

pMEMCPY: a simple, lightweight, and portable I/O library for storing data in persistent memory

Luke Logan, Jay Lofstead, Scott Levy, Patrick Widener, Xian-He Sun, Anthony Kougkas

Illinois Institute of Technology and Sandia National Labs

llogan@hawk.iit.edu, {glofst, sllevy, pwidene}@sandia.gov, {sun, akougkas}@iit.edu

Abstract—Persistent memory (PMEM) devices can achieve comparable performance to DRAM while providing significantly more capacity. This has made the technology compelling as an expansion to main memory. Rethinking PMEM as storage devices can offer a high performance buffering layer for HPC applications to temporarily, but safely store data. However, modern parallel I/O libraries, such as HDF5 and pNetCDF, are complicated and introduce significant software and metadata overheads when persisting data to these storage devices, wasting much of their potential. In this work, we explore the potential of PMEM as storage through pMEMCPY: a simple, lightweight, and portable I/O library for storing data in persistent memory. We demonstrate that our approach is up to 2x faster than other popular parallel I/O libraries under real workloads.

Index Terms—Persistent Memory, Parallel I/O, Libraries, Memory Management, Storage Management, Memory Mapped I/O, Data Layout

I. INTRODUCTION

Scientific applications generate massive amounts of data; however, storage performance lags behind CPU performance resulting in applications being bottlenecked by I/O both with other nodes as well as with storage. One approach to alleviate this problem is to expand the memory capacity of the nodes, enabling more local processing before requiring communication. PMEM (e.g., phase change memory and Intel DC Persistent Memory) offers an excellent solution that is cheaper than DRAM, but offers reasonably similar performance characteristics. This technology has driven considerable work into using DRAM as a working cache for an expanded PMEM main memory [1]. As compelling as this case is, it only addresses the inter-node portion of the bottleneck. For applications where communication with storage is a more serious concern, using that same PMEM technology as fast storage (instead of slower memory) offers a flexible resource that can address multiple kinds of workloads. For example, various works investigate the use of storage hierarchies in order to combat the I/O bottleneck [2]–[4]. In these works, storage such as PMEM, NVMe, SSD, and HDD are arranged in a hierarchy based on performance and capacity characteristics. Data is initially buffered in faster storage tiers and then asynchronously flushed to slower mass storage, which helps avoid costly data stalls. While there has been considerable work examining the use of node hosted storage technology with more favorable performance characteristics than hard drives, the interfaces for PMEM offer

another potential performance gain, but only if the software uses the devices with these more efficient interfaces.

Due to the DRAM-like performance of PMEM, software overheads are no longer negligible on the I/O path. For this reason, researchers have started rethinking the design of node-local storage stacks [5]–[9], which had previously been designed for slow storage technologies, such as hard drives. EXT4/XFS DAX [5] allows applications to directly store data in PMEM without first copying to DRAM using memory mapped I/O. SplitFS [6] aims to improve the performance of DAX by splitting metadata and storage operations between kernel space and user space respectively, allowing the majority of I/O operations to avoid the kernel entirely. NOVA [7] is a log-structured filesystem that aims to exploit the parallelism and random access properties of PMEM by storing logs per-inode as opposed to a global log. These works avoid many of the overheads introduced by the Linux kernel, such as context switching, splitting/merging, lock contention, and request reordering. However, improving node-local storage stacks is not enough. HPC applications typically use parallel I/O (PIO) libraries on top of node-local storage stacks to persist data. Fundamental changes in the design of PIO libraries must be made to gain the full benefits of PMEM for I/O.

Various PIO libraries exist, such as ADIOS [10], [11], HDF5 [12], and pNetCDF [13]. However, these libraries introduce significant programming burden, software overhead, and complex configuration spaces. In order to maximize the performance of these libraries and reduce the user’s burden, researchers have investigated the use of auto-tuning to identify optimal parameters specific to the characteristics of applications and parallel filesystems [14]–[17], with approaches such as genetic algorithms and Bayesian optimization. However, at a fundamental level, existing PIO libraries do not interact with PMEM efficiently, regardless of how well they are tuned. For example, all existing work depends on the use of MPI-IO and POSIX, which causes unnecessary networking communication and data copies that degrade the performance of I/O to PMEM [18]. Furthermore, PIO libraries tend to have complicated APIs, requiring many lines of code to store simple data structures, such as arrays. A simple *memcpy* interface is more desirable. PIO libraries should be designed with awareness of the underlying device characteristics in mind in addition to being more user-friendly.

In this work, we present pMEMCPY: a simple, lightweight,

and portable I/O library for storing data in persistent memory. Using the Persistent Memory Development Kit (PMDK) [19], applications have direct access to PMEM while maintaining consistency guarantees. Users can store data structures with a simple key-value store interface that adds the minimal metadata necessary to deserialize the data structures in addition to avoiding costly network communications and data copies that other PIO libraries introduce.

Our contribution offers an optimized approach for parallel I/O library design that can store application data structures in node-local PMEM directly with minimal overhead using a simple key-value store interface similar to *memcpy*. Through this style of I/O library, users can achieve the best possible PMEM performance for their storage operations and enjoy an API much closer to *memcpy*.

The rest of this paper is organized as follows. First in Section II is a deeper discussion of background and related work. Next in Section III we detail the reasoning and design decisions for our demonstration. Section IV presents a collection of evaluations comparing this approach against alternatives. Finally, in Section V we summarize the work.

II. BACKGROUND & RELATED WORK

There are various existing parallel I/O libraries, including HDF5, ADIOS, and pNetCDF. Furthermore, there are various libraries and APIs that exist to efficiently interact with PMEM. However, there has been no published approach, to our knowledge, that demonstrates how to optimize the I/O library for PMEM interfaces and simplify the API to a most basic *memcpy*-like approach.

A. Parallel I/O (PIO) Libraries

HDF5 [12] is a popular PIO library, and is used as the foundations for other popular PIO libraries, such as NetCDF4 [20]. HDF5 exposes a hierarchical namespace to users, where H5Groups are analogous to directories. HDF5 can store primitive types (ints, floats, doubles, etc.), compound data types (structures), and arrays (H5Datasets) of those types. Subsets of datasets can be taken using the Hyperslab APIs. HDF5 can store datasets using various data layout policies: contiguous, chunked, and compact. The contiguous layout stores arrays as a 1-D sequence of data, and is the default layout for HDF5. The chunked mode divides the array into fixed-size sub-arrays (i.e., chunks) where the dimensions of the sub-arrays are user-defined. In chunked mode, HDF5 also allows for the definition of filters, which are operations to perform on individual chunks, such as compression [21], [22]. Lastly, if the dataset is less than 64KB, the compact mode stores the dataset in its corresponding metadata entry. In order to persist data to storage, HDF5 allows multiple approaches: MPI Independent I/O, MPI Collective I/O, and POSIX I/O. In each of these cases, the final output of HDF5 is a single binary file. Furthermore, HDF5 introduced a multi-tiered buffer management system, Hermes [4], that allows users to manage the complexity of heterogeneous, multi-tiered storage environments without changing application code. While HDF5

is a feature-rich library that has specific functionality for buffering and prefetching, it has many limitations. The Neuroscience community, for example, has noted multiple flaws in the user-friendliness of this library [23]. HDF5 stores data in a single binary file, where metadata is not human readable. This also makes version control systems less efficient. Furthermore, HDF5 compound types do not support the nesting of compound types or dynamically sized arrays. Furthermore, MPI-IO relies on the underlying filesystem (for Linux, read/write) APIs in order to store data. However, read/write perform data copies which introduces unnecessary overhead [18].

An alternative to HDF5 and NetCDF4 is the pNetCDF [13] library. This was developed around the same time as NetCDF4 as an effective demonstration on how to maintain the NetCDF3 compatibility as much as possible while extending for 64-bit support. While the two libraries “compete”, the reality is that they co-exist peacefully and are widely supported as a pair rather than individually. For example, the NCAR PIO library [24] offers a single API that can switch to use either NetCDF4 or pNetCDF underneath. As with HDF5, pNetCDF is designed with MPI-IO as the primary IO interface for parallel IO and optimized for slow storage devices through additional work to prepare data to more efficiently be moved into storage. However, the performance gains of PMEM shifts the bottleneck of the storage device that afforded such optimizations without noticeably hurting performance to the I/O library itself. NVMe devices have had a similar effect [25], but PMEM offers additional performance exasperating the performance overhead the software layer imposes.

ADIOS is an alternative PIO library to HDF5, NetCDF and pNetCDF. ADIOS aims to encompass various I/O transport mechanisms (e.g., MPI-IO, POSIX, HDF5, and netCDF) under a simplified interface that is easily configurable and requires little change to application code to change which implementation is used. The original design of ADIOS was based on trying to reduce the code complexity of HDF5 and acknowledge that some of the performance optimizations employed by HDF5 and other libraries that use similar PIO techniques ultimately do not scale for writing or reading as well as hoped [26]. One approach ADIOS uses to address the performance gap is to use its own BP format whenever possible. BP offers delayed consistency, lightweight data characterization, and data resilience. Unlike HDF5, ADIOS stores data in the same format as it was produced on a process-by-process basis rather than constructing a global linearization of complex datastructures. For example, a 3D domain decomposition is stored as a single item in HDF5 with all three dimensions across all processes being linearized through a data rearrangement phase prior to hitting storage. This has the advantage of eliminating any potential artifacts from unusual process decompositions. ADIOS has each process write the data it owns with no coordination with other processes. This eliminates the data rearrangement phase, which can improving performance greatly. In particular, large 3D domain decompositions see radical performance improvements for both writing and reading [26]. ADIOS also supports transparent and custom

operators, similar to HDF5. Recently, ADIOS2 [10] was released, which provides a C++ interface that is more simplistic and extensible than that of the original ADIOS. Their recent revision includes a key-value store API for storing data. However, ADIOS2 suffers from the same drawbacks as the original when it comes to PMEM as it is storage device agnostic.

A more recent effort, Proactive Data Containers [27] offers a similar key-value store approach for data management. However, it is also designed to assume storage devices, such as SSDs with an assumption that non-volatile memory devices will have compatible interfaces to get the full performance benefits.

One system recognizing the need for a different interface to non-volatile memory is DStore [28]. However, DStore is intended as a way to store a log for an in-DRAM key-value store. Unlike other attempts to optimize key-value stores with PMEM, such as MongoDB-PMEM [29] and PMEM-RocksDB [30], DStore uses PMEM to store the logs rather than as the main store offering greater performance while still offering predictable consistency.

In all cases, effectively using PMEM using efficient interfaces is a relatively new endeavour that popular HPC I/O libraries have yet to embrace. While some progress has been made in the scale-out space, the recent DStore paper demonstrates that a simple “switch to the PMDK interface” may not be the most efficient nor optimal approach for achieving both performance and price/performance.

B. Accessing PMEM

PMEM can be exposed like any other storage device. Application developers can use traditional filesystem APIs such as POSIX, stdlib, and iostream in order to store data in PMEM. However, these interfaces introduce significant software overheads. For example, these interfaces will cause unnecessary data copies and memory allocations to occur. To avoid this, applications can access PMEM directly using memory mapped I/O (MMIO) and DAX. However, managing memory-mapped regions requires application developers to provide their own memory allocation functions and concurrency control mechanisms, which can cause data consistency and reliability concerns.

The Persistent Memory Development Kit (PMDK) [19] is a collection of libraries and tools for managing PMEM devices. It provides low-level primitives for interacting with PMEM and a transactional object store that utilizes memory mapping in order to interface with PMEM devices. What this really means is that PMEM is mapped directly in the memory space for a process enabling direct access. Unlike MPI-IO and POSIX I/O, this approach allows applications direct, zero-copy access to PMEM while providing consistency guarantees. PMDK provides optimized memory allocation functions, persistent locks, basic data structures (e.g., thread-safe lists), and transactions. This allows applications to have efficient and safe access to PMEM while reducing the complexity of managing memory-mapped files.




Fig. 1. Basic Machine Architecture

C. New Filesystems

In the Introduction, we covered many of the newer generation storage systems written from the ground up to take advantage of solid state, node local storage. However, these have all been written for NVMe devices, at best, and still assume a more traditional device interface. One major exception to this is DAOS [31]. The original design of DAOS [32] was to offer a new storage architecture, but still assuming non-PMEM storage devices. The current DAOS generations have been reimagined using Intel Optane PMEM devices as a core component. Using these devices, DAOS was able to achieve top marks on the IO500 list at sc19 [33]. More recent conversations with the DAOS team about Optane and DAOS or other storage use recommended at most 1% of the capacity using the PMEM devices as a way to ensure top performance for the most critical operations while keeping costs from spiraling out of control [34]. This makes DAOS a good potential candidate for using PMEM as a storage device, but it does not address the I/O library layer entirely. The plug-ins for HDF5 for speaking directly with DAOS and the DAOS native APIs may offer better support. However, the interfaces are still complex and focused on a container-like structure with POSIX-structures layered on top.

III. DESIGN & IMPLEMENTATION

This work offers pMEMCPY, a simplistic and portable I/O library for managing node-local PMEM. Our design assumes that the compute nodes running the application also contain PMEM. Data structures in memory are stored directly on PMEM without extra metadata, context switching, or data copies beyond what is necessary to reload the data during a different application run or for an analysis job. Our assumed, basic machine architecture is illustrated in Figure 1.

API: pMEMCPY exposes a key-value interface for storing and loading data from PMEM. Users can store primitive types, structured types, and arrays of these types using the templated load/store APIs. The C++ API is shown in Figure 2. In Figure 3, we demonstrate the usage of pMEMCPY for writing a 1-D array of data in parallel. In the example, each process writes 100 doubles to non-overlapping offsets in the array directly to PMEM. *alloc* is used to specify the final dimensions of the array, and *store* is used to persist pieces of the array generated by each process. In Figure 4, we show the equivalent HDF5 code.

```

1. #include <pmemcpy/pmemcpy.hpp>
2. pmemcpy::PMEM pmem;
3. pmem.mmap(std::string filename, int comm);
4. pmem.munmap();
5.
6. pmem.store<T>(std::string id, T &data);
7. pmem.alloc<T>(std::string id,
8.   int ndims, size_t *dims);
9. pmem.store<T>(std::string id, T *data,
10.  int ndims, size_t *offsets, size_t *dimspp);
11.
12. pmem.load<T>(std::string id);
13. pmem.load<T>(std::string id, T &num);
14. pmem.load<T>(std::string id, T *data,
15.  int ndims, size_t *offsets, size_t *dimspp);
16. pmem.load_dims(std::string id,
17.  int *ndims, size_t *dim);

```

Fig. 2. pMEMCPY API

```

1. #include <pmemcpy/pmemcpy.h>
2. int main(int argc, char** argv) {
3.   int rank, nprocs;
4.   MPI_Init(&argc, &argv);
5.   MPI_Comm_rank(MPI_COMM_WORLD, &rank);
6.   MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
7.   pmemcpy::PMEM pmem;
8.   size_t count = 100;
9.   size_t off = 100*rank;
10.  size_t dims = 100*nprocs;
11.  char *path = argv[1];
12.
13.  double data[100] = {0};
14.  pmem.mmap(path, MPI_COMM_WORLD);
15.  pmem.alloc<double>("A", 1, &dims);
16.  pmem.store<double>("A", data, 1, &off, &count);
17.  MPI_Finalize();
18. }

```

Fig. 3. pMEMCPY API Usage Example

HDF5 requires a user to create and free dataspace and dataset objects in addition to subsetting the dataset, and each of these interfaces contain many parameters. The dataspace defines the dimensions of the array, and the dataset represents the array within HDF5. The HDF5 version is 42 lines of code and 253 tokens, whereas our code is 16 lines and 132 tokens, which is a 92% reduction in the number of tokens. (Note, the number of lines calculation discounts lines where there is no code.) Similar to HDF5, NetCDF and pNetCDF requires users to define and allocate the dimensions of the array using special APIs, which adds unnecessary complexity. While ADIOS simplifies this, it still requires the user to store the dimensions of the array separately and then associate those variables with the array. pMEMCPY automatically stores the dimensions of the array and the per-process subarrays in the *store* API by appending “#dims” to the id; dimensions can be queried using *load_dims*. In Figure 5, we show the equivalent ADIOS code, which is 24 lines and 164 tokens. Overall, we see that pMEMCPY provides a more simplified and compact API than other libraries.

Data Transfer and Serialization: Unlike ADIOS, NetCDF, and pNetCDF which depend on POSIX and MPI-IO, pMEMCPY uses memory mapping and independent I/O to

```

1. #include <hdf5.h>
2. int main (int argc, char **argv) {
3.   int nprocs, rank;
4.   MPI_Init(&argc, &argv);
5.   MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
6.   MPI_Comm_rank(MPI_COMM_WORLD, &rank);
7.   hid_t file_id, dset_id;
8.   hid_t filespace, memspace;
9.   hsize_t count = 100;
10.  hsize_t offset = rank*100;
11.  hsize_t dims = nprocs*100;
12.  hid_t plist_id;
13.  herr_t status;
14.  char *path = argv[1];
15.  int data[100];
16.
17.  plist_id = H5Pcreate(H5P_FILE_ACCESS);
18.  H5Pset_fapl_mpio(plist_id,
19.   MPI_COMM_WORLD, MPI_INFO_NULL);
20.  file_id = H5Fcreate(path,
21.   H5F_ACC_TRUNC, H5P_DEFAULT, plist_id);
22.  H5Pclose(plist_id);
23.
24.  filespace = H5Screate_simple(1, &dims, NULL);
25.  dset_id = H5Dcreate(file_id, "dataset",
26.   H5T_NATIVE_INT, filespace, H5P_DEFAULT,
27.   H5P_DEFAULT, H5P_DEFAULT);
28.  H5Sclose(filespace);
29.  memspace = H5Screate_simple(1, &count, NULL);
30.  filespace = H5Dget_space(dset_id);
31.  H5Sselect_hyperslab(filespace,
32.   H5S_SELECT_SET, &offset,
33.   NULL, &count, NULL);
34.
35.  plist_id = H5Pcreate(H5P_DATASET_XFER);
36.  status = H5Dwrite(dset_id, H5T_NATIVE_INT,
37.   memspace, filespace, plist_id, data);
38.
39.  H5Dclose(dset_id);
40.  H5Sclose(filespace);
41.  H5Sclose(memspace);
42.  H5Pclose(plist_id);
43.  H5Fclose(file_id);
44.  MPI_Finalize();
45.  return 0;
46. }

```

Fig. 4. Equivalent HDF5 Example

store data in the node-local PMEM, which avoids unnecessary data copies, network/inter-process communications, and kernel interventions. When storing a data structure in PMEM, pMEMCPY serializes the data using well-known, portable serialization libraries, such as BP4 [10], CapnProto [35], and cereal [36]. By default, the BP4 serialization (same as ADIOS) is used; however, other serialization tools can be added, and serialization can be completely disabled. Unlike similar work which serializes data structures into an in-memory buffer and then copies to PMEM, pMEMCPY can serialize the data directly into PMEM without first placing it in DRAM, avoiding a significant data copying cost. Furthermore, we allow users to configure whether or not the MAP_SYNC flag is enabled when storing serialized data structures in a region of PMEM. The MAP_SYNC flag guarantees that, after a crash, a block that has been mapped into memory with write permissions will still be at the same offset within

```

1. #include <adios.h>
2. int main(int argc, char **argv) {
3.     int rank, nprocs;
4.     MPI_Init(&argc, &argv);
5.     MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
6.     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
7.     char *path = argv[1];
8.     char *config = argv[2];
9.     double data[100];
10.    int64_t adios_handle;
11.    size_t count = 100;
12.    size_t offset = 100*rank;
13.    size_t dimsf = 100*nprocs;
14.
15.    adios_init(config, MPI_COMM_WORLD);
16.    adios_open (&adios_handle, "dataset",
17.        path, "w", MPI_COMM_WORLD);
18.    adios_write (adios_handle, "count", &count);
19.    adios_write (adios_handle, "dimst", &dimst);
20.    adios_write (adios_handle, "offset", &offset);
21.    adios_write (adios_handle, "A", data);
22.    adios_close (adios_handle);
23.    adios_finalize (rank);
24.    MPI_Finalize ();
25.    return 0;
26. }

```

Fig. 5. Equivalent ADIOS Example. Note, there is a separate ADIOS config file that defines “A” in terms of count, off, and dimst.

the file [37]. While this improves crash consistency, this can introduce significant latency penalties that severely degrade performance, as shown in our evaluations. After serialization, a burst buffer, such as DataWarp [38], will then be triggered to asynchronously flush the buffered data to mass storage. The data will be stored in the same format as it was produced, similar to ADIOS, which avoids the network and inter-process communication required to restructure the data.

Data Layout: By default, pMEMCPY stores all application data in a single file similar to ADIOS, NetCDF, and pNetCDF. However, pMEMCPY uses the PMDK [19] to manage PMEM, which provides direct access to PMEM in addition to data consistency guarantees, concurrency control, and memory allocation policies. Metadata is stored in a flat namespace using a hashtable with chaining. This utilizes the high parallelism and random access characteristics of PMEM. Alternatively, unlike ADIOS, NetCDF, and pNetCDF, pMEMCPY can layout data hierarchically using the PMEM’s filesystem. In this approach, instead of writing to a single file, pMEMCPY stores the data structures in a directory and creates a file for each variable. Whenever a “/” is used in the id of the variable, a directory is created if it didn’t already exist.

IV. EVALUATIONS

Testbed: All tests were conducted in Chameleon Cloud using a Compute Skylake node. Compute Skylake nodes come with 192GB of RAM and 2x Intel(R) Xeon(R) Gold 6126 CPU @ 2.60GHz, for a total of 24 cores/48 threads. The OS used was Ubuntu 20.04 with kernel 5.4.0-70-generic. We used openmpi 3.1.6.

Emulating PMEM: Since we do not have access to PMEM, we emulate it using the approach presented in the Strata PMEM filesystem paper [8]. We utilize Linux’s PMEM

emulator to treat 80GB of DRAM as PMEM and format the resulting PMEM device using EXT4 with DAX enabled. We assume that PMEM has a read latency of 300ns, write latency of 125ns, read bandwidth of 30GB/s, and write bandwidth of 8GB/s [39]. We benchmarked DRAM bandwidth and latency using Intel’s MLC [40] and use nanosecond-accurate monotonic timers to add the additional latency and bandwidth constraints.

A. Real-App Evaluation

In this test, we demonstrate the performance impact of pMEMCPY over other popular PIO libraries using real workloads. In this evaluation, we use two workloads that were obtained through the help of scientists [26]. The first workload is a write-only 3-D domain decomposition problem where each process writes a rectangular region of data to storage. The second workload is a read-only workload that reads the regions from storage. For both tests, we use between 8 and 48 processes. This model represents a large memory regular stencil code common in compute models today. One example is the S3D combustion code [41] that was the inspiration for this configuration. This model has been previously used [26] to demonstrate potential I/O performance. In the write-only case, we generate 10 3-D rectangles. For each test, a total of 40GB of data is generated and the 40GB is divided equally among the processes. Each element in the rectangle is a double precision floating point value (8 bytes). The read workload is completely symmetrical to the write workload, where each process reads the same data that had been written. We measure the wall-clock time from the point at which the file is opened/mmapped to when the it is closed. We perform the I/O using ADIOS, NetCDF-4, pNetCDF, and pMEMCPY and compare the runtime between the different approaches. For NetCDF-4, we make sure to call `nc_def_var_fill()` with `NC_NOFILL` in order to prevent it from initializing variables with a default value, which causes significant overhead for write workloads. For pMEMCPY, we use BP4 serialization with the PMDK hashtable layout. We run each experiment 3 times and take the average of the runs. We only used a single node for this evaluation since pMEMCPY requires no network communication in order to store the data in local PMEM, and should therefore scale linearly with an increase in the number of nodes. The per-node data can be written to permanent storage asynchronously to mass storage using a burst buffer technology, as mentioned in Section III.

The results of the experiment are shown in Figures 6 and 7. From these figures, we see the effects of concurrency due to the CPU and PMEM wear off after 24 cores in the write case and for most of the reads, with the exception of PMCPY-B and NetCDF4. This makes sense considering the node has 24 physical CPU cores in total. For NetCDF, the performance differences were largely due to differences in the dimensions of the cube being read for the different process counts. For PMCPY-B, this was because the metadata updates were parallelized, which caused fewer stalls. Overall, we see that pMEMCPY outperforms ADIOS, NetCDF, and pNetCDF

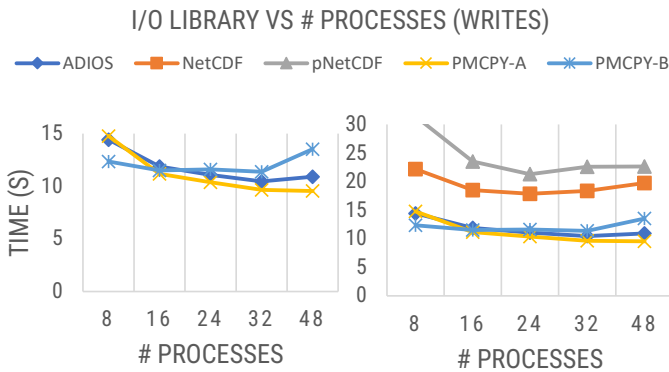


Fig. 6. Performance of writing a 40GB 3-D domain to PMEM for a varying number of processes. PMCPY-A has MAP_SYNC disabled, whereas PMCPY-B has it enabled. Each process writes an equal amount of data. pMEMCPY is 2.5x faster than pNetCDF and NetCDF by avoiding network communications and data copying costs. At 24 cores, pMEMCPY is faster than ADIOS by 15% when MAP_SYNC is disabled, and slightly slower when MAP_SYNC is enabled. Note, the left figure is a zoomed in version of the right figure.

in both workloads when MAP_SYNC is disabled. This is because pMEMCPY avoids unnecessary communications and data copies that other PIO libraries introduce. In the case of writes, all other PIO libraries first generate the cube in DRAM, serialize the cube into another DRAM buffer, and then copy the serialized cube to the PMEM whereas pMEMCPY generates the cube in DRAM and then serializes the cube directly into the PMEM, avoiding an entire copy of the cube. From these figures, we see ADIOS performs far better than NetCDF and pNetCDF in both read and write performance. This is because, similar to pMEMCPY, ADIOS stores data in the same format as it was produced, which avoids costly network communications and data copies during the write phase. Furthermore, since the workload is symmetrical, ADIOS does not need to realign any data, which mitigates data shuffling costs in the read phase. However, pNetCDF and NetCDF store data contiguously, which requires data to be shuffled during both reads and writes, incurring significant overhead. While ADIOS performs much better than pNetCDF and NetCDF, it still introduces data copying overheads that pMEMCPY avoids, causing its performance to be suboptimal. For example, in the case of reads, ADIOS requires the serialized data to be copied from PMEM into DRAM and then deserialized into another DRAM buffer. pMEMCPY deserializes the data directly from PMEM, avoiding the initial copy from PMEM to DRAM. Within pMEMCPY, we see that the choice of flags has a significant impact on performance. When MAP_SYNC is enabled, the performance benefit of serializing/deserializing directly from PMEM is completely lost, and can even cause performance to be worse than simply using POSIX read()/write(). Overall, we see that pMEMCPY can perform at least 15% better for writes and 2x better for reads depending on the level of security the user requires.

B. Discussion

While standard I/O libraries offer a familiar interface, that can come at a cost. ADIOS, with the design break from the previous generation demonstrates better performance, but is




Fig. 7. Performance of reading a 40GB 3-D domain from PMEM for a varying number of processes. PMCPY-A has MAP_SYNC disabled, whereas PMCPY-B has it enabled. Each process reads an equal amount of data. pMEMCPY is 5x faster than pNetCDF and NetCDF by avoiding network communications and data copying costs. pMEMCPY is 2x faster than ADIOS when MAP_SYNC is disabled. When enabled, pMEMCPY performs no better than ADIOS. Note, the left figure is a zoomed in version of the right figure.

still not optimal by a margin of 15% - 100%. Only by using an approach such as the one we demonstrate in pMEMCPY can the full potential of PMEM as a storage device be achieved.

V. CONCLUSIONS

Persistent memory (PMEM) is an extraordinarily fast persistent storage device typically thought of as an extension of DRAM main memory. However, using PMEM for storage requires revisiting the design of parallel I/O (PIO) libraries. With PMEM being integrated into compute nodes, PIO libraries should take full advantage of the characteristics of these devices. However, popular libraries, such as HDF5, ADIOS, and pNetCDF, introduce significant overheads when applications store and load data. Furthermore, they introduce complex interfaces and parameters that add unnecessary burden on programmers. In this paper, we introduced pMEMCPY: a simple, lightweight, and portable I/O library for storing data in persistent memory. We compared our design with ADIOS, NetCDF-4, and pNetCDF, and found that write speeds improved at least 15% and reads improved up to 2x.

ACKNOWLEDGEMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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.

REFERENCES

- [1] J. Ren, J. Luo, I. Peng, K. Wu, and D. Li, "Optimizing large-scale plasma simulations on persistent memory-based heterogeneous memory with effective data placement across memory hierarchy," in *Proceedings of the ACM International Conference on Supercomputing*, 2021, pp. 203–214.
- [2] K. Tang, P. Huang, X. He, T. Lu, S. S. Vazhkudai, and D. Tiwari, "Toward managing hpc burst buffers effectively: Draining strategy to regulate bursty i/o behavior," in *2017 IEEE 25th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE, 2017, pp. 87–98.

- [3] S. Byna, Q. Koziol, V. Vishwanath, J. Soumagne, H. Tang, J. Mu, B. Dong, R. A. Warren, F. Tessier, T. Wang *et al.*, "Proactive data containers (pdc): An object-centric data store for large-scale computing systems," in *AGU Fall Meeting Abstracts*, vol. 2018, 2018, pp. IN34B-09.
- [4] A. Kougkas, H. Devarajan, and X.-H. Sun, "Hermes: a heterogeneous-aware multi-tiered distributed i/o buffering system," in *Proceedings of the 27th International Symposium on High-Performance Parallel and Distributed Computing*, 2018, pp. 219-230.
- [5] "Direct access for files," 2014. [Online]. Available: <https://www.kernel.org/doc/Documentation/filesystems/dax.txt>
- [6] R. Kadekodi, S. K. Lee, S. Kashyap, T. Kim, A. Kolli, and V. Chidambaram, "Splitfs: Reducing software overhead in file systems for persistent memory," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019, pp. 494-508.
- [7] J. Xu and S. Swanson, "{NOVA}: A log-structured file system for hybrid volatile/non-volatile main memories," in *14th {USENIX} Conference on File and Storage Technologies ({FAST} 16)*, 2016, pp. 323-338.
- [8] Y. Kwon, H. Fingler, T. Hunt, S. Peter, E. Witchel, and T. Anderson, "Strata: A cross media file system," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 460-477.
- [9] S. Zheng, M. Hoseinzadeh, and S. Swanson, "Ziggurat: a tiered file system for non-volatile main memories and disks," in *17th {USENIX} Conference on File and Storage Technologies ({FAST} 19)*, 2019, pp. 207-219.
- [10] W. F. Godoy, N. Podhorszki, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Choi, K. Germaschewski, K. Huck *et al.*, "Adios 2: The adaptable input output system. a framework for high-performance data management," *SoftwareX*, vol. 12, p. 100561, 2020.
- [11] J. F. Lofstead, S. Klasky, K. Schwan, N. Podhorszki, and C. Jin, "Flexible io and integration for scientific codes through the adaptable io system (adios)," in *Proceedings of the 6th international workshop on Challenges of large applications in distributed environments*, 2008, pp. 15-24.
- [12] Q. Koziol, D. Robinson *et al.*, "Hdf5," Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), Tech. Rep., 2018.
- [13] J. Li, W.-k. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, and M. Zingale, "Parallel netcdf: A high-performance scientific i/o interface," in *SC'03: Proceedings of the 2003 ACM/IEEE Conference on Supercomputing*. IEEE, 2003, pp. 39-39.
- [14] B. Behzad, H. V. T. Luu, J. Huchette, S. Byna, R. Aydt, Q. Koziol, M. Snir *et al.*, "Taming parallel i/o complexity with auto-tuning," in *SC'13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2013, pp. 1-12.
- [15] B. Behzad, S. Byna, and M. Snir, "Optimizing i/o performance of hpc applications with autotuning," *ACM Transactions on Parallel Computing (TOPC)*, vol. 5, no. 4, pp. 1-27, 2019.
- [16] Z. Cao, "A practical, real-time auto-tuning framework for storage systems," Ph.D. dissertation, State University of New York at Stony Brook, 2019.
- [17] Z. Cao, V. Tarasov, S. Tiwari, and E. Zadok, "Towards better understanding of black-box auto-tuning: A comparative analysis for storage systems," in *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*, 2018, pp. 893-907.
- [18] J. Kim, Y. J. Soh, J. Izraelvitz, J. Zhao, and S. Swanson, "Subzero: zero-copy io for persistent main memory file systems," in *Proceedings of the 11th ACM SIGOPS Asia-Pacific Workshop on Systems*, 2020, pp. 1-8.
- [19] S. Scargall, "Pmdk internals: Important algorithms and data structures," in *Programming Persistent Memory*. Springer, 2020, pp. 313-331.
- [20] R. Rew and G. Davis, "Netcdf: an interface for scientific data access," *IEEE computer graphics and applications*, vol. 10, no. 4, pp. 76-82, 1990.
- [21] H. Devarajan, A. Kougkas, L. Logan, and X.-H. Sun, "Hcompress: Hierarchical data compression for multi-tiered storage environments," in *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2020, pp. 557-566.
- [22] H. Devarajan, A. Kougkas, and X.-H. Sun, "An intelligent, adaptive, and flexible data compression framework," in *2019 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. IEEE, 2019, pp. 82-91.
- [23] S.-A. Dragly, M. Hobbi Mobarhan, M. E. Lepperød, S. Tennøe, M. Fyhn, T. Hafting, and A. Malthe-Sørenssen, "Experimental directory structure (exdir): An alternative to hdf5 without introducing a new file format," *Frontiers in neuroinformatics*, vol. 12, p. 16, 2018.
- [24] J. M. Dennis, J. Edwards, R. Loy, R. Jacob, A. A. Mirin, A. P. Craig, and M. Vertenstein, "An application-level parallel i/o library for earth system models," *The International Journal of High Performance Computing Applications*, vol. 26, no. 1, pp. 43-53, 2012.
- [25] A. Aghayev, S. Weil, M. Kuchnik, M. Nelson, G. R. Ganger, and G. Amvrosiadis, "File systems unfit as distributed storage backends: lessons from 10 years of ceph evolution," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019, pp. 353-369.
- [26] J. Lofstead, M. Polte, G. Gibson, S. Klasky, K. Schwan, R. Oldfield, M. Wolf, and Q. Liu, "Six degrees of scientific data: Reading patterns for extreme scale science io," in *Proceedings of the 20th international symposium on High performance distributed computing*, 2011, pp. 49-60.
- [27] J. Soumagne, R. Warren, J. Mu, V. Vishwanath, F. Tessier, S. Byna, Q. Koziol, H. Tang, T. Wang, B. Dong, and J. Liu, "Final technical report - proactive data containers for scientific storage." [Online]. Available: <https://www.osti.gov/biblio/1577855>
- [28] S. Gugnani and X. Lu, "Dstore: A fast, tailless, and quiescent-free object store for pmem," in *Proceedings of the 30th International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '21. New York, NY, USA: Association for Computing Machinery, 2020, p. 31-43. [Online]. Available: <https://doi.org/10.1145/3431379.3460649>
- [29] Intel, "Mongodb persistent memory storage engine," Github, July 2021. [Online]. Available: <https://github.com/pmem/pmse>
- [30] J. Xu, J. Kim, A. Memaripour, and S. Swanson, "Finding and fixing performance pathologies in persistent memory software stacks," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 427-439.
- [31] Intel, "Daos: Revolutionizing high-performance storage with intel optane technology." [Online]. Available: <https://www.intel.com/content/dam/www/public/us/en/documents/solution-briefs/high-performance-storage-brief.pdf>
- [32] J. Lofstead, I. Jimenez, C. Maltzahn, Q. Koziol, J. Bent, and E. Barton, "Daos and friends: a proposal for an exascale storage system," in *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2016, pp. 585-596.
- [33] io500, "10 node challenge, io500-sc19," November. [Online]. Available: https://www.vi4io.org/io500/list/19-11/10node?fields=information_system,information_institution,information_storage_vendor,information_filesystem_type,information_client_nodes,information_client_total_procs,io500_score,io500_bw,io500_md,information_data,information_list_id&equation=&sort_asc=false&sort_by=io500_score&radarmax=6&query=
- [34] J. Lofstead, "Pmem for storage conversation in io500 general slack channel," Slack, February 2021. [Online]. Available: <https://io500workspace.slack.com/archives/C01BMTNT56K/p1612291357026500>
- [35] K. Varda, "Cap'n proto cerealization protocol," 2013. [Online]. Available: <https://capnproto.org/>
- [36] R. Voorhies, "cereal - a c++11 library for serialization," 2014. [Online]. Available: <https://usclab.github.io/cereal/>
- [37] J. Corbet, "Two more approaches to persistent-memory writes," 2017.
- [38] D. Henseler, B. Landsteiner, D. Petesch, C. Wright, and N. J. Wright, "Architecture and design of cray datawarp," *Cray User Group CUG*, 2016.
- [39] A. van Renen, L. Vogel, V. Leis, T. Neumann, and A. Kemper, "Persistent memory i/o primitives," in *Proceedings of the 15th International Workshop on Data Management on New Hardware*, 2019, pp. 1-7.
- [40] T. W. P. L. B. F. S. S. Vish Viswanathan, Karthik Kumar, "Memory latency checker," 2015. [Online]. Available: <https://software.intel.com/content/www/us/en/develop/articles/intel-memory-latency-checker.html>
- [41] D. Lignell, C. Yoo, J. Chen, R. Sankaran, and M. Fahey, "S3d: Petascale combustion science, performance, and optimization," in *Proceedings of the Cray Scaling Workshop, Oak Ridge National Laboratory, TN*, 2007.