

HiCCL: A Hierarchical Collective Communication Library

Mert Hidayetoglu
Stanford University
merth@cs.stanford.edu

Simon Garcia de Gonzalo
Sandia National Laboratories
simgarc@sandia.gov

Elliott Slaughter
SLAC National Accelerator
Laboratory
eslaught@slac.stanford.edu

Pinku Surana
SLAC National Accelerator
Laboratory
suranap@slack.stanford.edu

Wen-mei Hwu
University of Illinois at
Urbana-Champaign, Nvidia
w-hwu@illinois.edu

William Gropp
University of Illinois at
Urbana-Champaign
wgropp@illinois.edu

Alex Aiken
Stanford University
aaiken@stanford.edu

Abstract—HiCCL (Hierarchical Collective Communication Library) addresses the growing complexity and diversity in high-performance network architectures. As GPU systems have evolved into networks of GPUs with different multilevel communication hierarchies, optimizing each collective function for a specific system has become a challenging task. Consequently, many collective libraries struggle to adapt to different hardware and software, especially across systems from different vendors. HiCCL’s library design decouples the collective communication logic from network-specific optimizations through a compositional API. The communication logic is composed using multicast, reduction, and fence primitives, which are then factorized for a specified network hierarchy using only point-to-point operations within a level. Finally, striping and pipelining optimizations streamline execution.

Performance evaluation of HiCCL across four different machines—two with Nvidia GPUs, one with AMD GPUs, and one with Intel GPUs—demonstrates an average $17\times$ higher throughput than the collectives of highly specialized GPU-aware MPI implementations, and competitive throughput with those of vendor-specific libraries (NCCL, RCCL and OneCCL), while providing portability across all four machines.

I. INTRODUCTION

Communication libraries in high-performance computing (HPC) for scientific applications and machine learning offer *collectives*—standard functions such as Scatter, Reduce, and All-Reduce [30] that involve coordinated communication across processors. Our evaluation on leadership-class GPU systems revealed that the throughput of available library collectives is either not optimized in GPU-aware MPI implementations (e.g., MPICH [15] and OpenMPI [12]) or have limited cross-vendor portability (NCCL, RCCL and OneCCL). These performance and portability issues motivated us to develop a library design capable of producing throughput-optimized collectives portable across diverse network architectures and GPU programming models. However, there are many collective functions, each requiring a unique optimization strategy depending on the network architecture [31], making it unproductive to hand optimize a collective for one system and then repeat the process when porting to another system.

Modern systems typically feature a network comprised of multiple compute nodes each housing multiple GPUs [5], [9]–[11], which results in a nonuniform interconnect architecture among the GPU endpoints. Specifically, the network is hierarchical: GPUs within the same node communicate via a high-throughput intranode network, while GPUs on different nodes communicate through a lower-throughput network that links the nodes. Consequently, the bandwidth available for GPU communication is influenced by their placement within the hierarchical network architecture [28]. This hierarchy can extend to multiple levels, such as GPU dies in a device, devices across NUMA nodes, or nodes within racks, where communication bandwidth decreases at higher levels of the hierarchy, making data transfer between more “distant” GPUs costlier. Efficient hierarchical communication relies on having lower volume at higher levels of the hierarchy [3], [19], [24], which is currently achieved with carefully hand-designed and optimized implementations of collectives.

Our aim is to provide both high throughput and performance portability for collective communication operations across diverse hierarchical networks with GPUs of different vendors¹ while also automating much of the process of building optimized collectives. We present a method for constructing collective communication operations that achieves these goals in two steps: First, we specify the communication pattern for a collective operation as a composition of multicast, reduction, and fence primitives (Section III). Second, we identify and implement a number of optimizations that can be mechanically applied to the given composition (Section IV). These optimizations include striping communication across multiple network interface cards (NICs), pipelining communication in multiple stages, and whether to use a customized tree, ring, or hybrid virtual topology to best match the target network hierarchy. Combining these optimizations with a machine-specific description that fills in important constants (e.g., the number of levels in the hierarchy and the base communication

¹Today’s major GPU vendors are Nvidia, AMD, and Intel.

library to use for each level), users can create highly optimized collective operations tuned to a specific network. When porting between machines, only the machine description needs to change; the specification of the logic of the collective operation can be automatically optimized for the target network using the new machine description.

Our optimized collectives rely only on point-to-point communication functions of standard communication libraries and, when available, vendor-provided capabilities for a specific machine. This implementation strategy gives us both maximum flexibility to use the best point-to-point primitives for a specific level of a network hierarchy while also ensuring maximum portability (Section V).

We have implemented these ideas in a hierarchical collective communication library, HiCCL, which provides an API to build collective functions and apply hierarchical optimizations. HiCCL provides performance portability across systems of different shapes and sizes and of different vendors. We make the following contributions:

- We introduce a machine-agnostic specification of collective functions using multicast, reduction and fence primitives. We show these primitives are sufficient to express all collective functions in the MPI standard and their alternative implementations.
- We identify a set of unified hierarchical optimizations that are applicable to any collective function composed with the proposed primitives. We show how the optimizations adapt to diverse, modern GPU systems, and are sufficient to saturate the throughput of the various networks.
- We introduce HiCCL²: a hierarchical communication library which integrates multiple communication capabilities without relying on existing collective functions. We demonstrate HiCCL’s performance portability by matching or outperforming available MPI and vendor-provided libraries on different systems with Nvidia, AMD and Intel GPUs.

We evaluate HiCCL on four current HPC systems—Delta [11] and Perlmutter [5] with Nvidia GPUs, Frontier [10] with AMD GPUs, and Aurora [9] with Intel GPUs—and on eight collective functions that are listed in Table I. HiCCL achieves $17\times$ geomean speedup over the native MPI implementations across all platforms. Furthermore, it delivers competitive ($1.27\times$ against NCCL/RCCL), and in some cases superior ($12.1\times$ against OneCCL), throughput compared to the libraries provided by the GPU vendors.

II. BACKGROUND

This section provides an overview of collective functions and hierarchical communication on multi-GPU, multi-NIC node architectures.

A. Conventional Libraries and Collective Functions

MPI [29] is a well-established standard for message passing in distributed computing, with numerous implementations [12], [15], [20], [26]. GPU-aware implementations of

TABLE I: Collectives of MPI, NCCL, RCCL and OneCCL.

| Collective | MPI | NCCL / RCCL | OneCCL |
|----------------|----------------|---------------|-----------------|
| Scatter | MPI_Scatter | | |
| Broadcast | MPI_Bcast | ncclBcast | ccl::broadcast |
| Gather | MPI_Gather | | |
| Reduce | MPI_Reduce | ncclReduce | ccl::reduce |
| All-to-all | MPI_Alltoall | | ccl::alltoall |
| All-gather | MPI_Allgather | ncclAllgather | ccl::allgather |
| Reduce-scatter | MPI_Reduce_sc. | ncclReduceSc. | ccl::reduce_sc. |
| All-reduce | MPI_Allreduce | ncclAllreduce | ccl::allreduce |

MPI are available for almost all systems, including our test systems. NCCL [23] is a vendor-provided library, specifically developed for Nvidia GPUs and based on CUDA. RCCL [1] is an analogous library provided by AMD, mirroring NCCL’s API for compatibility but based on HIP. OneCCL [21] is Intel’s collective communication library. The NCCL, RCCL and OneCCL libraries offer the standard collective communication functions in Table I, while MPI’s API offers additional functions not shown [16].

B. Hierarchical Communication

Hierarchical GPU networks involve groups of GPUs organized into nested, multi-level structures, where groups within a level share the same number of GPUs. Normally inter-group communication bandwidth is lower than intra-group bandwidth on each level, so the aim of hierarchical communication strategies is to minimize data transfer volumes between groups at all levels.

As an example, we illustrate a simple two-level hierarchy with two groups of nodes, each with three GPUs. We describe a process for broadcasting d bytes from GPU 0 to all other GPUs. Figure 1(a) shows *direct* communication, where orange represents communication within the *root* node (the node with GPU 0) and blue represents communication across nodes. All five transfers move all d bytes, which is unnecessary. In contrast, *hierarchical* communication breaks the communication across nodes into two stages as shown in Figure 1(b): The first stage sends a single copy between the nodes (blue), and the second stage distributes the data within the sending (orange) and receiving (maroon) nodes. While direct communication has only a single stage where all transfers are done in parallel, in general inter-node bandwidth limitations will result in the direct strategy being much slower




Fig. 1: Broadcasting d bytes across six GPUs with (a) direct and (b) hierarchical methods. Each black dot corresponds to a GPU endpoint. Each set of three GPUs corresponds to a compute node. (a) Direct implementation redundantly moves three copies of data across nodes (in blue). (b) Hierarchical optimization moves a single copy across nodes (in blue), and distributes additional copies within each node (maroon).

²<https://github.com/merthidayetoglu/HiCCL>

than the hierarchical strategy. Furthermore, pipelining can be used to hide the overhead of multiple stages, a technique further discussed in Section IV-E.

C. Communications Across Multi-NIC Nodes

To provide adequate compute-to-communication ratios, multi-GPU nodes now incorporate multiple NICs [8], [32]. Our experiments with other communication libraries on our test systems show that each process is assigned to a single NIC statically throughout the lifetime of an application. As a result, in the common case where each process manages a single GPU, each GPU is implicitly bound to a single NIC. Consequently, both the direct and hierarchical implementations in Figure 1 utilize only a single NIC for communication across nodes. We address this limitation by offering multi-NIC striping with HiCCL, detailed in Section IV-C.

When using multiple NICs to handle the communication from a single GPU, it is important to understand exactly how that binding is done. Figure 2 shows g -to- k bindings within a hypothetical node of g GPUs and k NICs. When g is an exact multiple of k , we assign GPUs to NICs in a load-balanced way with (a) packed and (c) one-to-one mappings, which achieve the full bandwidth across nodes when all GPUs send/receive the same amount of data. However, when g is not a multiple of k , we use a (b) round-robin assignment that may lead to load imbalance across the NICs. Consequently, the utilization is only 75% of the theoretical bandwidth when all GPUs send/receive the same amount of data in Figure 2(b). This under-utilization will have implications in our evaluation in Section VI-C5.




Fig. 2: Various associations across g GPUs and k NICs per node ($k \leq g$). In our test systems, each GPU is logically bound to a single NIC via (a) packed, (b) round-robin, or (c) one-to-one associations.

III. COMPOSITION OF COLLECTIVES

HiCCL has a compositional design built on three primitives: *multicast*, *reduction*, and *fence*.

A. Collective Primitives

The simplest communication building block is point-to-point communication between two GPUs. Directly expressing collectives using point-to-point semantics often results in sub-optimal performance due to obscuring opportunities for the optimizations discussed in Section II-B. Therefore, we find it beneficial to use three higher-level primitives, which are ultimately implemented using point-to-point communication after optimizations have been applied.




Fig. 3: The (a) multicast and (b) reduction primitives form a simple tree structure with one root and multiple leaves.

The *multicast* primitive $M(i, j, d)$ expresses a one-to-many dependency between GPUs shown in Figure 3 (a), where the *root* GPU with *rank* i replicates d bytes of data to multiple *leaf* GPUs whose ranks are represented by the vector j . The leaf GPUs may be a sparse subset of all GPUs.

The *reduction* primitive $R(i, j, d, op)$ expresses many-to-one dependencies shown in Figure 3 (b), where data of size d at the leaf GPUs i are reduced to a single datum at the root GPU j . The reduction primitive additionally involves a computational operation such as sum, max, logical or, etc.

If the leaf set involves all GPUs, the multicast and reduction primitives correspond to the traditional Broadcast and Reduce collectives, respectively. On the other hand, when there is a single GPU in the leaf set, the primitives simplify naturally to point-to-point communication (with an omitted unary operation in the case of reduction).

Multiple multicast and reduction primitives can be composed by registering them in a persistent global communicator, `HiCCL::Comm<T>`; the semantics are that all registered primitives will be executed in parallel. The communicator is templated over the communicated data type T . Each primitive registration is made with an invocation of the C++ API shown in Listing 1. The registration functions accept pointers to the communication buffer (`sendbuf / recvbuf`), number of elements (`count`), ranks of the sender and receiver GPUs in scalar (i and j), and ranks of the receiver and sender GPUs in vector form (`j_vec` and `i_vec`) when registering a multicast or reduction primitive, respectively. The operation for a reduction is expressed by the last argument (`HiCCL::op`) of the registry function in Listing 1.

```
void Comm<T>::add_multicast(T *sendbuf, T *recvbuf, size_t
count, int i, std::vector<int> j_vec);
void Comm<T>::add_reduction(T *sendbuf, T *recvbuf, size_t
count, std::vector<int> i_vec, int j, HiCCL::op);
void Comm<T>::add_fence();
```

Listing 1: C++ API for registering primitives.

B. Single-Step Collectives

A single-step collective is composed of one or more primitives that are executed concurrently in a single step. If there are any race conditions between primitives, the result is undefined. Table II (Single) shows the realization of the standard collective communication functions using single-step collective functions. For consistency, the largest buffer size is chosen as dp , where p is the total number of GPUs. Summations Σ

TABLE II: Composition of Collective Functions on p Processes

| # Steps | Collective | Composition |
|----------|----------------|---|
| Single | Broadcast | $M(i, \mathbf{U}, dp)$ |
| | Reduce | $R(\mathbf{U}, j, dp, op)$ |
| | All-gather | $\sum_i M(i, \mathbf{U}, d)$ |
| | Reduce-scatter | $\sum_j R(\mathbf{U}, j, d, op)$ |
| | All-reduce | $\sum_j R(\mathbf{U}, j, dp, op)$ |
| | Scatter | $\sum_j R(i, j, d, op)$ |
| | Gather | $\sum_i M(i, j, d)$ |
| | All-to-all | $\sum_i \sum_j M(i, j, d)$ |
| Multiple | Broadcast | All-gather · Scatter |
| | Reduce | Gather · Reduce-scatter |
| | All-gather | Broadcast · Gather |
| | Reduce-scatter | Scatter · Reduce |
| | All-reduce | Broadcast · Reduce All-gather · Reduce-scatter |

express the parallel composition of multiple primitives into a single-step collective (the vector \mathbf{U} represents all participating GPUs). The range of the summations is $(0, p - 1)$. The summations (and so the registration of primitives) can be performed in any order. For example, Broadcast and Reduce can be expressed with a single primitive, which directly maps to a single line of code in HiCCL. All-to-all requires p^2 point-to-point primitives, and can be expressed with two nested loops and a single call to a primitive in three lines of HiCCL API code. The rest of the single-step collectives require p primitives that can be written using a loop with a single primitive in two lines of API code.

The top half of Figure 4(a) visualizes the composition of a Reduce-scatter on three processes, where the initial data is $3d$ bytes per GPU. It takes three primitives, R_0 , R_1 , and R_2 , to reduce the partial data with length d on each GPU.

Single-step collective design may yield redundant data movement. For example, the All-reduce in Table II (Single) is not efficient because the total data moved is dp^2 . This problem can be solved by formulating All-reduce as Reduce-scatter followed by an All-gather, which moves data of size dp (see Figure 4). The problem is that the All-gather depends on the result of the Reduce-scatter, violating the single-step design principle. We next describe the design of such multi-step collectives.

C. Multi-Step Collectives

A multi-step collective is a sequence of single-step collectives, where each step depends on the previous step. For composing multi-step collectives, HiCCL exposes a *fence* to express data dependencies. Table II (Multiple) shows example formulations of composite collectives as a sequence of two single-step collectives. In this algebraic formulation, the order of operations is from right to left, where \cdot represents the fence between the operations.

We use All-reduce³ as a motivating example because its multi-step form is more efficient than the single-step form. Figure 4 shows the composition of a Reduce-scatter followed

³All-reduce performance is critical in scientific simulation and machine learning applications.

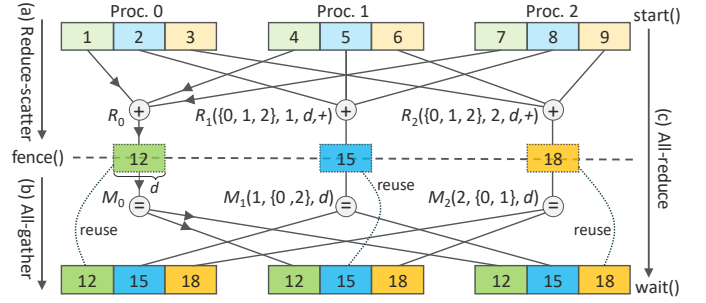


Fig. 4: Composition of (c) All-Reduce function as (a) Reduce-Scatter followed by an (b) All-Gather on three processes. The registration takes three reduction primitives, followed by a fence, and then followed by three multicast primitives. The dashed edges on the broadcasts can be omitted for in-place implementation.

```

1 using namespace HiCCL
2 // persistent communicator
3 Comm<float> comm();
4 // step 1) register Reduce-scatter using primitives
5 for(int j = 0; j < numproc; j++)
6     comm.add_reduction(sendbuf + j * count, count, all, j, op::sum);
7 // step 2) register fence to express data dependency
8 comm.add_fence();
9 // step 3) register All-gather using primitives
10 for(int i = 0; i < numproc; i++)
11     comm.add_multicast(recvbuf + i * count, recvbuf + i * count, count, i, others);
12 // optimization parameters for Aurora
13 std::vector<int> hierarchy{numproc / 12, 6, 2};
14 std::vector<Library> library{MPI, IPC, IPC};
15 int stripe(1); // number of stripes, default: 1
16 int ring(1); // node count in the ring, default: 1
17 int pipeline(16); // pipeline depth, default: 1
18 // initialization
19 comm.init(hierarchy, library, ring, stripe, pipeline);
20
21 comm.start(); // nonblocking start
22 // do other things...
23 comm.wait(); // blocking wait
    
```

Listing 2: Sample program for composing All-reduce as a Reduce-scatter followed by an All-gather.

by an All-gather that is functionally equivalent to the single-step All-reduce, but has higher throughput. In the multi-step algorithm, each GPU first reduces partial data from all GPUs and then broadcasts the result. To build this pattern, the user registers 1) p reduction primitives, 2) a fence, and 3) p multicast primitives in sequence, as shown in Listing 2 Lines, 4–11, where *all* is the vector $\{0, 1, 2, \dots, p-1\}$. Note that the receive buffer is reused for reducing partial data and therefore *others* represents the vector of all ranks but that of the root GPU.

The fence is a mechanism to express data dependencies between collections of primitives. HiCCL enforces the fine-grain dependencies between individual primitives in different steps without a global synchronization. For example, in Figure 4, M_0 depends on R_0 , M_1 depends on R_1 , and M_2 depends on R_2 . During execution, M_0 executes immediately after R_0 's completion, independent of M_1 and R_2 . However, it is inefficient for M_0 to wait until the output of R_0 is complete. Pipelined execution (Section IV-E) solves this problem by overlapping the execution of all primitives (on different data)

without violating data dependencies that are expressed with the logical fences.

Listing 2 presents registration of primitives into the communicator created in Line 3. Once the composition of primitives is registered, the communicator is initialized in line 19 of Listing 2 with the optimization parameters (Section IV-A). HiCCL provides `start()` and `wait()` functions for running the collective function. The former initiates the optimized communications from the CPU (although they take place on GPUs) and returns immediately. The latter blocks the CPU until the communication buffers on the corresponding GPU are safe to be reused.

IV. OPTIMIZATIONS

HiCCL applies a set of hierarchical optimizations to any collective pattern built with the primitives in Section III-A. These optimizations depend on a set of machine-specific parameters that are explained in this section.

A. Optimization Space

HiCCL’s optimization space is described with five parameters:

- 1) Integer factors of p for describing the network hierarchy (a vector).
- 2) The choice of implementation library for point-to-point communication for each level (a vector).
- 3) The striping factor for NICs (s).
- 4) The number of nodes for a ring (n).
- 5) Pipelining depth (m).

The parameters 1, 3 and 4 depend on the machine architecture, while 2 depends on the communication software stack and 5 depends on the message length. Note that HiCCL does not provide its own point-to-point communication operations, rather HiCCL can leverage the best available operations, including using different libraries for different levels of the communication hierarchy.

The optimizations are applied at the initialization step, after the collective composition (Section III) is defined. As an example, Listing 2 shows the parameters for a specific system in Lines 12–17. HiCCL does not automatically select these parameters, which are part of the input; however, in our experiments we found that we were able to reuse the same description of the network hierarchy for all collective communication operations on a particular machine. The parameters represent a virtual communication hierarchy that need not match the physical communication hierarchy, but of course the best performance will be achieved when the specified hierarchy matches the underlying machine.

B. Hierarchical Tree Structure

To exploit different network architectures, we parameterize the shape of the network hierarchy (see Figure 5) using a vector of integer factors (Listing 2, Line 13) that specifies a multilevel communication tree. Figure 5 (e) $\{2, 6, 2\}$ represents a three-level tree structure of 24 GPU endpoints. Assume there are two nodes, each node has six devices, and each device

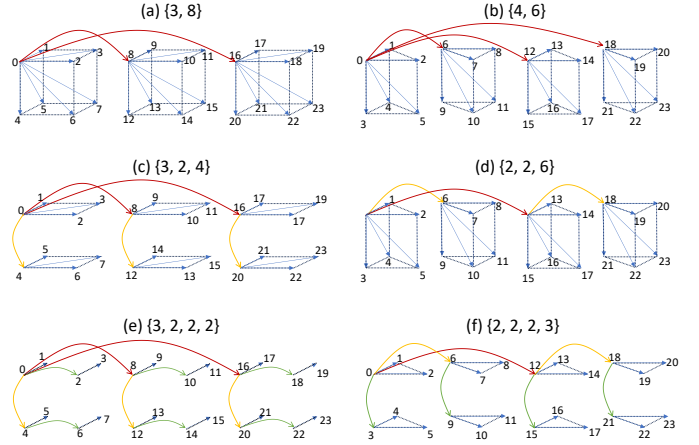


Fig. 5: Various tree structures and their notations across 24 GPUs. The examples shows (a)–(b) two, (c)–(d) three, and (e)–(f) four levels of hierarchies. The colors represents different communication links across: level 1 (red), level 2 (yellow), level 3 (green), and leaf (blue) levels. HiCCL implements each level with a the chosen communication library.

has two GPU dies. At Level 3, the first factor 2 partitions the endpoints into two groups of 12. At Level 2, the two nodes are sub-divided into six groups of two dies each. The leaf Level 1 specifies that the two dies within a device form a single group.

HiCCL assumes that the rank of each process/GPU is assigned in a way that reflects the network hierarchy; for example, Figure 5(c) represents a machine with 4 GPUs per node. Every sequence of four rank IDs, where the first is divisible by 4, represents the GPUs in a single node. Thus, the desired grouping of GPUs at each level is fully determined by the machine description and the numbering of ranks. The examples in Figure 5 are shown for a full set of leaves (GPUs). In case of custom collectives, the tree structure is pruned according to the sparsity of the leaf GPUs.

In practice, each level of the communication hierarchy has separate hardware and software, which means that different libraries may have very different performance at different levels of the hierarchy. As noted above, HiCCL allows specific libraries to be used at specific levels of the hierarchy to optimize communication for a specific system. For the mixed-library implementation of HiCCL, the library used for communication at each level of the hierarchy is specified as in Listing 2, Line 14. See Section V-A for options.

C. Multi-NIC Striping

Consider a broadcast factorized on 12 GPUs as $\{2, 2, 3\}$. The last level represents nodes with three GPUs each. The resulting tree structure has three levels, which are shown with dashed lines in Figure 6 (a). The message hops across nodes as ① orange dashed hop (g0-to-g6) and ② green dashed hops (g0-to-g3 and g6-to-g9). Then each staging GPU (g0, g3, g6, g9) multicasts the data within nodes with ③ red dashed hops. The problem is that each staging GPU is potentially bound to a single NIC (Section II-C), and therefore the hops across nodes underutilize the multi-NIC node architecture.

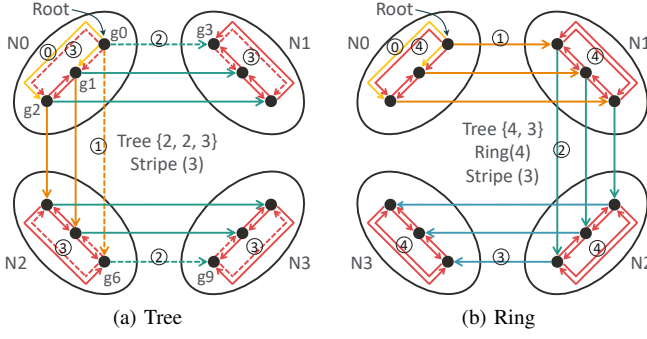


Fig. 6: (a) Tree and (b) ring factorizations on four nodes with three GPUs each. Striping forms multi-rail patterns that utilize all GPUs (and hence all NICs). With striping, (a) tree forms four stages ①–④ and (b) ring forms five stages ①–⑤, where each stage depends on the previous. The dashed route in (a) tree is shown for a single stripe originating from the root.

The aim of striping is to use all GPUs (and hence all NICs) in the inter-node communication (Section II-C). This heuristic splits each primitive into $s = 3$ parallel “stripes” in the corresponding root node, since there are $g = 3$ GPUs per node, which is achieved by an additional set of intra-node communication, as shown with ① solid golden hops. Striping is internally composed as $\sum_j^{g^1, g^2} R(0, j, d/s)$, followed by a fence (Section III-C). For preserving the original dependencies, each GPU in the root node (g_0, g_1, g_2) must multicast the corresponding partial data to all GPUs, i.e., $\sum_i^{g^0, g^1, g^2} M(i, \text{others}, d/s)$, where *others* represent all GPUs but i . The three parallel branches are further factorized recursively down to the point-to-point dependencies. As a result, striped factorization produces *multi-rail* communication patterns, automatically utilizing all NICs in the participating nodes. Figure 6 shows striping of (a) tree (making all striped lines solid) and (b) hybrid ring+tree (explained next).

D. Hybrid Ring+Tree Topology

A ring forms a communication chain across n conceptual nodes. Figure 6 (b) gives an example for $\text{ring}(4)$, where 4 is the number of nodes. To fully utilize the NICs, we employ striping by setting $\text{stripe}(3)$ because there are three GPUs per node, resulting in the extra ① golden hops. The striped ring starts from the root node and each stripe unfolds across ① (N_0 to N_1), ② (N_1 to N_2), and ③ (N_2 to N_3), until terminating in the last leaf node (N_3). Finally, the partial data is assembled on all leaf GPUs with ④ red hops, that are factorized using a tree within each node, resulting in a hybrid ring+tree pattern. A tree-only communication pattern is achieved using $\text{ring}(1)$, where there is only one conceptual node with all (p) GPUs and effectively no ring.

Overall, HiCCL factorizes each primitive with 1) striping, 2) ring, and 3) tree (in this order)—down to a dependency graph composed of multiple point-to-point communication stages. In the execution of this graph, each stage depends on the previous one, causing idle GPU time. To minimize idle time, HiCCL overlaps the communication stages with a generalized pipeline as we explain next.

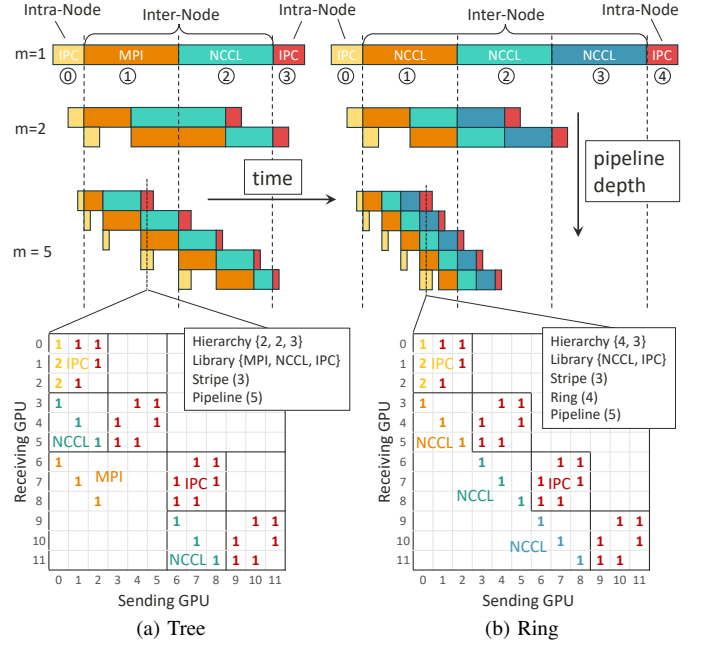


Fig. 7: Pipelining and the overlapped pattern of hierarchical communications for broadcast with (a) tree and (b) ring+tree virtual topologies. The machine models that produce these patterns are noted in the figure. In this case, HiCCL employs different implementation libraries (specifically, MPI, NCCL or IPC) for mixed communications across levels of the hierarchy.

E. Pipelining

HiCCL employs a pipelining optimization that overlaps communication across all steps of a multi-step collective operation. A sequential execution pipeline is made up of a series of stages that are marked with numbers as shown in Figure 7 ($m = 1$), where m is the pipeline depth. Each stage involves 1) a collection of point-to-point communications, or 2) the former followed by a set of computations (if a reduction is involved) to be executed on GPUs.

Figure 7(a) shows the pipeline for the tree example in Figure 6(a). When $m = 1$, stage ① corresponds to the intra-node striping, ①–② corresponds to a two-level binary tree across nodes, and ③ corresponds to the intra-node assembly. Similarly, Figure 7(b) shows the pipeline for the ring example in Figure 6(b). The original ring pipeline ($m = 1$) requires three stages across nodes that takes approximately 1.5 times longer to complete. By overlapping communications across nodes the ring is two times faster than a tree for this example.

To overlap the execution of multiple stages, without violating data dependencies, we partition the original payload (d/s) bytes across m channels in the pipeline, where each channel executes on $d/s/m$ bytes at a time. We overlap all channels in the final form of the execution pipeline as seen in Figure 7 ($m = 5$). In the final form, there are a few partially overlapped stages in the “warm-up” and “wind-down” of the pipeline, whereas the middle stages are fully overlapped. To overlap all stages, the pipeline must be deeper than the number of stages, which requires at least four channels for (a) tree and five channels for (b) ring in the examples in Figures 6–7.

The overlapped communication pattern changes across the stages of the pipeline. To illustrate, we represent the fully overlapped pattern with a hierarchical communication matrix in Figure 7 (bottom) for the (a) tree and (b) ring examples. Each non-zero entry in the matrices corresponds to a point-to-point communication. All entries are executed with the specified library for that level of the hierarchy. For demonstrating the mixed use of libraries, (a) tree uses ① MPI across the two groups of 6 GPUs, ② NCCL across nodes in the same group, and ③ IPC within nodes, i.e., 3×3 diagonal blocks in the matrices in Figure 7 (bottom). In (b) ring, we assume no hierarchy across nodes, and we use IPC in the diagonal blocks and NCCL elsewhere.

F. Performance Model

The communication cost for the Broadcast examples with pipelining are shown for (1) ring and (2) tree on n conceptual nodes with no physical hierarchy. The first and second terms correspond to inter-node and intra-node communication cost, respectively. Here, the variables represent α : communication latency, m : pipeline depth, k : number of NICs (per node), f : NIC bandwidth, and d : message length.

$$t_{\text{ring}} = \left(\alpha + \frac{d}{kf m} \right) (n + m - 2) + \mathcal{O}(d/m) \quad (1)$$

$$t_{\text{tree}} = \left(\alpha m + \frac{d}{kf} \right) \log n + \mathcal{O}(d/m) \quad (2)$$

We model the intra-node cost as $\mathcal{O}(d/m)$ because the details depend on the specific node architecture. Regardless, pipelining hides that cost and the residual intra-node overhead shrinks with $1/m$, as clearly seen in Figure 7 (red stages). Asymptotically, an infinitely deep pipeline ($m \rightarrow \infty$) with zero latency ($\alpha = 0$) hides the intra-node communication overhead completely. In practice, pipelining is most useful for large message sizes, where the latency term is negligible. Asymptotic pipelining yields $t_{\text{ring}} \sim d/kf$ that does not depend on the number of nodes, i.e., $\mathcal{O}(1)$. On the other hand, $t_{\text{tree}} \sim d \log n / kf = \mathcal{O}(\log n)$. On four nodes ($n = 4$) ring is theoretically two times faster than tree, and we will show this experimentally in Section VI-C.

TABLE III: Asymptotic Collective Throughput in GB/s

| Broadcast Reduce | Gather / All-Gather Scatter / Reduce-Scatter | All-Reduce | All-to-All |
|---|---|-----------------------|-----------------------|
| kf | $kf \frac{p}{p-g}$ | $kf \frac{p}{2(p-g)}$ | $kf \frac{p}{g(p-g)}$ |
| p : total # participated GPUs, g : # GPUs per node, k : # NICs per node, and f : rated NIC bandwidth in GB/s. | | | |

As our theoretical baseline of all collectives in Table I, we consider the upper bounds for throughput (GB/s) as shown in Table III. This upper bound solely depends on the total number of participating GPUs, collective pattern, the number of NICs and GPUs per node, and the rated bandwidth per NIC.

TABLE IV: Summary of Node Architecture of Test Systems

| System | CPUs | GPUs | NICs | B/W [†] |
|------------|--------------|---------------|---------|------------------|
| Delta | 1 AMD EPYC | 4 Nvidia A100 | 1 SS-11 | 25 GB/s |
| Perlmutter | 1 AMD EPYC | 4 Nvidia A100 | 4 SS-11 | 100 GB/s |
| Frontier | 1 AMD EPYC | 8 AMD MI250x* | 4 SS-11 | 100 GB/s |
| Aurora | 2 Intel Xeon | 12 Intel PVC* | 8 SS-11 | 200 GB/s |

*Each AMD MI250x and Intel PVC device involves two processor “dies” or “tiles”, which we count as separate GPUs.

[†] Rated unidirectional node bandwidth based on the number of NICs.

V. IMPLEMENTATION

HiCCL is implemented in C++ for distributed applications, libraries, and frameworks running on GPUs (and CPUs). HiCCL requires MPI for initialization, even if the user does not select MPI during execution.

A. Implementation Options

We implemented HiCCL by leveraging the point-to-point communication API of the chosen library for each level of the hierarchy. We integrated non-blocking point-to-point functions of MPI, NCCL, RCCL, and OneCCL to be used within and across nodes, and vendor-provided (CUDA, HIP or Level Zero) IPC put and get to be used within nodes. We also integrated CUDA, HIP and SYCL programming models for targeting Nvidia, AMD and Intel GPUs, respectively.

B. Persistent Design

HiCCL takes advantage of repetitive collective function calls by memoizing the optimized data movement and scheduling in internal data structures. On the second and subsequent uses of a communication operation, these structures are reused to avoid the cost of making on-the-fly decisions. Furthermore, HiCCL has no global synchronization either in the composition and synthesis nor in the execution.

VI. EVALUATION

We evaluate the performance portability of HiCCL across eight commonly used collective functions in Table I and four HPC systems with different hardware and software.

A. Experimental Setup

The node architectures of the test systems are summarized in Table IV, showing various numbers of CPUs, GPUs and NICs per node. We compare HiCCL collective throughput with that of a) corresponding native MPI implementations and b) vendor-provided collective communication libraries (NCCL, RCCL or OneCCL) on each system. The MPI implementations are based on OpenMPI (OMPI) for Delta and modified versions of MPICH for Perlmutter, Frontier and Sunspot. NCCL and RCCL use the AWS-ofi extension for portability to the Slingshot (SS-11) networks on these machines [2].

B. Measurements

We measure the peak throughput of each collective function on each system. We run the end-to-end collective function in multiple rounds: 5 warmup rounds and 10 measurement rounds. In each measurement round, we measure the elapsed time from a global synchronization to the moment that the

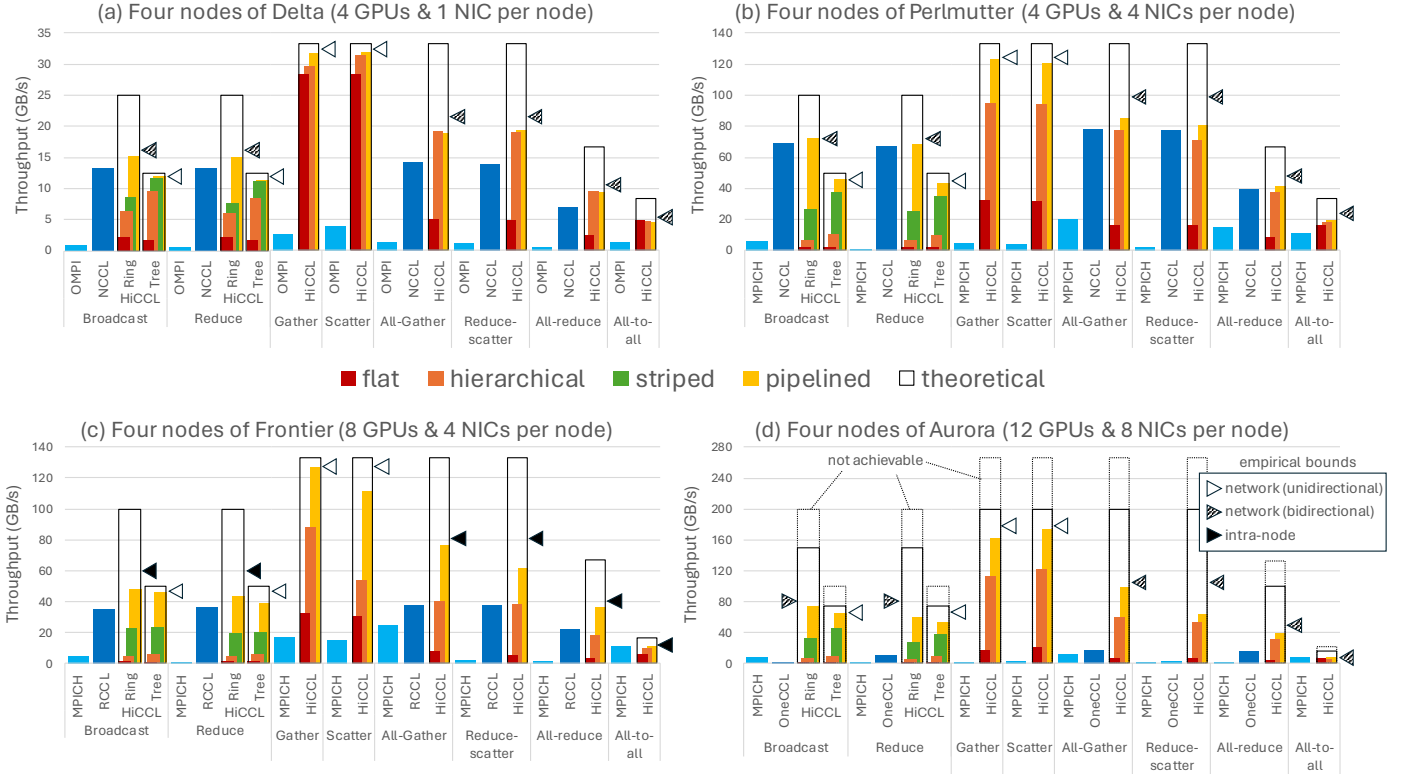


Fig. 8: Peak collective throughput (GB/s) on (a) Delta, (b) Perlmutter, (c) Frontier, and (d) Aurora. HiCCL optimizations are applied incrementally, where the frames around bars represents the theoretical limit based on the rated NIC performance. Triangle marks represents empirical bounds based on isolated measurements across two nodes. HiCCL throughput reaches to the empirical bounds on all systems, demonstrating performance portability across architectures.

communication buffers on all GPUs are safe to be reused. We run collectives with buffer sizes of pd bytes. For example, Scatter sends d bytes to each of the p processors. If a collective requires t seconds to execute, the throughput is dp/t (GB/s). We vary d across large message sizes (larger than a MB) until the throughput saturates the achievable bandwidth.

C. Collective Throughput

Figure 8 (a)–(d) shows the peak collective throughput on four nodes of each test system. We use available library implementations, presented with light blue (MPI) and dark blue (NCCL/RCCL/OneCCL) colors, as baselines. We confirmed these baseline results with administrators of each system, in some cases setting tuning flags when advised to do so.

1) *Overall Speedup*: In Figure 8, HiCCL results are shown with four bars, representing the incremental effect of optimizations. When all optimizations are applied, HiCCL’s geomean speedup over the MPI implementations is $12.52\times$, $14.22\times$, $9.76\times$, and $48.02\times$ on Delta, Perlmutter, Frontier, and Aurora, respectively. On the other hand, the speedup over vendor-provided libraries is $1.26\times$, $1.05\times$, $1.55\times$, and $12.01\times$ on the respective systems. The comparison suggests that while MPI implementations are not optimized for throughput, vendor-provided libraries generally are, with the exception of OneCCL on Aurora.

The rest of this subsection discusses the effect of individual optimizations of HiCCL on tested collective functions.

2) *Hierarchical Optimizations*: Red bars in Figure 8 represent direct implementations of collectives with non-blocking point-to-point functions, assuming there is no hierarchy across GPUs—i.e., the description of the network hierarchy for these experiments is just $\{p\}$, where p is the number of participating GPUs. Direct implementations use NCCL on Delta and Perlmutter, and MPI on Frontier and Aurora as they are the most performant options.

Orange bars in Figure 8 represent hierarchical optimizations with various factorizations (Section IV-B) that are specific to each system. These factorizations are shown in the third column of Table V, where the bold entries represent the hierarchies within nodes. Within nodes, Delta and Perlmutter are represented with a single level, e.g., $\{4\}$, due to the directly connected GPUs. On the other hand, Frontier and Aurora nodes consist two-level hierarchies, e.g., $\{4, 2\}$ and $\{6, 2\}$, where the lower level represents the dual-die devices and the upper level represents high-bandwidth links between these devices. Frontier has four devices and Aurora has six devices per node. On the other hand, the factorizations across nodes are to form virtual topologies with multi-rail tree (two levels) or ring (single-level) structures across nodes. Overall, hierarchical optimizations obtain a geometric average of $4.08\times$ speedup over the direct baseline on all systems and collectives.

3) *Multi-NIC Striping*: Green bars in Figure 8 represent the throughput with HiCCL’s multi-NIC striping (Section IV-C).

TABLE V: Hierarchical Factorizations and Libraries Used in Figures 8–9

| System | Topology | Hierarchy | Implement. Library |
|------------|-----------|--------------|----------------------|
| Delta / | Tree | {2, 2, 4} | {NCCL, NCCL, IPC} |
| Perlmutter | Ring+Tree | {4, 4} | {NCCL, IPC} |
| Frontier | Tree | {2, 2, 4, 2} | {MPI, MPI, IPC, IPC} |
| | Ring+Tree | {4, 4, 2} | {MPI, IPC, IPC} |
| Aurora | Tree | {2, 2, 6, 2} | {MPI, MPI, IPC, IPC} |
| | Ring+Tree | {4, 6, 2} | {MPI, IPC, IPC} |

This optimization is beneficial primarily for Broadcast and Reduce collectives, as these do not inherently utilize all NICs. On Delta, where each node has just one NIC, striping offers limited advantages, evidenced by a $1.29\times$ speedup. This improvement is attributed to four GPUs utilizing the single NIC more effectively than would a solitary GPU. In contrast, on multi-NIC nodes, such as Perlmutter, Frontier, and Aurora, striping yields speedups of $3.62\times$, $3.94\times$, and $4.76\times$, respectively, demonstrating significant throughput improvements.

4) *Pipelining*: Yellow bars in Figure 8 represent the pipelining optimization (Section IV-E). Pipelining hides intra-node communications with a tree topology and also provides asymptotic speedup with a ring+tree topology. To demonstrate, we test HiCCL’s Broadcast and Reduce collectives with both virtual topologies, denoted as “Tree” and “Ring” in Figure 8. We observe that the ring obtains up to $2.72\times$ speedup (on Perlmutter) yet does not saturate the throughput up to its theoretical limit. All other collectives use the tree topology, and pipelining is effective on systems with significant intra-node communication yet obtains no more than two times speedup. To explain these limitations, we use empirical bounds that are explained next.

5) *Upper Bounds*: The frames around our HiCCL results in Figure 8 represent the theoretical upper bounds that are given in Table III. Aurora is a special case because round-robin assignment (Section II-C) of 12 GPUs to 8 NICs. In this case, GPU i is assigned to NIC $i \bmod 8$, yielding GPUs 0–7 assigned to NICs 0–7 whereas GPUs 8–11 oversubscribe NICs 0–3. As a result, the first four NICs handle two GPUs each, whereas the remaining NICs handle a single GPU each, leading to load imbalance. Thus, the achievable bandwidth on Aurora is limited to 75% of the theoretical bandwidth with this strategy.

Even though we apply all optimizations aggressively and use large buffer sizes to saturate throughput, we converged to only a fraction of the theoretical limits. To understand why, we measured the unidirectional and bidirectional bandwidth across two nodes in isolation rather than using the numbers in the spec sheet. These empirical upper bounds are indicated by hollow (unidirectional) and striped (bidirectional) triangles marks in Figure 8. For example, the empirical bounds for Gather and Scatter are considered unidirectional, as the bottleneck in these operations is the root node, which either receives or sends messages in only one direction. On the other hand, other collectives send and receive messages at the same time and therefore their empirical bounds are the bidirectional utilization.

A surprising result is that the intra-node communication cost

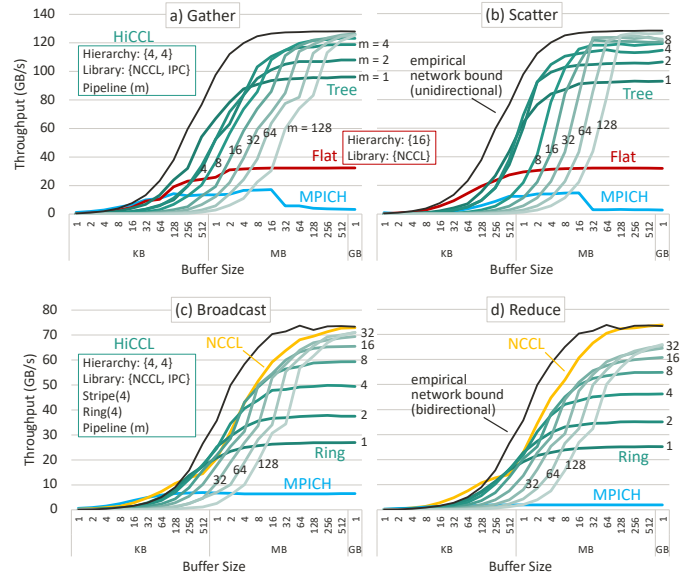


Fig. 9: Throughput with various pipeline depths and buffer sizes for tree implementation of (a) Gather and (b) Scatter, and ring+tree implementation of (c) Broadcast and (d) Reduce on four nodes of Perlmutter.

on Frontier is higher than that of inter-node communication, which prevents us from hiding the former with the latter, even when we align the virtual hierarchy with the node architecture. Therefore we also measured the intra-node cost in isolation and marked the corresponding empirical bound with dark triangles in Figure 8(c). Nevertheless, our results suggest that HiCCL almost always comes close to the maximum potential of the machine in practice.

D. Pipeline Depth

Although pipelining improves communication throughput, it must be used with caution. When applied too aggressively, message sizes become so small that the latency overhead between stages dominates, negatively impacting the overall throughput, as suggested in Equations (2)–(1). To illustrate this phenomenon, we vary the pipeline depth across four nodes of Perlmutter

Figure 9 shows the performance curves of (a) Gather and (b) Scatter with tree, and (c) Broadcast and (d) Reduce with ring across various pipeline depths (m) 1 to 128, where $m = 1$ means no pipelining and $m = 128$ means pipelining with 128 channels. Pipelining clearly improves throughput for large message lengths. Nevertheless, excessive use for smaller messages reduces the throughput.

Further implications of pipelining depend on the optimized communication pattern. In this case, we apply the tree algorithm for Gather / Scatter, and ring+tree algorithm for Broadcast / Reduce. Effectively, pipelining hides the intra-node communications with the tree implementation, and therefore converges to the empirical bound with a pipeline of only $k = 4$ stages. On the other hand, pipelining provides asymptotic speedup with the ring implementation, and it requires up to

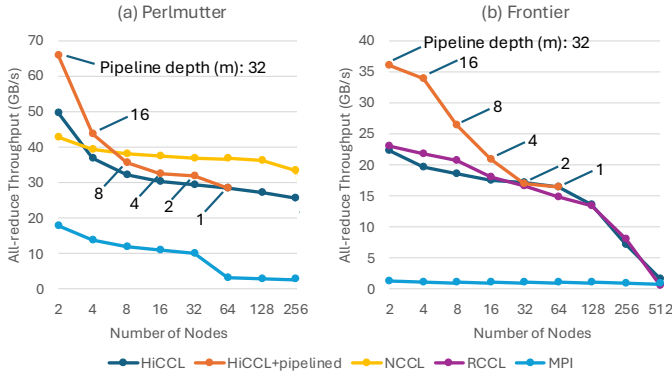


Fig. 10: Scaling on (a) Perlmutter and (b) Frontier. HiCCL applies throughput-oriented optimizations, e.g., pipelining, that are effective up to 256 nodes. The performance difference between NCCL and HiCCL are mainly due to the kernel call overhead involved in reductions, where NCCL is more optimized.

$k = 32$ levels in the pipeline to saturate the throughput. We observe similar behavior across all of our test systems.

As a reference, we take MPICH (light blue) and NCCL (gold) library implementations as well as the aforementioned empirical bounds. Since NCCL does not offer Gather and Scatter collectives, we implement them directly with NCCL’s point-to-point functions, as represented with the red curves. In (d) Reduce, even a deep pipeline falls short of NCCL’s throughput. Remember that Reduce involves an additional computation that requires a GPU kernel call. NCCL successfully hides the computational overhead with a CUDA streaming mechanism. We do not apply CUDA-specific optimizations for the sake of generality.

E. Scaling

Figure 10 demonstrates scaling on (a) Perlmutter and (b) Frontier. We use the All-reduce collective with the two-step formulation (Section III-C) on both machines: We compose the collective only once with the proposed API and change the virtual hierarchy across machines. HiCCL is compared with MPI, NCCL and RCCL, where available. To saturate the network in these experiments, buffer sizes were selected to be large (8.6 GB on Perlmutter and 17.2 GB on Frontier), which were determined based on the device memory capacity. Due to MPI’s limitations with large buffer sizes [17], a 1 GB buffer size was used for MPI in the experiments on both machines.

The scaling experiments on Frontier reveal the limitation of throughput-oriented optimizations we applied in this work. When scaled to more than 256 nodes (2,096 GPUs), the ratio of communication to computation drops significantly so that latency becomes the main bottleneck. In principle, latency-oriented collective design can be achieved with HiCCL’s API, however, it is not in the scope of this work. Additionally, MPI on Frontier is tuned to minimize latency at large scale, and therefore preferable to any other communication library on more than 256 nodes.

VII. RELATED WORK

To the best of our knowledge, HiCCL is the first compositional library that offers unified hierarchical communications across systems with Nvidia, AMD and Intel GPUs.

Previous work [6] implemented a multi-NIC striping algorithm with a lower-level communication API and for CPUs. For HiCCL, we applied the multi-NIC striping idea in the context of modern node architectures, where the endpoints are GPUs. We integrated the proposed striping algorithm with a high-level (GPU-aware) communication library, exploiting the fixed logical associations between GPUs and NICs as explained in Section II-C.

There is a body of previous work towards hierarchical collective communications. Nevertheless, they are either collective-specific [3], [24], system-specific [18], hardware-specific [13], [14], single-node [4], or CPU-only [25]. HiCCL optimizes all collective functions with unified optimizations, achieves portability across architectures, and also demonstrates high throughput across systems of different GPU vendors (Nvidia, AMD and Intel).

A similar line of research (MSCCL [4], [22], [27]) takes a different path on decomposition-based generalized optimization by building a larger programming system with a domain-specific language and code generation. However, the previous work’s scalability is limited up to 16 GPUs due to the cost of code synthesis based on an SMT solver [7]. In contrast, HiCCL is a standalone runtime library that strives to be minimal and is scalable up to hundreds of nodes. Specifically, the initialization cost of HiCCL does not take more than six seconds on a thousand GPUs.

VIII. CONCLUSION

HiCCL is a high-level communication library for optimizing collective functions. The collective pattern is built using the proposed compositional API (Section III). HiCCL applies hierarchical factorization on the original pattern, and optimizes it for a described machine (Section IV). The generalization of hierarchical optimizations across collective communications and machines is HiCCL’s highlight contribution. We implemented the library to be portable across Nvidia, AMD and Intel GPUs, and across various communication fabrics. In this way, the throughput of eight collective functions are improved on four distinct machines with different architectures and vendors. The geometric speedup of HiCCL is $17.0\times$ over the native GPU-aware MPI implementations on all systems, $1.15\times$ over NCCL, $1.55\times$ RCCL, and $12.1\times$ OneCCL where available. Our evaluation demonstrates scalability up to 256 nodes (1,024–2,048 GPUs). HiCCL not only addresses the immediate challenges posed by the current diversity in hierarchical networks, but also sets a foundation for future-proofing collective communication optimizations against the backdrop of evolving HPC landscapes.

ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S.

Department of Energy Office of Science and the National Nuclear Security Administration. This work was done on a pre-production supercomputer with early versions of the Aurora software development kit. This research used resources of the Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory and is based on research supported by the U.S. DOE Office of Science Advanced Scientific Computing Research Program, under Contract No. DE-AC02-06CH11357. This research used the Delta advanced computing and data resource which is supported by the National Science Foundation (award OAC 2005572) and the State of Illinois. Delta is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award ASCR-ERCAP0029675.

REFERENCES

- [1] Advanced Micro Devices (AMD). Radeon Collective Communications Library (RCCL). <https://rocm.docs.amd.com/projects/rccl/>, 2023. Accessed: 2023-04-01.
- [2] Amazon Web Services (AWS). AWS OFI RCCL. <https://github.com/ROCm/aws-ofi-rccl>, 2023. Accessed: 2023-04-01.
- [3] A. Bienz, L. N. Olson, W. D. Gropp, and S. Lockhart. Modeling data movement performance on heterogeneous architectures. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2021.
- [4] Z. Cai, Z. Liu, S. Maleki, M. Musuvathi, T. Mytkowicz, J. Nelson, and O. Saarikivi. Synthesizing optimal collective algorithms. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 62–75, 2021.
- [5] N. E. R. S. C. Center. Perlmutter Architecture, 2022.
- [6] S. Coll, E. Frachtenberg, F. Petrini, A. Hoesie, and L. Gurvits. Using multirail networks in high-performance clusters. *Concurrency and Computation: Practice and Experience*, 15(7-8):625–651, 2003.
- [7] L. De Moura and N. Björner. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.
- [8] D. De Sensi, S. Di Girolamo, K. H. McMahon, D. Roweth, and T. Hoefler. An in-depth analysis of the Slingshot interconnect. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–14. IEEE, 2020.
- [9] A. L. C. Facility. Aurora, 2022.
- [10] O. R. L. C. Facility. Frontier User Guide - System Overview, 2022.
- [11] N. C. for Supercomputing Applications. Delta, 2022.
- [12] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, et al. Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface: 11th European PVM/MPI Users' Group Meeting Budapest, Hungary, September 19-22, 2004. Proceedings 11*, pages 97–104. Springer, 2004.
- [13] R. L. Graham, D. Bureddy, P. Lui, H. Rosenstock, G. Shainer, G. Bloch, D. Goldenberg, M. Dubman, S. Kotchubievsky, V. Koushnr, et al. Scalable hierarchical aggregation protocol (SHARP): A hardware architecture for efficient data reduction. In *2016 First International Workshop on Communication Optimizations in HPC (COMHPC)*, pages 1–10. IEEE, 2016.
- [14] R. L. Graham, L. Levi, D. Bureddy, G. Bloch, G. Shainer, D. Cho, G. Elias, D. Klein, J. Ladd, O. Maor, et al. Scalable hierarchical aggregation and reduction protocol (SHARP) streaming-aggregation hardware design and evaluation. In *High Performance Computing: 35th International Conference, ISC High Performance 2020, Frankfurt/Main, Germany, June 22–25, 2020, Proceedings 35*, pages 41–59. Springer, 2020.
- [15] W. Gropp and E. Lusk. User's guide for MPICH, a portable implementation of MPI, 1996.
- [16] W. Gropp, R. Thakur, and E. Lusk. *Using MPI-2: Advanced features of the message passing interface*. MIT press, 1999.
- [17] J. R. Hammond, A. Schäfer, and R. Latham. To INT_MAX... and beyond! Exploring large-count support in MPI. In *2014 Workshop on Exascale MPI at Supercomputing Conference*, pages 1–8. IEEE, 2014.
- [18] M. Hidayetoğlu, T. Bicer, S. G. De Gonzalo, B. Ren, V. De Andrade, D. Gursoy, R. Kettimuthu, I. T. Foster, and W. H. Wen-mei. Petascale XCT: 3D image reconstruction with hierarchical communications on multi-GPU nodes. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–13. IEEE, 2020.
- [19] M. Hidayetoğlu, T. Biçer, S. G. De Gonzalo, B. Ren, D. Gürsoy, R. Kettimuthu, I. T. Foster, and W.-m. W. Hwu. MemXCT: Memory-centric X-ray CT reconstruction with massive parallelization. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–56, 2019.
- [20] IBM. IBM Spectrum MPI - Overview, 2017.
- [21] Intel Corporation. OneAPI Collective Communications Library (OneCCL). <https://oneapi-src.github.io/oneCCL/>. Accessed: 2024-03-29.
- [22] A. Jangda, J. Huang, G. Liu, A. H. N. Sabet, S. Maleki, Y. Miao, M. Musuvathi, T. Mytkowicz, and O. Saarikivi. Breaking the computation and communication abstraction barrier in distributed machine learning workloads. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 402–416, 2022.
- [23] S. Jeaugey. Nccl 2.0. In *GPU Technology Conference (GTC)*, volume 2, page 23, 2017.
- [24] S. Lockhart, A. Bienz, W. D. Gropp, and L. N. Olson. Characterizing the performance of node-aware strategies for irregular point-to-point communication on heterogeneous architectures. *Parallel Computing*, 116:103021, 2023.
- [25] X. Luo, W. Wu, G. Bosilca, Y. Pei, Q. Cao, T. Patinyasakdikul, D. Zhong, and J. Dongarra. Han: A hierarchical autotuned collective communication framework. In *2020 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 23–34. IEEE, 2020.
- [26] D. K. Panda, H. Subramoni, C.-H. Chu, and M. Bayatpour. The MVA-PICH project: Transforming research into high-performance MPI library for HPC community. *Journal of Computational Science*, 52:101208, 2021.
- [27] A. Shah, V. Chidambaram, M. Cowan, S. Maleki, M. Musuvathi, T. Mytkowicz, J. Nelson, O. Saarikivi, and R. Singh. TACCL: Guiding collective algorithm synthesis using communication sketches. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 593–612, 2023.
- [28] C. M. Siefert, C. Pearson, S. L. Olivier, A. Prokopenko, J. Hu, and T. J. Fuller. Latency and bandwidth microbenchmarks of US Department of Energy systems in the June 2023 Top 500 list. In *Proceedings of the SC'23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, pages 1298–1305, 2023.
- [29] M. Snir. *MPI—the Complete Reference: the MPI core*, volume 1. MIT press, 1998.
- [30] R. Thakur, R. Rabenseifner, and W. Gropp. Optimization of collective communication operations in mpich. *The International Journal of High Performance Computing Applications*, 19(1):49–66, 2005.
- [31] M. Wilkins, H. Wang, P. Liu, B. Pham, Y. Guo, R. Thakur, P. Dinda, and N. Hardavellas. Generalized collective algorithms for the exascale era. In *2023 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 60–71. IEEE, 2023.

- [32] C. Zimmer, S. Atchley, R. Pankajakshan, B. E. Smith, I. Karlin, M. L. Leininger, A. Bertsch, B. S. Ryujin, J. Burmark, A. Walker-Loud, et al. An evaluation of the CORAL interconnects. In *SCI9: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–18, 2019.