

# Unifying UPC and MPI Runtimes: Experience with MVAPICH

Jithin Jose    Miao Luo    Sayantan Sur    Dhabaleswar K. Panda

Department of Computer Science and Engineering, The Ohio State University

{jose, luom, surs, panda}@cse.ohio-state.edu

## Abstract

Unified Parallel C (UPC) is an emerging parallel programming language that is based on a shared memory paradigm. MPI has been a widely ported and dominant parallel programming model for the past couple of decades. Real life scientific applications require a lot of investment by domain scientists. As a result, many scientists choose MPI as a programming model as it is considered low-risk. It is unlikely that entire applications will be re-programmed using emerging UPC language (or PGAS paradigm) in the near future. It is more likely that parts of these applications will be converted to newer models. This requires that underlying implementation of system software be able to support both UPC and MPI simultaneously. Unfortunately, the current state-of-the-art of UPC and MPI interoperability leaves much to be desired both in terms of performance and ease-of-use.

In this paper, we propose an Integrated Native Communication Runtime (INCR) for MPI and UPC communications on InfiniBand clusters. Our library is capable of supporting both UPC and MPI communications simultaneously. This runtime is based on widely used MVAPICH (MPI over InfiniBand) Aptus runtime, which is known to scale to tens-of-thousands of cores. Our evaluation reveals that INCR is able to deliver equal or better performance than existing UPC runtime - GASNet on InfiniBand verbs. We observe that with UPC NAS benchmarks CG and MG (class B) at 128 processes, we outperform current GASNet implementation by 36% and 25%, respectively.

## 1. Introduction

Modern high-end computing (HEC) systems allow scientists and engineers to tackle grand challenge problems in their respective domains and make significant contributions to their fields. Examples of such problems include astro-physics, earthquake analysis, weather prediction, nanoscience modeling, multiscale and multiphysics modeling, biological computations, computational fluid dynamics, etc. The Message Passing Interface (MPI) is a very widely used parallel programming model. Scientific applications rely heavily on the performance and portability offered by it. Although MPI has been widely used for the past couple of decades, recently, there is an effort to improve programmer productivity for parallel applications. The hypothesis is that for several classes of applications, especially irregular applications, shared memory programming is much easier than explicit message passing. Addition-

ally, language based parallelism constructs provide greater flexibility of expression, analysis by software tools and optimization opportunities. For example, unneeded barriers in an application can be syntactically analyzed by tools and compiled away, eliminating the need for programmers to manually analyze dependencies and removing barriers manually. Often programmers tend to insert more barriers than strictly required for safety.

Unified Parallel C (UPC) [29] is an emerging language for parallel programming, that is based on extensions to ISO C99. It provides parallel iterations and remote memory access among a host of other features. UPC is a member of the Partitioned Global Address Space (PGAS) family of languages. Several flavors of UPC are available, along with different implementations of UPC compiler and runtimes on a variety of architectures including commodity clusters and leadership class machines such as the Blue Gene/P. The Berkeley UPC implementation [23] is one of the open-source and popular implementations that support high-performance interconnects such as Myrinet, InfiniBand and Quadrics through the portable GASNet runtime [14]. Even though UPC is an emerging language, there are not many real life scientific applications using it as their primary parallel programming model. Real life scientific applications require a lot of investment by domain scientists for design and verification, and their life cycles are in decades. Since MPI is standardized, adopted and portable, it provides a low-risk implementation choice. It is unlikely that entire applications will be re-designed using UPC in the near future. It is more likely that parts of these applications (primarily written in MPI) will be converted to UPC, resulting in a “hybrid” application.

InfiniBand [4] is an industry standard interconnection technology that has gained wide acceptance and is rapidly becoming one of the dominant interconnection technologies in HEC domain. Accordingly, the size of InfiniBand clusters is growing as well, with the Nebulae cluster in China having 120,640 cores. Design and implementation of highly scalable InfiniBand runtimes is not a trivial task as the software interface provided by *verbs* is at a very low-level. Using verbs, various buffer management, connection management, message coalescing, and reliability mechanisms need to be explored [19–22, 28].

Thus, there are two emerging trends: that of the need of hybrid programming support and need for extreme scalability on InfiniBand clusters. In order to realize the vision of highly scalable hybrid computing, the following questions need to be answered:

1. Can a communication library be designed for UPC that provides excellent performance and scalability on very large InfiniBand clusters?
2. Can one unified communication library support both UPC and MPI communications in a manner that does not degrade performance for either?

In this paper, we aim to answer the above questions. We have extended the MVAPICH-Aptus [21] runtime, which is known to be

one of the most scalable InfiniBand runtimes, to support UPC communications. MVAPICH/MVAPICH2 [25] are very widely used MPI libraries, supporting MPI-1 and MPI-2 standards specifically designed for InfiniBand. It is currently being used by more than 1,185 organizations in 59 countries. This extended MVAPICH library is capable of acting as a new GASNet communication conduit. Since these extensions are made to an MPI library, they do not interfere in any way with existing MPI communication infrastructure or affect MPI communication performance (MPI applications simply ignore the extensions). Our evaluation reveals that this new extended MVAPICH conduit is able to not only match the existing GASNet InfiniBand conduit in both microbenchmarks and application-level evaluation; but also to improve scalability significantly. We observe that with UPC NAS benchmarks CG and MG (class B) at 128 processes, the new design outperforms current GASNet implementation by 36% and 25%, respectively.

The rest of the paper is organized as follows. In Section 2, we describe background material for this work along with a discussion of related work. We contrast our work with existing software stack and research work. Then, in Section 3, we describe our design approach in detail. We present our evaluations in Section 4. Finally, we conclude and illustrate future directions in Section 5.

## 2. Background and Related Work

In this Section, we describe associated background work that is relevant to the contributions of this paper. We also cover other recent and related work in this area. We first start with the lowest InfiniBand layer. Then, we describe recent work in designing scalable communication techniques for MPI on InfiniBand. After that, we recount work in UPC and the GASNet communication system. Finally, we compare and contrast our work with related recent efforts to develop a common communication subsystem for multiple programming models.

### 2.1 InfiniBand Overview

InfiniBand [4] is an industry standard switched fabric that is designed for interconnecting nodes in HEC clusters. It is a high-speed, general purpose I/O interconnect that is widely used by scientific computing centers world-wide. The recently released TOP 500 rankings in June 2010 reveal that more than 40% of the computing systems use InfiniBand as their primary interconnect. The yearly growth rate of InfiniBand in the TOP 500 systems is pegged at 30%, indicating a strong momentum in adoption. It is also gaining ground in the commercial domain with the recent convergence around RDMA (Remote Direct Memory Access) over Converged Enhanced Ethernet (RoCE).

The InfiniBand specification clearly demarcates the duties of hardware (such as Host Channel Adapters (HCAs)) and software. Upper level software uses an interface called *verbs* to access the functionality provided by HCAs and other network equipment (such as switches). The verbs interface is a low-level communication interface that follows the Queue Pair model. Communication end-points are required to establish a queue pair between themselves. Each queue pair has a certain number of work queue elements. Upper level software places a work request on the queue pair that is then processed by the HCA. When a work element is completed, it is placed in the completion queue. Upper level software can detect completion by polling the completion queue.

Additionally, there are different types of Queue Pairs based on the type of transport used. There are Reliably Connected (RC) queue pairs that provide reliable transmission (retransmissions after packet losses are performed by the HCA). These RC queue pairs need to be established uniquely for each communicating pair. This implies an  $O(n^2)$  memory usage (for a system with  $N$  processes). Another type of queue pair is the Unreliable Datagram (UD). This

queue pair type does not provide reliable transmission, although it has a significant memory advantage – only one UD QP is capable of communicating with all remote processes. Thus, the memory usage of UD QP is  $O(n)$  (for a system with  $N$  processes).

Yet another feature of InfiniBand is Remote Direct Memory Access (RDMA). This feature allows software to remotely read memory contents of another remote process without any software involvement at the remote side. This feature is very powerful and can be used to implement high-performance communication protocols.

As indicated above, the verbs software layer is very low-level and the responsibility of communication protocols, buffer management mechanisms, connection management techniques etc. are left to the upper level software. The following Section 2.2 describes some recent work in designing scalable mechanisms using these verbs. Not all features are mentioned in this section and readers are encouraged to peruse the InfiniBand specification [4] for more details.

### 2.2 Scalable MPI Design over InfiniBand: MVAPICH-Aptus

MPI has been the dominant parallel programming model for the past couple of decades. It has been widely ported and several open-source implementations have been available. It has also achieved very good performance and scalability. As a result, all modern super-computers support it. During the past several years, InfiniBand has made great inroads into the HPC domain. Clusters using InfiniBand have been growing in size rapidly. One of the largest clusters is the Nebulae cluster in China with 120,640 cores. As indicated in the previous section, InfiniBand offers a low-level *verbs* interface with several types of Queue Pairs, with varying levels of services. This enables upper-level software, such as those in MPI implementations, to design flexible and high-performance connection management, buffer management, coalescing strategies, etc.

MVAPICH/MVAPICH2 [25] are high-performance implementations of MPI-1 and MPI-2 interfaces on InfiniBand, iWARP and RoCE (RDMA over Converged Ethernet). The internal design of MVAPICH has been systematically designed to achieve very good scalability by exploiting various InfiniBand features, such as Unreliable Datagrams [20], Shared Receive Queues (SRQ) [28], eXtended Reliable Connections (XRC) [22] along with connection management strategies such as On-demand connections and buffering strategies for message coalescing [19] and to improve memory efficiency. All of these optimizations have been combined into one unified runtime, called MVAPICH-Aptus [21]. To the best of our knowledge, this is the most scalable runtime on InfiniBand that offers high-performance and is open-source. Our team is in the process of bringing this framework into the MVAPICH2 library, so that MPI-2 applications can also leverage this scalable runtime.

### 2.3 UPC, PGAS and GASNet Communication System

Unified Parallel C (UPC) [29] is an emerging parallel programming language that aims to increase programmer productivity and application performance by introducing parallel programming and remote memory access constructs in the language. UPC is based on the Partitioned Global Address Space (PGAS) programming model. The PGAS programming model allows programmers to view a distributed memory supercomputer as a global address space, that may be partitioned to improve performance. There are several other PGAS programming languages, namely X10 [18], Chapel [10] and HPF [17], along with Global Address Space libraries such as Global Arrays [15].

The runtime implementations of UPC have been demonstrated to be scalable and provide very good performance to end applications through fine-grained remote memory accesses [9] and improved communication overlap [26]. In particular, the Blue Gene implementation of UPC, developed by IBM, has been demonstrated

to be highly scalable [7]. In this paper, we focus on the InfiniBand design and implementation of UPC through the popular GASNet communication library [14]. GASNet for InfiniBand clusters has only been mildly optimized with Firehose [8], which is a strategy to manage remote keys (required for RDMA). It is to be noted that as compared to the exhaustive scalability and optimization work carried out to scale MVAPICH/MVAPICH2 up to the level of 81,920 cores (of Pleiades at NASA) and 62,976-core (of Ranger at TACC), not much recent work has taken place in the InfiniBand implementation of GASNet. In particular, even on-demand connections and shared receive queues, which are considered to be fundamental requirements to scale InfiniBand runtimes, are missing in the current GASNet implementation. In this paper, we present our strategy to unify the runtimes of MVAPICH and GASNet in such a way that the InfiniBand communication schemes are unified between them in an efficient manner. Thus, our work *adds* major scalability features to the GASNet implementation on InfiniBand.

## 2.4 Common Communication Subsystems for Multiple Programming Model Support

High-performance communication libraries for HEC systems have historically been very closely tied with their corresponding programming models and their implementations. This has been due to the fact that each communication layer has been optimized heavily to suit the particular programming model and its requirements. Some libraries, such as ARMCI [16] are deployed along with MPI in a side-by-side manner to enable Global Arrays applications that occasionally use utility functions from MPI. There have been some efforts to unify the communication subsystems, such as CCS by Buntinas, et. al [12]. It is to be noted that even though the CCS layer had the potential to support multiple programming models, it did not offer the best performance as compared to GASNet (please refer to Fig. 2 of [12]), especially for short messages. Also, the authors of [12] did not evaluate whether CCS will be suitable for MPI itself as a communication layer or not. Additionally, there was no evaluation of UPC or GASNet on top of CCS. In this paper, we describe a truly unified communication library which provides the best performance to *both* UPC and MPI libraries in conjunction with thorough evaluation of the entire UPC stack through our new communication runtime.

The Low-level Application Programming Interface (LAPI) [3], is a unified communication layer that is known to support both MPI and pure LAPI communications. The LAPI layer is available only on IBM systems and is not available on general commodity cluster systems with InfiniBand. In this respect, it can potentially support GASNet implementation over LAPI along with MPI over LAPI. We have not yet found any publicly available document describing the feasibility of this approach. Additionally, publically available documents, such as [1] describes the limitations of LAPI supporting both MPI and LAPI communications. In particular, it mentions that in the event of concurrent communications, the DMA Receive Buffer may be too small to contain packets from both LAPI and MPI, and packets may be dropped, impairing performance. Further, the `MP_CSS_INTERRUPT` variable, that allows incoming messages to cause interrupts and message progress, is available to MPI API only, therefore, not progressing LAPI communications. We believe that the MVAPICH-Aptus design is not constrained by these limitations. Firstly, DMA buffer space is shared among all communication contexts. Secondly, MVAPICH-Aptus follows a “low-watermark” based flow control mechanism. In this mode, interrupts do not need to be generated for each incoming message, rather only *one* interrupt is generated when DMA buffer space is low. Thus, we believe that this is a much more flexible buffer management strategy compared to the one described for LAPI.

The MPI-2 RMA interface has been criticized before for its inadequacy as a compiler target [11]. The MPI community is actively discussing the upcoming MPI-3 standard. The MPI-3 RMA Working Group [27] is attempting to completely revamp the MPI-2 RMA semantics, in such a way that supporting PGAS languages becomes simpler at the MPI level. However, these discussions are a long way from inclusion in the formal standard. Our work is intended to influence these discussions by providing critical data about a high-performance UPC implementation in an existing MPI stack.

## 3. INCR: Design Requirements and Solutions

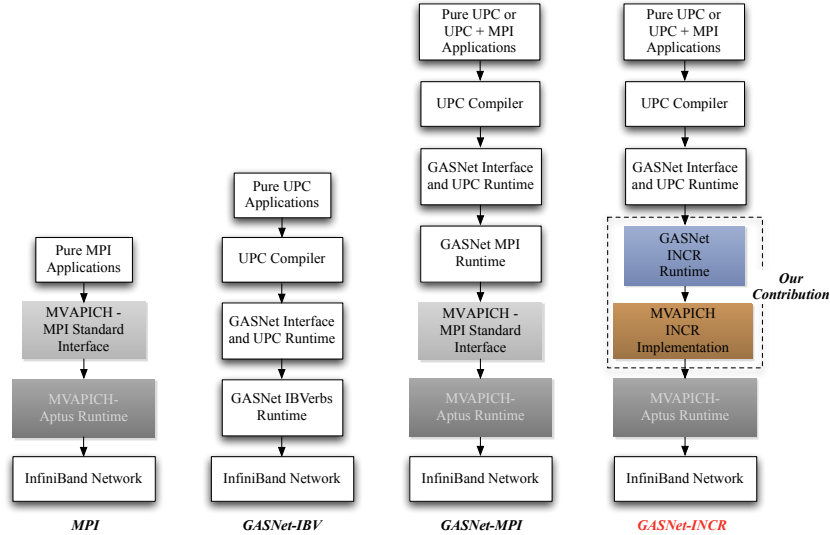
In this section, we discuss the design requirements and solutions for the implementation of the Integrated Communication Runtime (INCR) for InfiniBand clusters.

### 3.1 Design Requirements

The major design requirement of our work is to enable *simultaneous* MPI and UPC communications without imposing any performance penalties on either MPI or UPC, i.e. using our runtime, UPC programs should get identical or better performance than currently available software (and the same for MPI applications). Another design requirement is to eliminate buffer wastage inside the communication library. For example, we envision that the solution do not require separate communication resources, such as connections, buffers etc. for supporting simultaneous communications.

### 3.2 Design Approach

Currently, there exist several approaches that can be deployed on commodity InfiniBand clusters. We have outlined some of these approaches in Figure 1. Each implementation option is referred to the name given below the stack. The left-most alternative demonstrates that pure-MPI operations can work through the MPI standard implementations of MVAPICH on InfiniBand. The second alternative, marked “GASNet-IBV” indicates currently available GASNet implementation on OpenFabrics verbs. This implementation is on the lowest software layer made available by InfiniBand vendors. Currently, the two approaches “MPI” and “GASNet-IBV” can be combined together to support simultaneous communications of hybrid MPI and UPC applications. However, this approach suffers from two main drawbacks: (a) Communication progress of both UPC and MPI are separate, and it is possible to deadlock UPC and MPI by not progressing in their respective communications [5] and (b) It wastes communication resources, since both MPI and UPC allocate their own connection and buffer resources. Another approach is to use the “GASNet-MPI” stack for both MPI and UPC communications. This mode has the added advantage of progressing both MPI and UPC communications simultaneously. However, the disadvantage of this mode is that there is a mismatch of Active Messages (a fundamental design point of GASNet), and the MPI-1 point-to-point primitives. Due to this mismatch, there is a performance penalty imposed on most operations. As discussed in Section 2, MPI-2 RMA is not a viable target for compilers [11]. Thus, even though simultaneous communication is possible with “GASNet-MPI” mode, there is a performance penalty associated with it. Our design approach, presented in the right-most stack, “GASNet-INCR”, we extend the existing MVAPICH-Aptus runtime to support native active messages. In addition, we design a new communication conduit for GASNet that supports this interface. Using this approach, both UPC and MPI communication resources are shared. Further, we have the benefit of utilizing several *common* communication related optimizations that have been designed for MPI over the past several years (described in Section 2) for UPC communications as well.



**Figure 1.** An overview of various communication stack options available for MPI and UPC applications and our contribution

### 3.3 Internal Implementation Details

Our design requirements are to enable the highest possible performance for both GASNet and MPI communications. We made the following changes to implement MVAPICH-INCR.

**Packet Headers:** MPI communications stacks are required to carry message matching information triplet:  $\{\text{src}, \text{tag}, \text{context}\}$  to enable tag matching semantics of MPI. This represents an overhead of 12 bytes per packet. Additionally, packets are required to carry information for credits (such as last received packet), for flow-control. Overall, the minimum MPI packet header is 20 bytes. We have decided that this overhead is too heavy for GASNet. Thus, we have added extra packet types in the MVAPICH-Aptus implementation. Accordingly, the packets used by GASNet messages are handled completely separately from the MPI packets. This approach allows us to share common optimizations from the MPI communications layer, while providing least overhead for GASNet communications. Our GASNet-INCR has only four bytes of extra header on top of GASNet specific fields. We believe that in the future this can be improved upon as well.

**Extended Interfaces:** As indicated earlier in this Section, we have extended the communication implementation of the MVAPICH-Aptus runtime to support native Active Messages. In addition, we have also implemented the *extended* GASNet interfaces for natively supporting RDMA. Below is a short description of the interface.

*Active Messaging Interfaces:* These are the new active messaging functions that are implemented inside MVAPICH-Aptus.

- 1) `int incr_send_short_am_no_args(uint16_t dest_grank, uint16_t token, uint8_t handler)`: This function implements sending very short messages without any arguments to a remote destination. This is an optimization over the short message functions which need to send arguments. Since the handler at the remote end expects no arguments, the space for arguments in the message packet can be optimized out.
- 2) `int incr_send_short_am_with_args(uint16_t dest_grank, uint16_t token, uint8_t handler, va_list argptr, uint8_t numargs)`: This function implements sending

short messages to remote destinations. Short active messages do not carry any data payload.

- 3) `int incr_send_medium_msg(uint16_t dest_grank, uint16_t token, uint8_t handler, void * source_addr, uint32_t nbytes, va_list argptr, uint8_t numargs)`: This function implements sending medium size messages with data payload. The data payload can be sent over a bounce buffer implementation that is implemented using native RDMA [24]. It is to be noted that using MVAPICH-Aptus, RDMA connections are only opened in an on-demand basis and not an all-to-all manner during communication startup.

- 4) `int incr_send_long_msg(uint16_t dest_grank, uint16_t token, uint8_t handler, void * source_addr, uint32_t nbytes, uint32_t lkey, va_list argptr, uint8_t numargs, void * dest_addr, uint32_t rkey, incr_handle_t *handle_ptr)`: This function implements sending large messages. Since large messages may take a long time to send, this interface call supports returning a handle to the communication request. Using this handle, this call may be used in a non-blocking manner. Long messages are always sent using InfiniBand RDMA, with reliable connections (RC) established in an on-demand fashion.

*Extended Remote Memory Interfaces:* These interfaces implement the GASNet extended interface. These focus on efficient remote memory access. These functions currently do not return any error status, i.e. all errors in remote memory access are considered fatal.

- 1) `void incr_inline_put (uint16_t dest, void *rem_addr, void *local_addr, size_t nbytes, uint32_t rkey, incr_handle_t *handle_ptr)`: This function sends a very small ( $< 128$  bytes) data payloads directly to remote memory using RDMA. This utilizes the InfiniBand “inline” send operation, which reduces sender side overhead of DMAs. As long as send work queue elements (WQEs) are available, this function completes the put on returning.
- 2) `void incr_put (uint16_t dest, void *rem_addr, void *local_addr, size_t nbytes, uint32_t rkey, incr_handle_t *handle_ptr, uint32_t lkey)`: This function is used for remote put. The payload data is put directly using

RDMA write. If no send work queue elements (WQEs) are available, then this data can be buffered internally (decision made based on message size), and the call returns immediately. To test for completion of the data transfer, the handle needs to be checked upon. Since the call is implemented using RDMA put, RC connection is used (set up in an on-demand fashion).

3) void `incr_put_bulk` (uint16\_t `dest`, void \*`rem_addr`, void \*`local_addr`, size\_t `nbytes`, uint32\_t `rkey`, `incr_handle_t` \*`handle_ptr`, uint32\_t `lkey`): This function is also used for remote put. This is for supporting the GASNet put\_bulk interface, in which the local memory is guaranteed to be untouched until the put operation is completed, so that data need not be buffered internally if the work queue elements (WQEs) are not available.

4) void `incr_get` (uint16\_t `dest`, void \*`rem_addr`, void \*`local_addr`, size\_t `nbytes`, uint32\_t `rkey`, uint32\_t `lkey`, `incr_handle_t` \*`handle_ptr`): This function is used for remote get operations. This maps directly to RDMA Read operation, which uses RC connection (set up in an on-demand fashion).

## 4. Experimental Results

In this section, we compare performance evaluation results of the three design approaches: GASNet-IBV, GASNet-MPI and GASNet-INCR described in section 3.2. The existing implementation of GASNet implementation on OpenFabrics verbs is called GASNet-IBV and the implementation on MPI is called GASNet-MPI. GASNet-INCR, which is our contribution in this paper, is the GASNet implementation over extended MVAPICH-Aptus runtime.

GASNet interface consists of core APIs and extended APIs. The core API interface is a narrow interface based on Active Message paradigm. Extended APIs provide a richly expressive and flexible interface that provides medium and high-level operations on remote memory and collective operations [14]. GASNet-MPI has only the core APIs implemented where as GASNet-IBV and GASNet-INCR have the extended APIs for one sided put and get operations. As mentioned in Section 2, the GASNet-MPI implementation cannot utilize the MPI2-RMA interface for put and get due to several restrictions and is limited to only using the core APIs.

We compared these three design approaches from different angles. These include microbenchmark level performance evaluation, scalability analysis based on memory footprint, and performance analysis of different NAS benchmarks using these conduits. Detailed analysis of these performance evaluation is presented in the following sections. Section 4.2 covers the microbenchmark level performance analysis, and sections 4.3 and 4.4 covers the scalability analysis and NAS benchmark results respectively.

### 4.1 Experimental Platform

We used three different clusters for our experimental evaluations. Cluster A consists of four Intel Nehalem machines equipped with ConnectX QDR HCAs. Each node has eight Intel Xeon 5500 processors organized into two sockets of four cores each clocked at 2.40 GHz with 12 GB of main memory. Cluster B consists of 32 Intel Clovertown based systems equipped with ConnectX DDR HCAs. Each node in this cluster has eight Intel Xeon processors, organized into two sockets of four cores each clocked at 2.33 GHz with 6 GB of main memory. Cluster C consists of eight AMD Barcelona hosts. Each node has four sockets each with a Quad-Core AMD Opteron 8350 2GHz Processor with 512KB L2 cache and 2 MB L3 cache per core. Each node has a Mellanox MT25418 dual-port ConnectX HCA. RedHat Enterprise Linux Server 5 was used on all machines along with OFED version 1.4.2. We used Cluster A for the microbenchmark level experiments, Cluster B for scalability analysis and Cluster C for application level evaluations.

### 4.2 Microbenchmark Level Evaluation

We chose three representative benchmarks for the microbenchmark level performance analysis. These include the performance evaluation of UPC calls, `upc_memput`, `upc_memget` and `bupc_memput_signal`. The `upc_memput` call writes specified amount of data bytes to the remote side. Similarly, `upc_memget` fetches specified amount of data bytes from remote side. The `bupc_memput_signal` is part of the UPC extensions proposed by Berkeley Lab. It performs the same data movement semantics as that of `upc_memput`, and updates the specified semaphore in the remote side. The update of semaphore in the remote side signals the global completion of data movement. We used Cluster A for microbenchmark performance evaluation.

In `upc_memput` and `upc_memget` microbenchmarks, `upc_memput` and `upc_memget` were called in the sender thread for message sizes varying from one byte to two mega bytes. The receiver thread waits on a barrier. Time taken for the upc call for each of these message sizes is reported. We have split the latency results into two graphs, one showing the latency results for small payload sizes (1byte to 2K bytes) and other one for large payload sizes (4KB to 2MB). Micro-benchmark results are shown in Figures 2 and 3. Results indicate that the performance of GASNet-INCR is very much similar to that of the high performance GASNet-IBV. Since GASNet-MPI does not have the GASNet extended APIs, `upc_memget` and `upc_memput` calls are translated into active messages. These active messages are exchanged using MPI send/recv calls. This involves not only the overhead of translating into active messages, but also the extra MPI headers that these messages have to carry. But for GASNet-IBV and GASNet-INCR, `upc_memput` and `upc_memget` calls are translated into one sided RDMA put/get operations. This explains the huge performance difference between GASNet-MPI and the other two GASNet implementations.


We used ping-pong test in `bupc_memput_signal` microbenchmark. The sender thread calls `bupc_memput_signal` and then calls `bupc_sem_wait`. The receiver thread calls `bupc_sem_wait` and then calls `bupc_memput_signal`. The `bupc_sem_wait` call causes the thread to wait until the semaphore gets incremented. The `bupc_sem_wait` call in the receiver thread gets unblocked when the data movement initiated by the sender thread is over. Similarly, the semaphore in the sender thread gets unblocked when the data movement initiated by receiver thread is over. The test was also done for message sizes varying from one byte to two mega bytes. The one-way latency numbers are reported. Here also, the latency numbers are shown in two different graphs for fine grain analysis. The observed performance results of `bupc_memput_signal` benchmark is similar to that of `upc_memget` and `upc_memput` benchmarks. The results are shown in Figure 4.

The microbenchmark performance results indicate that the GASNet-INCR delivers similar performance as that of GASNet-IBV conduit, and it even outperforms GASNet-IBV conduit for small payload sizes in `bupc_memput_signal` microbenchmark. GASNet-MPI performs worse in all microbenchmarks.


### 4.3 Scalability Evaluation

In order to analyze the scalability aspect of the GASNet-INCR design, we conducted some scalability experiments as well. We used a simple UPC hello world program in this experiment. We measured memory foot-print of the process, with number of connections (processes) ranging from 16 to 256. The memory foot-print analysis results are shown in Figure 5. We used Cluster B for scalability analysis.


It was observed that for a 256 process UPC hello world execution, each process consumed a memory of about 227 MB in case of GASNet-INCR conduit, 265 MB in case of GASNet-IBV conduit and 237 MB in case of GASNet-MPI conduit. It can be noticed



(a) Mempup latency: Small




(b) Mempup latency: Large




(c) Mempup bandwidth


**Figure 2.** Mempup performance



(a) Memget latency: Small




(b) Memget latency: Large




(c) Memget bandwidth


**Figure 3.** Memget performance



(a) bupc Mempup latency: Small



(b) bupc Mempup latency: Large



(c) bupc Mempup bandwidth

**Figure 4.** Berkeley UPC (signal) mempup performance

that, with increase in number of connections, memory foot print increases almost linearly for GASNet-IBV conduit, where as memory foot print of GASNet-INCR and GASNet-MPI conduits remain almost constant. The low memory foot-print of GASNet-INCR and GASNet-MPI conduits is because it uses the MVAPICH-Aptus runtime, which is known to scale to tens-of-thousands of cores.

Analyzing the slope of the scalability performance results, it can be noted that, for GASNet-IBV, the memory foot-print increases by around 128KB with every additional process, where as the memory foot-print of GASNet-INCR and GASNet-MPI conduit remains constant. If we extrapolate the memory foot print for a 10,000 process run, then memory foot-print of GASNet-IBV conduit process will be around 1.4 GB, where as the memory foot-print of a GASNet-INCR conduit process will be around 250-300 MB. This shows that our design is highly scalable.

#### 4.4 Application Level Evaluation

To study the impact of our GASNet-INCR design, we evaluated its performance with a few NAS benchmarks [2]. We chose the Conjugate Gradient benchmark (CG), 3-D FFT PDE benchmark (FT) and the Multi-Grid benchmark (MG), so as to analyze GASNet designs from different angles. These benchmarks were ran for different problem sizes (class B and C) with 64 and 128 processes. We chose not to present experimental data for the NAS EP benchmark as it has very little communication. We used cluster C for these experiments.

CG benchmark has the smallest problem size (75,000 for Class B and 150,000 for Class C) among the three benchmarks and has a relatively more frequent communication pattern. The communication pattern in CG benchmark is mainly point-to-point operations; FT benchmark is dominant in collective operations; and MG benchmark is mainly utilizing point-to-point operations with 1% collective operations [13] [6]. The execution times for these different NAS benchmarks are shown in Figure 7.

The performance results observed for CG and MG benchmarks are similar as indicated in Figures 7(a) and 7(b). For 128 process CG benchmark, GASNet-INCR performs 1.36 times and 1.12 times faster as compared to GASNet-IBV conduit for classes B and C, respectively. For 128 process MG benchmark, GASNet-INCR performs 1.25 times and 1.20 times faster as compared to GASNet-IBV conduit for classes B and C, respectively. These results were a little surprising, because in the microbenchmark evaluation, we got similar performance results for GASNet-IBV and GASNet-INCR conduits.

In order to investigate further, we analyzed the communication pattern in these benchmarks. We timed the different sections of the CG benchmark. The results showed that much of the difference came from the barrier performance, in 'reduce-all' section of the benchmark. For a 64 process CG class B benchmark, the barrier time for GASNet-INCR conduit was just 3.63 seconds, where as for the GASNet-IBV conduit, it was 5.42 seconds. So we plotted the `upc_barrier` performance of the three conduits, for number of processes ranging from 2 to 128. Results indicate that the barrier performance of GASNet-IBV is not good, when the number of processes increase. For a 128 process barrier, GASNet-INCR takes 3.48 seconds where as GASNet-IBV and GASNet-MPI conduit takes 9.65 seconds and 4.95 seconds, respectively. The `upc_barrier` is implemented using active messages in the GASNet, so we suspect the active message implementation in GASNet-IBV for the poor barrier performance. We plotted the `MPI_Barrier` performance also along with `upc_barrier`. Results indicate that `MPI_Barrier` performs better than `upc_barrier` of all three GASNet implementations. This result underlines the need for hybrid programming models. Hybrid programs can make use of the highly-optimized MPI collectives and other constructs, thereby

avoiding the need for re-inventing these in new programming models.

Another surprising result was that the GASNet-MPI conduit was giving almost similar performance results in FT benchmark and it even out-performed other two conduits for 64-process FT(C) benchmark. In order to investigate this, we took the FT benchmark and timed the different sections. Here the major difference was identified in the 'all-to-all' section of the benchmark. For a 64 process class B FT benchmark, the all-to-all section took 4.67 seconds using GASNet-IBV conduit, but it just took 2.67 seconds using the GASNet-MPI conduit. In the all-to-all section of FT benchmark, all the processes fetch a chunk of around 130K bytes from the shared memory area, which is actually a `upc_memget`. The `upc_memget` implementation in the GASNet-MPI conduit is different as compared to that of GASNet-IBV and GASNet-INCR. Since the later two conduits implement GASNet Extended API for Get (`gasnet-get-bulk`), the GASNET layer translates the `upc_memget` to `gasnet_get_bulk`. Thus, in these conduits, the `upc_memget` gets translated as an RDMA Get of 130K.

But for the GASNet-MPI conduit, the Extended APIs interfaces can not be implemented (due to MPI-2 RMA inadequacies). Accordingly, GASNet translates the `upc_memget` to active messages (Long Active Message) [14]. In GASNet AMMPI implementation, the maximum long active message limit is set to 65K bytes. So in case of the GASNet-MPI conduit, the 130K `upc_memget` is translated into two 65K long active messages. Therefore, the Get operations do not start immediately, rather consist of the request-and-reply set of operations. This results in some delay between the processes, and not all Get operations start at the same time. This may result in better network utilization, as congestion is avoided. In order to verify this, we modified the GASNet-IBV conduit implementation to disable the Extended API interface and set the maximum long message limit as 65K bytes. With these modifications, we observed similar performance for GASNet-MPI conduit and GASNet-IBV conduit for the FT benchmarks. Thus, it is our understanding that the all-to-all implemented by point-to-point memgets is not a very efficient mechanism, as it robs the underlying communication layer of optimization opportunities. We also note that `MPI_Alltoall` algorithms in MPI stacks (including MVAPICH and MVAPICH2), have congestion management by using hierarchical algorithms along with pair-wise send-recv (to avoid flooding the network with numerous very large messages). We are continuing to investigate this and will include a more detailed description of this result in the camera ready version of the paper.

The application level performance results indicate that the GASNet-INCR design performs better than the GASNet-IBV conduit. It also highlights that hybrid programming models can provide better performance along with improved programmer productivity.

## 5. Conclusions and Future Work

During the past few years, there has been an increasing interest in improving programmer productivity for scientific applications. UPC and other PGAS languages have emerged as alternatives to the still popular MPI programming model. At the same time, InfiniBand technology is being widely deployed and very large clusters, consisting of many tens-of-thousands of cores are being deployed for conducting ground-breaking science. In this context, there is a need for a communication library that supports both UPC and MPI programming models, provides excellent performance for them and is highly scalable. In this paper, we have presented our designs for Integrated Native Communication Runtime (INCR) for MPI and PGAS on InfiniBand clusters. Our design and evaluation reveal that we are able to achieve our primary design objectives.






Figure 5. Memory Consumption

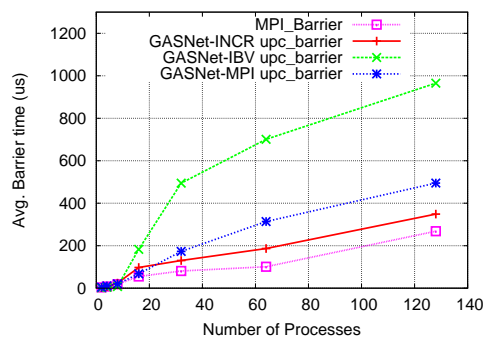



Figure 6. Barrier Performance



(a) CG Class B and C (64 procs and 128 procs)

(b) MG Class B and C (64 procs and 128 procs)

(c) FT Class B and C (64 procs and 128 procs)

Figure 7. NAS benchmarks performance

For UPC NAS benchmarks CG and MG (class B) at 128 processes, GASNet-INCR outperforms the GASNet-IBV runtime by 36% and 25%, respectively. Also the memory scalability analysis reveals that GASNet-INCR is highly scalable.

In the future, we aim to continue working in this direction. In this paper we presented performance results for UPC using our communication runtime. We plan to extend our evaluation to hybrid MPI and UPC applications. Our team is involved in optimizations of several well-known applications from the San Diego Supercomputer Center and Texas Advanced Computing Center. We plan to investigate optimization opportunities by using hybrid programming models in those applications. Additionally, we are also working towards more generalization of our proposed runtime. Our aim is to support PGAS in general in the future.

## Acknowledgments

This research is supported in part by U.S. Department of Energy grants #DE-FC02-06ER25749 and #DE-FC02-06ER25755; National Science Foundation grants #CNS-0403342, #CCF-0702675, #CCF-0833169 and #CCF-0916302. grants from Intel, Mellanox, Cisco, QLogic, and Sun Microsystems; Equipment donations from Intel, Mellanox, AMD, Appro, Chelsio, Dell, Fujitsu, Fulcrum, Microway, Obsidian, QLogic, and Sun Microsystems. We would like to thank Mr. Jonathan Perkins of The Ohio State University for

extending help in setting up the experimental systems on a short notice.

## References

- [1] Parallel environment for aix 5l mpi programming guide.
- [2] NAS Parallel Benchmarks. <http://www.nas.nasa.gov/Software/NPB/>.
- [3] Message Passing Interface and Low-Level Application Programming Interface (LAPI). <http://www-03.ibm.com/systems/software/parallel/index.html>.
- [4] InfiniBand Industry Standard Specification. <http://www.infinibandta.org/>.
- [5] Berkeley UPC User's Guide: Mixing UPC and C/C++/MPI/FORTRA. <http://upc.lbl.gov/docs/user/interoperability.shtml>.
- [6] P. B. Ahmad, A. Faraj, and X. Yuan. Communication characteristics in the nas. In *In Fourteenth IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2002)*, pages 729–734, 2002.
- [7] C. Barton, C. Casçaval, G. Almási, Y. Zheng, M. Farreras, S. Chatterje, and J. N. Amaral. Shared memory programming for large scale machines. In *PLDI '06: Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation*, pages 108–117, New York, NY, USA, 2006. ACM. ISBN 1-59593-320-4. doi: <http://doi.acm.org/10.1145/1133981.1133995>.



- [8] C. Bell and D. Bonachea. A New DMA Registration Strategy for Pinning-Based High Performance Networks. In *Workshop on Communication Architecture for Clusters (CAC), in conjunction with Int'l Symposium on Parallel and Distributed Computing (IPDPS)*, 2003.
- [9] C. Bell and D. Bonachea and R. Nishtala and K. Yelick. Optimizing Bandwidth Limited Problems Using One-Sided Communication and Overlap. In *Int'l Symposium on Parallel and Distributed Computing (IPDPS)*, 2006.
- [10] Cray, Inc. The Chapel Programming Language. <http://chapel.cray.com/>.
- [11] D. Bonachea and J. Duell. Problems with using MPI 1.1 and 2.0 as Compilation Targets for Parallel Language Implementations. *Intl. J. High Performance Computing and Networking*, 1(1/2/3):91–99, 2004.
- [12] D. Buntinas and W. Gropp. Designing a Common Communication Subsystem. In *Euro PVM/MPI Conference*, 2005.
- [13] D. H. Bailey, T. Harris, R. Van der Wiggart, W. Saphir, A. Woo and M. Yarrow. The NAS Parallel Benchmarks 2.0. In *Technical Report NAS-95-010*. NASA Ames Research Center, 1995.
- [14] Editor: Dan Bonachea. GASNet specification v1.1. Technical Report UCB/CSD-02-1207, U. C. Berkeley, 2008.
- [15] Environmental Molecular Sciences Laboratory and Pacific Northwest National Laboratory. The GA Toolkit. <http://www.emsl.pnl.gov/docs/global/>.
- [16] Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory. Aggregate Remote Memory Copy Interface (ARMCI). <http://www.emsl.pnl.gov/docs/parsoft/armci/>.
- [17] High Performance Fortran Forum (HPFF). High Performance Fortran. <http://hpff.rice.edu/>.
- [18] IBM Research. X10 Programming Language. <http://x10-lang.org/>.
- [19] M. Koop, T. Jones, and D. K. Panda. Reducing Connection Memory Requirements of MPI for InfiniBand Clusters: A Message Coalescing Approach. In *7th IEEE Int'l Symposium on Cluster Computing and the Grid (CCGrid07)*, May 2007.
- [20] M. Koop, S. Sur, Q. Gao, and D. K. Panda. High Performance MPI Design using Unreliable Datagram for Ultra-Scale InfiniBand Clusters. In *21st ACM International Conference on Supercomputing (ICS07)*, June 2007.
- [21] M. Koop, T. Jones, and D. K. Panda. MVAPICH-Aptus: Scalable High-Performance Multi-Transport MPI over InfiniBand. In *IEEE Int'l Parallel and Distributed Processing Symposium (IPDPS 2008)*, April 2008.
- [22] M. Koop, J. Sridhar, and D. K. Panda. Scalable MPI Design over InfiniBand using eXtended Reliable Connection. In *IEEE Int'l Conference on Cluster Computing (Cluster 2008)*, September 2008.
- [23] Lawrence Berkeley National Laboratory and University of California at Berkeley. Berkeley UPC - Unified Parallel C. <http://upc.lbl.gov/>.
- [24] J. Liu, J. Wu, and D. K. Panda. High Performance RDMA-Based MPI Implementation over InfiniBand. In *In 17th Annual ACM International Conference on Supercomputing (ICS)*, 2003.
- [25] Network-Based Computing Laboratory. MVAPICH: MPI over InfiniBand, 10GigE/iWARP and RoCE. <http://mvapich.cse.ohio-state.edu/>.
- [26] R. Nishtala and P. Hargrove and D. Bonachea and K. Yelick. Scaling Communication-Intensive Applications on BlueGene/P Using One-Sided Communication and Overlap. In *Int'l Symposium on Parallel and Distributed Computing (IPDPS)*, 2009.
- [27] Remote Memory Access Working Group. MPI3 RMA Wiki. <https://svn.mpi-forum.org/trac/mpi-forum-web/wiki/RmaWikiPage>.
- [28] S. Sur, L. Chai, H.-W. Jin, and D. K. Panda. Shared Receive Queue Based Scalable MPI Design for InfiniBand Clusters. In *IEEE Int'l Parallel and Distributed Processing Symposium (IPDPS 2006)*, April 2006.
- [29] UPC Consortium. UPC Language Specifications, v1.2. Technical Report LBNL-59208, Lawrence Berkeley National Lab, 2005.