Hermes: A Heterogeneous-Aware Multi-Tiered Distributed I/O Buffering System

Anthony Kougkas, Hariharan Devarajan, Xian-He Sun
Illinois Institute of Technology, Department of Computer Science, Chicago, IL
akougkas@hawk.iit.edu, hdevarahan@hawk.iit.edu, sun@iit.edu

ABSTRACT

Modern High-Performance Computing (HPC) systems are adding extra layers to the memory and storage hierarchy, named deep memory and storage hierarchy (DMSH), to increase I/O performance. New hardware technologies, such as NVMe and SSD, have been introduced in burst buffer installations to reduce the pressure for external storage and boost the burstiness of modern I/O systems. The DMSH has demonstrated its strength and potential in practice. However, each layer of DMSH is an independent heterogeneous system and data movement among more layers is significantly more complex even without considering heterogeneity. How to efficiently utilize the DMSH is a subject of research facing the HPC community. In this paper, we present the design and implementation of *Hermes*: a new, heterogeneous-aware, multi-tiered, dynamic, and distributed I/O buffering system. *Hermes* enables, manages, supervises, and, in some sense, extends I/O buffering to fully integrate into the DMSH. We introduce three novel data placement policies to efficiently utilize all layers and we present three novel techniques to perform memory, metadata, and communication management in hierarchical buffering systems. Our evaluation shows that, in addition to automatic data movement through the hierarchy, *Hermes* can significantly accelerate I/O and outperforms by more than 2x state-of-the-art buffering platforms.

CCS CONCEPTS

• **Information systems** → **Distributed storage**; *Record and buffer management*; *Main memory engines*; *Storage class memory*; *Cloud based storage*; *Hierarchical storage management*; • **Hardware** → *External storage*; *Emerging architectures*; *Memory and dense storage*;

KEYWORDS

I/O buffering, Heterogeneous buffering, Layered buffering, Deep memory hierarchy, Burst buffers

ACM Reference format:

Anthony Kougkas, Hariharan Devarajan, Xian-He Sun. 2018. Hermes: A Heterogeneous-Aware Multi-Tiered Distributed I/O Buffering System. In *Proceedings of HPDC '18: International Symposium on High-Performance Parallel and Distributed Computing*, Tempe, AZ, USA, June 11–15, 2018 (HPDC '18), 12 pages.
DOI: 10.1145/3208040.3208059

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HPDC '18, Tempe, AZ, USA

© 2018 ACM. 978-1-4503-5785-2/18/06...\$15.00

DOI: 10.1145/3208040.3208059

1 INTRODUCTION

Data-driven science is a reality and in fact, is now driving scientific discovery [28]. An International Data Corp. (IDC) report [44] predicts that by 2025, the global data volume will grow to 163 zettabytes, ten times the 16.1ZB of data generated in 2016. The evolution of modern storage technologies is driven by the increasing ability of powerful High-Performance Computing (HPC) systems to run data-intensive problems at larger scale and resolution. In addition, larger scientific instruments and sensor networks collect extreme amounts of data and push for more capable storage systems [23]. Modern I/O systems have been developed and highly optimized through the years. Popular interfaces and standards such as POSIX I/O, MPI-IO [51], and HDF5 [22] expose data to the applications and allow users to interact with the underlying file system through extensive APIs. In a large scale environment, the underlying file system is usually a parallel file system (PFS) with Lustre [41], GPFS [47], PVFS2 [45] being some popular examples. However, as we move towards the exascale era, most of these storage systems face significant challenges in performance, scalability, complexity, and limited metadata services [7, 19], creating the so called *I/O bottleneck* which will lead to less scientific productivity [43, 48].

To reduce the I/O performance gap, modern storage subsystems are going through extensive changes, by adding additional levels of memory and storage in a hierarchy [5]. Newly emerging hardware technologies such as High-Bandwidth Memory (HBM), Non-Volatile RAM (NVRAM), Solid-State Drives (SSD), and dedicated buffering nodes (e.g., burst buffers) have been introduced to alleviate the performance gap between main memory and the remote disk-based PFS. Modern supercomputer designs employ such hardware technologies in a heterogeneous layered memory and storage hierarchy, we call *Deep Memory and Storage Hierarchy (DMSH)* [12, 26]. For example, Cori system at the National Energy Research Scientific Computing Center (NERSC) [38], uses CRAY's Datawarp technology [16]. Los Alamos National Laboratory Trinity supercomputer [34] uses burst buffers with a 3.7 PB capacity and 3.3 TB/s bandwidth. Summit in Oak Ridge National Lab is also projected to employ fast local NVMe storage for buffering [54].

As multiple layers of storage are added into HPC systems, the complexity of data movement among the layers increases significantly, making it harder to take advantage of the high-speed and low-latency storage systems [10]. Additionally, each layer of DMSH is an independent system that requires expertise to manage, and the lack of automated data movement between tiers is a significant burden currently left to the users [32]. Furthermore, popular I/O middleware, such as HDF5, PnetCDF [31], and ADIOS [33], are configured to operating with the traditional memory-to-disk I/O endpoints. This middleware provides great value by isolating users

from the complex effort to extract peak performance from the underlying storage system, but it will need to be updated to handle the transition to a multi-tiered I/O configuration [32]. There is a need to seamlessly and transparently support access to DMSH.

In this paper, we present the design and implementation of Hermes: a new, heterogeneous-aware, multi-tiered, dynamic, and distributed I/O buffering system. Hermes enables, manages, and supervises I/O buffering into DMSH and offers: a) *vertical and horizontal distributed buffering in DMSH* (i.e., access data to/from different levels locally and across remote nodes), b) *selective layered data placement* (i.e., buffer data partially or entirely in various levels of the hierarchy), c) *dynamic buffering via system profiling* (i.e., change the buffering schema dynamically by monitoring the system status such as capacity of buffers, messaging traffic, etc.). Hermes accelerates applications' I/O access by transparently buffering data in DMSH. Data can be moved through the hierarchy effortlessly and therefore, applications have a capable, scalable, and reliable middleware software to navigate the I/O challenges towards the exascale era. Lastly, by supporting both POSIX and HDF5 interfaces, Hermes offers ease-of-use to a wide-range of scientific applications.

The contributions of this work include:

- presenting the design and implementation of Hermes: a new, heterogeneous-aware, multi-tiered, dynamic, and distributed I/O buffering system (Section 3.1).
- introducing three novel data placement policies to efficiently utilize all layers of the new memory and storage hierarchy (Section 3.2.2).
- presenting the design and implementation of three novel techniques to perform *memory*, *metadata*, and *communication* management in hierarchical buffering systems (Section 3.3.2).
- evaluating Hermes' design and technical innovations showing that our solution can grant better performance compared to the state-of-the-art buffering platforms (Section 4).

2 BACKGROUND

2.1 Modern Application I/O Characteristics

Modern HPC applications are required to process large volume, velocity and variety of data, leading to an explosion of data requirements and complexity [15]. Many applications spend significant time of the overall execution in performing I/O making storage a vital component in performance [56]. Furthermore, scientific applications often demonstrate bursty I/O behavior [27, 37]. Typically, in HPC workloads, short, intensive, phases of I/O activities, such as checkpointing and restart, periodically occur between longer computation phases [1, 8]. The intense and periodic nature of I/O operations stresses the underlying parallel file system and thus, stalls the application. To appreciate how important and challenging the I/O performance of a system is, one needs to deeply understand the I/O behavior of modern scientific applications. More and more scientific applications generate very large datasets, and the development of several disciplines greatly relies on the analysis of massive data. We highlight some scientific domains that are increasingly relying on High-Performance Data Analytics (HPDA), the new generation of data-intensive applications, which involve sufficient data volumes and algorithmic complexity to require HPC

resources: *Computational Biology*: The National Center for Biotechnology Innovation maintains the GenBank database of nucleotide sequences, which doubles in size every 10 months. The database contains over 250 billion nucleotide bases from more than 150,000 distinct organisms. *Astronomy*: Square Kilometre Array project run by an international consortium operates the largest radio telescope in the world which produces staggering data as presented in the keynote speech during the 2017 SC conference. As highlighted, the incoming images are of 10 PBs and the produced 3D image is 1 PB each. *High-Energy Physics*: The Atlas experiment for the Large Hadron Collider at the Center for European Nuclear Research generates raw data at a rate of 2 PBs per second and stores approximately 100 PBs per year of processed data.

2.2 A New Memory and Storage Hierarchy

Accessing, storing, and processing data is of the utmost importance for the above applications which expect a certain set of features from the underlying storage systems: a) high I/O bandwidth, b) low latency, c) reliability, d) consistency, e) portability, and f) ease of use. New system designs that incorporate non-volatile buffers between the main memory and the disks are of particular relevance in mitigating the periodic burstiness of I/O. The new DMSH promises to offer a solution that can efficiently support scientific discovery in many ways: improved application reliability through faster checkpoint-restart, accelerated I/O performance for small transfers and analysis, fast temporary space for out-of-core computations and in-transit visualization and analysis. Building hierarchical storage systems is a cost-effective strategy to reduce the I/O latency of HPC applications. However, while DMSH systems offer higher I/O performance, data movement between the layers of the hierarchy is complex and significantly challenging to manage. Moreover, there is no software yet that addresses the challenges of DMSH.

Middleware layers, like MPI-IO and parallel HDF5, try to hide the complexity by performing coordinated I/O to shared files while encapsulating general purpose optimizations. However, the actual optimization strategy of these middleware layers is dependent on the underlying file system software and hardware implementation. More importantly, these middleware libraries are designed with memory-to-disk endpoints and are not ready to handle I/O access through a DMSH system, which is ultimately left to the user. Ideally, the presence of multiple layers of storage should be transparent to applications without having to sacrifice performance or increase programming difficulty. System software and a new middleware solution to manage these intermediate layers can help obtain superior I/O performance. Ultimately, the goal is to ensure that developers have a high-performance I/O solution that minimizes changes to their existing software stack, regardless of the underlying storage.

Deep memory and storage hierarchies require a scalable, reliable, and high-performance software to efficiently and transparently manage data movement. New data placement and flushing policies, memory and metadata management, and an efficient I/O communication fabric is required to address DMSH complexity and realize its potential. We believe that a radical departure from the existing software stack for the scientific communities is not realistic. Therefore, we propose to raise the level of abstraction by introducing a new middleware solution, Hermes, and make it easier for the user to perform I/O on top of a DMSH system. In fact, Hermes

Figure 1: Software stack and Hermes internal design.

supports existing widely popular I/O libraries such as MPI-IO and HDF5 which makes our solution highly flexible and production-ready. We envision a buffering platform that can be application- and system-aware, and thus, hide lower level details allowing the user to focus on his/her algorithms. We strive for maximizing productivity, increasing resource utilization, abstracting data movement, maximizing performance, and supporting a wide range of scientific applications and domains.

3 DESIGN AND IMPLEMENTATION

3.1 Hermes Architecture

3.1.1 Design overview. Hermes is designed as a middleware layer - sitting between applications and DMSH as shown in Figure 1. As a middleware library, Hermes captures I/O calls, both POSIX and HDF5 (i.e., `fopen`, `fread`, `fwrite`, and `H5Fcreate`, `H5Dread` etc.) and redirects them to different layers of DMSH. Legacy applications can easily connect to Hermes by simple linking (i.e., `LD.PRELOAD`) or recompiling the code with our library. There are no changes to user code and there is no need to upgrade to a different workflow. We design Hermes to easily work with existing software. Our goal is to maximize user productivity by making I/O buffering transparent. Furthermore, Hermes also provides a new buffering API for users who want to explicitly take control of the data movement between layers of DMSH. This mode also allows Hermes to perform active buffering where data is shipped to the buffer nodes along with specific instructions or operations to be performed on them. For example, a user can pass a set of integers to Hermes instructing it to first store them to the buffer nodes, then sort them, compress the sorted list and lastly persist the final result to the remote PFS. This flow can be easily executed by a series of hinting mechanisms (i.e., flags) that Hermes provides to the user. Our hinting mechanism is a simple bit encryption which indicates predetermined operations like sorting, compression/decompression, deduplication and others. For user defined operations, Hermes provides a bootstrapping mechanism in which the user can submit his/her functions. The library will then compile and place the executables to a registry of operations to be handled by the buffering nodes. Reserved bits are used for user-defined operations. The high-level architecture of Hermes can be seen in Figure 2. In DMSH systems, besides the main memory, every compute node might be equipped with an NVMe device or even an SSD. Additionally, shared buffering nodes, such as burst buffers, will most likely

Figure 2: Hermes internal design.

be present and positioned close to the compute nodes. Finally, a remote PFS supports all compute nodes with persistence and fault tolerance as important features. Hermes is a platform that aims to enable efficient access to the layers of DMSH and as such we distinguish two data paths: *a vertical and a horizontal hierarchy*. Vertical hierarchy refers to data movement within a compute node and all the way down to the burst buffers and PFS. Horizontal hierarchy refers to sending data to another compute node's RAM or NVMe device. The horizontal data movement is greatly optimized if there is an RDMA-capable network but Hermes can also support systems with no RDMA. Therefore, a DMSH system could consist of several layers, performance-wise, such as local RAM, remote RAM, local NVMe, remote NVMe, burst buffers, and PFS (numbered in fig. 2).

3.1.2 Internal components. Figure 1 demonstrates the design of Hermes library and all the internal components that work together to achieve an efficient, transparent, and easy-to-use data access in all layers of a DMSH (i.e., both vertically and horizontally). The main Hermes library is complemented by a set of tools and services that help achieve broader goals such as multi-tenancy, adaptability, etc. Brief description of each component's responsibilities:

API: The API is responsible to intercept all I/O calls from the applications. It also calculates the operations to be carried out by the buffering nodes in case of an active buffering scenario.

Data Placement Engine: This engine is responsible to map data onto DMSH. In other words, the data placement engine calculates the data destination, where in the hierarchy should the data be redirected. It maps data according to various data placement policies.

Data Organizer: The main responsibility of this component is to move data between the layers of DMSH. It is triggered by other components according to certain criteria which makes it an event-based component. For instance, if there is no space left in NVMe, data organizer is triggered to move data down to the burst buffers and thus freeing space in NVMe. This component is responsible to carry out all data movement either for prefetching reasons, evictions, lack of space, or hotness of data etc.

Metadata Manager: The MDM maintains two types of metadata information: user's and Hermes library's internal metadata. Since Hermes can transparently buffer data by intercepting I/O calls, MDM keeps track of user's metadata operations (i.e., files, directories, permissions etc.) while consulting the underlying PFS. Additionally, since data can be buffered anywhere in the hierarchy, MDM

tracks the locations of all buffered data and internal temporary files that contain user files.

Cache manager: This component is responsible to handle all buffers inside Hermes. It is equipped with several cache replacement policies such as least recently used (LRU) and least frequently used (LFU). It works in conjunction with the prefetcher. It can be configured to hold "hot" data for better I/O latency. It is also responsible to implement application-aware caching schemas.

Prefetcher: This component is performance-driven. It implements several typical prefetching algorithms such as sequential data access, strided access, and random access. Hermes also supports user defined prefetching. In a way, the prefetcher becomes Hermes' client for reading operations much like application cores are when writing data in DMSH.

Messaging Service: This component is used to pass small messages across the cluster of compute nodes. This component does not involve any data movement which is actually done by either the application cores or other Hermes components such as the data organizer and prefetcher. Instead, this component provides an infrastructure to pass instructions to other nodes to perform operations on data or facilitate its movement. For example, a typical type of message in Hermes is to flush buffered data of a certain file to the next layer or to PFS.

I/O Clients: These clients refer to simple calls using the appropriate API based on the layer of the hierarchy. For instance, if Hermes data placement engine maps some data to the burst buffers, then the respective I/O client will be called and perform the `fwrite()` call. Internally, Hermes can use POSIX, MPI-IO, or HDF5 to perform the I/O. An important feature of Hermes is that user's data structures are mapped to Hermes' internal structures at each layer of DMSH. For example, an original dataset of an HDF5 file could be mapped into a temporary POSIX file in NVMe. The I/O clients give Hermes the flexibility to "talk" to several data destinations and manage the independent systems (e.g., `memcpy` for RAM, `fwrite()` for NVMe, `MPI_File_write()` for burst buffers).

System Profiler: This component is a service outside the main library. It is designed to run once during the initialization. It performs a profiling of the underlying system in terms of hardware resources. It tries to detect the availability of DMSH and measure each layer's respective performance. It is crucial to identify the parameters that Hermes needs to be configured with. Using this information, the data placement engine can do a better job when mapping data to different layers. Each system will have different hierarchy. Additionally, each hierarchy will demonstrate different performance characteristics. In our prototype implementation this component is external and results are manually injected to the configuration of the library. We plan to automate this process.

Schema Parser: This component accepts a user-defined buffering schema and embeds it into the library. This schema is passed in a XML format and Hermes is configured accordingly. For instance, if user chooses to aggressively buffer a certain dataset or file, then Hermes will prioritize this data higher up in the hierarchy and also the cache manager will get informed not to evict this specific buffered dataset. All this is possible because Hermes will use the user's instructions to offer the best buffering performance. In our prototype implementation schema parser is external and is planned to be automated in future versions of Hermes.

Applications Coordinator: This component is designed to offer support in a multiple-application environment. It manages the access to the shared layers of the hierarchy such as the burst buffers. Its goal is to minimize interference between different applications sharing this layer. Additionally, it coordinates the flushing of the buffers to achieve maximum I/O performance. More information on this component can be found in [29].

All the above components allow Hermes to offer a high performance I/O buffering platform which is highly configurable, easily pluggable to several applications, adaptable to certain system architectures, and feature-rich yet lightweight.

3.2 Hermes Buffering Modes and Policies

3.2.1 Buffering modes. Similar to other buffering systems, Hermes offers several buffering modes (i.e., configurable by the user) to cover a wide range of different application needs such as I/O latency, fault tolerance, and data sharing:

A. Persistent: in this mode, data buffered in Hermes is also written to the PFS for permanent storage. We have designed two configurations for this mode. 1) Synchronous: directs write I/O onto DMSH and also to the underlying permanent storage before confirming I/O completion to the client. This configuration is designed for use cases such as *write-though* cache or *stage-in* for read operations. Since all data also exist in the PFS, synchronous-persistent mode is highly fault-tolerant, offers strong data consistency, is ideal for data sharing between processes, and supports read-after-write workloads. However, it demonstrates the highest latency and lowest bandwidth for write operations since data directed to the buffers also need to be written in the PFS. 2) Asynchronous: directs write I/O onto DMSH and completion is immediately confirmed to the client. The contents of buffers are eventually written down to the permanent storage system. The trigger to flush buffered data is configurable and can be: i) per-operation, flushing is triggered at the end of current `fwrite()`, it also flushes all outstanding previous operations, ii) per-file, flushing is triggered upon calling `fclose()` of a given file (this is similar to Data Elevator approach), iii) on-exit, flushing is triggered upon application exit (this is similar to Datawarp approach), and iv) periodic, flushing is periodically triggered in the background (this is the default Hermes setting). This configuration is designed for use cases such as *write-back* cache and *stage-out* for read operations. It provides low-latency and high bandwidth to the application since processes return immediately after writing to the buffers. It also offers eventual consistency since data are flushed down eventually. It is ideal for write-heavy workloads and out-of-core computations.

B. Non-persistent: in this mode, I/O is directed to DMSH and is never written down to the permanent storage. It is designed to offer a scratch space for fast temporary I/O. Upon application exit, Hermes deletes all buffered data. This mode can be used for scenarios such as quickly storing intermediate results, communication between processes, in-situ analysis and visualization. In case of buffering node failures, application must restart. This mode offers high bandwidth and low latency. Lastly, applications can reserve a specific allocation (i.e., capacity on buffers) for which data preservation is guaranteed by Hermes (similar to Datawarp reservations). These allocations expire with the application lifetime. In case of buffer overflow, Hermes will transparently swap buffer contents

to the PFS much like memory pages are swapped to the disk by the OS. The mechanism was designed to offer some extra degree of flexibility to Hermes. For example, let us assume that an application writes simulation results every 5 minutes. These results are directly read from the buffers by an analysis kernel which writes the final result to the PFS for permanent storage. Simulation data can be deleted or overwritten after the analysis is done. Hermes can utilize this periodic and bursty I/O behavior and write the next iteration on top of the previous one instead of wasting extra buffer space. To achieve this conditional overwriting of data, Hermes utilizes a flagging system to define the lifetime of buffered data.

C. Bypass: in this mode, as the name suggests, I/O is performed directly against the PFS effectively bypassing Hermes. This mode resembles *write-around* cache designs.

3.2.2 Data placement policies. In DMSH systems, I/O can be buffered to one or more layers of the hierarchy. There are two main challenges: i) how and where in the hierarchy data are placed, ii) how and when do buffers get flushed either in the next layer or all the way down to PFS. In Hermes, the first challenge is addressed by the data placement engine (DPE) component and the second by the data organizer. We designed four different data placement policies to cover a wide variety of applications' I/O access patterns. Each policy is described by a dynamic programming optimization¹ and follows the flow of Algorithm 1. The general idea of the algorithm is as follows. First, if the incoming data can fit in the current layer's remaining capacity, it places the data there (i.e., *PlaceData()*). In case it does not fit, based on the constraint of each policy, it tries one of the following: a) solve again for next layer (i.e., *skip()*), b) place as much data as possible in the current layer and the rest in next (i.e., *split()*), and c) flush current layer and then place new incoming I/O (i.e., *flush()*). We implemented the DP algorithm using memoization techniques to minimize the overhead of the solution. We further provide a configuration knob to tune the granularity of triggering the optimization code for data placement.

A. Maximum Application Bandwidth (MaxBW): this policy aims to maximize the bandwidth applications experience when accessing Hermes. The DPE places data in the highest possible layer of DMSH in a top-down approach, starting from RAM, while balancing bandwidth, latency, and the capacity of each layer. The approach applies to all layers making the solution recursively optimal in nature. The above data placement policy is expressed as an optimization problem where DPE minimizes the time taken to write the I/O in the current layer and the access latency to serve the request, effectively maximizing the bandwidth. The data organizer moves data down periodically (or when triggered) to increase the available space in upper layers for future incoming I/O. Data movement between layers is performed asynchronously. This policy is the default Hermes configuration.

B. Maximum Data Locality: this policy aims to maximize buffer utilization by simultaneously directing I/O to the entire DMSH. The DPE divides and places data to all layers of the hierarchy based on a data dispersion unit (e.g., chunks in HDF5, files in POSIX and independent MPI-IO, and portions of a file in collective MPI-IO). Furthermore, Hermes maintains a threshold based on the capacity ratio between the layers of the hierarchy. This ratio reflects on the

Algorithm 1: Hermes algorithm to calculate data placement in DMSH (pseudo code)

```

1  Hermes-DPE(data request, DMSH layer);
2  if data can fit in current layer then
3    | PlaceData();           // buffer data in this layer
4  else
5    | MaxConstraint() // based on selected policy
6    | - skip();           // buffer in next layer
7    | - split();         // buffer in both current and next layers
8    | - flush();        // buffer in current layer after flushing
9    | );
10 end
```

relationship between each layer (e.g., system equipped with 32GB RAM, 512GB NVMe, and 2TB burst buffers creates a capacity ratio of 1-16-64). The data placement in this policy accounts for both layer's capacity and data's spatial locality. The above process is recursive and can be expressed as an optimization problem. DPE minimizes the time taken to write the I/O in the current layer and the degree of data dispersion (i.e., how many layers data are placed to) effectively maximizing the buffer utilization. Data movement between layers is performed asynchronously. This policy is ideal for workflows that encapsulate partitioned I/O. For instance, one could prioritize a certain group of MPI ranks over another (e.g., aggregator ranks) or one type of file over another (e.g., metadata files over data files).

C. Hot-data: this policy aims to offer applications a fast cache for frequently accessed data (i.e., hot-data). The DPE places data in the hierarchy based on a hotness score that Hermes maintains for each file. This score encapsulates the access frequency of a file. Highest scored files will be placed higher up in DMSH since they are expected to be accessed more often. This ensures that layers with lower latency and higher bandwidth will serve critical data such as metadata, index files, etc. The DPE also considers the overall file size to efficiently map data to each layer (i.e., smaller files buffered in RAM whereas larger files in burst buffers). The data placement policy can be expressed as an optimization problem where DPE minimizes the time taken to write the I/O in the current layer considering both hotness and capacity of layers. The data organizer demotes or promotes data based on the hotness score and the data movement is performed asynchronously. This policy is ideal for workflows that demonstrate a spectrum of hot-cold data.

D. User-defined: this policy aims to support user-defined buffering schemas. Users are expected to submit an XML file with their preferred buffering requirements. This file is parsed during initialization by the schema parser component and used by the DPE to make data placement decisions. For instance, user can define certain files to always be in RAM (i.e., never get evicted), or which HDF5 chunks to get buffered in NVMe etc.

3.3 Implementation Details

3.3.1 Node design. The new DMSH system architecture suggests that compute nodes may be equipped with one or more non-volatile storage device and share access to a burst buffer deployment. Hermes is designed to support all the new trends in system design. Figure 3 demonstrates Hermes node design. Each application core uses an I/O API (i.e., POSIX, MPI-IO, HDF5 etc.) which in turn

¹Full mathematical formulation of each policy can be found in the Appendix.

Figure 3: Compute node design in Hermes.

is captured by Hermes. A dedicated core per node, called *Node Manager*, is exclusively used by Hermes services. Specifically, this multi-threaded core is responsible for metadata management, data organization and movement between layers, messaging services between compute nodes (horizontal hierarchy), local memory management such as placement of data in buckets, eviction policies, and finally prefetching. The ratio between application cores and the Hermes node manager is configurable and is suggested to be around 64-to-1 (i.e., similar to I/O forwarding layer present in several supercomputing sites). If an I/O forwarding layer exists, Hermes can utilize the I/O cores there. However, our design is not limited only to such systems and can be widely deployed.

3.3.2 Critical components. During I/O buffering into DMSH, there are three critical operations: memory, metadata, and communication management. To achieve high-performance in each of these critical operations, Hermes incorporates several novel technical innovations. As it can be seen in Figure 3, RAM is split into application memory and Hermes memory, which is further divided in bucket pool, MDM, and message queue.

A. RAM management. We have designed a new memory management system to offer fast and efficient use of main memory, a very crucial resource in any buffering platform. Hermes stores data in buckets, an abstract notion of a data holder. Buckets have a configurable fixed size and consist of a collection of memory pages. All buckets are allocated during the bootstrapping of the system, creating a bucket pool. This allows Hermes to avoid the cost of per-request memory allocation (i.e., only pay the cost in the beginning before application starts), to better control memory usage by avoiding expensive garbage collection, and to define the lifetime of memory allocations per application (i.e., re-use the same buckets after data have been flushed down). Bucket pools are organized in four regions: available buckets, RAM cache, NVMe cache, burst buffers cache. The bucket pool is managed by the bucket manager who is responsible to keep track of the status of each bucket (e.g., full - available). The bucket, as a unit of buffering, is extremely critical to achieve high performance, low latency, and increases design flexibility (e.g., better eviction policies, hot data cache etc.).

Figure 4: RAM operations throughput.

We implemented Hermes' memory management using MPI one-sided operations. Specifically, buckets are placed in a shared dynamic Remote Memory Access (RMA) window. This allows easier access to the buckets from any compute node and a better global memory management. MPI-RMA implementations support RDMA-capable networks which further diminishes the CPU overhead. Access to buckets occurs using *MPLPut()* and *MPLGet()*. Update operations are atomic with exclusive locking only on the bucket being updated. To support fast querying (e.g., location of a bucket, list of available buckets, etc.) the bucket manager indexes the RMA window and bucket relationships much like how *inode* tables work. The structure of a bucket includes an identifier (uint32), a data pointer (void*), and a pointer (uint32) to the next bucket. Hermes' buckets are perfectly aligned with RAM's memory pages which optimizes performance especially for applications with unaligned accesses. Finally, to ensure data consistency and fault tolerance, Hermes maps (via *mmap()*) the entire MPI-RMA window and the index structure to a file stored in a non-volatile layer of the hierarchy (configured by user). We suggest placing this special file to the burst buffers since if a compute node fails, the local NVMe device will become unavailable till the node is fixed.

Figure 4 motivates our design for Hermes' memory management. In this test, we issued a million fwrites of various sizes (from 64KB to 2MB) and measured the achieved memory operations per second. The test was conducted on our development machine that runs CentOS 7.1. In the test's baseline, we intercept each *fwrite()*, allocate a memory buffer (i.e., *malloc()*), copy data from user's buffer to the newly allocated space (i.e., *memcpy()*), and finally flush the buffer (i.e., *free()*) once the data are written to the disk. As a slightly optimized baseline case we used Google's TC Malloc. In contrast, Hermes intercepts each *fwrite()*, calculates how many buckets are required to store the data and asks the bucket manager for them, and copies data from user's buffer to the acquired buckets. Once data are written to the disk, buckets are marked by the data organizer as available and no freeing is performed. As it can be seen in figure 4, Hermes outperforms Linux's *Malloc* by 3x and *TCMalloc* by 2x. Hermes managed to sustain more than 3 million memory ops/sec, whereas the baselines, 1 and 2 million ops/sec respectively. Interestingly, as the allocation size grows, Linux's *Malloc* struggles in performance compared to *TCMalloc*. The pre-allocation and efficient management of the buckets and the lack of freeing of buffers helped Hermes to maintain stable high performance.

B. Metadata management. Any metadata service in distributed systems is subject to scalability and performance issues. Metadata in a buffering platform like Hermes consist of data distribution information (e.g., which node, which layer in DMSH, which bucket,

Figure 5: Metadata Manager throughput.

etc.) and maintenance of both user’s and internal file namespaces. Hermes’ metadata manager is distributed and aims to offer highly concurrent and asynchronous operations. To achieve this, Hermes employs a novel distributed hashmap design, implemented using RMA windows and MPI one-sided operations. A hashmap consists of keys that correspond to specific values. Our design uses two RMA windows: i) *key window*, which is indexed to support efficient querying and ii) *value window*, for data values. This practically allows any process to simply *MPI_Get()* a specific key and then fetch its respective value. We use a 2-way hashing: first, the key is hashed to a specific node and then into a value that resides on that node. The MPI one-sided operations allow Hermes to perform metadata operations without interrupting the destination node. RDMA-capable machines will be able to perform even faster by using the RDMA controller for any data movement. Additionally, the RMA windows are dynamic effectively allowing the metadata to grow in size as required, similarly with rehashing in traditional hashmap containers. Lastly, our hashmap design liberates us to use complex structures, such as objects and nested custom datatypes, to describe a certain file and its metadata information. In contrast, popular in-memory key-value such as Redis or MemCached use simple datatypes for keys and values (e.g., strings or integers) which can be a limiting factor to metadata services. Additionally, these key-value stores offer features that are not useful in our use case such as replication, timestamps, and other features that only add overhead if one does not need or intend to use them.

Hermes’ MDM uses several maps: i) file handler to file: maintains file handlers of opened files, $\{fh, filename\}$, ii) file to metadata properties: maintains all typical file properties (e.g., permissions, ownership, timestamps etc.), $\{filename, \{filestat\}\}$, iii) files to location in DMSH: maintains data distribution information, $\{filename, \{offset, size\}, \{node, layer, type, identifier, freq\}\}$, and iv) node to current status: maintains information for each node’s current status such as remaining capacity, hot data access frequencies, etc., $\{node, \{layer, size, \dots\}\}$. These maps allow fast queries and $O(1)$ read/write MDM operations without the need to execute separate services (e.g., a memcached server). Creation and update of metadata information is performed by using *MPI_EXCLUSIVE* locks which ensures FIFO consistency. Read operations use a shared lock which offers higher performance and concurrency. Finally, Hermes’ MDM exposes a simple and clean API to access its structures (e.g., *mdm_update_on_open()*, *mdm_get_file_stat()*, *mdm_sync_meta()*, etc.).

In Figure 5 we compare Hermes’ MDM performance with a custom MPI-based solution, Memcached, and Redis. In this test, we issue a million metadata operations and we measure the MDM throughput in operations per second. First, we implemented a

custom MPI-based solution where one process per node is the MDM and answers queries from other processes. Upon receiving one, it queues the operation, it spawns a thread to serve the operation, and it goes back to listening. The spawned thread removes the operation from the queue and performs the operation. While this approach is feasible, it uses a dedicated core per node. Another approach is to use an in-memory key-value store. We implemented the MDM using Memcached and Redis, two of the most popular solutions. In this approach, one memcached or Redis server per node is always running and awaits for any metadata operations. There is no explicit queuing but its implementation uses multi-threaded servers with locks and internal queues to support concurrent operations. Again, a dedicated core is required to run the server. Lastly, Hermes is using our own hashmap to perform metadata operations. Each processes accesses the shared RMA window to get or put metadata. There is no dedicated core used. As it can be seen in Figure 5, our solution outperforms by more than 7x the MPI-based custom solution and by more than 2x the Memcached and Redis versions. Update operations are more expensive since clients first need to retrieve the metadata, update them, and then push them back.

C. Messaging service. Many operations in Hermes involve communication between different compute nodes, buffering nodes, and several other components. The messaging service does not involve in data movement but instead provides the infrastructure to pass instructions between nodes. For instance, horizontal access to the deep memory hierarchy involves sending data across the network to a remote RAM or NVMe. Another example is when the prefetcher gets triggered by one process it will fetch data to a layer of the hierarchy for subsequent read operations. Finally, when the buffers are flushed to the remote parallel file system for persistence, a system-wide coordination is required. All the above cases, require a high-performance and low latency messaging service to be in place. Hermes implements such messaging service by utilizing our own distributed queue via MPI one-sided operations. We designed a scalable messaging service by leveraging the asynchronicity of MPI RMA operations. When a process needs to communicate with another process across the compute nodes, it simply puts a message into the distributed queue that is hosted by all compute nodes. A shared dynamic RMA window is used to hold the queue messages. Each message has a type (i.e., an instruction to be carried out), its associated attributes, and a priority. As with the distributed hashmap above, if there is an RDMA controller it will be used to avoid interrupting the destination core. There is no need to employ listeners or other always-on services such as Apache ActiveMQ [49] or Kafka [30] leading to better resource utilization. Additionally, we define our own bit encoding to keep the messages small and avoid costly serializations/transformations and therefore lead to lower latencies and higher throughput. Hermes messaging service aims to offer higher overall performance avoiding network bottlenecks and communication storms.

In Figure 6 we compare Hermes’ performance with a custom MPI-based solution, Memcached, and NATS. In this test, we issue a million queue operations (e.g., publish - subscribe) and we measure the messaging rate in messages per second. As described above, we implemented a custom MPI-based solution where one process per node accepts messages from other processes. We also implemented a distributed queue using Memcached where each

Figure 6: Messaging Service throughput.

message becomes a key-value pair (i.e., ID-message). Furthermore, we explored NATS, a popular, in-memory, high-performance, and open source messaging system. In both latter options, a dedicated core needs to run server code. Lastly, Hermes is using our own distributed priority queue to execute the messaging service. Each processes puts or gets messages from the shared RMA window while no dedicated core is used. As it can be seen in figure 6, Hermes outperforms the custom MPI-based messaging implementation by more than 12x. This is expected since the server process gets saturated from the overwhelming rate of incoming messages. As a result, client processes needs to wait blocked for the server to accept their message. The handler thread cannot match the rate of new messages. A similar picture is evident in the memcached solution where Hermes performs more than 8x faster. However, in memcached, up to 4 handler threads are spawned which possibly leads to better performance compared to the custom MPI-based one. Finally, NATS performance is really good with more than 300000 published messages per second. However, Hermes outperforms NATS by more than 2x for publishing and more than 3x for subscribe operations.

3.4 Design Considerations

In this subsection, we briefly discuss concerns regarding the design and features of any buffering platform, especially one that supports a DMSH system such as Hermes. The goal is to present some of our ideas and to generate discussion for future directions.

A. High-performance:

Concern 1: How to support and manage heterogeneous hardware? Hermes is aware of the heterogeneity of the underlying resources via the system profiler component which identifies and benchmarks all layers present in the system. Hermes aims to utilize each hardware resource to its best of its capabilities by avoiding hurtful workloads. Instead, Hermes' I/O clients generate access patterns favorable to the each medium.

Concern 2: How to avoid excessive network traffic?

Hermes' messaging service is carefully designed to operate with small-sized messages with bit encoding. Furthermore, by using asynchronicity and RDMA capable hardware our solution ensures the low network overhead.

Concern 3: How to support low-latency applications?

The several data placement policies of Hermes' DPE provide tunable performance guarantees for a variety of workloads. For low latency applications, Hermes can leverage the performance characteristics of each layer by placing data to the fastest possible layer. Additionally, our novel memory management ensures that data can be efficiently cached in RAM before ending up to their buffer.

Concern 4: How to avoid possible buffer overflow?

Hermes' Data Organizer component manages the capacities of the layers and moves data up and down the hierarchy (i.e., between the layers). In corner cases of overflow, Hermes provides explicit triggers to the data organizer to re-balance the layers and move data based on the buffer capacity on each layer.

Concern 5: How to scale the buffer capacity?

Hermes' DPE can place data in remote RAM and NVMe devices, and thus, scaling is horizontal by adding more compute nodes. Additionally, Hermes can support RAM Area Network (RAN) deployments [57] to further extend the buffer capacity.

B. Fault tolerance:

Fault tolerance guarantees are based on the buffering mode selected (i.e., sync, async). In case of asynchronous buffering mode, buffered data are written to a fault tolerant layer such as a PFS eventually which means for a small window of time buffer contents are susceptible to failures. In our prototype implementation, buffers are flushed based on an event-driven architecture and also periodically to decrease the possibilities of losing critical data. As a future step, we want to investigate the following options: i) Checkpointing with configurable frequency. ii) Random replication per write operation. iii) DPE skips the failing component for incoming I/O.

C. Data consistency:

Concern 1: Data consistency model?

Hermes supports strong consistency for the application since our design avoids having the same buffered data in multiple locations and copies. Once a write is complete, any other process can read the data via either a local or a remote call. Excessive locking is avoided by using MPI RMA operations and memory windows. The model supported is single-writer, multiple-readers.

Concern 2: Support of highly concurrent metadata operations?

Upon opening a file, metadata are loaded from the PFS to the local RAM of the process that opened it. Then, Hermes randomly selects two other nodes and replicates metadata there. We do this to increase the availability of the metadata info and avoid saturation of one node's RAM. When another process wants to access the metadata, it randomly selects one of the replica copies and performs the get. If it needs to update the metadata, Hermes propagates the update to all replicas. This is synchronous to ensure consistency.

D. Hermes limitations: Hermes' DPE component implements our data placement policies based on the assumption that the user knows exactly what his/her workload involve, and thus, selecting the appropriate policy is not trivial. As a suggestion, the user can first profile his/her application using typical monitoring and profiling tools, such as Darshan [9], extract knowledge regarding the I/O behavior, and make the right policy choice.

4 EVALUATION

4.1 Methodology

Overview: To evaluate Hermes, we have conducted two set of experiments. We first explored how Hermes' data placement policies handle different workloads and application characteristics using synthetic benchmarks. We then compare Hermes with state-of-the-art buffering platforms, namely Data Elevator and Cray's DataWarp, using real applications. As performance metric, we use the overall execution time in seconds which we further divide to: i) time to write/read to/from buffers, and ii) time to flush buffers to PFS. Computation time is excluded since it is the same among all systems.

Device	RAM	NVMe	SSD	HDD
Model	M386A4G40DM0	Intel DC P3700	Intel DC S3610	ST9250610NS
Connection	DDR4 2133Mhz	PCIe Gen3 x8	SATA 6Gb/s	SATA 7200rpm
Capacity	128 GB(8GBx16)	1.2 TB	1.6 TB	2.4 TB
Latency	13.5 ns	20 μ s	55-66 μ s	4.16 ms
Max Read BW	13000 MB/s	2800 MB/s	550 MB/s	115 MB/s
Max Write BW	10000 MB/s	1900 MB/s	500 MB/s	95 MB/s
Test Config	32x client nodes	RamFS emulated	8x burst buffers	16x PFS servers
ReadBW tested	92647 MB/s	38674 MB/s	3326 MB/s	883 MB/s
WriteBW tested	86496 MB/s	33103 MB/s	2762 MB/s	735 MB/s

Figure 7: Testbed specifications.

As reference, we include a baseline of *no buffering* in which data are written/read directly to/from the PFS. We run all tests ten times and we report the average time.

Hardware: All experiments were conducted on Chameleon [13]. More specifically, we used the bare metal configuration with 32 client nodes (i.e., up to 1024 MPI ranks), 8 burst buffer nodes, and 16 PFS storage nodes. Each node has a dual Intel(R) Xeon(R) CPU E5-2670 v3 running at 2.30GHz with a total of 48 cores, and 128 GB RAM. Each burst buffer node is equipped with an SSD drive and each PFS node with an HDD. We emulated one NVMe device per client node by deploying a DRAM-based file system (i.e., RAMDISK) and imposing latency and bandwidth penalties to match the actual NVMe performance [20, 52, 55]. In order to correctly calculate the added latency and lowered bandwidth, we captured the performance characteristics of real NVMe devices present in the hierarchy appliances of Chameleon. Figure 7 lists all the hardware specifications and performance measurements. Lastly, to better capture the architecture of a modern supercomputer, we setup our cluster topology as follows: all 32 client nodes and 8 burst buffers are interconnected with 56Gbps Infiniband network and the 16 storage nodes are connected to the rest via a 10Gbps Ethernet network.

Software: The operating system of the cluster is CentOS 7.1, the MPI version is Mpich 3.2, the PFS we used is OrangeFS 2.9.6, the in-memory key-value stores are Memcached 1.4.36 and Redis 4.0.6, and lastly the distributed queue we used is NATS Server 1.0.4.

Applications: We evaluate Hermes using our own synthetic benchmark that emulates common scientific application workloads such as alternation between computation - I/O phases, read-after-write, read-once, read-many etc. It uses POSIX-IO to issue requests to the file system and operates in a typical file-per-process pattern. We also use two real science applications: Vector Particle-In-Cell (VPIC), a general purpose simulation code for modeling kinetic plasmas in spatial multi-dimensions, and Hardware Accelerated Cosmology Code (HACC), a cosmological simulation that studies the formation of structure in collisionless fluids under the influence of gravity in an expanding universe. Both of these simulations perform computations and produce output files periodically that need to be persisted in PFS. Also, both demonstrate a periodic behavior with time steps (i.e., iterations) that include the checkpoint and restart as well as the analysis outputs produced by the simulations. At the end of each step, VPIC writes a single HDF5 file containing properties of 8 million particles. VPIC tends to be extremely I/O intensive (i.e., write-only, write-heavy), since the portion of computation is small. In contrast, HACC has read-after-write workload where, at every step, simulation writes out a single shared file (i.e.,

Figure 8: Benchmark: Alternating Compute-I/O phases.

MPI Collective I/O) that various analysis modules read back. We used 16 time steps for both simulations resulting to total I/O of 1TB.

4.2 Experimental Results

4.2.1 Synthetic Benchmarks. Our synthetic benchmark is highly tunable to generate workloads that can stress the buffering system under various use-cases. We designed two test-cases to evaluate Hermes' data placement policies.

Alternating Compute-I/O phases: In this test, each process first performs some computations (emulated by *sleep()* calls) and then writes 64MB in a file-per-process fashion. We repeat this pattern 16 times with 1024 processes resulting in 1TB total I/O size. We vary the ratio of computation over I/O time to emulate three distinct types of applications: *data-intensive*, *compute-intensive*, and *balanced*. We assume that all data written to the buffers need to be also written to the disk-based remote PFS. Therefore, Hermes is configured in *persistent asynchronous* mode. We measure the overall time spent in I/O, in seconds, which consists of write-time and flush-time. Figure 8 shows the results. As it can be seen, the baseline writes directly to PFS (i.e., no flush-time) and maintains stable write performance regardless of the computation-I/O ratio. In Data Elevator and DataWarp, data are written to the burst buffers resulting to similar write-time between them. The difference in performance comes from data flushing. Data Elevator overlaps flushing with computation phases, and thus, as the computation-I/O ratio increases, flush-time decreases (i.e., flushing is hidden behind computation). On the other hand, DataWarp flushes data only once the application finishes and demonstrates stable flush-time regardless of the computation-I/O ratio. In Hermes, data are written in all layers of the DMSH (i.e., RAM, NVMe, and burst buffers in our system). We evaluate both *MaxBW* and *MaxLocality* data placement policies since they buffer data differently. *MaxBW* places data in a top-down fashion. It starts with RAM for the first iterations of the test, and once this layer is full, it first moves data down to NVMe to create space in RAM and then places the incoming iteration in RAM. On the other hand, *MaxLocality* uses layers concurrently. It writes the first iterations in RAM and once this layer is full it goes on to the next without any data movement between layers. It is clear that for data-intensive applications where the rate of incoming I/O is high, *MaxBW*'s data movement between layers imposes some performance losses, and thus, *MaxLocality*'s write performance is slightly higher. As the computation-I/O ratio increases however, *MaxBW* can overlap data movement between layers with computations. Therefore, for compute-intensive workloads, *MaxBW* outperforms *MaxLocality* by 4x in write-time since it ensures that incoming I/O can be written in RAM. For flushing, both policies leverage any

Figure 9: Benchmark: Repetitive Read operations.

computation time available to asynchronously flush buffer contents to PFS, similarly with Data Elevator. However, Hermes flushes all layers of the DMSH concurrently which decreases flush-time significantly. In summary, in this test Hermes offers **8x** and **2x** higher write performance when compared to No Buffering baseline and state-of-the-art buffering platforms respectively.

Repetitive Read operations: In this test, the benchmark is configured to create a write-once, read-many workload. Each process first writes 32MB in a file-per-process approach and then reads back 32MB of data (not necessarily the same data). We have 16 phases of this pattern with 1024 processes aggregating the I/O to 1TB. We vary the repetition of read operations as follows: i) *Read-once*, where 32MB of data is read only once, ii) *Read-many x4*, where 8MB of data is read 4 times (i.e., still 32MB in total), and iii) *Read-many x16*, where 2MB of data is read 16 times. This pattern resembles workloads where portions of data such as metadata information, indices of files, etc., are frequently accessed creating a data hotness spectrum. In this test, we assume that buffers are used as scratch space (i.e., temporary I/O), and thus, Hermes is configured in *non-persistent* mode. The total time, in seconds, is divided into write-time and read-time. As it can be seen in Figure 9, the baseline writes and reads directly from the PFS and maintains a stable performance irrespective of the workload type. In Data Elevator and DataWarp, data are written/read to/from the burst buffers respectively. This results to a considerable performance improvement over the baseline. Since repetitive read operations are treated as new, it shows stable performance across different workloads. In contrast, Hermes implements a *HotData* data placement policy to offer higher performance for this type of workloads. Since HotData will promote frequently accessed data in upper layers, repetitive read operations access data always from RAM resulting in significant performance boost for Read-many x4 and x16. On the other hand, MaxBW, while offering a competitive performance across the tested workloads, does not cache frequent accessed data in RAM and demonstrates a stable performance across the tested workloads. In summary, in this test Hermes offers **38x** and **11x** higher read performance when compared to No Buffering baseline and state-of-the-art buffering platforms respectively.

4.2.2 Real Applications. To test our system under real applications workload, we configured Hermes in *persistent asynchronous* mode since data need to be stored in the PFS for future access and selected the default data placement policy, *MaxBW*.

VPIC: This application demonstrates a write-only I/O access pattern where at the end of each time step, each process writes data to an HDF5 file. During this evaluation we executed the application

Figure 10: I/O Buffering performance with VPIC-IO.

Figure 11: I/O Buffering performance with HACC-IO.

for 16 time steps. We strong scaled the application from 256 to 1024 total ranks and we measured the total time. In Figure 10 we report only the I/O time which consists of write-time (i.e., what the application experiences) and flush-time (i.e., persisting the data asynchronously). As it can be seen, all tested solutions scale linearly with the number of MPI ranks. In the largest tested scale of 1024 ranks, the baseline completed the test in 1192 seconds. Both Data Elevator and DataWarp wrote the entire dataset in 438 seconds. This is approximately a 2.5x improvement over the baseline. However, due to the higher bandwidth of the DMSH, Hermes' write performance is **5x** and **2x** higher than the baseline and the two buffering platforms we tested, respectively. When considering data flushing, Data Elevator overlaps small computations between each time step and flushes the contents of burst buffers in 1115 seconds whereas DataWarp flushes everything at the end in 1274 seconds. In contrast, Hermes leverages the computations but also the concurrency of the DMSH to flush all buffered data to PFS in 637 seconds. In summary, in this test, Hermes outperformed the baseline and state-of-the-art buffering platforms by **40%** and **85%** respectively.

HACC: This application demonstrates a read-after-write I/O access pattern where during each time step, each process reads back data previously written using MPI-Collective IO. During this evaluation we executed the application for 16 time steps. We strong scaled the application from 256 to 1024 total ranks and we measured the total time. In Figure 11 we report only the I/O time which consists of write-time, read-time, and flush-time. As it can be seen, all tested solutions scale linearly with the number of MPI ranks. In the largest tested scale of 1024 ranks, the baseline completed the test in 1313 seconds. Both Data Elevator and DataWarp performed I/O in 348 seconds. This is approximately a 3.7x improvement over the baseline. However, when considering data flushing, Data Elevator completed the test in 773 and DataWarp in 985 seconds effectively reducing the total improvement to 1.6x and 1.3x respectively. In contrast, Hermes completed the entire test in 494 seconds showcasing the potential of a DMSH system. The performance improvement

is substantial when compared to No Buffering baseline with 7.5x faster I/O operations. Hermes outperformed Data Elevator and DataWarp by 2x due to higher bandwidth of the DMSH.

5 RELATED WORK

New hardware technologies have been developed and can be used to build new memory and storage hierarchies using non-volatile memory (NVRAM) such as phase-change memory (PCM) [42], memristors [50], and Flash memory [12]. Flash-based SSD technology has been widely studied [24], characterized [21], and evaluated for different application types [3, 14]. Researchers also advocate the use of shared buffer technologies, such as burst buffers [6], to accelerate I/O. Existing work has considered NVMe devices as a viable solution for I/O staging [25, 26]. Caulfield proposed Moneta [11], an architecture with NVRAM as an I/O device for HPC applications. Ekel extended Moneta with a real PCM device to understand the performance implications of using NVRAM [2]. Dong studied NVRAM for HPC application checkpointing [18]. Kannan studied NVRAM for I/O intensive benchmarks in Cloud environments [26]. Wang proposed BurstMem [53], a technology for optimizing I/O using burst buffers. Sato et al., show how the burst buffers can boost performance of checkpointing tasks by 20x [46].

Active Buffers [35, 36] exploits one-sided communication for I/O processors to fetch data from compute processors' buffers and performs actual writing in the background while computation continues. IOLite [40], proposes a single shared memory per-node for leveraging inter-process communication and buffering of I/O. Such an approach led to 40% boost in performance. Nitzberg [39] proposes collective buffering algorithms for improving I/O performance by 100x on IBM SP2 at NASA Ames Research Center. PLFS [4] remaps an application's preferred data layout into one which is optimized for the underlying file system.

While all the above work emphasizes the benefits of using each technology individually, none introduced a complete I/O buffering platform that leverages the DMSH. The closest work to Hermes is Data Elevator [17], a new system that transparently moves data in a hierarchical system. The authors focused on systems equipped with burst buffers and demonstrated a 4x improvement over other state-of-the-art burst buffer management systems such as Cray's DataWarp [16]. However, they did not address local memory and local non-volatile devices such as NVMe. Hermes considers both local resources and shared resources like burst buffers. Furthermore, Hermes extends buffering into remote resources and tackles data movement to a more complicated landscape of I/O-capable devices.

6 CONCLUSIONS

To increase I/O performance, modern storage systems are presented in a new memory and storage hierarchy, called Deep Memory and Storage Hierarchy. However, data movement among the layers is significantly complex, making it harder to take advantage of the high-speed and low-latency storage systems. Additionally, each layer of the DMSH is an independent system that requires expertise to manage, and the lack of automated data movement between tiers is a significant burden currently left to the users.

In this paper, we present the design and implementation of Hermes: a new, heterogeneous-aware, multi-tiered, dynamic, and distributed I/O buffering system. Hermes enables, manages, and

supervises I/O buffering into the DMSH and offers a buffering platform that can be application- and system-aware, and thus, hide lower level details allowing the user to focus on his/her algorithms. Hermes aims to maximizing productivity, increasing resource utilization, abstracting data movement, maximizing performance, and supporting a wide range of scientific applications and domains. We have presented three novel data placement policies to efficiently utilize all layers of the new memory and storage hierarchy as well as three novel techniques to perform *memory*, *metadata*, and *communication* management in hierarchical buffering systems. Our evaluation results prove Hermes' sound design and show a 8x improvement compared to systems without I/O buffering support. Additionally, Hermes outperforms by more than 2x state-of-the-art buffering platforms such as Data Elevator and Cray's DataWarp.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grants no. CCF-1744317, CNS-1526887, and CNS-0751200.

REFERENCES

- [1] Sean Ahern, Sadaf R Alam, Mark R Fahey, Rebecca J Hartman-Baker, Richard F Barrett, Ricky A Kendall, Douglas B Kothe, Richard T Mills, Ramanan Sankaran, Arnold N Tharrington, et al. 2007. *Scientific application requirements for leadership computing at the exascale*. Technical Report. Oak Ridge National Laboratory (ORNL); Center for Computational Sciences.
- [2] Ameen Akel, Adrian M Caulfield, Todor I Mollov, Rajesh K Gupta, and Steven Swanson. 2011. Onyx: A Prototype Phase Change Memory Storage Array. *HotStorage* 1 (2011), 1.
- [3] David G Andersen, Jason Franklin, Michael Kaminsky, Amar Phanishayee, Lawrence Tan, and Vijay Vasudevan. 2009. FAWN: A fast array of wimpy nodes. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. ACM, 1–14.
- [4] John Bent, Garth Gibson, Gary Grider, Ben McClelland, Paul Nowoczynski, James Nunez, Milo Polte, and Meghan Wingate. 2009. PLFS: a checkpoint filesystem for parallel applications. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. ACM, 21.
- [5] John Bent, Gary Grider, Brett Kettering, Adam Manzanara, Meghan McClelland, Aaron Torres, and Alfred Torre. 2012. Storage challenges at Los Alamos National Lab. In *Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on*. IEEE, 1–5.
- [6] Wahid Bhimji, Debbie Bard, Melissa Romanus, David Paul, Andrey Ovsiyannikov, Brian Friesen, Matt Bryson, Joaquin Correa, Glenn K Lockwood, Vakho Tsulaia, et al. 2016. Accelerating science with the NERSC burst buffer early user program. *CUG2016 Proceedings* (2016).
- [7] D Brown, Paul Messina, D Keyes, J Morrison, R Lucas, J Shalf, P Beckman, R Brightwell, A Geist, J Vetter, et al. 2010. Scientific grand challenges: Crosscutting technologies for computing at the exascale. *Office of Science, US Department of Energy, February* (2010), 2–4.
- [8] Philip Carns, Kevin Harms, William Allcock, Charles Bacon, Samuel Lang, Robert Latham, and Robert Ross. 2011. Understanding and improving computational science storage access through continuous characterization. *ACM Transactions on Storage (TOS)* 7, 3 (2011), 8.
- [9] Philip Carns, Robert Latham, Robert Ross, Kamil Iskara, Samuel Lang, and Katherine Riley. 2009. 24/7 characterization of petascale I/O workloads. In *Cluster Computing and Workshops, 2009. CLUSTER'09. IEEE International Conference on*. IEEE, 1–10.
- [10] Adrian M Caulfield, Joel Coburn, Todor Mollov, Arup De, Ameen Akel, Jiahua He, Arun Jagatheesan, Rajesh K Gupta, Allan Snaveley, and Steven Swanson. 2010. Understanding the impact of emerging non-volatile memories on high-performance, I/O-intensive computing. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society, 1–11.
- [11] Adrian M Caulfield, Arup De, Joel Coburn, Todor I Mollov, Rajesh K Gupta, and Steven Swanson. 2010. Moneta: A high-performance storage array architecture for next-generation, non-volatile memories. In *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 385–395.
- [12] Adrian M Caulfield, Laura M Grupp, and Steven Swanson. 2009. Gordon: using flash memory to build fast, power-efficient clusters for data-intensive applications. *ACM Sigplan Notices* 44, 3 (2009), 217–228.
- [13] Chameleon.org. 2017. Chameleon system. (2017). <https://www.chameleoncloud.org/about/chameleon/>
- [14] Shimin Chen. 2009. FlashLogging: exploiting flash devices for synchronous logging performance. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*. ACM, 73–86.
- [15] S Conway and C Dekate. [n. d.]. High-Performance Data Analysis: HPC Meets Big Data. ([n. d.]). <http://www.hpcuserforum.com/presentations/tuscon2013/IDCHPDABigDataHPC.pdf>
- [16] CRAY Inc. 2017. DataWarp technology. (2017). <http://www.cray.com/sites/default/files/resources/CrayXC40-DataWarp.pdf>
- [17] Bin Dong, Suren Byna, Kesheng Wu, Hans Johansen, Jeffrey N Johnson, Noel Keen, et al. 2016. Data Elevator: Low-Contention Data Movement in Hierarchical Storage System. In *High Performance Computing (HiPC), 2016 IEEE 23rd International Conference on*. IEEE, 152–161.

- [18] Xiangyu Dong, Naveen Muralimanohar, Norm Jouppi, Richard Kaufmann, and Yuan Xie. 2009. Leveraging 3D PCRAM technologies to reduce checkpoint overhead for future exascale systems. In *High Performance Computing Networking, Storage and Analysis, Proceedings of the Conference on. IEEE*, 1–12.
- [19] Jack Dongarra, Pete Beckman, Terry Moore, Patrick Aerts, Giovanni Aloisio, Jean-Claude Andre, David Barkai, Jean-Yves Berthou, Taisuke Boku, Bertrand Braunschweig, et al. 2011. The international exascale software project roadmap. *The international journal of high performance computing applications* 25, 1 (2011), 3–60.
- [20] Subramanya R Dulloor, Sanjay Kumar, Anil Keshavamurthy, Philip Lantz, Dheeraj Reddy, Rajesh Sankaran, and Jeff Jackson. 2014. System software for persistent memory. In *Proceedings of the Ninth European Conference on Computer Systems*. ACM, 15.
- [21] Kaoutar El Maghraoui, Gokul Kandiraju, Jeofon Jann, and Pratap Pattnaik. 2010. Modeling and simulating flash based solid-state disks for operating systems. In *Proceedings of the first joint WOSP/SIPEW international conference on Performance engineering*. ACM, 15–26.
- [22] Mike Folk, Albert Cheng, and Kim Yates. 1999. HDF5: A file format and I/O library for high performance computing applications. In *Proceedings of Supercomputing*, Vol. 99. 5–33.
- [23] Tony Hey, Stewart Tansley, Kristin M Tolle, et al. 2009. *The fourth paradigm: data-intensive scientific discovery*. Vol. 1. Microsoft research Redmond, WA.
- [24] Yongsoo Joo, Junhee Ryu, Sangsoo Park, and Kang G Shin. 2011. FAST: Quick Application Launch on Solid-State Drives. In *FAST*. 259–272.
- [25] Sooyong Kang, Sungmin Park, Hoyoung Jung, Hyoki Shim, and Jaehyuk Cha. 2009. Performance trade-offs in using nvram write buffer for flash memory-based storage devices. *IEEE Trans. Comput.* 58, 6 (2009), 744–758.
- [26] Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, Dejan Milojicic, and Vanish Talwar. 2011. Using active NVRAM for I/O staging. In *Proceedings of the 2nd international workshop on Petascale data analytics: challenges and opportunities*. ACM, 15–22.
- [27] Youngjae Kim, Raghul Gunasekaran, Galen M Shipman, David A Dillow, Zhe Zhang, and Bradley W Settlemyer. 2010. Workload characterization of a leadership class storage cluster. In *Petascale Data Storage Workshop (PDSW), 2010 5th. IEEE*, 1–5.
- [28] Rob Kitchin. 2014. Big Data, new epistemologies and paradigm shifts. *Big Data & Society* 1, 1 (2014), 2053951714528481.
- [29] Anthony Kougkas, Hariharan Devarajan, and Xian-He Sun. 2017. *Harmonia: An Interference-Aware Dynamic I/O Scheduler for Shared Non-Volatile Burst Buffers*. Technical Report. Illinois Institute of Technology.
- [30] Jay Kreps, Neha Narkhede, Jun Rao, et al. 2011. Kafka: A distributed messaging system for log processing. In *Proceedings of the NetDB*. 1–7.
- [31] Jianwei Li, Wei-keng Liao, Alok Choudhary, Robert Ross, Rajeev Thakur, William Gropp, Robert Latham, Andrew Siegel, Brad Gallagher, and Michael Zingale. 2003. Parallel netCDF: A high-performance scientific I/O interface. In *Supercomputing, 2003 ACM/IEEE Conference. IEEE*, 39–39.
- [32] Glenn K Lockwood, Damian Hazen, Quincey Koziol, RS Canon, Katie Antypas, Jan Balewski, Nicholas Balthaser, Wahid Bhimji, James Botts, Jeff Broughton, et al. 2017. *Storage 2020: A Vision for the Future of HPC Storage*. Technical Report. NERSC.
- [33] Jay F Lofstead, Scott Klasky, Karsten Schwan, Norbert Podhorski, and Chen Jin. 2008. 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*. ACM, 15–24.
- [34] Los Alamos National Lab. [n. d.]. Trinity specs. ([n. d.]). <http://www.lanl.gov/projects/trinity/specifications.php>
- [35] Xiaosong Ma, Marianne Winslett, Jonghyun Lee, and Shengke Yu. 2001. Faster collective output through active buffering. In *Parallel and Distributed Processing Symposium, Proceedings International, IPDPS 2002, Abstracts and CD-ROM. IEEE*, 8–pp.
- [36] Xiaosong Ma, Marianne Winslett, Jonghyun Lee, and Shengke Yu. 2003. Improving MPI-IO output performance with active buffering plus threads. In *Parallel and Distributed Processing Symposium, 2003. Proceedings. International. IEEE*, 10–pp.
- [37] Ningfang Mi, Alma Riska, Qi Zhang, Evgenia Smirni, and Erik Riedel. 2009. Efficient management of idleness in storage systems. *ACM Transactions on Storage (TOS)* 5, 2 (2009), 4.
- [38] NERSC. [n. d.]. Cori system burst buffer design. ([n. d.]). <https://www.nersc.gov/users/computational-systems/cori/burst-buffer/>
- [39] Bill Nitzberg and Virginia Lo. 1997. Collective buffering: Improving parallel I/O performance. In *High Performance Distributed Computing, 1997. Proceedings. The Sixth IEEE International Symposium on. IEEE*, 148–157.
- [40] Vivek S Pai, Peter Druschel, and Willy Zwaenepoel. 2000. IO-Lite: a unified I/O buffering and caching system. *ACM Transactions on Computer Systems (TOCS)* 18, 1 (2000), 37–66.
- [41] JB Peter. 2004. The Lustre storage architecture. *Cluster File Systems, Inc* (2004).
- [42] Moinuddin K Qureshi, Vijayalakshmi Srinivasan, and Jude A Rivers. 2009. Scalable high performance main memory system using phase-change memory technology. *ACM SIGARCH Computer Architecture News* 37, 3 (2009), 24–33.
- [43] Daniel A Reed and Jack Dongarra. 2015. Exascale computing and big data. *Commun. ACM* 58, 7 (2015), 56–68.
- [44] David Reinsel, John Gantz, and John Rydning. 2017. Data Age 2025: The Evolution of Data to Life-Critical. *Donfit Focus on Big Data* (2017).
- [45] Robert B Ross, Rajeev Thakur, et al. 2000. PVFS: A parallel file system for Linux clusters. In *Proceedings of the 4th annual Linux showcase and conference*. 391–430.
- [46] Kento Sato, Kathryn Mohr, Adam Moody, Todd Gamblin, Bronis R De Supinski, Naoya Maruyama, and Satoshi Matsuoka. 2014. A user-level infiniband-based file system and checkpoint strategy for burst buffers. In *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on. IEEE*, 21–30.
- [47] Frank B Schmuck and Roger L Haskin. 2002. GPFS: A Shared-Disk File System for Large Computing Clusters. In *FAST*, Vol. 2.
- [48] John Shalf, Sudip Dosanjh, and John Morrison. 2010. Exascale computing technology challenges. In *International Conference on High Performance Computing for Computational Science*. Springer, 1–25.
- [49] Bruce Snyder, Dejan Bosanac, and Rob Davies. 2017. Introduction to apache ActiveMQ. *Active MQ in Action* (2017), 6–16.
- [50] Dmitri B Strukov, Gregory S Snider, Duncan R Stewart, and R Stanley Williams. 2008. The missing memristor found. *nature* 453, 7191 (2008), 80–83.
- [51] Rajeev Thakur, William Gropp, and Ewing Lusk. 1999. Data sieving and collective I/O in ROMIO. In *Frontiers of Massively Parallel Computation, 1999. Frontiers' 99. The Seventh Symposium on the. IEEE*, 182–189.
- [52] Haris Volos, Guilherme Magalhaes, Ludmila Cherkasova, and Jun Li. 2015. Quartz: A lightweight performance emulator for persistent memory software. In *Proceedings of the 16th Annual Middleware Conference*. ACM, 37–49.
- [53] Teng Wang, Sarp Oral, Yandong Wang, Brad Settlemyer, Scott Atchley, and Weikuan Yu. 2014. Burstmem: A high-performance burst buffer system for scientific applications. In *Big Data (Big Data), 2014 IEEE International Conference on. IEEE*, 71–79.
- [54] Whitt, Justin L. 2017. Oak Ridge Leadership Computing Facility: Summit and Beyond. (2017). [https://indico.cern.ch/event/618513/contributions/2527318/attachments/1437236/2210560/SummitProjectOverview\[...\].pdf](https://indico.cern.ch/event/618513/contributions/2527318/attachments/1437236/2210560/SummitProjectOverview[...].pdf)
- [55] Kai Wu, Yingchao Huang, and Dong Li. 2017. Unimem: Runtime Data Management on Non-Volatile Memory-based Heterogeneous Main Memory. *arXiv preprint arXiv:1705.00249* (2017).
- [56] Bing Xie, Yezhou Huang, Jeffrey S Chase, Jong Youl Choi, Scott Klasky, Jay Lofstead, and Sarp Oral. 2017. Predicting output performance of a petascale supercomputer. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*. ACM, 181–192.
- [57] Dawid Zawislak, Brian Toonen, William Allcock, Silvio Rizzi, Joseph Insley, Venkatram Vishwanath, and Michael E Papka. 2016. Early investigations into using a remote ram pool with the V13 visualization framework. In *In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV), Workshop on. IEEE*, 23–28.

APPENDIX

A. Maximum Application Bandwidth (MaxBW):

$$DPE_{MaxBW}(s, C_i) = \begin{cases} (s/BW_i) * A_i & , s \leq C_i \\ \min \left(\frac{DPE(s, C_{i+1})}{DPE(C_i, C_i) + DPE(s - C_i, C_{i+1})} \right) & , s > C_i \end{cases} \quad (1)$$

where s is the request size, C is a layer's remaining capacity in MBs, i is the current layer, BW is the bandwidth in MB/s, A is the access latency in ms, and $Move(min_size, dest)$ triggers data organizer to recursively move at least min_size data to $dest$ layer.

B. Maximum Data Locality:

$$DPE_{MaxLocality}(s, d, L_i, R_i) = \begin{cases} (s/BW_i) * d & , L_i \& s \leq R_i \\ \min \left(\frac{(s/BW_i) * (d + 1)}{DPE(s, d, L_{i+1}, R_{i+1})} \right) & , !L_i \& s \leq R_i \\ \min \left(\frac{DPE(s, d, L_{i+1}, R_{i+1})}{DPE(R_i, d, L_i, R_i) + DPE(s - R_i, d, L_{i+1}, R_{i+1})} \right) & , s > R_i \end{cases} \quad (2)$$

where s is the request size, d is the degree of data dispersion into DMSH, L is the locality of a dispersion unit in layer (i.e., if it exists in this layer or not), R is a layer's capacity threshold, i is the current layer, BW is the bandwidth in MB/s, and $ReOrganize(min_size)$ is a function that triggers data organizer to recursively move at least min_size data to maintain the locality of a dispersion unit.

C. Hot-data:

$$DPE_{HotData}(s, h, H_i, C_i) = \begin{cases} (s/C_i)/BW & , h \geq H_i \& s \leq C_i \\ \min \left(\frac{DPE(s, h - 1, H_{i+1}, C_{i+1})}{DPE(C_i, h, H_i, C_i) + DPE(s - C_i, h - 1, H_{i+1}, C_{i+1})} \right) & , h \geq H_i \& s > C_i \\ \min \left(\frac{Evict(s - C_i, h, i + 1) + DPE(s, h, H_i, C_i)}{DPE(s, h + 1, H_i, C_i)} \right) & , h < H_i \& s \leq C_i \\ \min \left(\frac{DPE(s, h, H_{i+1}, C_{i+1})}{DPE(C_i, h + 1, H_i, C_i) + DPE(s - C_i, h, H_{i+1}, C_{i+1})} \right) & , h < H_i \& s > C_i \end{cases} \quad (3)$$

where s is the request size, h is the file's hotness score, H is the minimum hotness score present in a layer, C is a layer's remaining capacity in MBs, i is the current layer, BW is the bandwidth in MB/s, and $Evict(min_size, score, dest)$ is a function that triggers data organizer to recursively move at least min_size data to the $dest$ layer with $score$ hotness.