13] LAMMPS: Molecular Dynamics Simulator, Available at http://www.cs.sandia.gov/ sjplimp/lammps.html.

[14] Matrix Transposition example, http://www.sara.nl/userinfo/reservoir/mpi/mpi-intro.

[15] MPICH - A Portable Implementation of MPI, Available at http://www.mcs.anl.gov/mpi/mpich.

[16] NASA Parallel Benchmarks, Available at http://www.nas.nasa.gov/NAS/NPB.

[17] Parallel NBody Simulations, Available at http://www.cs.cmu.edu/ scandal/alg/nbody.html.

[18] ParaDyn: Parallel Molecular Dynamics With the Embedded Atom Method, Available at http://www.cs.sandia.gov/ sjplimp/download.html.

[19] F. Petrini, D. J. Kerbyson, and S. Pakin, "The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8192 Processors of ASCI Q", *IEEE/ACM SC2003 Conference*, 2003.

[20] Pittsburg Supercomputing Center, Available at http://www.psc.edu/machines/tcs/lemieux.html.

[21] R. Rabenseinfner, "Automatic MPI counter profiling of all users: First results on CRAY T3E900-512," In *Proceedings of the Message Passing Interface Developer's and User's Conference*, pages 77-85, 1999.

[22] I. Rosenblum, J. Adler, and S. Brandon, "Multi-processor molecular dynamics using the Brenner potential: Parallelization of an implicit multi-body potential," *International Journal of Modern Physics*, 10(1): 189-203, February, 1999.

[23] P. Patarasuk, A. Faraj, and X. Yuan, "Pipelined Broadcast on Ethernet Switched Clusters." The *20th IEEE International Parallel & Distributed Processing Symposium* (IPDPS), April 25-29, 2006.

[24] T.B. Tabe, J.P. Hardwick, and Q.F. Stout, "Statistical analysis of communication time on the IBM SP2," *Computing Science and Statistics*, 27: 347-351, 1995.

[25] T. Tabe and Q. Stout, "The use of the MPI communication library in the NAS Parallel Benchmark," *Technical Report CSE-TR-386-99*, Department of Computer Science, University of Michigan, Nov 1999.

[26] R. Thakur, R. Rabenseifner, and W. Gropp, "Optimizing of Collective Communication Operations in MPICH," *ANL/MCS-P1140-0304*, Mathematics and Computer Science Division, Argonne National Laboratory, March 2004.

[27] Two-D FFT, http://www.mhpcc.edu/training/workshop/parallel_develop.

[28] The UC/ANL Teragrid cluster, http://www.uc.teragrid.org/tg-docs/user-guide.html.

[29] The Virginia Numerical Bull Session Ideal Hydrodynamics, http://wonka.physics.ncsu.edu/pub/VH-1.

[30] J. S. Vetter and A. Yoo, "An empirical performance evaluation of scalable scientific applications," In *Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, pages 1-18, 2002.