

Mapping Communication Layouts to Network Hardware Characteristics on Massive-Scale Blue Gene Systems

Pavan Balaji¹ · Rinku Gupta¹ · Abhinav Vishnu² · Pete Beckman³

Abstract For parallel applications running on high-end computing systems, which processes of an application get launched on which processing cores is typically determined at application launch time without any information about the application characteristics. As high-end computing systems continue to grow in scale, however, this approach is becoming increasingly infeasible for achieving the best performance. For example, for systems such as IBM Blue Gene and Cray XT that rely on flat 3D torus networks, process communication often involves network sharing, even for highly scalable applications. This causes the overall application performance to depend heavily on how processes are mapped on the network. In this paper, we first analyze the impact of different process mappings on application performance on a massive Blue Gene/P system. Then, we match this analysis with application communication patterns that we allow applications to describe prior to being launched. The underlying process management system can use this combined information in conjunction with the hardware characteristics of the system to determine the best mapping for the application. Our experiments study the performance of different communication patterns, including 2D and 3D nearest-neighbor communication and structured Cartesian grid communication. Our studies, that scale up to 131,072 cores of the largest BG/P system in the United States (using 80% of the total system size), demonstrate that different process mappings can show significant difference in overall performance, especially on scale. For example, we show that this difference can be as much as 30% for P3DFFT and up to twofold for HALO. Through our proposed model, however, such differences in performance can be avoided so that the best possible performance is always achieved.

1 Introduction

The massive-scale systems being deployed throughout the world will utilize a large amount of shared hardware including shared caches, memory, and network infrastructure. For example, systems such as the IBM Blue Gene (BG) [1]

This work was supported in part by the National Science Foundation Grant #0702182 and by Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

¹Mathematics and Computer Science Division, Argonne National Laboratory

E-mail: {balaji, rgupta}@mcs.anl.gov

²High Performance Computing Group, Pacific Northwest National Laboratory E-mail: abhinav.vishnu@pnl.gov

³Argonne Leadership Computing Facility, Argonne National Laboratory E-mail: beckman@mcs.anl.gov

and Cray XT [2] have started utilizing flat 3D torus networks where each node connects directly to its six neighbors. Therefore, unless each node communicates with only its physically nearest neighbors, it will be forced to share network links with other communication. Unless the physical placement of application processes matches the communication characteristics of the application, such sharing can result in significant communication contention and performance loss.

Communication libraries such as the Message Passing Interface (MPI) optimize application communication and process mapping at runtime using virtual topology functionality (such as Cartesian topologies or graph topologies). The idea of such functionality is that the application can specify the layout of the data it wants to process, and the runtime system can reorder process ranks while matching the data layout to the network hardware topology. A number of applications do not take full advantage of such functionality, however, mainly because it is not descriptive enough. That is, it does not allow the application to specify its exact communication pattern within the data layout.

For example, computational kernels such as partial differential equation solvers [3], molecular dynamics simulations [4], climate/ocean-modeling systems [5], and ray-tracing applications with domain-based parallelism [6] all use Cartesian grids of different dimensions and structure (2D, 3D, unstructured). Depending on the application and what it is trying to solve, however, the communication pattern might be completely different. For instance, partial differential equation solvers and climate/ocean-modeling systems perform nearest-neighbor communication along the edges or diagonals or both. Fast Fourier transforms (FFTs), on the other hand, communicate with all processes in each dimension. Ray-tracing applications with domain-based parallelism communicate with processes along their diagonals. With the current model of communication libraries, the application can specify that it want its processes to be laid out as a 2D or 3D

Cartesian grid, but it cannot specify that it will communicate with its neighbors along the edges, or neighbors along the diagonals, or all processes in its row/column, or any other communication pattern.

Thus, in practice, many applications try to manually map processes to the network topology. While this is a reasonable approach for small to medium scale topologies, it is not a feasible solution for massive-scale systems with hundreds of thousands of processors. The important aspect is that application developers know their applications communication pattern while the application is being launched. But, there currently exists no good way to provide this information to the process management system in order to optimize the process layout while taking network hardware characteristics into account.

In previous work [7, 8], we noticed the variance of communication performance with process mapping, and we studied the network congestion characteristics of flat torus networks that cause this behavior. In this paper, we extend that research by first performing a detailed analysis of the mapping of logical process layouts on the physical network topology. Then, we take advantage of this analysis to allow the application to describe its communication pattern, estimate the expected network contention through a simple first-order approximation model, and use this estimate to identify the right mapping to use. Moreover, we study the performance benefits such a model would provide, using evaluations with various micro-benchmarks and real application kernels including P3DFFT [9] and the HALO ocean-modeling kernel [10]. Our experiments, which scale up to 131,072 cores of the largest BG/P system in the United States (using 80% of the total system size), demonstrate that different process mappings can show significant difference in overall performance, especially on large-scale systems. For example, we show that this difference can be as much as 30% for P3DFFT and up to twofold for HALO. Through our proposed model, however, such differences in performance can be avoided, so that the best possible performance is always achieved.

2 Complexity in Process Mapping

In this section, we describe the complexity of process mapping on large-scale systems. Specifically, in Section 2.1, we describe the application viewpoint of application processes and their logical layout. In Section 2.2, we tie such logical mapping to the physical system topology for 3D torus-based systems.

2.1 Application Logical Process Layout

Most communication libraries, including MPI, do not expose the physical layout of the underlying system to the application in order to improve portability across different systems. Thus, applications form logical topologies of the

processes available to match the problem they are trying to solve (often depending on the data layout corresponding to the application science). For example, if we consider ocean/climate modeling and other computational fluid dynamics applications, the data representation is often a 2D plane, 3D volume, or even a multidimensional unstructured grid [11]. Thus, these applications form logical grids of the available processes, with each process having a part of the overall data. The computation on the local data, however, depends on the partially evaluated results from neighboring processes on the logical process grid formed by the application.

In other computational kernels such as FFTs [12, 9], the interaction is based on the Cartesian dimension. Processes are divided into a logical multidimensional grid (e.g., 2D grid). Each process initially has data corresponding to one dimension of the data grid. Once all processes are done computing their parts, the entire data grid is transposed, resulting in each process locally getting a different dimension of the data grid. Thus, a process interacts only with other processes in each dimension it is a part of. For example, in a 2D grid, each process interacts only with \sqrt{N} processes in each of the two dimensions.

Similarly, in ray-tracing applications with domain-based parallelism [6], while working on one particular ray, each process interacts only with other processes that share data corresponding to that ray. This is several times a diagonal group of processes in a 2D logical grid.

The common aspect in these and many other applications is that all of them rely on a logical layout of the processes without any information about the physical placement or mapping of the actual processes.

2.2 Process Mapping for BG/P

Optimal mapping of logical communication patterns onto the actual physical topology of the system is an intricate task. As the complexity of the communication patterns and system topology increases, this task becomes even more challenging. This section describes the mapping complexity for applications with different logical process layouts while being mapped onto a 3D physical system grid. The complexity is similar for other layouts such as Cartesian dimension-based communication and diagonal communication patterns, but these are not discussed here.

One-dimensional communication is one of the simpler patterns among nearest-neighbor communication patterns—every process communicates with its two logically neighboring processes. Consider mapping such an application having 512 processes onto 512 nodes connected in a $8 \times 8 \times 8$ 3D physical grid topology. Such a grid will have 8 nodes along each of the X, Y, and Z axes. The process management framework on BG/P allows application processes to be mapped onto the system in different ways, as represented

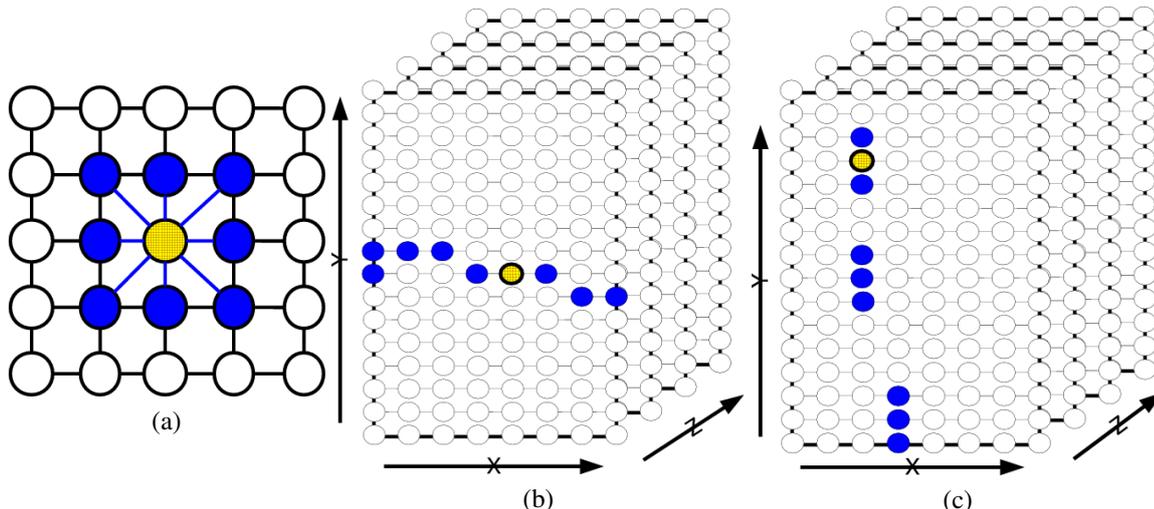


Fig. 1 2D Mapping for Nearest-Neighbor Communication: (a) Logical Mapping, (b) XYZ Physical Mapping, (c) YXZ Physical Mapping

by the mapping string; for example, with an XYZ mapping, processes are mapped to nodes along X-axis first, followed by Y-axis and Z-axis respectively. Thus, for a 1D logical nearest-neighbor communication, most of the application processes will get mapped next to each other in a sequential order.

Two-dimensional and three-dimensional communication patterns are more challenging to map optimally. Every process communicates with up to 8 (in 2D mapping) and 26 (in 3D mapping) logically neighboring processes.¹ Figure 1 illustrates such process mapping for 2D nearest-neighbor communication, while Figure 2 shows 3D mapping. Figure 2(a) shows a portion of the 3D logical grid as viewed by the application. Figures 2(b) and 2(c) show the process mapping of this logical grid on to a $8 \times 16 \times 8$ physical grid with XYZ and XZY mappings, respectively. As can be seen from these figures, communicating groups of processes can get dispersed throughout the physical network, resulting in increased communication overlap between process groups and performance degradation. More important, the dispersal of processes depends heavily on the process mapping (e.g., an XYZ mapping in these figures shows a different, and slightly smaller, dispersal as compared with the YXZ mapping), which causes different mappings to have different performance characteristics.

Based on these observations, from the system perspective we can conclude that the physical layout of the application processes (and thus the communication overhead) is highly dependent on the mapping type and the dimensions of the physical grid/torus.

While in theory any mapping of processes to cores is possible, different systems have different restrictions on what mappings they allow. For example, the IBM Blue Gene/P

system allows different combinations of process mappings along the X, Y, Z dimensions of the 3D torus. We note that for simplicity of description we have discussed process mappings only along the X, Y, Z dimensions of the 3D torus. However, BG/P systems consist of four cores per compute node, which can be considered to be a fourth dimension represented as “T.” Thus, an application can be mapped by using various mappings such as TXYZ, XYZT, and TYXZ.² Other mappings that do not form a symmetric ordering of ranks in one of these orders are not supported. Thus, the “best performance” we can achieve is artificially restricted by this requirement.

Table 1 Argonne BG/P System Torus Dimensions

Nodes	Dimensions
512	$8 \times 8 \times 8$
1024	$8 \times 8 \times 16$
2048	$8 \times 8 \times 32$
4096	$8 \times 16 \times 32$
8192	$8 \times 32 \times 32$
16384	$16 \times 32 \times 32$
32768	$32 \times 32 \times 32$

Further complicating this issue, applications on the BG/P system run on a subset of nodes defined by a “partition.” Two different partitions can have the same number of cores while having completely different dimension sizes. For example, partitions of both $8 \times 16 \times 32$ and $16 \times 16 \times 16$ dimensions can be used for a 4,096-process job. The dimensions of different partition sizes of the torus on the BG/P system can be configured by the system administrator and would determine how much network sharing, and consequently contention, each mapping would cause for any given partition

¹ Some applications treat diagonally intersecting processes as neighbors as well (star stencils), while some applications treat only processes intersecting along the edges as neighbors (box stencils).

² TXYZ indicates that Message Passing Interface processes are ordered with priority given first to cores of a node, followed by nodes on the X-axis, Y-axis, and Z-axis, respectively.

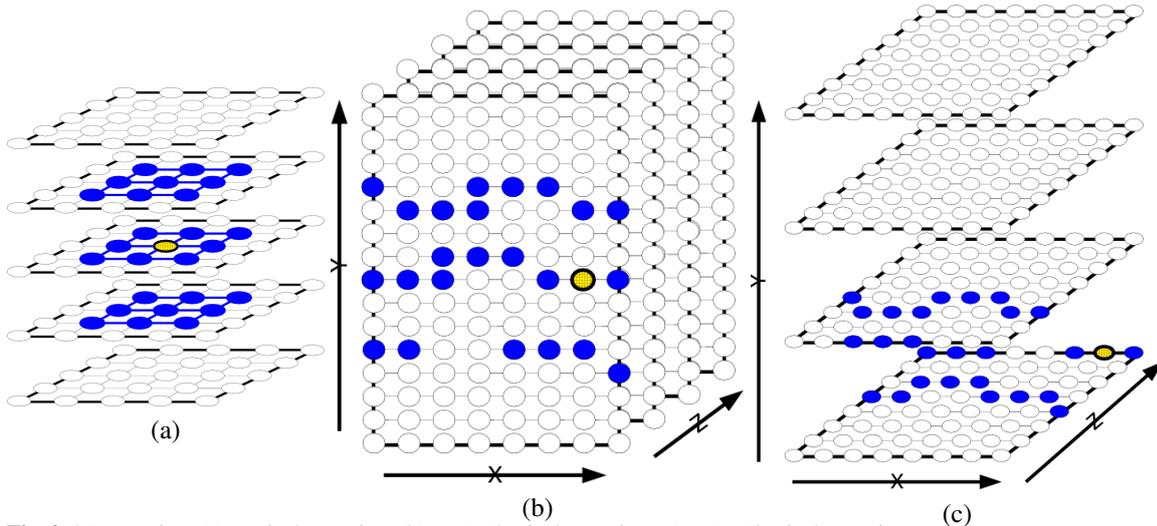


Fig. 2 3D Mapping: (a) Logical Mapping, (b) XYZ Physical Mapping, (c) XZY Physical Mapping

shape. Table 1 shows the dimensions of the torus for varying system sizes for the Argonne BG/P system.

3 Communication Contention Model

In this section, we describe our methodology for analyzing different communication patterns of applications and mapping them optimally.

The basic idea of our approach is to (1) understand the communication pattern of the application by allowing the application to describe it, (2) understand the physical platform network topology and the routing algorithm behavior on the platform, and (3) map these two categories of information to calculate the network contention for the given application pattern on the available network. The contention analysis model, as its output, predicts an optimal mapping that would give the best performance for a given application on a given platform. We have chosen the IBM BG/P system as an example platform for this paper. This information, consisting of application pattern and its optimal mapping, can then be integrated in the IBM BG/P job launching system. Applications can then specify their application pattern during job submission time, and the optimal mapping can be automatically and transparently picked up by the runtime system while launching the application.

Our current analysis is specific to symmetric communication patterns where all processes are involved in similar communication. Instances of such symmetric communication include 1D, 2D, and 3D nearest-neighbor communication patterns, which we analyze in this section.

3.1 Routing on Blue Gene/P

Our contention analysis model relies on the routing algorithm used by BG/P. Specifically, for small messages, BG/P uses a static routing algorithm in a dimension-wise order

(e.g., first along the X-axis, then the Y-axis, and finally the Z-axis). For large messages, BG/P uses a destination-based adaptive routing algorithm. To avoid livelocks, however, it always picks one of the minimum distance routes. In other words, at each hop, the adaptive routing algorithm considers only the outgoing links that reduce the hop count to the destination node. Among the outgoing links being considered, it selects the one with the least amount of other ongoing traffic.

As a first-order approximation, we assume that the data packets in a large message are split equally on all possible paths at each hop. This approximation allows us to estimate the amount of data traffic on each hop caused by the communication between any pair of processes.

We, currently, consider only symmetric communication patterns in our research. We note that every partition on the BG/P is fully symmetric and forms a torus on all dimensions. This symmetry allows us to model the contention caused by a single process for its communication with all its peers, and simply extend it to all the processes in the entire system.

3.2 Example Contention Analysis

To analyze the optimal topology for a given physical partition layout and application, our model needs to understand the contention that may arise because of shared network links during the application communication. We explain the logic used in our model through an example.

Consider a 128×128 2D process grid (16,384 cores). On BG/P, this will require 4,096 quad-core nodes. The partition dimensions for 4,096 nodes on the Argonne BG/P are $8 \times 16 \times 32$ along the X-, Y-, and Z-axis, respectively.

Let us analyze the TXYZ mapping of processes, with a row-major mapping of the 2D 128×128 processor logical grid. Each row of 128 processors of the logical grid

will occupy the four rows of the X-axis (8 nodes; i.e., 32 cores lie along the X-axis). Now consider a node communicating with its eight nearest neighbors. Its neighbors on either side of it, on the same row, will lie either on the same physical node or on a neighboring node (since all cores of a node are allocated before the next node), thereby incurring minimal communication overhead (no link contention). The three neighbors on the row above it will lie on nodes on the same X-axis at a distance of four hops (since one row of the 2D grid spans four rows of the physical torus) from its node. Thus, each link along these four hops will be used three times by this communicating center node. Each of these links along the four hops will be used by other sets of nearest-neighbor communicating nodes as well, in both directions bringing its average contention count to 24.³

In a TZXY mapping, on the other hand, the Z-axis has 32 nodes (i.e., 128 processors along its axis). Thus, a single row of the 2D grid can map along a single Z-axis plane. Analyzing this reveals that for every node, the neighbors above or below a process are mapped to a single node, one hop away. On average, the bidirectional contention count for every link is 6. Similarly an XYZT mapping would have an average bidirectional contention count for every link of about 10.5. Thus, contention seen by the TZXY mapping is the least, while that seen by TXYZ is the highest. This analysis is substantiated by the results presented in Section 4. Note that for completely symmetric communication patterns such as nearest neighbor, it is sufficient to evaluate the contention caused by one process.

We performed similar analyses for other communication patterns such as dimension-wise communication (used by 3DFFT libraries) and diagonal communication (used by ray-tracing applications). However, descriptions for those have been omitted because of space restrictions.

Contention counts will vary with the type of mapping, type of communication, partition size of the system, logical grid size of the application and order (row major, column major, etc.) in which logical grid rows are mapped onto the partition cores. In the rest of the paper, because of lack of space, we do not analyze contention count further. Rather, we demonstrate the results obtained from our experiments on the BG/P system.

4 Experiments and Analysis

In this section, we discuss the impact of various mappings on different micro-benchmarks and applications. Our experiments also demonstrate the performance achieved by the model described in Section 3 (called ‘Contention Detection Model’ and represented by “CDL” in the charts), which is automatically and transparently picked by the runtime system when launching the application.

³ “Contention count” can be considered to be an indication of how many communication streams are going over a link at the same time.

4.1 Microbenchmark-Based Evaluation

We first evaluate 2D and 3D logical nearest-neighbor communication. Both experiments use star stencils, so each process has $(3^d - 1)$ neighbors, where d is the logical process grid dimensionality.

2D Nearest-Neighbor Communication: In this experiment, each process does point-to-point communication exchanging some data with its logical neighbors. In real application kernels, this data would typically correspond to the *ghost cells* that form the bordering data points between the two processes. Figure 3(a) shows the communication performance for a system size of 16,384 processes, and Figure 3(b) shows the same for a system size of 65,536 processes.

From the figure, we see that the difference in performance between the various mappings increases with both message size (as shown in each subfigure) and system size (comparing the two subfigures). In fact, for a system size of 65,536 cores and a message size of 1 MB, we notice almost a *sevenfold* difference in performance between the different mappings. This behavior is expected; as the message size increases, each network link is used for longer periods of time, increasing the possibility of congestion and performance degradation. Similarly, as the system size increases, the number of communication flows increases proportionally with the number of *pairs* of processes available, that is, $O(N^2)$, where N is the number of processes in the system. But the number of network links increases only as $O(N)$; this situation leads to more congestion and performance degradation. The results in Figure 3(a) also show that CDL model can determine which layout has the least contention (in this case TZXY) and automatically pick that for the user. These results are similar to those reported in Section 3.

3D Nearest-Neighbor Communication: This experiment is similar to the 2D nearest-neighbor experiment except that processes are logically laid out in three dimensions. Thus, each process has 26 logically nearest neighbors with which it performs point-to-point communication. Because of space restrictions, we simply point out, as can be seen in Figure 4(a) and Figure 4(b), that the overall trend is similar to that of the 2D nearest-neighbor benchmark: (1) there is a significant performance difference between different mappings, (2) the performance difference increases with message size and system size, and (3) the CDL model allows the system to automatically pick the best mapping. In fact, for 65,536 processes and a message size of 1 MB, we notice an order-of-magnitude difference in performance.

4.2 Evaluation of Application Kernels

Next, we use two application kernels, P3DFFT and Halo, to analyze the impact of process mapping on overall performance.

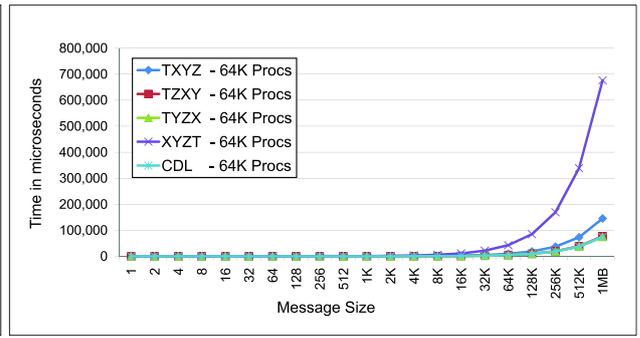
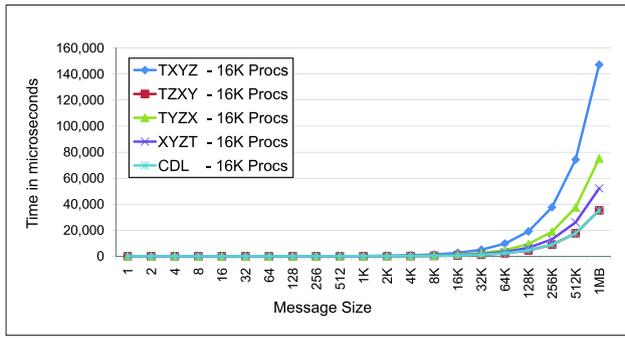


Fig. 3 2D Communication Benchmark Performance: (a) 16K Cores, (b) 64K Cores

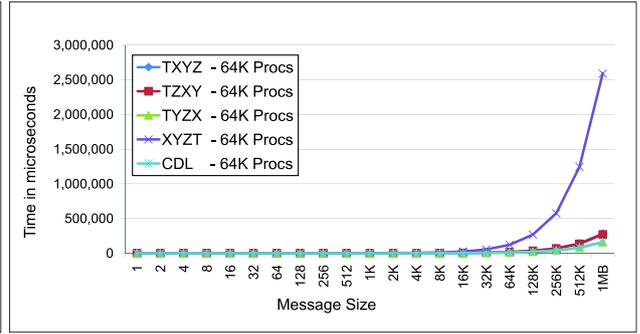


Fig. 4 3D Communication Benchmark Performance: (a) 16K Cores, (b) 64K Cores

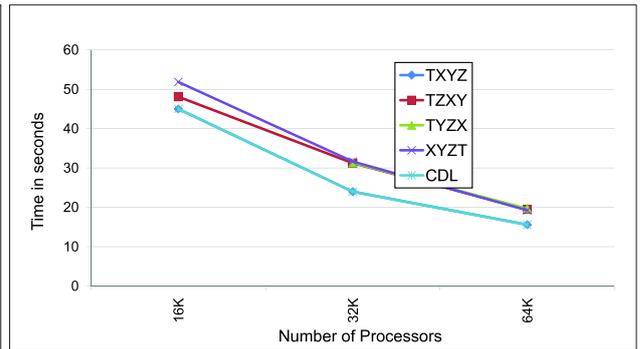
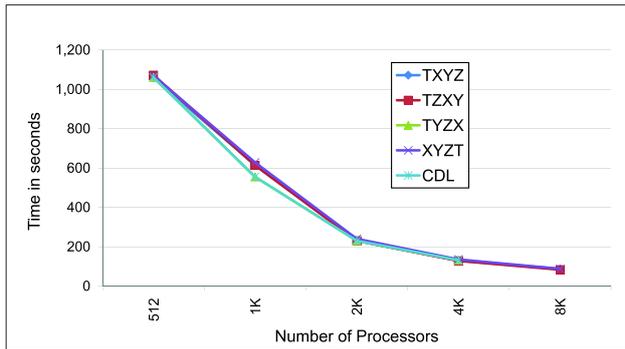


Fig. 5 (a) P3DFFT Performance: (a) 512 to 8K Cores, (b) Greater than 8K Cores

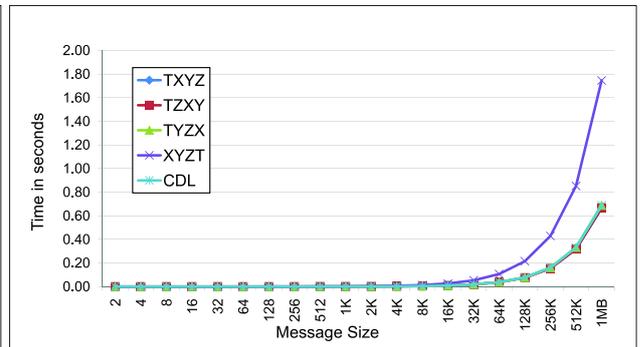
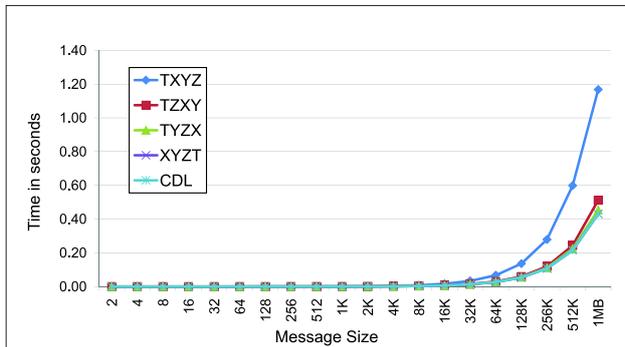


Fig. 6 Halo Application on (a) 32K Cores, (b) 128K Cores

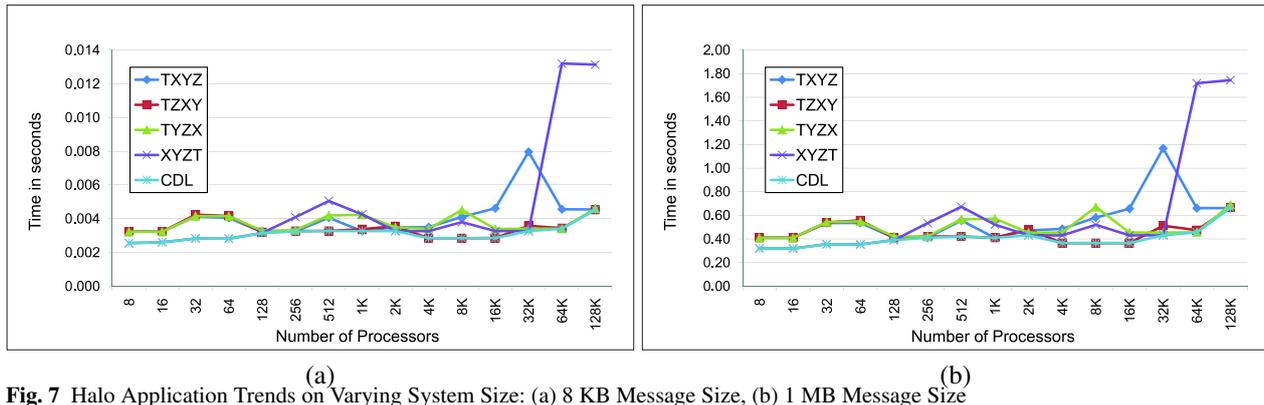


Fig. 7 Halo Application Trends on Varying System Size: (a) 8 KB Message Size, (b) 1 MB Message Size

P3DFFT Library: FFT [13] is an efficient algorithm to compute the discrete Fourier transform (DFT) and its inverse. FFT, as applied to a three-dimensional volume of space, is typically referred to as 3DFFT. The goal of 3DFFT is essentially to perform 1D Fourier transform on each of the three dimensions of the 3D data mesh. Most parallel 3DFFT implementations rely on a sequential version of 1DFFT (that performs the transform on one dimension at a time) and then transpose the data grid when needed. Multiple implementations of parallel 3D FFT are available today. P3DFFT [14, 15] is a popular implementation of 3DFFT used in various application domains, including digital speech and signal processing, solving partial differential equations, molecular dynamics simulations, and Monte Carlo simulations.

Figure 5 shows the performance of P3DFFT with increasing system size for different process mappings. For visibility, we split the data into two graphs—Figure 5(a) shows the performance for small system sizes (up to 8,192 cores), while Figure 5(b) shows the performance for large system sizes (more than 8,192 cores). As shown in the figures, for small systems, process mapping has no significant impact. When the system size increases beyond 8,192 cores, however, we do notice a performance difference between the different mappings. In fact, this difference is as high as 25% for 65,536 cores, which makes the Contention Detection Model’s (CDL) benefits substantial for large-scale systems.

The Halo Application Kernel: The NRL Layered Ocean Model (NLOM) [16] simulates semi-enclosed seas, major ocean basins, and global ocean. The current implementation of the model uses a tiled data-parallel programming style. Its general nature allows implementations in various programming models including MPI, OpenMP, Co-Array Fortran, and shared memory. This makes NLOM a good candidate for benchmarking both hardware and the associated communication software. The Halo benchmark [10] simulates an NLOM 2-D exchange for an $N \times N$ subdomain for different values of N . In other words it performs nearest-neighbor exchanges from a 2D array. Each of these

exchanges is performed by using several algorithms, different for each programming technique. When Halo is used on a smaller number of nodes, the performance is affected predominantly by the system latency, whereas for larger node counts the bandwidth becomes a dominating factor [10].

Figure 6 shows the performance of Halo with varying message size, for different process mappings (system sizes of 32,768 cores and 131,072 cores). Again, we notice a significant improvement in overall performance for the CDL model, depending on the mapping. For 131,072 cores, we notice more than *twofold* difference in performance. Figure 7 shows the performance of Halo with varying system sizes (message sizes of 8 KB and 1 MB). We notice three trends: (1) no single mapping performs the best in all cases, which is also true for all other results shown in the paper; (2) the overall performance difference between the best mapping and the worst mapping generally increases with system scale; and (3) the CDL model picks the best mapping in all cases.

5 Related Work

In the past few years, a lot of research has been devoted to studying techniques for mapping application topology (tasks) onto the system’s actual physical topology [17–20] and to optimizing these techniques [21–23]. Research [24–27] provides different methods for optimizing problem layout on BG/L. All this research assumes that applications can describe their communication patterns effectively to the runtime system; however, this is not currently the case. Most mapping techniques assume a possible communication pattern based on the data layout and optimize the process layout accordingly.

The NAS Parallel Benchmarks [28] were evaluated with different mappings such as Gray codes and permutations of the X, Y, Z coordinates on the BG/L systems spanning up to 256 processors. While that work is similar to a subset of research proposed here, it was done only with microbenchmarks and on very small system size. None of the prior work has shown these effects at the scales we have shown. Our ex-

periments are done on a scale nearly 100 times larger than the previous papers (80% of the largest BG/P system in the United States). Since network sharing is going to increase in future exascale systems, scale of the applications is key to understanding network saturation behavior.

Work also has been done in the context of STAR-MPI, which focuses on adaptive techniques for various collective operations and utilizes the most optimal communication algorithm [29]. While such tools can help choose the best communication algorithm for a given system process mapping, they do not explicitly try to understand the network congestion behavior caused by different mappings; thus, they are mostly complementary to the work done in this paper.

6 Concluding Remarks

We analyzed the impact of different process mappings on application performance on a massive BG/P system. Using this analysis, we propose a communication description model/language (CDL) that applications can use to specify their communication patterns prior to being launched. We demonstrated the performance benefits the CDL model can provide, using microbenchmarks and real application kernels including P3DFFT and the Halo ocean-modeling kernel. Our experiments, which scaled up to 131,072 cores of the Blue Gene/P system indicate that different process mappings can show significant difference in overall performance, especially on large-scale systems. Through our CDL model, however, one can avoid these differences in performance, thereby obtaining the best possible performance.

References

1. IBM Blue Gene Team. Overview of the IBM Blue Gene/P project. *IBM J. Res. Dev.*, 52(1-2):199–220, 2008.
2. Cray Research, Inc. *Cray T3D System Architecture Overview*, 1993.
3. Argonne National Laboratory. PETSc. <http://www.mcs.anl.gov/petsc>.
4. S. Kumar, C. Huang, G. Almasi, and L. V. Kale. Achieving strong scaling with NAMD on Blue Gene/L. In *IEEE International Parallel and Distributed Processing Symposium*, 2007.
5. Naval Research Laboratory. Naval Research Laboratory Layered Ocean Model (NLOM). <http://www.navo.hpc.mil/Navigator/Fall99.Feature.html>.
6. C. Rabiti, M. A. Smith, D. Kaushik, W. S. Yang, and G. Palmiotti. Parallel Method of Characteristics on Unstructured Meshes for the UNIC Code. In *PHYSOR*, Interlaken, Switzerland, Sept. 14–19 2008.
7. P. Balaji, A. Chan, R. Thakur, W. Gropp, and E. Lusk. Toward Message Passing for a Million Processes: Characterizing MPI on a Massive Scale Blue Gene/P. *Special edition of the Springer Journal of Computer Science on Research and Development (presented at the International Supercomputing Conference (ISC)); Best Paper Award*, 2009.
8. P. Balaji, H. Naik, and N. Desai. Understanding Network Saturation Behavior on Large-Scale Blue Gene/P Systems. In *Proceedings of the International Conference on Parallel and Distributed Systems (ICPADS)*, Shenzhen, China, Dec. 8–11 2009.
9. D. Pekurovsky. P3DFFT webpage, Feb. 2009. <http://www.sdsc.edu/us/resources/p3dfft/index.php>.
10. A. J. Wallcraft. The Halo Benchmark. <http://www.navo.hpc.mil/Navigator/PDFS/Fall1999.pdf>.
11. P. Fischer, J. Lottes, D. Pointer, and A. Siegel. Petascale Algorithms for Reactor Hydrodynamics. *JPCS*, 125(1), 2008.
12. M. Frigo and S. G. Johnson. The Design and Implementation of FFTW3. *Proceedings of the IEEE*, pages 216–2331, 2005.
13. James W. Cooley and John W. Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. *MCOM*, 19(90):297–301, 1964.
14. San Diego Supercomputing Center. P3DFFT. <http://www.sdsc.edu/us/resources/p3dfft/>.
15. A. Chan, P. Balaji, W. Gropp, and R. Thakur. Communication Analysis of Parallel 3D FFT for Flat Cartesian Meshes on Large Blue Gene Systems. In *Proceedings of the IEEE/ACM International Conference on High Performance Computing (HiPC)*, Bangalore, India, Dec. 17–20 2008.
16. A. J. Wallcraft. The NRL Layered Ocean Model Users Guide. NOARL Report 35, Naval Research Laboratory, Stennis Space Center, MS, 1991.
17. J. Traff. Implementing the MPI Process Topology Mechanism. In *SC*, pages 1–14, 2002.
18. J. Hur. An Approach for Torus Embedding. In *ICPP*, page 301, Washington, DC, USA, 1999. IEEE Computer Society.
19. C. Ou and S. Ranka and G. Fox. Fast and Parallel Mapping Algorithms for Irregular Problems. *TJS*, 10(2):119–140, 1996.
20. S. Bokhari. On the Mapping Problem. *TC*, 30(3):207–214, 1981.
21. S. W. Bollinger and S. Midkiff. Heuristic Technique for Processor and Link Assignment in Multicomputers. *TC*, 40(3):325–333, 1991.
22. N. Mansour, R. Ponnusamy, A. Choudhary, and G. C. Fox. Graph Contraction for Physical Optimization Methods: A Quality-cost Tradeoff for Mapping Data on Parallel Computers. In *ISC*, pages 1–10, New York, NY, USA, 1993. ACM.
23. T. Chockalingam and S. Arunkumar. Randomized Heuristics for the Mapping Problem. The Genetic Approach. In *Parallel Computing*, pages 1157–1165, 1992.
24. G. Bhagnot, A. Gara, and P. Heidelberger et al. Optimizing Task Layout on the Blue Gene/L Supercomputer. In *IBM J. Res. Dev.*, volume 49, pages 489–500, 2005.
25. G. Almasi, C. Archer, J. Castanos, et al. Implementing MPI on the BlueGene/L Supercomputer. In *Euro-Par*, pages 833–845, 2004.
26. H. Yu, I. Chung, and J. Moreira. Topology Mapping for Blue Gene/L Supercomputer. In *SC*, page 116, New York, NY, USA, 2006. ACM.
27. T. Agarwal, A. Sharma, A. Laxmikant, and L. V. Kale. Topology-aware task mapping for reducing communication contention on large parallel machines. *IPDPS*, 0:122, 2006.
28. B. Smith and B. Bode. Performance Effects of Node Mappings on the IBM BlueGene/L Machine. In *Euro-Par*, pages 1005–1013, 2005.
29. A. Faraj, X. Yuan, and D. Lowenthal. STAR-MPI: Self Tuned Adaptive Routines for MPI Collective Operations. In *Proceedings of the 20th Annual International Conference on Supercomputing (ICS)*, pages 199–208, Cairns, Queensland, Australia, 2006.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.