



Kernel-assisted and topology-aware MPI collective communications on multicore/many-core platforms



Teng Ma^{a,*}, George Bosilca^a, Aurelien Bouteiller^a, Jack J. Dongarra^{a,b,c}

^a University of Tennessee, Knoxville, TN, USA

^b Oak Ridge National Laboratory, Oak Ridge, TN, USA

^c University of Manchester, Manchester, UK

ARTICLE INFO

Article history:

Received 10 September 2012

Received in revised form

28 January 2013

Accepted 29 January 2013

Available online 11 February 2013

Keywords:

MPI

Multicore

Cluster

HPC

Collective communication

Hierarchical

ABSTRACT

Multicore Clusters, which have become the most prominent form of High Performance Computing (HPC) systems, challenge the performance of MPI applications with non-uniform memory accesses and shared cache hierarchies. Recent advances in MPI collective communications have alleviated the performance issue exposed by deep memory hierarchies by carefully considering the mapping between the collective topology and the hardware topologies, as well as the use of single-copy kernel assisted mechanisms. However, on distributed environments, a single level approach cannot encompass the extreme variations not only in bandwidth and latency capabilities, but also in the capability to support duplex communications or operate multiple concurrent copies. This calls for a collaborative approach between multiple layers of collective algorithms, dedicated to extracting the maximum degree of parallelism from the collective algorithm by consolidating the intra- and inter-node communications.

In this work, we present HierKNEM, a kernel-assisted topology-aware collective framework, and the mechanisms deployed by this framework to orchestrate the collaboration between multiple layers of collective algorithms. The resulting scheme maximizes the overlap of intra- and inter-node communications. We demonstrate experimentally, by considering three of the most used collective operations (Broadcast, Allgather and Reduction), that (1) this approach is immune to modifications of the underlying process-core binding; (2) it outperforms state-of-art MPI libraries (Open MPI, MPICH2 and MVAPICH2) demonstrating up to a 30x speedup for synthetic benchmarks, and up to a 3x acceleration for a parallel graph application (ASP); (3) it furthermore demonstrates a linear speedup with the increase of the number of cores per compute node, a paramount requirement for scalability on future many-core hardware.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

While the insatiable demand for increasing computing power from the domain sciences motivates the deployment of powerful High Performance Computing (HPC) systems, thermal and power consumption concerns have curbed the growth of both compute node count and processor frequency. As an alternate source of processing power, multicore clusters have become the most prominent form of HPC systems, and exhibit a rapid increase in the number of cores per compute node. The top ranking machine in the August 2012 Top500 list, the LLNL Sequoia computer, uses more than one and a half million cores.¹ Processors with 8–12 cores are the norm today, and it is not uncommon to have such processors

deployed in multiple socket boards featuring 8–96 cores, with network-style interconnection between caches or to the memory banks (e.g. Intel QPI or AMD Hyper-transport). Unfortunately, this new hardware trend challenges the assumptions made by most current HPC programming models, directly threatening the performance efficiency of the machines. Namely, within compute nodes, non-uniform memory accesses (NUMA), memory and shared cache hierarchies, weaken the assumptions of regular load balance and uniform link bandwidth and latency.

In the era of the single-core cluster, the Message Passing Interface (MPI) standard has enjoyed a wide adoption in the HPC community, thanks to two key features: implementations could provide the highest level of performance while maintaining the application's portability. With respect to portability, not only an MPI code compiled on different machines, but it also exhibits an excellent efficiency, because network topologies and collective patterns of communications are accounted for by the MPI library rather than the application itself. With the introduction of multicore compute nodes, both of these features

* Corresponding author.

E-mail address: xiaok1981@gmail.com (T. Ma).

¹ <http://top500.org/lists/2012/06>.

have been threatened in MPI, most implementations treating multicore compute nodes as mere SMP units and ignoring their internal hierarchies. To alleviate these issues, some attempts have been made to use hybrid programming models, retaining MPI between computer nodes and a thread-centric approach (pthreads, OpenMP, TBB, . . .) between cores. The suitability of this approach is questionable, with research showing a similar number of applications that benefited from the approach compared with failures to reach any performance improvement. Moreover, from the productivity point of view several drawbacks are evident: it imposes a significant complexity on programmers, renders explicit the management of hierarchies which defeats performance portability, and imposes a major rewrite of legacy applications. We believe that the MPI standard is a competitive proposition for harnessing the power of multicore clusters, should the implementations of the standard use the proper techniques to account for core and memory link properties, especially in the area of collective communications.

Indeed, recent advances in MPI collective communications have already demonstrated that the performance issues incurred by multicore memory hierarchies can be solved on shared memory multicore computer nodes. The careful mapping between the collective topology and the core distance [16], and the use of single-copy kernel assisted mechanisms deep inside the collective algorithms [15] have been proven to greatly increase the shared memory communication efficiency. However, on distributed memory machines, like clusters of multicores, a single approach cannot encompass the extreme variations not only in the bandwidth and latency capabilities, but also in features such as the aptitude to operate multiple concurrent copies. Efficient multicore shared memory approaches are so specific, including kernel assisted copies, that they cannot apply to network communications; on the other hand, regular network approaches fail to extract performance off shared memory links. This calls for a collaborative approach between multiple layers of collective algorithms, dedicated to managing intra and inter computer node communications.

In this paper, which is an extension of our distinguished work [8], we present how HierKNEM, a kernel-assisted topology-aware collective framework, orchestrates the collaboration between multiple layers of collective algorithms. Leaders are selected among the core-centric collective algorithm, to participate in the inter-node collective topology. Intra-node communications are managed by offloading memory copies to non-leader processes, taking advantage of the kernel-assisted single-copy approach to balance the memory copy load among available cores. The resulting scheme enables perfect overlap of intra-node communication with inter-node communications, thanks to innovative hierarchical algorithms. We demonstrate experimentally, by considering three distinct collective patterns (one-to-many, many-to-many and many-to-one), that (1) this approach is immune to modifications of the underlying process-core binding; (2) it outperforms state-of-art MPI libraries (Open MPI, MPICH2 and MVAPICH2) demonstrating up to a 30x speedup for messages between 8 and 256 KB in synthetic benchmarks, and up to 3x speedup for a parallel graph application (ASP); (3) it demonstrates a linear speedup with the increase of the number of cores per computer node, a paramount requirement for scalability on future many-core hardware.

The rest of this paper is organized as follows: Section 2 introduces related work on current efforts to optimize MPI collective communication on multicore clusters and the application of a kernel-assisted approach into MPI libraries. Then, Section 3 describes the framework for kernel-assisted hierarchical collective communications on clusters of multicore and details three collective algorithms: one-to-all (Broadcast), all-to-one (Reduce), all-to-all (Allgather), and their corresponding implementations in a new

Open MPI collective component: HierKNEM. These algorithms are experimentally compared with state-of-the-art MPI implementations to assess the benefits of the hierarchical approach in Section 4. Finally, Section 5 concludes the paper with a discussion of the results.

2. Related work

The legacy approach to implement collective communication is to adopt one of many different communication topologies (linear, chain, split binary tree, binomial tree, etc.) [9]. These basic approaches can be refined by enabling parallel treatment through message pipelining, a technique in which large messages are split into smaller chunks to maximize steady-state bandwidth. Furthermore, a runtime decision module can be used to select the best algorithm and tuning parameters, according to message size, communicator size, and other input variables [6]. The Tuned collective module, in Open MPI, is iconic of such an approach; while MPICH2 and other MPI libraries feature a similar approach restricted at the compilation time. Unfortunately, while mapping communications to specific network topologies yield drastic performance improvement on single-core architectures, the increase in the number of cores renders this problem more complex, requiring additional parameters that could reflect the runtime processes' topology in view of physical distances. To further exacerbate this issue, intelligent process placements, used as a bridge between applications and MPI libraries, e.g. MPIPP [5], have the tendency to exacerbate the amount of point-to-point communications between specific processes as they consider the collective communication as a simple set of point-to-point communications. This triggers two issues: on one hand it imposes the preselected topology for all subsequent runs and computes the neighborhood relationship based on a communication pattern likely to change with the execution environment. This irregular mapping leads to a further mismatch between collective topologies and underneath hardware [16].

The conventional effort toward adapting collective communications to hierarchical hardware topologies is leader-based hierarchical algorithms [14,20,22,17,12,18]. Earlier attempts toward hierarchical approaches on collective communication on clusters of SMPs [14] or Grids [13] focused on reducing the amount of data or the number of messages crossing the low bandwidth or high latency links, respectively. Combining with the SMP-aware method, leader-based hierarchical algorithms were widely applied to all collective communications patterns, e.g., MPI on Quadrics networks [20], Open MPI's Hierarch collectives, or MVAPICH2 on Infiniband networks [17,12,18]. In these SMP-aware methods, multicore compute nodes are often treated as shared memory nodes without internal hierarchy. As a consequence, the layered collective components that handle inter and intra-node communications do not cooperate tightly, leading to suboptimal pipelining and sometimes contradictory tuning choices. This results in another difference with our proposed work: the intra-node communication is mainly implemented by a copy-in/copy-out approach using a shared memory segment.

The copy-in/copy-out approach requires two memory copies for each message, greatly wasting memory bandwidth and CPU cycles. When applying this approach into leader-based hierarchical algorithms, leader processes are heavily involved in intra-node data movement [15], resulting in serializing the inter- and intra-node communications. Most of the intra-node communication overhead accumulates and results in a significant overhead that cannot benefit from overlap by inter-node communications. Obviously, such overhead is bound to increase with the number of cores; the copy-in/copy-out at the leader process has to be sequentially executed once for each of the processes participating in the collective communication.

To reduce the overhead from double memory copies in the copy-in/copy-out approach, one-sided single-copy methods have been proposed. SMARTMAP [2,1] is an effort to make use of a simple page table management in catamount systems to implement single-copy intra-node communication. Another direction is the kernel-assisted approach such as LiMIC [11] in MVAPICH2 or KNEM [4] in MPICH2 and Open MPI. This kernel-assisted approach has been widely used to speed up large messages' point-to-point communication on shared memory machines [4]. Furthermore, an intra-node collective communication component, KNEM collective [15], is implemented into Open MPI, based directly on the KNEM copy and not implemented over KNEM-enabled point-to-point communication. The KNEM collective harnesses KNEM's single-copy and direction control techniques to offload memory copies to non-root processes, providing a significant performance boost [15]. Also, efforts have been made to take into account NUMA hierarchies in the process placement and to optimize intra-node collective algorithms to include architectural features [16]. However, these projects focus solely on improving communications within a single shared memory multicore compute node. The aspects regarding cooperation of these complex algorithms with the inter-node layer of the collective communication have not been addressed so far. There is an obvious need to develop algorithms that encompass both layers and account for all particularities and varieties of hardware.

3. Collective algorithm composition

3.1. Framework

As hinted previously, most existing approaches to develop hierarchical collective communications are based on a multi-level approach where the top level represents the largest area network, and each subsequent level is for a smaller area network. While they provide interesting performance compared with single-level approaches, they do not benefit from the entire overlapping potential of collective algorithms, as the transition processes (i.e. processes that are leafs in one level and become root on the next), are step by step blocked in a collective for a particular level. What has been missing in these attempts at providing hierarchical collective operations on clusters of multicore system was the ability to express a multi-level algorithm with a very tight level of interoperability between the levels. In the present effort, we want to enable an unprecedented level of integration between different algorithms, by dissolving the boundaries between the levels, and allowing the transition processes to overlap collective between the inter and intra levels.

From a technical point of view, in most of the hierarchical approaches including ours, collective communication is divided between inter- and intra-node communication. Each process has an intra-node communicator encompassing all processes hosted on the same physical compute node. Among these local processes, a leader process is selected to represent the compute node in the inter-node layer. All non-leader processes only communicate with the local leader process and then messages are forwarded by the leader process to remote leader processes on remote compute nodes. The advantage is that the messages carried through expensive inter-node links are explicit, giving leverage for the algorithm composition to minimize cross-traffic volume. From a technical standpoint, what differentiates our approach compared to previous attempts is the level of integration between the layers of the hierarchy, allowing multiple algorithms to coordinate their pipelining strategies at a very low level.

One major challenge for multi-level algorithms is to coordinate around the usage of common resources. In this particular instance, one should pay attention to the load imposed on the memory

bus. This load is two-fold: on one side sending/receiving data over the network translates into moving data across the PCI bus from the memory bus. On the other side, moving data inside the compute node generates memory bus traffic, and therefore collides with the network transfer (the data in Fig. 4 highlight this fact). Therefore, special care has been taken to minimize the number of memory transfers at the inter-node level. The approach chosen in this framework is to base all intra-node memory transfers on the KNEM collective components, described in [15]. KNEM's offloading capability is naturally matched up with leader-based hierarchical collectives: workloads of memory copies can be off-loaded onto non-leader processes. Non-leader processes can simultaneously read or write leader processes' memory through KNEM primitives; meanwhile, leader processes can dedicate themselves to inter-node forwarding, without sequentialization experienced by less integrated hierarchical approaches. For communication strictly within large NUMA compute nodes, different approaches yield varying performance. Our new hierarchical algorithms leverage the knowledge accumulated on a single compute node [15] to design sound algorithm compositions that can cope with a large number of cores within compute nodes. The experimental section demonstrates how well these approaches collaborate with another layer.

In this new context, we provide three improved versions of the most used collective communications: a one-to-many (Broadcast), a many-to-many (Allgather) and a many-to-one (Reduce).

3.2. Broadcast

Let us assume the intra-node communicator for each compute node is *lcomm*, the inter-node communicator for leader processes is *lcomm*, and process rank is *P*. Suppose a two-level broadcast algorithm, using a spanning tree-based approach for the inter-node level and a linear approach for the intra-node level. Our HierKNEM broadcast algorithm is adaptive enough to handle special cases, e.g. when all processes are allocated on a single compute node, our broadcast is transformed into a linear algorithm identical to the KNEM one; when each compute node has a single process in the communicator, our HierKNEM broadcast is automatically morphed into a spanning tree broadcast identical to the inter-node level.

Algorithm 1 presents the pseudo-code of the HierKNEM Broadcast. In order to save space, we trimmed the pseudo-code handling the special cases mentioned above and presented the algorithm processing a general case: each compute node has more than one process bound to different cores and all leader processes are organized into a spanning tree with more than two levels: a root node, intermediate nodes, and leaf nodes. At first, each leader process registers 'rbuf' into KNEM device and gets a 'cookie' back at step 2. This cookie is a unique identifier to point to an entry recording rbuf's physical memory address, and any other process in the compute node having this identifier can access (based on the granted right) this registered buffer via the KNEM module. This cookie will then be broadcast to all non-leader processes on the same compute node (steps 3 and 33). Afterward the message is divided into equal-sized fragments and forwarded in a pipelining fashion along the spanning tree composed of all leader processes (between steps 4 and 29). In this particular context, father and children mentioned in Algorithm 1 refer to the process up and down the spanning tree from the current process *P*. For intermediate and leaf nodes in the spanning tree, once the leader processes receive a segment from its father node, they will notify all non-leader processes on the same compute node to fetch the segment (steps 15 and 21). Upon receiving this notification at step 38, each non-leader process will fetch the segment by a KNEM *get* operation at step 39. This *get* operation is one-sided and will be offloaded to the non-leader processes. Therefore the overhead of

```

Input: MPI_Bcast(void *rbuf, int count, MPI_Datatype dtype, int root, MPI_Comm comm)
1 if  $P$  is leader process then
2   Register rbuf into KNEM device and get a cookie;
3   Broadcast this cookie to all non-leader processes on the same compute node;
4   if  $P$  is root process then
5     for  $i \leftarrow 1$  to  $seg\_num$  do
6       | Lsend segment  $i$  to its children in spanning tree;
7       | Wait for all lsend;
8     end
9   else if  $P$  is a leader process in an intermediate compute node then
10    Post lrecv for 1st segment from its father;
11    for  $i \leftarrow 1$  to  $seg\_num-1$  do
12      | Post lrecv for next segment(segment  $i+1$ ) from its father;
13      | Wait for previous lrecv(segment  $i$ );
14      | Lsend received segment(segment  $i$ ) to its children;
15      | Barrier in the lcomm communicator;
16      | Wait for all lsend;
17    end
18    if  $i \equiv seg\_num$  then
19      | Wait for previous lrecv(last segment);
20      | Lsend last segment to its children;
21      | Barrier in the lcomm communicator;
22      | Wait for lsend;
23    end
24  else
25    for  $i \leftarrow 1$  to  $seg\_num$  do
26      | Recv segment  $i$  from its father;
27      | Barrier in the lcomm communicator;
28    end
29  end
30  Barrier in the lcomm communicator;
31  Deregister buffer from KNEM device;
32 else
33  Get KNEM cookie from the leader process;
34  if  $P$  is on the same compute node with root process then
35    Fetch the whole data from root process by KNEM;
36  else
37    for  $i \leftarrow 1$  to  $seg\_num$  do
38      | Barrier in the lcomm communicator;
39      | Fetch segment  $i$  from leader process by KNEM;
40    end
41    Barrier in the lcomm communicator;
42  end
43 end

```

Algorithm 1: The HierKNEM Broadcast Algorithm.

intra-node data movement can be overlapped at the leader process with the forwarding between leader processes on the upper level (steps 12, 14, or 20).

This is the fundamental reason why our HierKNEM collective can outperform other collective components: intra-node communication is offloaded to non-leader processes and leader processes can dedicate themselves into inter-node message forwarding. In an ideal situation, the intra-node communication overhead can be completely hidden from the overall execution time and the entire collective communication execution time made close to the inter-node collective execution time (the collective on the leader processes communicator). In the event of a perfect overlap, a multicore broadcast operation can be made *number-of-compute-nodes* dependent instead of *number-of-cores* dependent.

3.3. Reduce

Similar to the Broadcast algorithm, the HierKNEM Reduce uses an inter-node communicator (llcomm) and intra-node communicator (lcomm). In addition to these two communicators, the HierKNEM Reduce creates another local communicator, a subset of the lcomm, to organize all non-leader processes on the same compute node (new_comm). This new_comm is used to isolate leader processes from the intra-node reduction. The HierKNEM Reduce is actually a double-leader algorithm: the 1st leader process participate in the upper level (inter-node) reduction while the 2nd leader process will be the root for an intra-node reduction on each of the new_comm communicators and

```

Input: MPI_Reduce(void *sbuf, void *rbuf, int count, MPI_Datatype dtype, MPI_Op op, int root, MPI_Comm comm)
1 if  $P$  is the 1st leader process then
2   for  $i \leftarrow 1$  to  $seg\_num$  do
3     | Wait notification from 2nd leader;
4     | Reduction in the llcomm for segment  $i$ ;
5   end
6 else if  $P$  is the 2nd leader process then
7   for  $i \leftarrow 1$  to  $seg\_num$  do
8     | Fetch segment  $i$  from 1st leader;
9     | Reduction between two leaders' segment  $i$ ;
10    | Reduction in the new_comm for segment  $i$ ;
11    | Push reduction result of segment  $i$  to 1st leader's tmpbuf;
12    | Notify 1st leader that pushing operation is done;
13  end
14 else if non-leader processes exist inside the compute node then
15  for  $i \leftarrow 1$  to  $seg\_num$  do
16    | Reduction in the new_comm for segment  $i$ ;
17  end
18 end

```

Algorithm 2: The HierKNEM Reduce Algorithm.

responsible for updating the 1st leader with the local contribution to the upper level reduction. Algorithm 2 describes the HierKNEM Reduce for a general case: each compute node has more than two processes participating in a reduction operation: one leader for the inter-node reduction and another leader for the intra-node reduction. In order to save space, we trimmed the algorithm of the handling of special cases, the internal management and distribution of KNEM registrations.

The 2nd leader fetches segment i from the 1st leader's sbuf by a KNEM get operation (step 8), and applies the reduction operation between their sbuf's segment i (step 9). As a root process, the 2nd leader calls an intra-node reduction for segment i in the new_comm with the result from step 9. After finishing this reduction, the 2nd leader will push the reduction result of segment i to the 1st leader by a KNEM writing. After getting the notification from the 2nd leader (step 3), the 1st leader will trigger an inter-node reduction between leader processes with pushed results for segment i . The intra-node reduction for segment $i + 1$ can be overlapped with inter-node reduction for segment i thanks to KNEM's one-sided operation and the pipelining reduction algorithm between hierarchical communicators.

3.4. Allgather

The HierKNEM provides two algorithms for Allgather: a leader-based algorithm for clusters of small compute nodes (2–6 cores per compute node) and a ring algorithm for large compute nodes. The leader-based algorithm has three steps: (1) gathering messages into leader processes; (2) exchanging data between leader processes; and (3) broadcasting data from leader processes to non-leader processes. Steps 1 and 3 happen inside a compute node while step 2 exchanges data using inter-node communications. At the inter-node level (step 2), the leader processes are organized into a logical ring and each leader process communicates only with the left and right neighbors in this ring. Once leader processes get a message from step 1 or step 2, they will notify non-leader processes to fetch data by KNEM copy. Because KNEM copy in step 1 or 3 is one-sided and always offloaded onto non-leader processes, leader processes only synchronize with non-leader processes before or after non-leader processes write or read data into or from leaders. This synchronization overhead is minimal compared with the cost of intra-node data movement. As a result, the leader processes can dedicate themselves to inter-node data exchanging and steps 1–3 can be totally overlapped. The critical path of our algorithm depends on the overhead of inter-node exchanging or intra-node gather (step 1) and broadcast (step 3). When intra-node communication cost exceeds inter-node exchanging time (more

cores per compute node or faster network), the leaders' memory bandwidth is overloaded by this ad-hoc memory access pattern. Thus, the overall throughput is seriously restricted by such a simple combination of Gather and Broadcast operations. So in clusters of large NUMA compute nodes, the HierKNEM Allgather adopts a ring-based algorithm to distribute data both at the inter-node and intra-node levels in order to avoid such hot-spots on leader processes. The HierKNEM Allgather ring algorithm is similar to the MPICH Allgather ring algorithm [21]: all processes are organized into a logical ring and each process receives messages only from its left neighbor and sends messages only to its right neighbor. This send and receive will be executed number of `comm_size-1` times and a local memory copy will be executed at the beginning. A notable improvement over the ordinary ring algorithm, the construction of the HierKNEM's logical ring is not based on the order of MPI ranks but adheres to the physical process distance in terms of sockets and NUMA compute nodes. Thus, processes physically close are clustered together into a set. Only processes on edges between sets communicate through slow links: inter-node links or inter-socket links.

4. Experimental evaluation

We used two clusters of the Grid5000 experimental platform: Stremi and Paraplue. The Stremi cluster features 32 compute nodes, each with two AMD Opteron 6164 HE twelve-core CPUs (24 cores per compute node). Each socket has 10 MB L3 caches and two NUMA memory nodes, and 6 cores in each socket share one NUMA memory node with 12 GB of memory (48 GB of memory per compute node). These 32 compute nodes are interconnected by Gigabit Ethernet. The Paraplue cluster is identical to Stremi, except that the 32 compute nodes are interconnected through a 20 G Infiniband network.

Our HierKNEM collective is based on Open MPI version 1.5.3. We compared the HierKNEM collective with Open MPI's (1.5.3) Tuned, Hierarch collective, MPICH2 version 1.4.1 on the Ethernet cluster (Stremi) and MVAPICH2 version 1.7 on the Infiniband cluster (Paraplue). All implementations that support kernel assisted memory operations use KNEM version 0.9.6 [4] except MVAPICH2. MVAPICH2 uses LIMIC2 0.5.5 [11] as the kernel assisted memory copy module.

For intra-node communications, HierKNEM, Tuned, and Hierarch collective modules are configured to use the SM/KNEM BTL (byte transfer layer) as underneath point-to-point communication helper. SM/KNEM BTL uses KNEM copy to speed up point-to-point communication; for performance reasons, the copy-in/copy-out approach is still used for messages smaller than 4 KB. The same configuration is applied to MPICH2 or MVAPICH2: KNEM/LIMIC copy is enabled for large message transfer (LMT). For inter-node communications, the appropriate low level point-to-point transport module is used, depending on the underlying hardware (Open IB, TCP). For all MPI libraries, the process/core binding is the default uniform "by-core" strategy, except when explicitly mentioned. In this default strategy, sequential MPI ranks are bound into adjacent processor cores until all slots of a compute node have been used, then the same process is applied for the next compute node in the list. To summarize, the underlying technology used by our HierKNEM algorithm and all other collective components to perform point-to-point operations is similar and uses KNEM copies, similarly process placement is comparable; therefore any performance difference roots solely in the proposed collective operation innovations.

The Intel MPI benchmark suite IMB-3.2 [10] is used to assess the difference between the collective components on a variety of collective operations. The ASP [19] problem is a typical example of a parallel graph shortest path search algorithm. It is used to illustrate how performance differences in micro-benchmarks translate into application improvement.




Fig. 1. Non-uniform I/O effect on point-to-point communication (Pingpong Test) execution time on Paraplue cluster; runtime is normalized to the result for Pingpong test between two compute nodes' core 0 (the smaller, the better).

4.1. Leader selection

Similar with Non-Uniform memory access (NUMA), Non-Uniform Input/Output access (NUIO) phenomenon was found in some platforms because some I/O devices are closer to some processors and memory banks than to the others [7]. The latency or bandwidth of inter-node communication between leader processes may be affected by the selection of leader processes due to this Non-Uniform Input/Output access. We did four pingpong tests between the Paraplue's two compute nodes to inspect the NUIO impact on our experiment platforms. To simplify the results in the graph, core 0, 6, 12, and 18 are selected in the experiments to represent other cores on four NUMA memory nodes. And Pingpong tests are only executed between two cores with the same id on two compute nodes. Fig. 1 shows the execution time of pingpong tests between pairs of cores normalized to the runtime between core 0s. In most cases, the selection of processes on core 0 or core 6 as leaders will give out the best inter-node communication performance. But the difference is very trivial, about 1%–2% for message sizes selected as pipeline sizes, so even a wrong leader selection here will not lead to a performance disaster of HierKNEM collectives. In the following experiments, HierKNEM keeps selecting processes with the minimum core id as leader processes in collective communications. In a future release, we will add a module to help HierKNEM select leaders according to the hardware locality [3].

Another issue related with the leader selection is how many leader processes were selected in the collectives. HierKNEM can be configured to select one leader process from each compute node, each board, each NUMA memory node, or each socket. The run-time configuration is decided by the ratio between inter-node and intra-node communication. For example in a scenario of large NUMA compute nodes connected by fast interconnections, e.g. a 20 G or 40 G Infiniband network, the runtime of intra-node collective communication (broadcast) possibly greatly exceed the runtime of inter-node forwarding (sending/receiving). HierKNEM will adjust the whole hierarchical topology by selecting leaders from each NUMA memory node instead of from each compute node. The HierKNEM Broadcast on the Paraplue Cluster follows this selection. This flexible selection will lead to a better overlapping between intra- and inter-node communication, and meanwhile intra-node memory accesses will be spread evenly over NUMA memory nodes. Except this special case, other leader selections still follow one leader process per compute node in HierKNEM collectives.

4.2. Pipeline size

In the HierKNEM collective component, both the Broadcast and the Reduce operations are pipelining algorithms, in which

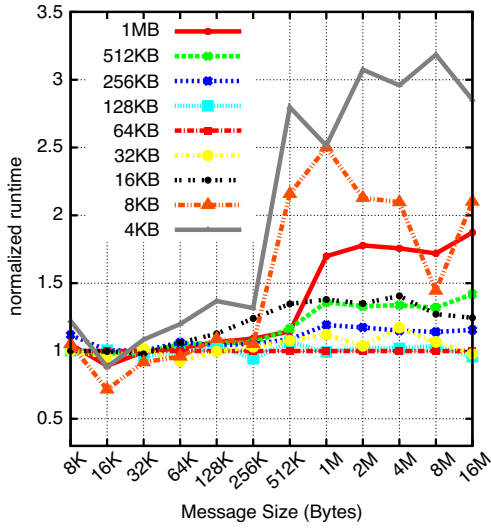


Fig. 2. Effect of pipeline size on HierKNEM Broadcast execution time (Paraplue cluster: 768 processes, 32 compute nodes, Infiniband 20 G); runtime is normalized to the result for 64 KB pipeline size (the smaller, the better).

Table 1 Best pipeline size for broadcast and reduce for differing network capacities.

Operation	Paraplue (IB20 G)		Stremi (Ethernet)	
	Message size	Pipeline size	Message size	Pipeline size
Broadcast	[8 KB, 64 KB]	8 KB	[8 KB, 512 KB]	16 KB
	[64 KB, ∞)	64 KB	[512 KB, ∞)	32 KB
Reduce	[2 K, 16 MB]	64 KB	[2 K, 16 MB]	64 KB
	[16 MB, ∞)	1 MB	[16 MB, ∞)	1 MB

messages are split into several smaller chunks. Tuning an optimal size of a chunk is a key criterion of every pipeline algorithm. Fig. 2 presents the effect of the pipeline size on the HierKNEM Broadcast execution time. In this Broadcast test, 768 processes are spawned on the Paraplue cluster. To ease figure clarity, the execution time for all pipeline sizes is normalized to the runtime obtained with a pipeline size of 64 KB (t_z/t_{64}). One can see that the pipeline size is indeed critical to the HierKNEM collective performance, and a wrong selection of pipeline sizes leads to significant penalty. On one hand, a too small pipeline size results in inefficient inter-node communication, as the small message latency comes to dominate, preventing the full point-to-point bandwidth from being leveraged; as an example, the Broadcast with a pipeline size of 4 KB is more than 3 times slower than with 64 KB. On the other hand, a too large pipeline size results in long pipeline fan-in and fan-out phases, where the pipeline algorithm is not at steady-state efficiency. Experimentally, 8 KB or 64 KB is the ideal pipeline size for the Broadcast operation on the Paraplue cluster for messages smaller or larger than 64 KB. We did similar experiments for HierKNEM's Broadcast and Reduce on both the Paraplue and Stremi clusters. Table 1 shows the best pipeline size for each operation on each type of cluster. Both HierKNEM's Broadcast and Reduce algorithms use the pipeline size in Table 1 in the following tests.

4.3. Impact of Kernel-assisted approaches

Intra-node point-to-point operations can be implemented using shared memory or kernel-assisted approaches. Kernel-assisted approaches have the benefit of decreasing the number of memory copies, a factor critical in memory intensive code sections such as the collective communications. Fig. 3 shows a comparison of the broadcast collective bandwidth on the Stremi

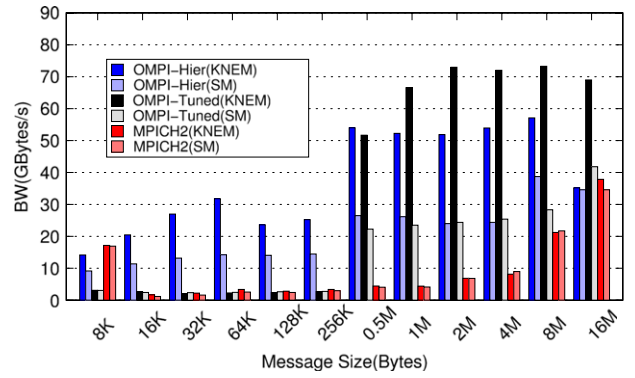


Fig. 3. Comparison between Kernel-assisted and shared memory approach: Broadcast bandwidth on Stremi (Ethernet, 768 processes, 24 cores/compute node).

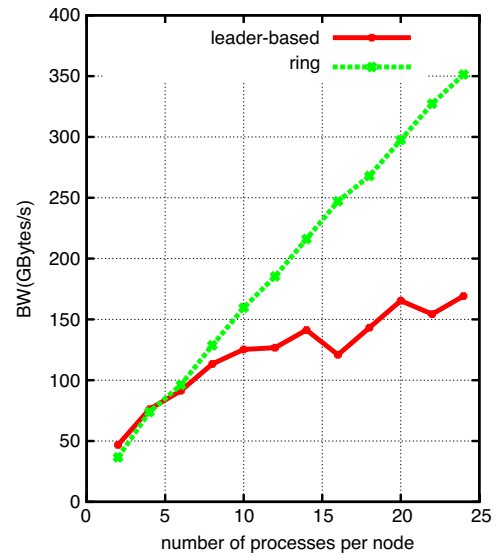


Fig. 4. Bandwidth comparison between Leader-based and Ring Allgather algorithms, when increasing the number of processes per compute node (from 2 to 24), on Paraplue's 32 compute nodes.

cluster between collective modules over the kernel-assisted or shared memory point-to-point communication. Kernel-assisted approach shows a huge performance boost in Open MPI's collective modules: Tuned and Hierarch collectives, up to 3× for some message sizes. On the opposite side, the performance difference regarding MPICH2 is less flagrant. Other collective operations, such as Reduce and Allgather, exhibit similar speedup when kernel-assisted approach are replacing shared memory point-to-point communication. Moreover, this performance improvement on collective communications is directly inherited by the applications using them, as highlighted in Section 4.8. Thus, in the remaining of this paper we focus on kernel-assisted approaches, discarding the suboptimal approaches built on top of shared memory point-to-point communications.

4.4. Allgather algorithm selection

Although the two levels of algorithms are tightly integrated, there are still a variety of combinations that are possible, whose performance greatly varies depending on hardware features and properties. In the case of the Allgather algorithm, we identified two combinations of interest: both use the pipelined Tuned collective module between compute nodes, but the internal Tuned operation differs depending on the number of cores between compute nodes. Between cores, the algorithm can rely on the leader

originating all messages simultaneously (referred to as “leader-based” algorithm), but for large core counts, this approach has the potential to result in heavy traffic contention on the memory bus of the core hosting the leader. For a larger number of cores per compute node, the ring algorithm has more potential to even out the load on all cores. Fig. 4 shows the aggregate bandwidth for the two algorithms combinations, for a 512 KB message’s Allgather operation on Paraplue’s 32 compute nodes, when increasing the number of processes per compute node from 2 to 24. The leader-based algorithm has a slight performance advantage in dual-core or quad-core compute nodes, as the parallel KNEM accesses overlap one another. For larger setups, the bandwidth contention on the leader core prevents aggregate bandwidth to scale, while the ring algorithm, which proves more scalable thanks to evenly distributing data access load across all memory links, dominates. Results (not presented here) are similar for other message sizes, and when using different inter-node networks on Stremi and Paraplue clusters. In the following tests, we use the ring algorithm as we mainly target large multicore compute nodes.

4.5. Collective communication performance

In this Section we will analyze the performance of three of the most used collective communication operations, namely Broadcast, Reduce and Allgather. They cover all of the regular collective patterns available in MPI, one-to-many, many-to-one and many-to-many, providing a quite extensive view of all the potential performance improvement in collective communications.

4.5.1. Broadcast performance

Fig. 5 presents the aggregate Broadcast bandwidth for HierKNEM, Open MPI’s Hierarch and Tuned modules, and MPICH2 or MVAPICH2 on, respectively, the Ethernet Stremi cluster or the Infiniband Paraplue cluster. On Stremi (Fig. 5(a)), for message size between 8 and 256 KB, HierKNEM Broadcast provides a significant speedup, sometimes up to 30x, when compared with MPICH2 and Open MPI. Compared with Open MPI’s Hierarch, our HierKNEM version provides more than twice the aggregate bandwidth in this message size range.

For larger message sizes (superior to 512 KB), the most important tuning factors are process mapping and proper pipelining to evenly spread the workload across cores and links. In the Tuned module, the “by core” binding luckily happens, in this experiment, to match the hardware topology; and the pipeline size selected by the tuned module to optimize the network communications is suitable for core communications. The Hierarchical module of Open MPI is not as successful for large messages, because the intra-node and inter-node layers do not cooperate to evenly spread the load of the pipelining algorithm. The leader processes are unavailable for long periods of time when they take part in the shared memory local operation, resulting in effectively sequentializing the local and remote collective operations without opportunity for overlap. With such a large core count, the large intra-node overhead offsets the benefits of the standard hierarchical algorithm. In contrast, the HierKNEM algorithm obtains better performance in all cases, thanks to explicitly taking into account process mapping and using directional KNEM control to offload parts of the operations onto the leaf processes, hence enabling intra and inter-communications to overlap.

Similarly, on the Infiniband cluster (Fig. 5(b)), in most cases, the HierKNEM Broadcast still outperforms other collective components. One major difference in the results, when compared with the Ethernet case, is that the performance of the classical hierarchical algorithm is much better for small message sizes. On the Infiniband network, the tuning parameters selected by the two non-cooperating algorithms forming the hierarchical collective are

matching better. However, as one can see, the tuning parameters for larger message size are not as lucky; the performance for a large messages drops, with the notable exception of 512 KB messages, for which the pipeline length matches the balance for 32 processes and 24 cores. This discrepancy illustrates the difficulty of tuning the behavior of separate collective algorithms cooperating in a hierarchical manner. Even with expert knowledge, it is unrealistic to tune Open MPI’s Tuned collectives on such a complex system with so many hierarchies and diverse networks, when a small variation in message size results in unexpected and dramatic performance consequences. Although the HierKNEM module is not immune to the challenges of unpredictable and unstable performance on varying hardware, the fact that both algorithms select compatible tuning parameters, that the outer collective operation can overlap imperfection on the inner operation and that the collective topology is constructed to match core hierarchies greatly alleviates this difficulty, as illustrated by more stable results across the message size range.

4.5.2. Reduction performance

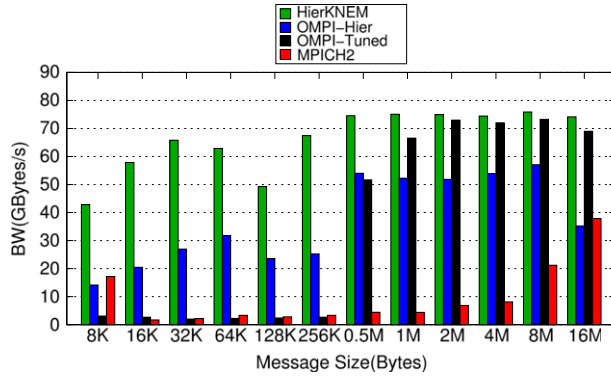
Fig. 6 presents the aggregate Reduce bandwidth on the Ethernet cluster (Fig. 6(a)). For message sizes between 2 and 32 KB, the HierKNEM Reduce competes closely with Open MPI’s Hierarch Reduce. After 64 KB, the HierKNEM Reduce dominates other collective components, thanks to a good overlapping between inter-node Reduce and intra-node Reduce. Similarly with the Broadcast, the Hierarch Reduce worsens for large messages due to the increased intra-node Reduce overhead which cannot be dodged by overlap. Again, the performance of the Tuned Reduce improves for messages larger than 4 MB, but is still 19%–28% slower than the HierKNEM Reduce.

On the Infiniband cluster (Fig. 6(b)), the HierKNEM Reduce clearly dominates for message size in the range of [2B, 32 KB] and (1, 16 MB]. When message size is between 64 KB and 1 MB, although HierKNEM Reduce still achieves significant speedup when compared with Open MPI’s Hierarch and Tuned Reduce, it is a little behind MVAPICH2 performance, about 20% in the worst case. By profiling a 64 KB message’s Reduction operation with 32 processes on Paraplue’s 32 compute nodes (no multicore or hierarchies), we discovered that the Open MPI Tuned Reduction suffers from a serious performance limitation on the Infiniband network; meanwhile MVAPICH2 enjoys very good performance (366 μ s for Open MPI compared to 281 μ s for MVAPICH2). As our HierKNEM composite algorithm reuses the original Tuned module for inter-node communications, it suffers from the same defect and cannot compete with MVAPICH2, until the Open MPI community addresses this issue.

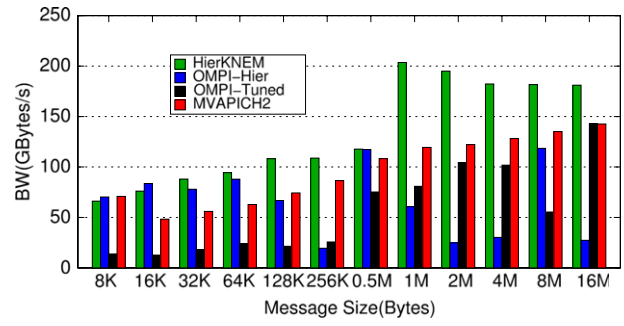
4.5.3. Allgather performance

Fig. 7 presents the aggregate Allgather bandwidth. The HierKNEM Allgather is enabled only when the message size is larger than 8 KB. The biggest message size is 1 MB, because of the large amount of memory required for this all-to-all operation between 768 processes exhausting available system memory for larger sizes. On both clusters, the HierKNEM Allgather adopts a ring algorithm, as described in Section 3.4. The Open MPI Hierarch module is not presented for this collective operation, as it has not been implemented.

On the Infiniband cluster (Fig. 7(b)), both MVAPICH2 and Tuned Allgather operations slightly outperform HierKNEM’s Allgather. In this message range, Open MPI’s Tuned Allgather adopts a similar but more aggressive algorithm: neighbor exchange. Different than the ordinary ring algorithm which exchanges rcount data in each round, $2 \times$ rcount data are exchanged in each round in the neighbor exchange algorithm (except the first round). Neighbor exchange algorithm arranges processes continuous in MPI ranks

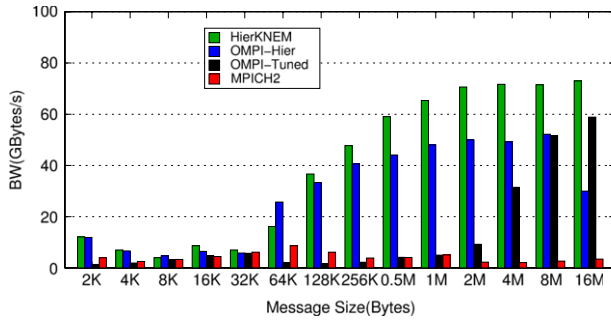


(a) Stremi (Ethernet).

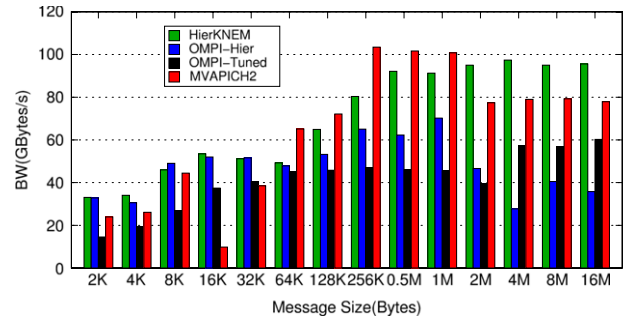


(b) Paraplue (IB20 G).

Fig. 5. Aggregate Broadcast bandwidth of collective modules on multicore clusters (768 processes, 24 cores/compute node).




(a) Stremi (Ethernet).




(a) Paraplue (IB20 G).

Fig. 6. Aggregate Reduce bandwidth of collective modules on multicore clusters (768 processes, 24 cores/compute node).



(a) Stremi (Ethernet).



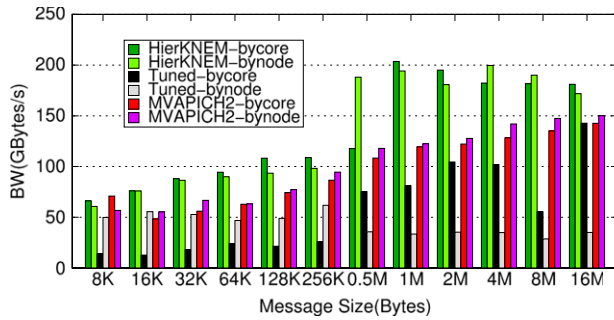
(a) Paraplue (IB20 G).

Fig. 7. Aggregate Allgather bandwidth of collective modules on multicore clusters (768 processes, 24 cores/compute node).

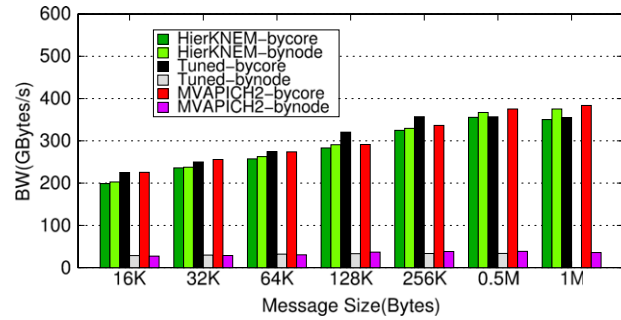
together in a ring and the “by core” binding strategy used in this test coincidentally maps the logical ring of the Tuned Allgather correctly to the underlying hardware topology. As a consequence, Tuned and HierKNEM are actually running on the same underlying hardware topology, but Tuned does not have to pay for the extra cost of detecting the physical distance between processes, and fast networks like 20 G Infiniband show better performance in a more aggressive policy due to more bandwidth capacity available. On the Ethernet cluster (Fig. 7(a)), the HierKNEM Allgather outperforms all other collective components for all message sizes. While adopting a similar ring topology for large messages, the Tuned Allgather on this Ethernet cluster suffers up to 50% in performance loss because the aggressive policy in the neighbor exchange algorithm overloads slow networks decreasing the overall network throughputs. Clearly, adopting an aggressive strategy, like the neighbor exchange algorithm, should depend on the capacity of the networks.

4.6. Impact of process placement

It is well known that process placement can have a major impact on collective operations performance. Approaches such as MPIPP [5] have been designed to detect communication patterns during a “tuning run”, whose result is used to hint process placement to decrease long distance communication volume during subsequent production runs. However, this approach is not practical in many cases, as it requires being able to run smaller problems that exhibits similar communication patterns; and the collective algorithm underlying communication topology might depend on the message and communicator size. Another difficulty is that oftentimes, one might want to optimize for the pattern of point-to-point operations explicitly realized at the application level (such as the typical hypercube topology found in many CG implementations), which means that the process placement may or may not fit the expectations of the collective modules. As a consequence, the default deployment approach is usually less



(a) Broadcast.



(a) Allgather.

Fig. 8. Impact of process mapping: aggregate Broadcast and Allgather bandwidth of the collective modules for two different process-core bindings: by core and by compute node (Paraplue cluster, IB20 G, 768 processes, 24 cores/compute node).

elaborate and simply allocates ranks sequentially on the available resources.

Fig. 8 shows the impact of two typical process placements on the performance of the Broadcast and Allgather operations. The goal of this experiment set is to investigate the sensitivity of the hierarchical approaches to variations in the process placement. As such, more than raw performance, it is the difference between the same algorithm on different mappings that is of interest here. Hierarch has been trimmed from the figure, because it does not feature an Allgather operation, and uses a similar topology as HierKNEM for the Broadcast (hence similar performance trends). Considering the Broadcast (Fig. 8(a)), one can witness that hierarchical approaches (HierKNEM and MVAPICH2 both feature a hierarchical algorithm) reach more stable performance. The Tuned algorithm exhibit very unstable performance trends, for some message sizes the by node binding reaches better performance, while it is the contrary for larger messages.

Fig. 8(b) further displays the importance of considering hierarchical features to enable portability of performance across varied process mappings. In this algorithm, the HierKNEM algorithm demonstrates very stable performance when changing from by core to by node process mappings. The performance variation between two bindings is less than 10%, which is very small when compared to the tremendous performance penalty suffered by non hierarchical algorithms, commonly more than 6 \times and sometimes up to 14 \times increased communication time. In the “by node” binding, the Tuned Allgather uses a ring algorithm for large messages; every edge of the logical ring (768 edges in this case) passes through inter-node links (Infiniband), causing a serious traffic congestion on the Infiniband network. This clearly illustrates the penalty suffered by topology-unaware algorithms when considering irregular process-core bindings. Although our HierKNEM collective pays an overhead due to constructing the internal topology, it provides stable performance independently of process placement. Such a flexible process placement is a desirable feature to enable deeper optimization of the hard-coded point-to-point communication patterns and ensure maximum performance with default settings on complex architectures.

4.7. Core per node scalability

In the next experiment, we investigate the trend of aggregate bandwidth when varying the number of cores per compute node. The total number of compute nodes is left unchanged (32 compute nodes), but the number of processes per compute node is increasing for each experiment, reaching up to the maximum of 24 processes per compute node. The message size is kept constant at 2 MB. Processes on each compute node are bound to cores sequentially.

Table 2

Kernel-assisted approach comparison: ASP application execution runtime execution breakdown on Stremi (Ethernet, 768 processes, 24 cores/compute node). Using KNEM and SM.

Problem size (K)	HierKNEM		Tuned		Hierarch		MPICH2	
	Bcast (s)	Total (s)	Bcast (s)	Total (s)	Bcast (s)	Total (s)	Bcast (s)	Total (s)
16	20.3	97.4	229	308	31.7	109	128	204
32	79	711	929	1560	173	806	417	1020
			Tuned-SM		Hierarch-SM		MPICH2-SM	
16			229	309	66	145	124	201
32			929	1574	245	899	429	1040

On both clusters (Fig. 9(a) and (b)), the aggregate bandwidth of HierKNEM Broadcast achieves a linear speedup when more cores per compute node are involved, because our HierKNEM Broadcast dodge the intra-node communication overhead by overlapping it with the inter-node message forwarding. Increasing processes (cores) per compute node does not increase the overall Broadcast completion time on these two platforms. This linear speedup can be maintained until the time necessary to perform the entire intra-node communication (a KNEM Broadcast) exceeds the inter-node forwarding time of the network.

4.8. Application performance

We have asserted the maximum possible performance improvement by solely executing synthetic benchmarks over the modified operations. It is now needed to evaluate how much of this improvement results in improved performance for applications. To evaluate the impact of the HierKNEM collective algorithms on real application performance, we consider a typical parallel graph application: ASP [19]. This application unfolds the parallel Floyd–Warshall algorithm to solve the all pairs shortest path problem. At the beginning of each iteration the master process broadcasts a row of the square matrix representing edges weight to all peers in the communicator, in order to distribute the workload. The outer loop of the algorithm iterates on rows, until the entire matrix is treated. Overall, for a matrix of size N , the algorithm performs N broadcasts, with a message size of $\text{column_num} \times \text{type_size}$. As a consequence, MPI_Bcast contributes to the majority of the runtime of the ASP’s MPI usage.

Table 2 compares the overall execution time and communication time (mostly MPI_Bcast) of the ASP application on the Stremi cluster when using different collective modules. By subtracting the communication time from the overall execution time, one can assert that ASP’s computational part remains generally constant for a given problem size, independently of the communication setup. The major performance difference between these four setups comes from the communication overhead (MPI_Bcast). The




Fig. 9. Core per compute node scalability: aggregate bandwidth of Broadcast for 2 MB messages on multicore clusters (32 compute nodes).

cost of communications occupies 21% of the overall application runtime for the HierKNEM collective, while it rises to 74% when using Open MPI's default collective. Even considering the hierarchical broadcast, the HierKNEM's ability to overlap between inter and intra communications shows a significant improvement in this application.

The second part of the Table 2 shows the same application (ASP) on the same platform using shared memory point-to-point communication instead of kernel-assisted approaches. Similarly with the prior results in Section 4.3 most collective modules over shared memory point-to-point communication lose performance compared with kernel-assisted point-to-point communication due to double memory copies and more memory usage, sometimes up to 30%. The impact on collective performance is directly translated to waste time for the applications.

4.9. Experiments accuracy

In this experimental evaluation, we preferred using widely accepted benchmarks, such as the Intel Message Passing (IMB) benchmark, for they are specifically designed to eliminate noise and artifacts: for small and intermediate messages, experiments are executed thousand of times with different roots in order to eliminate any hot caching and pipelining effect, while the experiments are executed hundreds of times for large messages. To further emphasize reproducibility, for each experiment, we reserved the whole clusters to avoid sharing network switches with other users. In addition to these precautions, every experiment was realized multiple times, and we took in account not only the minimum, maximum and average values, but the standard deviation as well. For the message size of 512 KB, the IMB tests resulted in a bandwidth between 25.6 and 26.2 GB, with an average at 25.9 GB, and a standard deviation of 0.17. Similarly, particular care has been taken to the standard deviation of the ASP application performance. Among 18 times of repeated experiments, the maximum ASP execution time was 146 s, the minimum time 144 s, the average execution time is 145.1, and the standard deviation is 0.6. These small standard deviation values are indicative that the design of the IMB benchmark, ASP application, and the precautions we have taken guarantee the accuracy and stability of our experimental results.

5. Conclusion

In this paper, we described a kernel-assisted topology-aware collective framework: HierKNEM, which enables efficient combinations of multiple layers of collective algorithms, to tackle collective communication on clusters of many-core compute nodes. The algorithms are built reusing modular combinations of existing collective algorithms (such as the Tuned and the

KNEM components in Open MPI). The main contributions of this paper are: (1) propose an adaptive hierarchical collective framework to enable tight collaboration between the collective algorithms pertaining to different layers of the hierarchy, (2) combine offloading and pipelining techniques into the hierarchical framework to release leader processes from intra-node data movement, hence maximizing the overlap between inter- and intra-node communications, and (3) build internal collective topologies to form a mapping between the runtime process-core binding and the hardware features, which means stable collective performance independently of process placement.

We demonstrated the benefits of this approach by devising three hierarchical integrated collective algorithms, one of the most useful for each major type of collective communication (one-to-many: Broadcast, many-to-one: Reduce, and many-to-many: Allgather). Experimental results demonstrate that our approach outperforms not only non hierarchy aware state-of-art MPI implementations (MPICH2 and Tuned Open MPI), although these setups benefit from kernel assisted memory copies as well (KNEM), but also significantly outperforms approaches that account for the hierarchy (MVAPICH2 Broadcast, Hierarch component in Open MPI). A simple leader based algorithm that does not enable pipeline coordination, and intra-node copies offloading, under-performs compared to our HierKNEM approach that introduces these features. The performance improvement is visible not only in synthetic benchmarks, but also results in up to a ten-fold performance improvement when compared to the default non hierarchy aware strategy, and still features two-fold improvements when compared to other hierarchical strategies.


Acknowledgments

Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>).

References

- [1] R. Brightwell, Exploiting direct access shared memory for MPI On multi-core processors, *International Journal of High Performance Computing Applications* (2010).
- [2] R. Brightwell, K. Pedretti, T. Hudson, SMARTMAP: operating system support for efficient data sharing among processes on a multi-core processor, in: *Proc. of the ACM/IEEE Conference on Supercomputing*, 2008.
- [3] F. Broquedis, J. Clet-Ortega, S. Moreaud, N. Furmento, B. Goglin, G. Mercier, S. Thibault, R. Namyst, HWLOC: a generic framework for managing hardware affinities in HPC applications, in: *The Euromicro International Conference on Parallel, Distributed and Network-Based Computing*, 2010.
- [4] D. Buntinas, B. Goglin, D. Goodell, G. Mercier, S. Moreaud, Cache-efficient, intranode large-message MPI communication with MPICH2-Nemesis, in: *International Conference on Parallel Processing*, 2009.

- [5] H. Chen, W. Chen, J. Huang, B. Robert, H. Kuhn, MPIPP: an automatic profile-guided parallel process placement toolset for SMP clusters and multiclusters, in: Proc. of International Conference on Supercomputing, 2006.
- [6] G.E. Fagg, G. Bosilca, J. Pješivac-Grbović, T. Angskun, J.J. Dongarra, Tuned: a flexible high performance collective communication component developed for Open MPI, in: Proceedings of DAPSYS, 2006.
- [7] B. Goglin, S. Moreaud, Dodging non-uniform I/O access in hierarchical collective operations for multicore clusters, in: Workshop on Communication Architecture for Scalable Systems (CASS) Held in Conjunction with International Parallel and Distributed Processing Symposium, 2011.
- [8] HierKNEM: an adaptive framework for kernel-assisted and topology-aware collective communications on many-core clusters, Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS'12, IEEE, 2012.
- [9] L. Huse, Collective communication on dedicated clusters of workstations, in: Recent Advances in Parallel Virtual Machine and Message Passing Interface, 1999.
- [10] Intel MPI benchmarks 3.2, <http://software.intel.com/en-us/articles/intel-mpi-benchmarks/>.
- [11] H.W. Jin, S. Sur, L. Chai, D.K. Panda, LiMIC: support for high-performance MPI intra-node communication on Linux cluster, in: International Conference on Parallel Processing, 2005.
- [12] K. Kandalla, H. Subramoni, A. Vishnu, D.K. Panda, Designing topology-aware collective communication algorithms for large scale infiniband clusters: case studies with scatter and gather, in: IEEE International Symposium on Parallel Distributed Processing, Workshops and Ph.D. Forum, IPDPSW, 2010.
- [13] N.T. Karonis, B.R. De Supinski, I.T. Foster, W. Gropp, E.L. Lusk, A multilevel approach to topology-aware collective operations in computational grids, in: Computing Research Repository, 2002.
- [14] N.T. Karonis, B.R. de Supinski, I. Foster, W. Gropp, E. Lusk, J. Bresnahan, Exploiting hierarchy in parallel computer networks to optimize collective operation performance, in: The International Parallel and Distributed Processing Symposium, 2000.
- [15] T. Ma, G. Bosilca, A. Bouteiller, B. Goglin, J.M. Squyres, J.J. Dongarra, Kernel assisted collective intra-node MPI communication among multi-core and many-core CPUs, in: Proc. of International Conference on Parallel Processing, 2011.
- [16] T. Ma, T. Herault, G. Bosilca, J.J. Dongarra, Process distance-aware adaptive MPI collective communications, in: Proc. of IEEE International Conference on Cluster Computing, 2011.
- [17] A.R. Mamidala, L. Chai, H.W. Jin, D.K. Panda, Efficient SMP-aware MPI-level broadcast over infiniband's hardware multicast, in: Workshop on Communication Architecture for Clusters (CAC) Held in Conjunction with International Parallel and Distributed Processing Symposium, 2006.
- [18] A. Mamidala, A. Vishnu, D.K. Panda, Efficient shared memory and RDMA based design for MPI_Allgather over infiniband, in: Recent Advances in Parallel Virtual Machine and Message Passing Interface, 2006.
- [19] A. Plaat, H.E. Bal, R.F.H. Hofman, T. Kielmann, Sensitivity of parallel applications to large differences in bandwidth and latency in two-layer interconnects, in: Future Generation Computer Systems, 2001.
- [20] Y. Qian, A. Afsahi, RDMA-based and SMP-aware multi-port all-gather on multi-rail QsNet II SMP clusters, in: Proc. of International Conference on Parallel Processing, 2007.
- [21] R. Thakur, W. Gropp, Improving the performance of collective operations in MPICH, in: Recent Advances in Parallel Virtual Machine and Message Passing Interface, 2003.
- [22] H. Zhu, D. Goodell, W. Gropp, R. Thakur, Hierarchical collectives in MPICH2, in: Recent Advances in Parallel Virtual Machine and Message Passing Interface, 2009.




Teng Ma is a Ph.D. student in EECS Department of University of Tennessee, Knoxville. His research interest focused on design, implementation, and modeling of parallel communication libraries, parallel computer architectures; cluster and grid computing; modeling of parallel benchmarks; clusters with hardware accelerators; parallel programming models; parallel I/O and distributed file system.



George Bosilca is a Research Associate Professor of Electrical Engineering and Computer Science, University of Tennessee, Knoxville. He obtained his Ph.D. from the University of Paris XI with a background in computer architecture and parallel computing. Active member of several large scale projects (Open MPI, MPICH-V, CCI, DAGuE), his research encompass a large area in the distributed computing world. From low level network protocols to algorithmic based fault tolerance, Dr. Bosilca research target to reduce the gap between peak and sustained performance on large scale execution environments, effectively taking

advantage of the heterogeneous capabilities of current and future computing platforms.



Aurelien Bouteiller received his Ph.D. from University of Paris in 2006, under the direction of Franck Cappello. His research is focused on improving performance and reliability of distributed memory systems. Toward that goal, he investigated automatic (message logging based) checkpointing approaches in MPI, Algorithm Based fault tolerance approaches and their runtime support, mechanisms to improve communication speed and balance within nodes of many-core clusters, and employing emerging data flow programming models to negate the raise of jitter on large scale systems (DAGuE project). These works

resulted in over twenty-five publications in international conferences and journals and three distinguished paper awards from IPDPS and EuroPar. He his also a contributor to Open MPI and participates to the MPI-3 Forum.



Jack J. Dongarra holds an appointment at the University of Tennessee, Oak Ridge National Laboratory, and the University of Manchester. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. He was awarded the IEEE Sid Fernbach Award in 2004 for his contributions in the application of high performance computers using innovative approaches; in 2008 he was the recipient of the first IEEE Medal of Excellence in Scalable Computing; in 2010 he was the first recipient of the SIAM Special Interest Group on Supercomputing's award for Career Achievement; and in 2011 he was the recipient of the IEEE IPDPS 2011 Charles Babbage Award. He is a Fellow of the AAAS, ACM, IEEE, and SIAM and a member of the National Academy of Engineering.