SpCCL: A Sparsity-Aware Collective Communication Library for GPU Platforms

Lannie Dalton Hough, Emir Gencer, Hoffmann Muki, Abhinav Bhatele

Department of Computer Science, University of Maryland, College Park, MD, USA
E-mail: {ldhough, egencer, hoffmuki}@umd.edu, bhatele@cs.umd.edu

***Abstract*—Collective communication is essential to high performance computing and machine learning workloads, yet libraries such as NCCL do not exploit sparsity in message payloads. Sending only nonzero values can reduce network traffic, but explicitly handling sparsity introduces challenges such as compression and decompression overheads. We address these challenges with sparsity-exploiting versions of all-gather, reduce-scatter, and all-reduce collectives. Our implementations use a new bitvector-based format, Pici, designed for low space overhead and fast GPU-based compression and decompression. Further, our collective algorithms adapt to the degree of sparsity in data, modifying data representations during the course of the collective. At 99% input sparsity, our collectives achieve up to 5.25×, 2.5×, and 2.66× speedups over NCCL for all-gather, reduce-scatter, and all-reduce, respectively. Integrating our collectives into a representative deep learning application, we achieve a 26% end-to-end speedup.**

***Index Terms*—sparse data, collective algorithms, compression**

I. Introduction

Sparsity, i.e., the presence of many zero values in data, is common in certain parallel workloads, including high performance computing (HPC) [1]–[4], and even more so in distributed deep learning (DL) [5]–[8]. In practice, the sparse data often takes the form of an array or matrix of N elements, of which nnz are nonzero. The degree of sparsity is expressed by the fraction of zeros as $\sigma = (N - nnz)/N$, often reported as a percentage. When working with sparse data, it is often useful to quantify the degree of sparsity with terms like "highly sparse" or "low density." However, such terms are not standardized, and what counts as highly sparse may vary by problem domain or researcher. In this paper, we primarily contrast two sparsity regimes. The high-sparsity (HS) regime comprises >99% sparse data, while the non-high-sparsity (NHS) regime comprises data with 1–99% sparsity, inclusive. We consider data that is less than 1% sparse to be near-dense.

Exploiting sparsity in data (both application input data and large arrays in memory) to improve the performance of computation and communication is common [9]–[11]. Communication optimizations center on sending only important information (e.g., nonzero values), often in collective operations such as all-gather, reduce-scatter, and all-reduce. As communication overheads in parallel applications can represent a significant fraction of execution time, particularly as applications scale to more GPUs and nodes, such optimizations can be particularly helpful for large messages.

Historically, sparsity-aware kernels and communication collectives have often targeted the HS regime or specific sparsity patterns. However, workloads such as the training of pruned deep neural networks often exhibit unstructured sparsity patterns at lower degrees of sparsity [5], [7]. In this paper, we introduce communication collectives that exploit unstructured sparsity across a wider range of degrees of sparsity and significantly improve communication performance.

The task of implementing efficient sparsity-aware collectives (sparse collectives, in short) is challenging because of several features of problems with sparse data (sparse problems, in short):

- **Irregularity**: Unlike many dense problems, sparse problems are highly variable in their characteristics. For example, both the degree of sparsity and the sparsity structure (where the nonzeros occur) can vary and impact which algorithms, kernels, and sparse formats are effective.
- **Dynamic changes**: In the case of collectives that include a reduction such as reduce-scatter and all-reduce, the degree of sparsity can decrease across reduction steps during the collective ("dynamic sparsity"). If processes contribute sparse data with differing nonzero indices, the sparsity of intermediate and final outputs decreases ("fill-in"). Thus, better support for non-high-sparsity data is required in such cases.
- **Overheads**: Conversion between dense and sparse formats (compression and decompression) can incur significant overheads. We must ensure that these overheads do not offset the benefits of communicating less data.

Additionally, even when data has a large number of zeros, operations surrounding the collective may still assume a dense representation. These concerns indicate that we must consider not only how space-efficient a sparse format is but also the performance of its compression and decompression kernels.

We observe that traditional formats for representing sparse data, such as Coordinate (COO) [12] and Compressed Sparse Row (CSR) [13], are poorly suited to collectives where fill-in quickly drives data outside of the HS regime. These formats use a large amount of metadata per nonzero element. While acceptable for data in the HS regime, this becomes increasingly undesirable as sparsity decreases. We introduce a novel bitvector-based format, Pici, which instead uses a small constant amount of metadata per input element.

SC26, November 15-20, 2026, Chicago, Illinois, USA
979-8-3195-4789-7/26/$31.00 ©2026 IEEE

We also observe that below a certain sparsity threshold, which depends on many factors, including compression and decompression speeds, compression and decompression frequency, and the bandwidth of the communication path, the overhead of compression and decompression outweighs the benefits of communicating less data. Due to fill-in in collectives with reductions, this may occur after some intermediate steps in the collective algorithm that do benefit from exploiting sparsity initially. To address this concern, we design our collectives to adapt to dynamic sparsity, changing the representation of data from sparse to dense as needed during the course of the collective. To leverage prior collective optimization work, we implement our sparse collectives within a fork of the NCCLX collective communication library [14] and name it SpCCL (**Sp**arsity-Aware **C**ollective **C**ommunication **L**ibrary).

We evaluate SpCCL's collectives on NERSC's flagship Perlmutter supercomputer [15]. For a range of input sparsities, we achieve improved performance compared to both other sparse collective implementations and NCCL [16] and NCCLX [14]. We also incorporate our collectives into a representative deep learning (DL) application and demonstrate end-to-end performance speedups.

This paper makes the following key contributions:

- We introduce Pici (pronounced "peachy"), a novel bitvector-based sparse representation that is space-efficient across a wide range of sparsity levels, along with supporting CUDA kernels. This format imposes a low, sparsity-independent metadata overhead of approximately 3.15% for fp32 data, while also supporting efficient parallel compression and decompression.
- We develop optimized implementations of three sparse collectives that use Pici: all-gather, reduce-scatter, and all-reduce. We focus on effective support for unstructured data in the NHS regime, but SpCCL collectives are also effective in the HS regime.
- We mitigate the problem of fill-in by introducing a sparsity-aware algorithm that adapts the data representation depending on the degree of sparsity in the data, the current stage of the collective, and the links that messages traverse.
- Using SpCCL with Pici-based collectives, we achieve substantial speedups over state-of-the-art dense communication libraries and a prior sparsity-aware all-reduce implementation. We also demonstrate gains in a representative DL application. At 99% sparsity, SpCCL outperforms a dense NCCL baseline for all-gather, reduce-scatter, and all-reduce by up to 5.25×, 2.5×, and 2.66×, respectively.

II. Background

In this section, we provide the background necessary to understand requirements and tradeoffs involved in representing sparse data and implementing sparse collectives.

A. Modeling Performance of Collective Routines

Our work focuses on three collectives: all-gather (AG), reduce-scatter (RS), and all-reduce (AR). Libraries frequently implement all-reduce as RS+AG [17]. Before designing a sparse collective, we consider a cost model to determine under what conditions we can realize speedups. The Hockney model is a simple and effective way to model communication [18] and is sufficient to reason about the ideas this paper introduces, although more sophisticated models do exist (such as LogGP [19] and max-rate [20]). Under the Hockney model, the α term represents a message's startup latency. The β term represents the inverse of a link's peer-to-peer bandwidth. n describes the message size in bytes. Taken together:

$$T_{\text{comm}} = \alpha + n\beta \tag{1}$$

We use the Hockney model to characterize how different collective communication algorithms affect each term. For example, with the ring algorithm, each process communicates with immediate neighbors in a circular fashion, forwarding messages around the ring in $p-1$ steps. With n as per-process output payload, ring all-gather can be modeled as:

$$T_{\text{ring}} = (p-1) \times \left(\alpha + \frac{n}{p} \times \beta\right) \tag{2}$$

While bandwidth-optimal, ring can suffer at very large process counts due to the latency term, which scales linearly. In contrast, some recursive algorithms divide the communication task into $\log_2(p)$ steps, either doubling (AG) or halving (RS) the message volume at each step.

Compression yields speedups by reducing transmitted message volume. This implies that if messages are already small enough for α to be dominant, further reducing communication volume is unlikely to achieve significant speedups. As a consequence, we focus on the large message regime (densely formatted inputs of hundreds of MiB). Large messages are typical of distributed training, for example, when synchronizing gradients using distributed data parallelism (DDP) [21].

Although the Hockney model is useful for modeling dense collectives, it is insufficient for sparse collectives, where compression and decompression introduce additional overhead. We extend the model with two additional terms, T_{comp} and T_{decomp}:

$$T_{\text{comm_sparse}} = \alpha + n\beta + T_{\text{comp}} + T_{\text{decomp}} \tag{3}$$

The impact of these additional terms varies depending on the collective in question. Additionally, n in $T_{\text{comm_sparse}}$ represents the sparse payload of the modeled message, which depends on format choices and may vary across algorithm steps. While we defer a more detailed discussion of these considerations to Section IV, minimizing both T_{comp} and T_{decomp} is clearly beneficial. This necessitates a sparse format with high-performance compression and decompression kernels.

B. Overview of Sparse Data Formats

Existing literature describes numerous sparse representations [22], [23]. In order to reconstruct (decompress) densely formatted data, sparse formats should store nonzero values and metadata that can be used to identify the positions of the nonzero values. Two of the most common and general-purpose

formats are Coordinate (COO) and Compressed Sparse Row (CSR) [12], [13]. COO represents sparse matrices by storing row and column indices for each nonzero in the matrix, as well as the value itself. CSR uses less metadata than conventional COO by omitting per-nonzero row indices and instead storing a per-row prefix sum of the number of nonzeros (nnz) encountered. This generally results in more complex compression and decompression routines.

A common variant of COO treats inputs as one-dimensional, storing only a single flat index per nonzero. This cuts index metadata in half. For the rest of this work, COO refers to this variant. Prior work uses COO in sparse collectives, such as SparCML's all-reduce implementations [10]. Alongside Pici, we use COO in this work. Space overhead calculations assume 32-bit indices and 32-bit values.

Observe that in "index + value" formats, such as COO and CSR, metadata overhead scales linearly with nnz. Relative to formats that impose a sparsity-independent overhead of m metadata bytes, COO uses less space when $4 \times nnz < m$. While often requiring less metadata than alternatives at higher sparsities, COO becomes less competitive with formats featuring sparsity-independent metadata overhead as sparsity decreases. At 50% sparsity, COO requires the same amount of space as the dense representation, and so space savings cannot be realized.

Other formats exploit structure in sparse data [24]. Block Sparse Row (BSR) generalizes CSR to *blocks* of nonzeros and explicitly stores zeros within a nonzero block. Diagonal (DIA) stores nonzero diagonals. Structured formats such as BSR can be highly space-efficient when the sparsity pattern closely matches the structure being exploited by the format. However, structured formats can be highly inefficient for unstructured sparse data. Consider BSR with 16×16 blocks, where each element has an independent 99% chance of being zero. The chance of an entire block being zero is only $0.99^{256} \approx 0.076 = 7.6\%$, and all of the zero values in a nonzero block must still be stored.

Finally, bitvector-based formats use bitvectors to implicitly encode the positions of nonzeros. The simplest variant of this is Bit-Vector (BV), which stores a single bit per input element and each nonzero value in a compacted values array [25]. Bits are set to 1 for corresponding nonzero elements or 0 otherwise. Compressed Bit-Vector (CBV) reduces space by encoding runs of zeros and nonzeros. Observe that for BV, metadata size does not depend on the sparsity pattern. Its low, sparsity-independent metadata overhead makes BV appealing across much of the non-high-sparsity regime. For 32-bit values, BV has 3.125% metadata overhead relative to the dense representation. Thus, below 96.875% sparsity, BV is more space-efficient than COO. For a comprehensive overview of sparse formats, we refer the reader to Langr et al. [22].

III. Pici: A Space-Efficient Sparse Representation

We now describe Pici, the design of its compression and decompression algorithms, and the implementation of corresponding CUDA kernels.

A. Key Design Considerations for Pici

A sparse representation suitable for supporting arbitrary sparsity patterns for data in the NHS regime and for mitigating the impact of the fill-in problem should (1) have low space overhead for data of varying degrees of sparsity, and (2) support efficient parallel compression and decompression. Recall that fill-in is the densification of data across reduction steps in multi-stage reducing collectives.

Among the formats described in Section II-B, Bit-Vector (BV) satisfies the first design constraint particularly well. BV's metadata overhead is independent of how nonzeros are distributed, so it is suitable for unstructured sparse data. The overhead is also sparsity-independent, and as a result, it does not increase as sparsity decreases due to fill-in. Hence, we use a bitvector-based sparse representation as the foundation for designing Pici.

Unfortunately, raw bitvectors are not ideal for satisfying the second design requirement. Consider parallel decompression with an N-element bitvector B and an array of nonzero values A (the components of the BV format). Naïvely, threads start scanning from different segments of the bitvector in order to parallelize scans that identify the positions of nonzero values. For a thread starting the scan at an arbitrary position i, if $B[i] = 1$, the algorithm knows that a nonzero value should be placed at position i in the decompressed dense buffer. However, unless $i = 0$, the algorithm does not know which value from A corresponds to this position. The algorithm needs to know how many bits in B are set to 1 before position i.

Given an N-element bitvector, we could solve the value-lookup problem with access to a metadata array I storing, for each bit, the number of preceding 1-bits in the bitvector (a prefix sum). This naïve approach violates the low space overhead requirement because of the high storage requirements of I. In practice, bitvector-based formats may mitigate this problem by dividing the bitvector into spans of length k and storing only $\lceil N/k \rceil$ prefix sums. A thread can then initiate a k-element scan from the start of any span. With one thread available per span, each thread scans at most k elements.

Observe that increasing k reduces index overhead but decreases the number of immediate offsets into A. This saves space but reduces the number of spans that can be decompressed independently without additional prefix computation.

To satisfy our design requirements, we ideally want to keep index overhead low while facilitating high parallelism. To achieve this, we take inspiration from the field of succinct data structures. Works in this area demonstrate that in exchange for a modest amount of span-level computation, data structures that can answer rank queries on bitvectors may achieve significant index space savings [26], [27]. Informally, rank denotes the number of 1s or 0s in a bitvector, up to some position i. Answering rank_1 queries is isomorphic to our problem of determining the nonzero count up to the start of an arbitrary span.

Rank-enabling data structures such as poppy and pasta introduce many ideas to keep space overhead and query time low [26], [27]. They feature hierarchical indices of decreasing

width covering increasingly small spans of the bitvector. The indices are interleaved to minimize cache misses. For their most granular indexed span of 512 bits, the `POPCNT` instruction enables faster scans, as it can return the number of 1-bits in a 64-bit word while taking only a few cycles.

The space overhead of poppy and pasta is exceptionally low, at 3.125% relative to the bitvector. However, these structures are still not ideal for supporting parallel compression and decompression on GPUs. They target low-latency random queries on CPUs, but within a thread, compression and decompression scan many adjacent bits or elements in series. Rather than adopting these structures directly, we leverage the key insight that modest span-local computation can reduce index space and design a data structure to better fit our access patterns and the CUDA programming model.

B. Description of the Pici Sparse Data Format

The Pici data format consists of three parts: two metadata arrays and an array A of nonzero values from the dense representation of the original data. The first metadata array is a bitvector, B, used to identify nonzero positions, similar to the BV format. The second metadata array is an index array, I, which stores prefix sums of nonzero counts to provide immediate offsets into the values array, A.

To construct the three Pici components, the compression process receives an N-element densely formatted input (a contiguous buffer of zero and nonzero elements). Pici then views the densely formatted input as being composed of 64×64 row-major "virtual" tiles (see Figure 1). Each tile corresponds to a 4096-bit span of B and one exclusive prefix sum in I. Because I contains one prefix sum per 4096-bit span of B, $k = 4096$. A single additional entry at the end of I stores the total number of nonzeros (nnz) in the input.

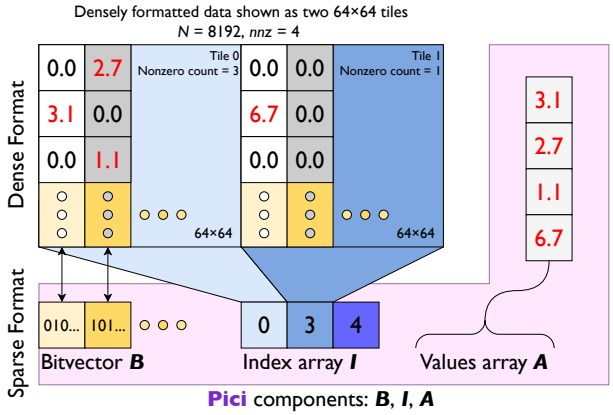

Fig. 1. Overview of how a densely formatted input of $N = 8192$ elements covered by two virtual tiles is represented in the Pici sparse format.

Although the layout of tile elements is row-major, each 64-bit word in B corresponds to a single tile *column*. Because of Pici's tile- and column-based view into the input, the bitvectors of Pici and BV are generally not identical. The layout of nonzeros in Pici's A also follows the view hierarchy, and compacted nonzeros are ordered first by their parent tile, next by their parent column, and finally by their intra-column position. The choices of $k = 4096$ and the column-first tile view are informed by the CUDA programming model to facilitate fast compression and decompression, discussed in more detail in Section III-C.

Observe that while Pici features supporting indices in I, it uses a flat indexing scheme instead of poppy-style hierarchical and packed indices, avoiding overhead from packing variable-width indices. Additionally, with a substantially larger k than poppy and pasta, Pici's index overhead is even lower, at only $\approx 0.78\%$ relative to the bitvector. Pici's total overhead relative to a 32-bit dense representation is only $\approx 3.15\%$. Figure 1 shows how Pici components map to the tile-based view of the densely formatted source.

C. GPU Kernels for Compression and Decompression

Next, we describe the implementation of high-performance compression and decompression kernels in CUDA. We begin with compression. Data in Pici tiles is accessed by GPU threads in a warp. In CUDA, a "warp" is a group of threads that executes instructions using the SIMT model. When a warp is processing data, the ideal access pattern is for threads to access adjacent locations in memory. These accesses are referred to as "coalesced" memory accesses. For compression, we assign each tile to be processed cooperatively by a single warp of 32 threads, such that each thread is responsible for two columns. Now our decision to organize row-major tiles into columns makes sense. As threads access elements within adjacent columns, their memory accesses are coalesced. If there are more tiles than available warps, warps process multiple tiles. Algorithm 1 presents a high-level, simplified view of compression, which proceeds in three phases.

In **Phase 1** of Algorithm 1, each thread in the warp scans the two columns it is responsible for (Lines 5–6). When encountering nonzero values, threads flip a bit in a local word corresponding to the column and increment local column counters. The kernel batches writes to B in 64-bit words, avoiding per-bitflip writes (Lines 7–8). To obtain the full tile nonzero count, the warp reduces the local column counts in parallel using the `__shfl_down_sync` warp-shuffle intrinsic (Line 9). CUDA warp shuffles allow for the direct register-to-register exchange of values across a warp's threads, bypassing slower DRAM and shared memory. The reduction requires only $\log_2 32 = 5$ shuffles, as one reduction step uses the two thread-local column counts. After reduction, the first lane in the warp writes the tile count to the index array I (Line 11).

In **Phase 2** of Algorithm 1, we prefix-sum the tile counts to complete the prefix sums in I (Lines 15–17). Typically, this operation would require synchronization across CUDA blocks/streaming multiprocessors (SMs). However, in SpCCL, we ultimately compress and decompress independent data segments within each CUDA block used by collectives. We discuss the segmentation regime in more detail in Section V. The segmentation allows us to instead implement a simple

Algorithm 1: High-level, simplified Pici compression

```
   Input: Densely formatted input array D of length N
   Output: Bitvector B, index array I, values array A
   // Phase 1: Populate B, partial I
1  foreach tile_idx assigned to warp in parallel across warps do
2      t ← TileView(D, tile_idx);
3      foreach lane_idx ∈ {0, ..., 31} in parallel across threads do
           // Assigned cols are strided by 32
4          (c_1, c_2) ← (lane_idx, lane_idx + 32);
           // (col nnz, local 64-bit col bitvector)
5          (nnz_c1, bv_c1) ← ScanColumn(t, c_1);
6          (nnz_c2, bv_c2) ← ScanColumn(t, c_2);
7          B[tile_idx × 64 + lane_idx] ← bv_c1;
8          B[tile_idx × 64 + lane_idx + 32] ← bv_c2;
9          nnz_tile ← WarpShuffleReduce(nnz_c1 + nnz_c2);
10         if lane_idx = 0 then
11             I[tile_idx] ← nnz_tile;
12         end
13     end
14 end
   // Phase 2: Complete I
15 SyncWorkers();
16 ExclusivePrefixSumAndStoreTotalNNZ(I);
17 SyncWorkers();
   // Phase 3: Compact values array A
18 foreach tile_idx assigned to warp in parallel across warps do
19     t ← TileView(D, tile_idx);
20     foreach lane_idx ∈ {0, ..., 31} in parallel across threads do
           // Thread-assigned cols are adjacent
21         (c_1, c_2) ← (2 × lane_idx, 2 × lane_idx + 1);
22         bv_c1 ← B[tile_idx × 64 + c_1];
23         bv_c2 ← B[tile_idx × 64 + c_2];
24         (nnz_c1, nnz_c2) ← Popcounts(bv_c1, bv_c2);
25         pair_off ← WarpExPrefixSum(nnz_c1 + nnz_c2);
26         offset_c1 ← I[tile_idx] + pair_off;
27         offset_c2 ← offset_c1 + nnz_c1;
28         CompactColumnIntoA(t, c_1, bv_c1, A, offset_c1);
29         CompactColumnIntoA(t, c_2, bv_c2, A, offset_c2);
30     end
31 end
```

block-local prefix sum and fuse phases 1-3 together in a single device function.

In **Phase 3** of Algorithm 1, we compact nonzero values into A using the now fully constructed B and I. Threads still process two columns in each tile (two batches of columns per tile). Except for column zero, the number of nonzeros preceding each column is not immediately available. We build per-column offsets into A on the fly using a Hillis-Steele-style prefix sum implemented via warp shuffles on per-column nonzero counts obtained by the `__popcll` popcount intrinsic (two per thread, Lines 24–27) [28]. To quickly identify nonzeros within a column, we apply to bitvector words a combination of masking and the `__ffsll` intrinsic, which identifies the location of the word's least significant 1-bit. These operations occur within the CompactColumnIntoA calls on Lines 28–29. As per-column nonzero counts may vary, warp divergence can occur here, but the maximum per-column nonzero count in the column batch bounds loop iterations. Depending on data sparsity, the number of iterations may be significantly fewer than 64.

Decompression follows essentially the same procedure as Phase 3 of compression but moves values from the compacted values array into a dense buffer. We also implement a scatter-add variant of "decompression," facilitating reductions in all-reduce and reduce-scatter. Note how the popcount intrinsic and warp shuffles allow Pici to use a relatively large k value of 4096 while still supporting intra-tile parallelism. Using popcounts executed in parallel across a warp's threads, followed by a short warp-shuffle prefix sum, we can produce temporary prefix sums for each column in a tile. This lets Pici retain the low index overhead of $k = 4096$ while still efficiently computing temporary column-granularity offsets into A and using them to distribute a tile's work across 32 threads.

IV. Sparse Collective Algorithms in SpCCL

Next, we describe the design of our collectives. We focus on the algorithmic and format-selection decisions that we make, deferring implementation details to Section V.

A. Selection of Collective Algorithms

We must select a base algorithm for each of our communication collectives. For example, we could use a ring algorithm or recursive doubling/halving as discussed in Section II-A. However, we are limited by which algorithms NCCLX supports. For reduce-scatter and all-gather, NCCLX supports only Ring on Perlmutter. Some algorithms such as CollNet require specialized hardware, and NCCLX only implements Parallel Aggregated Trees for inter-node communication [29]. For all-reduce, NCCLX also implements and supports the Tree algorithm on Perlmutter. Similar to the recursive algorithms described in Section II-A, Tree has a logarithmic number of communication steps.

To choose between Ring and Tree for all-reduce, we consider empirical results from several collective communication studies [30]–[32]. These studies find that for large messages, Ring often performs better empirically than recursive algorithms, despite having a worse latency term and communicating the same data volume. The authors attribute this to the ring communication pattern, which can support effective pipelining and bandwidth utilization. We validate this intuition by comparing NCCLX's Ring and Tree algorithms for a large collective input of 512 MiB and a smaller input of 50 MiB. Ring outperforms Tree at both sizes, suggesting that it is an appropriate choice for our regime. All of our sparse collective implementations utilize Ring.

B. Where and When SpCCL Compresses and Decompresses

Compression and decompression are necessary for sparse communication to occur, but their placement and frequency within the algorithm differ by collective. All-gather is the simplest and most forgiving case. Data is unmodified during the course of the collective, so compression can occur once before communication. In our implementation, each process compresses its contribution to the output before the first ring communication step and forwards the sparse representation around the ring unmodified. As a process receives shards of data from peers at each step, it immediately decompresses them into the output buffer. Consider an all-gather on p processes with an expected output buffer of size n, data

sparsity σ, and $\omega = 0.0315$ representing Pici's metadata overhead. Our extended, idealized cost model is:

$$n' = n \times (1 - \sigma + \omega) \tag{4}$$

$$T_{\text{cd}} = T_{\text{comp}}(\frac{n}{p}) + T_{\text{decomp}}(n' \times \frac{p-1}{p}) \tag{5}$$

$$T_{\text{sparse_ring}} = (p-1) \times \left(\alpha + \beta \times \frac{n'}{p}\right) + T_{\text{cd}} \tag{6}$$

On the other hand, reduce-scatter must accumulate data after receipt from peers at each ring step. We accumulate into a dense buffer using our Pici scatter-add kernel. For the next communication step to benefit from sparsity, this accumulated dense buffer must be recompressed. Thus, unlike all-gather, reduce-scatter must compress data at each of the $p-1$ ring steps. Because reduce-scatter has additional computational overhead and densifies data, we expect smaller speedups for reduce-scatter than for all-gather. Our ring all-reduce is implemented as a two-phase RS+AG, and each phase mirrors the behavior of our standalone RS/AG collectives.

C. An Adaptive Algorithm for Dynamic Sparsity

One purpose of a sparse representation such as COO or Pici is to reduce storage requirements and communication volume relative to a dense representation. However, metadata requirements result in the size of the sparse representation matching or exceeding the dense one for low sparsities (for example, at 50% sparsity with COO). In these cases, communicating the sparse representation would be wasteful, especially when compression and decompression overheads are considered. A simple heuristic for communication collectives is to swap data to a dense format when the size of the sparse representation exceeds the dense representation in size. This observation has been made previously in other work, such as SparCML (using COO), which uses it to swap to a dense format during all-reduce reduction steps [10].

However, the format-size heuristic does not adequately account for the holistic performance of sparse collectives. Swapping formats (whether COO or Pici) based only on representation sizes is insufficiently conservative unless $T_{\text{comp}} = T_{\text{decomp}} \approx 0$, which is unrealistic.

For our collectives, we add an adjustable threshold τ that switches to communicating a dense representation when sparsity σ is $\leq \tau$. This decision is typically informed by the degree of sparsity computed as a byproduct of an earlier compression. However, the reduction process does not naturally compute the post-reduction sparsity after receiving a peer's data, so the decision of whether to compress data before sending is made based on an *extrapolated* post-reduction sparsity. We assume uniform sparsity and independent positions across processes and extrapolate according to the following formula, derived from the inclusion-exclusion principle:

$$\sigma_{\text{next}} = \sigma_{\text{prev}}\sigma_0 \tag{7}$$

where σ_{next} is the extrapolated sparsity, σ_{prev} is the previously computed sparsity, and σ_0 is the initially computed sparsity, measured at step 0. Step 0 always executes compression, and the first threshold-informed send decision is based directly on the computed sparsity. For example, if we set the threshold to 0.6 and sparsity is extrapolated to be $\leq$60% during an intermediate step, the collective sends densely formatted data.

We also consider that in modern GPU clusters, communication typically occurs over heterogeneous networks. For example, Perlmutter nodes feature four 25 GB/s/direction NVLink links connecting each GPU pair, while each GPU node is equipped with four 25 GB/s/direction Slingshot 11 NICs. Thus, nominal peak aggregate off-node injection bandwidth is equivalent to the one-way bandwidth between a single pair of GPUs. Consequently, inter-node sparse sends should more readily provide speedups. To exploit this, we introduce **link-aware** adaptivity, where the collective makes decisions on send format based on whether it is sending data over inter- or intra-node links. We support this with a second threshold: $intra_thresh$. This threshold controls intra-node send decisions, while our first threshold, now $inter_thresh$, controls inter-node decisions. These are only relevant to reduce-scatter and all-reduce, as all-gather does not densify data.

Notably, $inter_thresh$ and $intra_thresh$ are not independent. Consider a process that will post an inter-node send after receiving a sparse representation from an intra-node send. It must scatter-add this data into a dense buffer, requiring a more irregular memory access pattern and more writes to memory. Thus, the decision made at the prior intra-node send step affects the runtime of the inter-node send step. To identify appropriate thresholds, we sweep combinations of threshold pairs at intervals of 0.1 across various GPU counts and input/output buffer sizes for SpCCL using Pici. Figure 2 illustrates a subset of one of these sweeps. We observe that for both thresholds, thresholding either too aggressively or too conservatively results in slowdowns. While the exact best thresholds are mildly sensitive to collective input size, we find that for fp32 reduce-scatter, $intra_thresh = 0.6$ and $inter_thresh = 0.5$ are generally strong choices.

Finally, for all-reduce, the expected ideal threshold varies by algorithm phase (RS/AG). If, as suggested by our cost model, the sparse all-gather phase can achieve speedups at lower sparsities, it does not make sense to forward data in a dense format through all-gather just because the reduce-scatter phase terminated with a dense representation. We propose **phase-aware** adaptivity and introduce an all-gather threshold: ag_thresh. This determines at the RS-AG transition step in all-reduce whether the AG phase should proceed with a sparse or dense representation. We find that a value of 0.1 is appropriate for all-gather, which is, as expected, much lower than the RS thresholds.

Although ideally we could compute thresholds analytically, the thresholds depend on too many interacting factors for this to be practical (compression and decompression speeds at different sparsities and message sizes, heterogeneous link bandwidths, densification rate, etc.). Empirical tuning is a practical choice and incurs only a one-time offline cost.

Observe that our thresholding algorithm permits one or

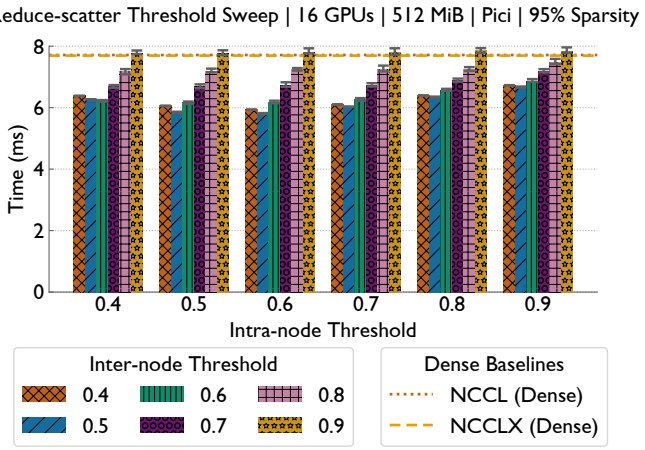

Fig. 2. Impact of intra- and inter-node threshold values on reduce-scatter performance with Pici using a 512 MiB input size. Initial input sparsity is 95%, resulting in substantial densification on 16 GPUs. Observe that swapping to dense too soon (high thresholds) or too late (low thresholds) limits the performance improvements that can be realized over dense.

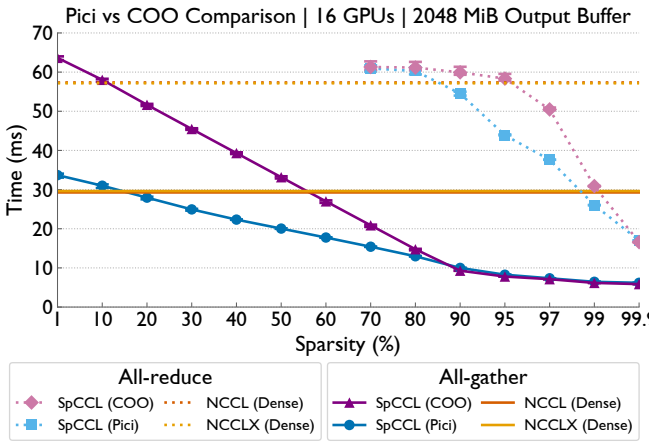

Fig. 3. Comparison of Pici and COO in SpCCL with NCCL (Dense) and NCCLX (Dense) for all-gather and all-reduce on 16 GPUs. While COO can perform comparably or slightly better than Pici at higher sparsities, it becomes less effective as sparsity decreases, especially for all-reduce. Results are similar for other tested buffer sizes and GPU counts.

more dense-send ring steps to be sandwiched between sparse-send ring steps. After a dense-send step (where the nonzero count is not computed), measured data sparsity is unavailable at the next ring step. This necessitates that we extrapolate densification forward through (potentially multiple) dense-send ring steps, using Equation 7, where σ_{prev} may now represent a prior-step computed sparsity or a previous extrapolation. When compressions occur after extrapolations, they calibrate the tracked sparsity using the computed nonzero count.

D. COO Support in SpCCL

In addition to Pici, we implement support for COO in our collectives and perform additional sweeps to identify appropriate COO thresholds. These thresholds are higher than for Pici, at $intra_thresh = 0.8$, $inter_thresh = 0.7$, and $ag_thresh = 0.5$, respectively. Higher thresholds are expected because COO becomes less space-efficient relative to Pici as sparsity decreases.

Figure 3 compares SpCCL using Pici and COO with dense NCCL and NCCLX baselines for sparse all-gather and all-reduce. While COO achieves comparable or slightly superior performance to Pici at higher sparsities, it falls off significantly as sparsity decreases, especially for all-reduce, where it is comparable only within the HS regime above 99% sparsity. For all-reduce, Pici is $1.19\times$ faster than COO at 99% sparsity and $1.33\times$ faster at 95% (where COO is slower than dense). For all-gather, COO remains competitive at slightly lower sparsities (down to around 90%) than would be expected based on space requirements alone. We attribute this to its simpler decompression routine, because decompression dominates all-gather overhead. While compression is similar to Pici in structure (identify nonzeros, build an index for stream compaction, compact nonzeros), decompression only requires reading index-value pairs.

V. Implementing Sparse Collectives in NCCLX

As previously mentioned, SpCCL collectives are implemented as extensions to NCCLX in order to exploit preexisting collective optimizations rather than reimplementing them in a redundant engineering effort [14]. Our implementation uses NCCLX's underlying communication components that are largely unchanged from stock NCCL, so our descriptions of how NCCLX operates apply to NCCL as well.

A. Overview of the NCCLX Collective Library

We now provide an overview of NCCLX concepts that are relevant to our implementation and optimizations. For a more comprehensive analysis of NCCL, we refer the reader to the work of Hu et al. [33]. Because NCCLX collectives are implemented in CUDA kernels, they use streaming multiprocessors (SMs) to handle GPU work. Without sufficient SMs, a collective can fail to post enough work to saturate links. To address this, NCCLX splits collectives into **channels**. Each channel corresponds to a single CUDA block that executes on a single SM, and each channel has an independent progress engine driving progress on disjoint segments of the data.

In the Simple communication protocol, after dividing the work that each process is responsible for into channels, NCCLX further subdivides work into chunks of **slices**. Slices are the granularity of individual network transmissions, and they are pipelined using send and receive FIFO buffers, with pointers into these buffers controlling producer-consumer synchronization. For intra-node receives with our configuration, the receive FIFO is the peer-mapped memory comprising a peer's send FIFO (enabled by GPUDirect P2P over NVLink [34]). For inter-node receives, a host-side proxy thread interacts with NICs, which write data directly to GPU DRAM without going through the host. The "receive + reduce (AR/RS) + send" slice lifecycle is handled by fused NCCLX primitives that instantiate different versions of the templated `genericOp`

device function. Other communication protocols would complicate compression and decompression and are designed for low-latency regimes.

B. Extending NCCLX to Support SpCCL Algorithms

To support compression in NCCLX, we must decide at which granularity to apply it. The most natural solution is to compress data at the slice level immediately before sends are posted and to decompress immediately upon receipt. Compressing data at any other granularity would clash with NCCLX's architecture and interfere with pipelining, as the slice is the unit of pipelining. As a consequence, we implement compression and decompression within `genericOp`.

When sending dense-format messages, the collective determines message sizes statically based on input data size. When sending messages in a sparse format with SpCCL, the collective cannot determine message sizes or message format in advance for any individual receipt. To communicate message formats and sizes to peers, we extend the Simple protocol by prepending a 48-byte header to all messages. The header contains format and message size metadata, as well as offsets indicating where different components of our sparse formats begin.

Sparse representations can exceed the size of dense representations as sparsity decreases, so we also enlarge the FIFO buffers to ensure that even a worst-case compressed slice fits within its FIFO reservation. For Pici, format overhead is only 3.15% relative to the dense representation, a significant advantage over COO's worst-case, in which one index is needed per nonzero value. For reduction, NCCLX accumulates directly into the send FIFO. For sparse steps, we cannot do this, as we must compress into the send FIFO. We reuse the application send/input buffer as accumulation scratch and then compress into the send FIFO buffer. This allows us to avoid allocating additional scratch space and preserves dense collective signatures. For dense reductions, we reuse heavily vectorized NCCLX primitives such as `reduceCopy`.

C. Optimization 1: Increasing Parallelism with Channels

NCCLX heuristically adjusts the number of channels a collective uses (capped at 64). When relying on these heuristics, we observe that our collectives significantly underperform dense baselines. We are incurring additional compute overhead for each (generally smaller) message, so the default heuristics about how many channels are necessary to saturate the links no longer hold. We experiment with raising the number of channels our collectives utilize, sweeping 4–64 channels across multiple GPU counts and input/output buffer sizes. Figure 4 illustrates the results of one of our channel sweeps.

We observe that while dense collectives generally do not benefit from using a large number of channels, sparse collectives benefit substantially from elevated channel counts. For example, 64 channels is best 72.5% of the time for sparse all-gather (with 32 channels being best in remaining cases). Cases where 32 channels win tend to be those with the smallest collective inputs. The maximum size of a slice is fixed by the implementation and runtime configuration. Thus, with small enough inputs, adding more channels can result in multiple channels each sending the same number of smaller slices rather than reducing the number of per-channel slices.

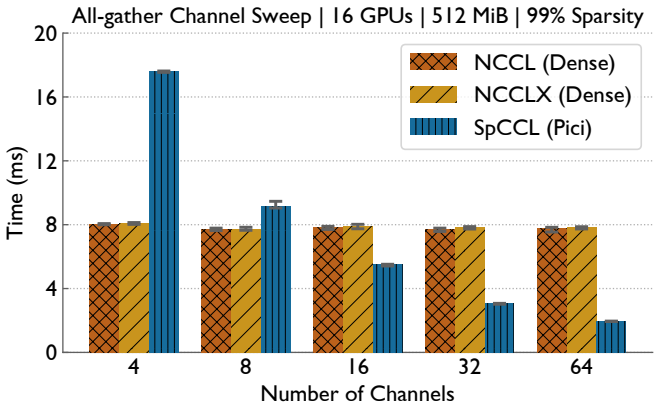

Fig. 4. Impact of the number of channels used on all-gather performance on 16 GPUs (512 MiB output, 99% sparsity). Observe that sparse all-gather (Pici) benefits substantially from additional channels.

Using more channels can reduce the potential for computation-communication overlap due to the collective utilizing more SMs and placing more concurrent demand on shared memory-system resources. Thus, SpCCL's current implementation may be less well-suited to workloads with heavy computation-communication overlap. The overlap and contention concerns may be partially mitigated in cases where sparsity provides speedups, as the SMs remain occupied by the collective for less time.

D. Optimization 2: Avoiding Uncoalesced Reads over NVLink

For intra-node receives, data is available via the FIFO buffer (peer-mapped memory). For dense accumulation, the memory access pattern is perfectly coalesced. However, scatter-adds from compressed data result in uncoalesced reads over NVLink. As NVLink has much lower bandwidth than local HBM, when receiving Pici-formatted data, we first copy the received data into a staging buffer and perform the scatter-add from local memory. To avoid allocating an additional staging buffer, we reuse the send FIFO. This is safe, as the scatter-add must complete before we compress back into the send FIFO.

VI. Experimental Setup

Below, we provide an overview of our methodology for evaluating SpCCL collectives on Perlmutter [15].

A. Collective Microbenchmarks

First, using Pici, we compare the standalone performance of SpCCL all-gather, reduce-scatter, and all-reduce with NCCL/NCCLX baselines. For all-gather, we report the per-GPU output size; for reduce-scatter, the input size; for all-reduce, the shared input/output size. We sweep a range of sparsity levels. Each element is independently selected to be nonzero with probability equal to $1 - \sigma$. Each process uses a

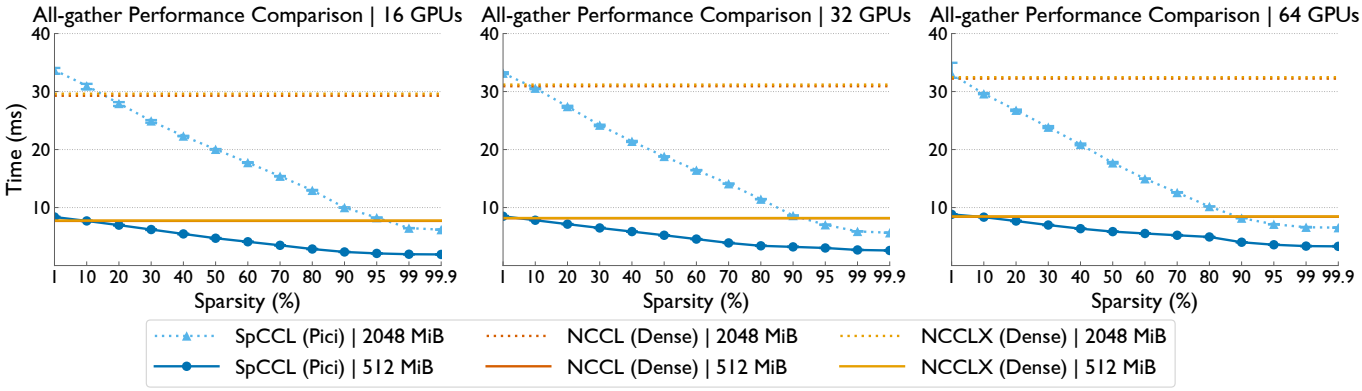

Fig. 5. Performance comparison of SpCCL all-gather (Pici) versus NCCL (Dense) and NCCLX (Dense) baselines on several GPU counts for 512 and 2048 MiB output sizes. SpCCL all-gather significantly outperforms dense baselines, only falling behind dense at 1–10% input sparsity.

distinct pseudorandom seed. With respect to densification, this is a pessimistic non-adversarial scenario.

For all-reduce, we also compare against SparCML sparse all-reduce implementations [10]. SparCML supports four algorithms for all-reduce, including ring and recursive variants. OmniReduce is a potential additional sparse baseline, but its reliance on InfiniBand RDMA verbs prevents us from using it on Perlmutter without porting to CXI libfabric [11].

We conduct all experiments on Perlmutter [15]. GPU nodes have either four 40 GB or four 80 GB A100 GPUs. We use 40 GB nodes in all microbenchmarks. Inter-node communication occurs over the Slingshot 11 interconnect fabric, while intra-node communication uses NVLink connections.

We compile SpCCL and NCCLX against CUDA 12.9. Our dense NCCL baseline uses Perlmutter's stock nccl/2.24.3 module. We obtain NCCL runtimes with NVIDIA's nccl-tests software [35]. We implement a microbenchmarking harness based on nccl-tests and osu-micro-benchmarks for evaluating NCCLX, SpCCL, and SparCML (but make no modification to SparCML collectives) [36], [37].

To account for network variability [38]–[40], we run each tested configuration across at least five separate Slurm jobs. Within each non-tuning job, each configuration runs five warmups and 1000 timed iterations. We report runtimes as the average across jobs, processes, and iterations at the best tested channel count, with an untimed barrier before each iteration and a timed barrier at the end of each iteration (consistent with nccl-tests behavior). For Pici, we set $inter_thresh = 0.5$, $intra_thresh = 0.6$, and $ag_thresh = 0.1$. We remove obvious outlier data points and additionally plot min/max times across jobs (never < 4 after filtering). SparCML runs use 50 timed iterations due to significantly elevated runtimes relative to NCCL and SpCCL. For each SparCML data point, we report results from the best-performing algorithm.

B. Application Case Study: Gradient-Pruned DDP Training

We evaluate SpCCL (Pici) on gradient-pruned Distributed Data Parallel (DDP) training [21]. We implement gradient pruning using a sampling-based top-k approach [6], imposing 99% collective input sparsity and determining the pruning threshold by sampling 0.1% of gradients. We further implement a leaky error feedback (EF) mechanism [6], [41] in a Triton kernel. Although collective input sparsity is 99%, process-dependent nonzero positions cause sparsity to decrease during the all-reduce.

We experiment with a 1.5B GPT-2 XL-based [42] model and a 3.3B StarCoder2-3B-based [43] model on 40 GB and 80 GB GPU nodes, respectively. We reduce the latter to 28 layers due to the additional memory overhead of EF. We set the sequence length to 512 and the micro-batch size to two per GPU. We train for 100 iterations on the WikiText-103 dataset [44] using Megatron-LM as our training framework [45]. We run each configuration three times, reporting mean runtimes of the last 80 iterations. We discard the first 20 iterations as warmup. The breakdowns in Figure 12 use only the run with the median mean iteration time for each configuration.

VII. Results

We now discuss and interpret our results from benchmarking individual collectives and from the training case study.

A. Microbenchmark Results

NCCL and NCCLX dense baselines achieve near-identical performance (Figures 5, 9). Based on our cost model, we expect to observe the largest speedups over dense baselines at the highest sparsities, with greater speedups for all-gather than for reduce-scatter and all-reduce. This is consistent with our results. Using Pici, SpCCL all-gather achieves speedups over dense at sparsities as low as 10%, highlighting Pici's fast decompression and ability to save space even at 10% sparsity (Figure 5, 64 GPUs, 2048 MiB). At 99% sparsity, SpCCL all-gather achieves speedups of up to 5.25× (Figure 6). Additionally, when we hold per-GPU input size constant for all-gather, we observe that speedups over NCCL actually *increase* with GPU count. This can be observed in the bottom-left-to-top-right diagonals in the heatmaps of Figure 6 (rows refer to the output buffer size). This is likely due in part to the step-0-only compression cost being amortized over more ring steps.

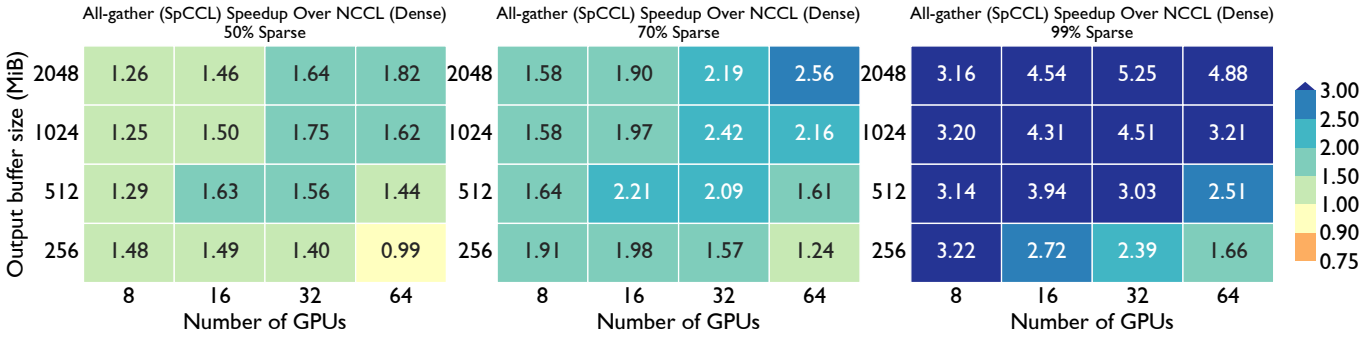

Fig. 6. SpCCL all-gather (Pici) speedups over NCCL all-gather (Dense). Speedups are greatest for large output buffer sizes and at higher sparsities, but substantial speedups are achieved even at 50% sparsity. SpCCL all-gather achieves speedups of up to $5.25\times$ at 99% sparsity.

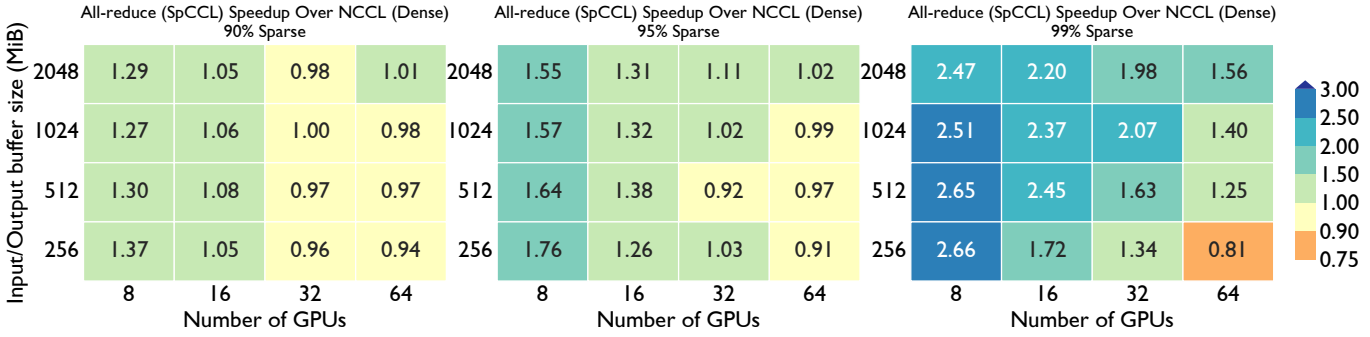

Fig. 7. SpCCL all-reduce (Pici) speedups over NCCL (Dense). At 99% sparsity, substantial speedups can be achieved even at 64 GPUs, while for lower sparsities of 90% and 95%, speedups are only consistently observed on lower GPU counts due to densification. SpCCL all-reduce achieves speedups of up to $2.66\times$ at 99% sparsity.

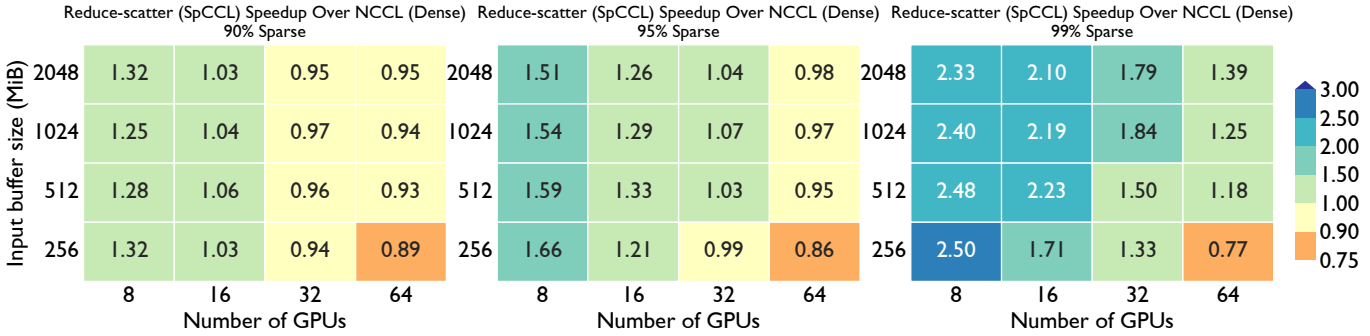

Fig. 8. SpCCL reduce-scatter (Pici) speedups over NCCL (Dense). At 99% sparsity, substantial speedups can be achieved even at 64 GPUs, while for lower sparsities of 90% and 95%, speedups are only consistently observed on lower GPU counts due to densification. SpCCL reduce-scatter achieves speedups of up to $2.5\times$ at 99% sparsity.

For sparse all-reduce and reduce-scatter, speedups over NCCL at a given input sparsity decrease as process count rises (Figures 7 and 8). This is expected due to densification, as each additional ring step decreases sparsity. Maximum speedups at 99% sparsity are $2.66\times$ (AR) and $2.5\times$ (RS). At 99% sparsity, worthwhile speedups can be observed even up to 64 GPUs. The exception is the smallest 256 MiB input.

The per-process per-channel contribution at 64 GPUs with 32 active channels while using stock NCCLX is only 128 KiB for a 256 MiB input buffer. With the default FIFO buffer size, this is split over two slices, yielding 64 KiB individual message sizes. Prior work establishes MPI RMA put latencies on Perlmutter over Slingshot 11 to be about 2.49 μs [46]. The crossover point between latency-bound and bandwidth-bound messages in the Hockney model is $n^* \approx \alpha/\beta$, which for Perlmutter inter-node messages is 60.8 KiB $\approx \frac{2.49\ \mu s}{1/25\ \text{GB/s}}$. Although MPI RMA put latencies are not perfectly representative of our NCCLX send path and the exact slice size may vary by FIFO size, this is consistent with densely formatted slices being near the latency/bandwidth crossover, before accounting

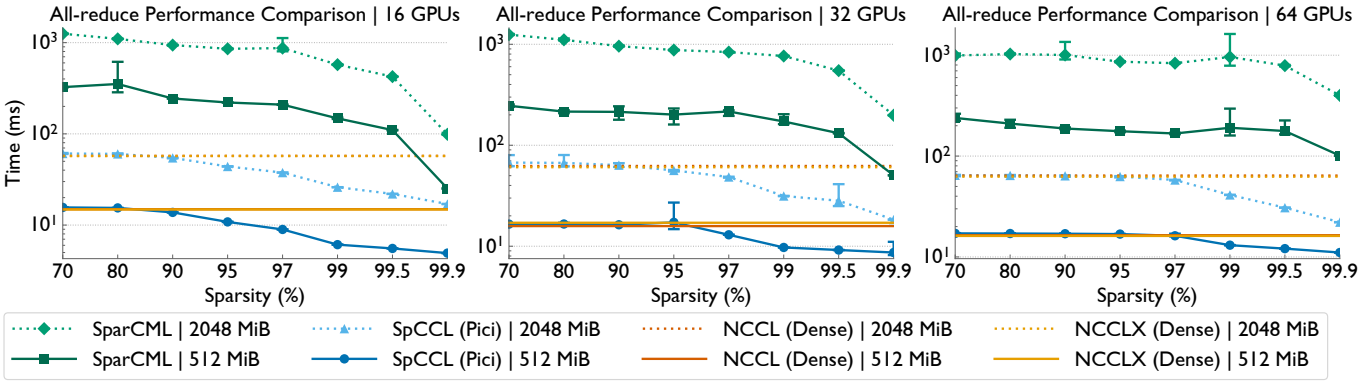

Fig. 9. Comparison of SpCCL all-reduce (Pici) with dense (NCCL and NCCLX) baselines and SparCML on several GPU counts and input buffer sizes. Speedups are highest at higher input sparsities and decrease (due to densification) as sparsity decreases or as GPU counts increase.

for compression and decompression overheads. Thus, the slowdown at 256 MiB inputs on 64 GPUs is not surprising. For such messages, our threshold choices based on the bandwidth-bound regime are no longer reasonable, and it would be better to set very aggressive thresholds or fall back to dense NCCLX.

Figure 9 demonstrates that SparCML is not competitive, even against dense NCCL/NCCLX. In most cases, it is at least an order of magnitude slower. The SparCML authors originally report speedups against dense MPI. However, SparCML is not GPU-aware, and the original experiments used much slower networks. Our results highlight the utility in starting a sparse collective implementation from a strong dense baseline on modern networks.

Figure 10 isolates the impact of SpCCL's most significant design choices. Observe that each design decision or optimization produces a measurable speedup, and only the fully optimized configuration outperforms NCCLX in this experiment configuration. Increasing the channel count provides the largest individual improvement ($3.80\times$).

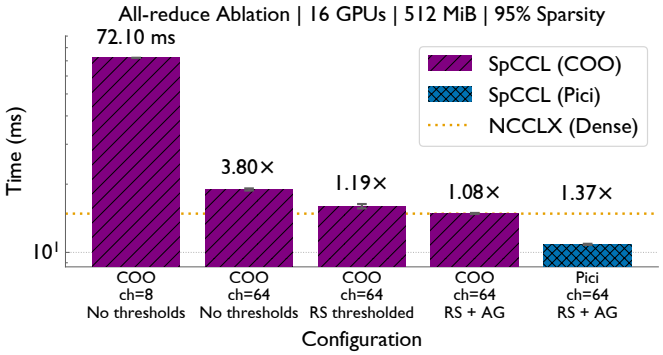

Fig. 10. Ablation demonstrating the effect of increasing channels, applying thresholds during the reduce-scatter (RS) and all-gather (AG) phases, and exchanging COO for Pici for 95% sparse all-reduce inputs on 16 GPUs. Annotations represent speedups relative to the preceding configuration.

B. Case Study: Gradient-Pruned DDP Training

Next, we evaluate SpCCL in an end-to-end gradient-pruned DDP training application. Figure 11 shows weak scaling

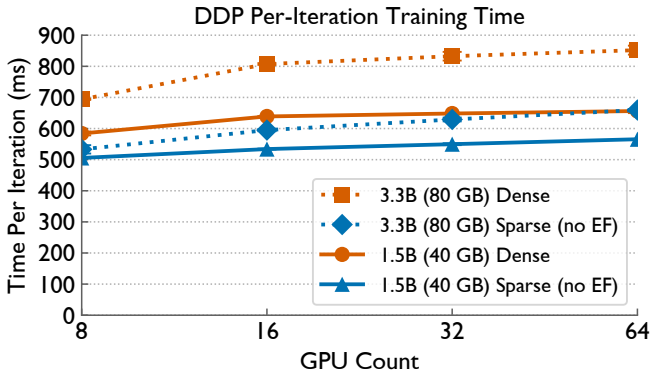

Fig. 11. Performance impact of using SpCCL in DDP training for 40 GB and 80 GB GPUs. Pruning improves end-to-end time for all GPU counts.

performance across GPU counts for 40 GB and 80 GB GPUs. End-to-end performance with SpCCL is at least 13% faster than the dense baseline on eight 40 GB GPUs and 22% faster on 64 80 GB GPUs. The improvement peaks at 16 GPUs, reaching 16% and 26% when running on 40 GB and 80 GB GPUs, respectively.

Figure 12 illustrates the time spent in all-reduce and pruning over 80 iterations on 32 GPUs with and without EF. The 80 GB runs spend more time in all-reduce than 40 GB runs, which explains the speedup differences. Without EF, in both configurations, SpCCL achieves over $2\times$ speedup on the all-reduce portion, while pruning introduces 1.5 s and 3 s of overhead for 40 GB and 80 GB runs, respectively. Pruning time scales linearly with gradient size, as it compares each tensor element against the threshold once and possibly discards it. Since we compute the threshold over a sampled subset of gradients, the sorting cost is minimal.

Error feedback introduces overhead in both pruning and all-reduce. In scaling tests, we do not observe EF overhead increasing with GPU count. During pruning, additional error buffer memory traffic accounts for this cost. The source of all-reduce overhead is less apparent. We hypothesize that as

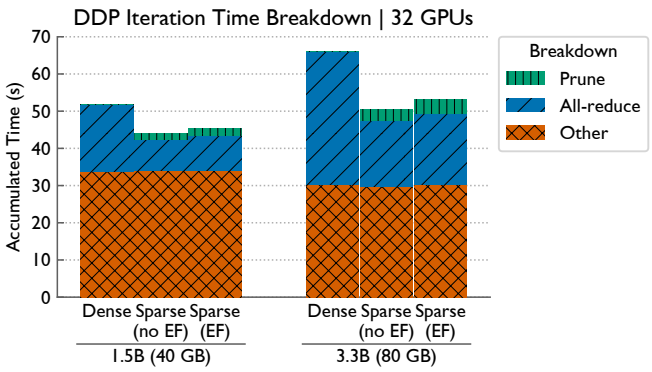

Fig. 12. Breakdown of all-reduce and pruning time on 32 GPUs. All-reduce constitutes a larger fraction of total time in 80 GB runs than in 40 GB runs. In all cases, sparse collectives reduce all-reduce time by over 40%.

error accumulates, the combined error and gradients become more uniform in magnitude during certain iterations. This may reduce compression efficiency, degrading all-reduce performance. SpCCL achieves significant speedups with and without EF, demonstrating that sparse collectives can effectively accelerate gradient synchronization in distributed training.

We note that although our application study is intended to be a realistic test of SpCCL, it is not exhaustive of all sparsification strategies and model scales. Results may be better or worse on different platforms or training scenarios, and we do not claim that our sparsification strategy is optimal. Rather, we highlight the efficacy of SpCCL as a tool for conducting sparse training research.

VIII. Related Work

Although prior work explores sparse communication collectives, few systems implement exact collectives designed to work with arbitrary sparsity patterns (particularly outside of all-reduce). SparCML is the work that is conceptually most similar to SpCCL. It supports multiple sparse all-reduce implementations but operates on host buffers [10]. Unlike SpCCL, it supports only a COO-style representation of sparse data. OmniReduce implements sparse all-reduce as a streaming aggregation system using a block-based sparse representation. Its microbenchmarks apply block-wise sparsity rather than element-wise sparsity. Block-wise sparsity is highly favorable to OmniReduce's approach [11], [47]. Wijerathne et al. describe block-based collectives implemented in RCCL (and thus exclusive to AMD GPUs) and demonstrate speedups on block-wise sparse inputs [48]. OmNICCL implements sparse all-reduce but relies on BlueField-2 SmartNICs for in-network aggregation [49]. D-DOSA similarly relies on DPUs/SmartNICs [50].

ZEN is a holistic sparse gradient synchronization system with all-reduce semantics and load balancing, whereas SpCCL is a drop-in NCCL replacement for supported collectives that works with GPU-resident unstructured sparse inputs [51]. ZEN also demonstrates speedups over baselines, albeit on lower-bandwidth networks and with more support overlap (less fill-in) than in the SpCCL microbenchmarks, making the results not directly comparable.

Several works implement *lossy* sparse collectives, introducing sparsity during the course of the collective. This may be tolerable for some applications, but lossy collectives are out of scope for this work: SpCCL collectives are exact. O*k*-Top*k* and gTopKAllReduce exemplify this line of research [52], [53].

Hecate implements sparse collectives, but they are sparse in their communication pattern (which processes participate) rather than in the data communicated [54]. Similarly, Hoefler et al. also describe collectives that are sparse in their communication pattern [55].

Coruscant introduces a tile- and bitvector-based sparse matrix representation that is similar to Pici [9]. However, the authors target SpMM rather than communication or standalone whole-input CUDA compression and decompression. It uses one index per 64-element column in a tile, resulting in 50% index space overhead relative to the bitvector, compared to Pici's $\approx 0.78\%$.

Our sparse collectives are exact with respect to their inputs. Thus, specific strategies for gradient pruning such as Deep Gradient Compression or Gaussian$_k$ are orthogonal to our work [6], [56]. Such methods may benefit from our work so long as they do not require additional sparsification during collectives or alter communication schedules.

IX. Conclusion

In this work, we introduce improvements for the collective communication of sparse data in all-gather, all-reduce, and reduce-scatter across a range of sparsities. With Pici, we reduce the representation overhead of communicating data in the non-high-sparsity regime without making assumptions about sparsity structure. Optimized Pici kernels enable efficient compression and decompression. Through our phase- and link-aware adaptive algorithm, we avoid compression and decompression when their overhead is predicted to outweigh the benefits of sending less data. By implementing SpCCL as an extension to NCCLX, we build our library on a highly optimized foundation, and we highlight the integration and optimization decisions that make the integration effective. We demonstrate substantial speedups over NCCL of up to $5.25\times$ for all-gather, $2.5\times$ for reduce-scatter, and $2.66\times$ for all-reduce at 99% sparsity. These gains translate to end-to-end performance improvements of up to 26% during gradient-pruned training of 1.5B and 3.3B LLMs on Perlmutter, emphasizing the utility of our approach in a representative application.

Acknowledgments

The authors wish to thank Siddharth Singh, Prajwal Singhania, Ashley Adames Perez, Bahar Asgari, Donghyeon Joo, and Alan Zaoxing Liu for valuable advice and feedback on this work. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility, operated under Contract No. DE-AC02-05CH11231 using NERSC

awards DDR-ERCAP0034262, ALCC-ERCAP0034775, and DDR-ERCAP0037864.

AI usage: Claude Code (Opus 4.X) and Codex (GPT 5.X) assisted with implementing code for portions of the methodology and experiments, plotting scripts, and minor text and grammar polishing throughout the paper. The authors developed the primary conceptual contributions, and the core GPU kernels were predominantly human-authored.

REFERENCES

[1] J. D. Trotter, S. Ekmekçibaşı, D. Sağbili, J. Langguth, X. Cai, and D. Unat, "Cpu- and gpu-initiated communication strategies for conjugate gradient methods on large gpu clusters," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 298–315. [Online]. Available: https://doi.org/10.1145/3712285.3759774

[2] E. Bavier, M. Hoemmen, S. Rajamanickam, and H. Thornquist, "Amesos2 and belos: Direct and iterative solvers for large sparse linear systems," *Scientific Programming*, vol. 20, no. 3, p. 243875, 2012. [Online]. Available: https://doi.org/10.3233/SPR-2012-0352

[3] X. S. Li, P. Lin, Y. Liu, and P. Sao, "Newly released capabilities in the distributed-memory superlu sparse direct solver," *ACM Trans. Math. Softw.*, vol. 49, no. 1, Mar. 2023. [Online]. Available: https://doi.org/10.1145/3577197

[4] Y. Qiu, *Scalable and Efficient Material Point Methods on Modern Computational Platforms*. University of California, Los Angeles, 2024.

[5] S. Singh and A. Bhatele, "Exploiting sparsity in pruned neural networks to optimize large model training," in *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2023, pp. 245–255. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/IPDPS54959.2023.00033

[6] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," *arXiv preprint arXiv:1712.01887*, 2017. [Online]. Available: https://doi.org/10.48550/arXiv.1712.01887

[7] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," *arXiv preprint arXiv:1803.03635*, 2018. [Online]. Available: https://doi.org/10.48550/arXiv.1803.03635

[8] A. K. Ranjan, S. Singh, C. Wei, and A. Bhatele, "Plexus: Taming billion-edge graphs with 3D parallel full-graph GNN training," in *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. ACM, Nov. 2025. [Online]. Available: https://doi.acm.org/10.1145/3712285.3759890

[9] D. Joo, H. Hosseini, R. Hadidi, and B. Asgari, "Coruscant: Co-designing gpu kernel and sparse tensor core to advocate unstructured sparsity in efficient llm inference," in *Proceedings of the 58th IEEE/ACM International Symposium on Microarchitecture®*, 2025, pp. 232–245. [Online]. Available: https://doi.org/10.1145/3725843.3756065

[10] C. Renggli, S. Ashkboos, M. Aghagolzadeh, D. Alistarh, and T. Hoefler, "Sparcml: High-performance sparse communication for machine learning," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–15. [Online]. Available: https://doi.org/10.1145/3295500.3356222

[11] J. Fei, C.-Y. Ho, A. N. Sahu, M. Canini, and A. Sapio, "Efficient sparse collective communication and its application to accelerate distributed deep learning," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 676–691. [Online]. Available: https://doi.org/10.1145/3452296.3472904

[12] N. Bell and M. Garland, "Efficient sparse matrix-vector multiplication on cuda," Nvidia Technical Report NVR-2008-004, Nvidia Corporation, Tech. Rep., 2008.

[13] S. Eisenstat, M. Gursky, M. Schultz, and A. Sherman, "Yale sparse matrix package. ii. the nonsymmetric codes," Tech. Rep., 1977.

[14] M. Si, P. Balaji, Y. Chen, C.-H. Chu, A. Gangidi, S. Hasan, S. Iyengar, D. Johnson, B. Liu, J. Ren *et al.*, "Collective communication for 100k+ gpus," *arXiv preprint arXiv:2510.20171*, 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2510.20171

[15] NERSC, "Perlmutter system architecture," https://docs.nersc.gov/systems/perlmutter/architecture/.

[16] "Nccl," 2020. [Online]. Available: https://docs.nvidia.com/deeplearning/nccl/user-guide/docs/overview.html

[17] R. Rabenseifner, "Optimization of collective reduction operations," in *International Conference on Computational Science*. Springer, 2004, pp. 1–9. [Online]. Available: https://doi.org/10.1007/978-3-540-24685-5_1

[18] R. W. Hockney, "The communication challenge for mpp: Intel paragon and meiko cs-2," *Parallel Comput.*, vol. 20, no. 3, p. 389–398, Mar. 1994. [Online]. Available: https://doi.org/10.1016/S0167-8191(06)80021-9

[19] A. Alexandrov, M. F. Ionescu, K. E. Schauser, and C. Scheiman, "Loggp: incorporating long messages into the logp model—one step closer towards a realistic model for parallel computation," in *Proceedings of the Seventh Annual ACM Symposium on Parallel Algorithms and Architectures*, ser. SPAA '95. New York, NY, USA: Association for Computing Machinery, 1995, p. 95–105. [Online]. Available: https://doi.org/10.1145/215399.215427

[20] W. Gropp, L. N. Olson, and P. Samfass, "Modeling mpi communication performance on smp nodes: Is it time to retire the ping pong test," in *Proceedings of the 23rd European MPI Users' Group Meeting*, ser. EuroMPI '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 41–50. [Online]. Available: https://doi.org/10.1145/2966884.2966919

[21] S. Li, Y. Zhao, R. Varma, O. Salpekar, P. Noordhuis, T. Li, A. Paszke, J. Smith, B. Vaughan, P. Damania, and S. Chintala, "Pytorch distributed: Experiences on accelerating data parallel training," *Proc. VLDB Endow.*, vol. 13, no. 12, p. 3005–3018, Aug. 2020. [Online]. Available: https://doi.org/10.14778/3415478.3415530

[22] D. Langr and P. Tvrdik, "Evaluation criteria for sparse matrix storage formats," *IEEE Transactions on parallel and distributed systems*, vol. 27, no. 2, pp. 428–440, 2015. [Online]. Available: https://doi.org/10.1109/TPDS.2015.2401575

[23] J. Gao, W. Ji, F. Chang, S. Han, B. Wei, Z. Liu, and Y. Wang, "A systematic survey of general sparse matrix-matrix multiplication," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–36, 2023. [Online]. Available: https://doi.org/10.1145/3571157

[24] Y. Saad, "Sparskit: A basic tool kit for sparse matrix computations," Tech. Rep., 1990.

[25] S. Kestur, J. D. Davis, and E. S. Chung, "Towards a universal fpga matrix-vector multiplication architecture," in *2012 IEEE 20th International Symposium on Field-Programmable Custom Computing Machines*. IEEE, 2012, pp. 9–16. [Online]. Available: https://doi.org/10.1109/FCCM.2012.12

[26] D. Zhou, D. G. Andersen, and M. Kaminsky, "Space-efficient, high-performance rank and select structures on uncompressed bit sequences," in *International Symposium on Experimental Algorithms*. Springer, 2013, pp. 151–163. [Online]. Available: https://doi.org/10.1007/978-3-642-38527-8_15

[27] F. Kurpicz, "Engineering compact data structures for rank and select queries on bit vectors," in *International Symposium on String Processing and Information Retrieval*. Springer, 2022, pp. 257–272. [Online]. Available: https://doi.org/10.1007/978-3-031-20643-6_19

[28] W. D. Hillis and G. L. S. Jr., "Data parallel algorithms," *Communications of the ACM*, vol. 29, no. 12, pp. 1170–1183, 1986. [Online]. Available: https://doi.org/10.1145/7902.7903

[29] S. Jeaugey, "Pat: a new algorithm for all-gather and reduce-scatter operations at scale," 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2506.20252

[30] J. Huang, P. Majumder, S. Kim, A. Muzahid, K. H. Yum, and E. J. Kim, "Communication algorithm-architecture co-design for distributed deep learning," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2021, pp. 181–194. [Online]. Available: https://doi.org/10.1109/ISCA52012.2021.00023

[31] R. Thakur, R. Rabenseifner, and W. Gropp, "Optimization of collective communication operations in mpich," *The International Journal of High Performance Computing Applications*, vol. 19, no. 1, pp. 49–66, 2005. [Online]. Available: https://doi.org/10.1177/1094342005051521

[32] S. Singh, K. Pradeep, M. Singh, C. Wei, and A. Bhatele, "The big send-off: Scalable and performant collectives for deep learning," in *2026 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2026, pp. 1314–1328. [Online]. Available: https://doi.org/10.1109/IPDPS65963.2026.00108

[33] Z. Hu, S. Shen, T. Bonato, S. Jeaugey, C. Alexander, E. Spada, J. Dinan, J. Hammond, and T. Hoefler, "Demystifying nccl: An in-depth analysis of gpu communication protocols and algorithms," in *2025 IEEE Symposium on High-Performance Interconnects (HOTI)*, 2025, pp. 48–59. [Online]. Available: https://doi.org/10.1109/HOTI66940.2025.00024

[34] NVIDIA, "Gpudirect," https://docs.nvidia.com/cuda/gpudirect-rdma/index.html.

[35] ——, "Nccl tests," https://github.com/NVIDIA/nccl-tests, 2017.

[36] D. K. Panda, H. Subramoni, C.-H. Chu, and M. Bayatpour, "The mvapich project: Transforming research into high-performance mpi library for hpc community," *Journal of Computational Science*, vol. 52, p. 101208, 2021. [Online]. Available: https://doi.org/10.1016/j.jocs.2020.101208

[37] Network-Based Computing Laboratory, The Ohio State University, "OSU Micro-Benchmarks," accessed 2026-04-04. [Online]. Available: https://mvapich.cse.ohio-state.edu/benchmarks/

[38] S. A. Smith, C. Cromey, D. K. Lowenthal, J. Domke, N. Jain, J. J. Thiagarajan, and A. Bhatele, "Mitigating inter-job interference using adaptive flow-aware routing," in *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '18. IEEE Computer Society, Nov. 2018. [Online]. Available: https://doi.org/10.1109/SC.2018.00030

[39] A. Bhatele, J. J. Thiagarajan, T. Groves, R. Anirudh, S. A. Smith, B. Cook, and D. K. Lowenthal, "The case of performance variability on dragonfly-based systems," in *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, ser. IPDPS '20. IEEE Computer Society, May 2020. [Online]. Available: https://doi.org/10.1109/IPDPS47924.2020.00096

[40] C. Wei, K. Pradeep, and A. Bhatele, "The case of the elusive application performance on production GPU supercomputers," in *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, ser. IPDPS '26. IEEE Computer Society, May 2026. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/IPDPS65963.2026.00079

[41] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu, "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs," in *Interspeech 2014*, 2014, pp. 1058–1062. [Online]. Available: https://doi.org/10.21437/Interspeech.2014-274

[42] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," Tech. Rep., 2019.

[43] A. Lozhkov, R. Li, L. B. Allal, F. Cassano, J. Lamy-Poirier, N. Tazi, A. Tang, D. Pykhtar, J. Liu, Y. Wei, T. Liu, M. Tian, D. Kocetkov, A. Zucker, Y. Belkada, Z. Wang, Q. Liu, D. Abulkhanov, I. Paul, Z. Li, W.-D. Li, M. Risdal, J. Li, J. Zhu, T. Y. Zhuo, E. Zheltonozhskii, N. O. O. Dade, W. Yu, L. Krauß, N. Jain, Y. Su, X. He, M. Dey, E. Abati, Y. Chai, N. Muennighoff, X. Tang, M. Oblokulov, C. Akiki, M. Marone, C. Mou, M. Mishra, A. Gu, B. Hui, T. Dao, A. Zebaze, O. Dehaene, N. Patry, C. Xu, J. McAuley, H. Hu, T. Scholak, S. Paquet, J. Robinson, C. J. Anderson, N. Chapados, M. Patwary, N. Tajbakhsh, Y. Jernite, C. M. Ferrandis, L. Zhang, S. Hughes, T. Wolf, A. Guha, L. von Werra, and H. de Vries, "Starcoder 2 and the stack v2: The next generation," 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2402.19173

[44] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer sentinel mixture models," *CoRR*, vol. abs/1609.07843, 2016. [Online]. Available: https://doi.org/10.48550/arXiv.1609.07843

[45] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-lm: Training multi-billion parameter language models using model parallelism," Tech. Rep., 2020. [Online]. Available: https://doi.org/10.48550/arXiv.1909.08053

[46] P. H. Hargrove and D. Bonachea, "Gasnet-ex rma communication performance on recent supercomputing systems," 2022. [Online]. Available: https://doi.org/10.25344/S40C7D

[47] SANDS Lab, "omnireduce-experiments," gitHub repository, commit a2bdc8082c805a1cb86499ce1eae7a9b48fbf8c4, accessed 2026-04-04. [Online]. Available: https://github.com/sands-lab/omnireduce-experiments/commit/a2bdc8082c805a1cb86499ce1eae7a9b48fbf8c4

[48] D. Wijerathne, H. Javaid, G. Zhong, D. Wu, X. Y. Kom, and M. Baldi, "Sparse collectives: Exploiting data sparsity to improve communication efficiency," in *Proceedings of the 2nd Workshop on Networks for AI Computing*, 2025, pp. 64–66. [Online]. Available: https://doi.org/10.1145/3748273.3749206

[49] T. Gu, J. Fei, and M. Canini, "Omniccl: Zero-cost sparse allreduce with direct cache access and smartnics," in *Proceedings of the 2024 SIGCOMM Workshop on Networks for AI Computing*, 2024, pp. 75–83. [Online]. Available: https://doi.org/10.1145/3672198.3673804

[50] Z. Yu, W. Li, S. Guo, Q. Li, F. Qi, and J. Xiu, "D-dosa: Dpu-based dataflow offloading and sparse allreduce framework for distributed training," *IEEE Transactions on Cloud Computing*, 2025. [Online]. Available: https://doi.org/10.1109/TCC.2025.3634564

[51] Z. Wang, Z. Xu, J. Xi, Y. Wang, A. Shrivastava, and T. E. Ng, "{ZEN}: Empowering distributed training with sparsity-driven data synchronization," in *19th USENIX Symposium on Operating Systems Design and Implementation (OSDI 25)*, 2025, pp. 537–556.

[52] S. Li and T. Hoefler, "Near-optimal sparse allreduce for distributed deep learning," in *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2022, pp. 135–149. [Online]. Available: https://doi.org/10.1145/3503221.3508399

[53] S. Shi, Q. Wang, K. Zhao, Z. Tang, Y. Wang, X. Huang, and X. Chu, "A distributed synchronous sgd algorithm with global top-k sparsification for low bandwidth networks," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2019, pp. 2238–2247. [Online]. Available: https://doi.org/10.1109/ICDCS.2019.00220

[54] Y. Qing, G. Zhu, F. Li, L. Lei, Z. Sun, X. Guan, S. Zhao, X. Chen, D. Huang, S. Wang *et al.*, "Hecate: Unlocking efficient sparse model training via fully sharded sparse data parallelism," *arXiv preprint arXiv:2502.02581*, 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2502.02581

[55] T. Hoefler and J. L. Traff, "Sparse collective operations for mpi," in *2009 IEEE International Symposium on Parallel & Distributed Processing*. IEEE, 2009, pp. 1–8. [Online]. Available: https://doi.org/10.1109/IPDPS.2009.5160935

[56] S. Shi, X. Chu, K. C. Cheung, and S. See, "Understanding top-k sparsification in distributed deep learning," *arXiv preprint arXiv:1911.08772*, 2019. [Online]. Available: https://doi.org/10.48550/arXiv.1911.08772