

...And eat it too: High read performance in write-optimized HPC I/O middleware file formats

Milo Polte*, Jay Lofstead†, John Bent‡, Garth Gibson*, Scott A. Klasky§, Qing Liu§, Manish Parashar¶, Norbert Podhorszki§, Karsten Schwan†, Meghan Wingate‡, Matthew Wolf†
*Carnegie Mellon University †Georgia Institute of Technology ‡Los Alamos National Lab
§Oak Ridge National Lab ¶Rutgers University

1. INTRODUCTION

As HPC applications run on increasingly high process counts on larger and larger machines, both the frequency of checkpoints needed for fault tolerance [14] and the resolution and size of Data Analysis Dumps are expected to increase proportionally. In order to maintain an acceptable ratio of time spent performing useful computation work to time spent performing I/O, write bandwidth to the underlying storage system must increase proportionally to this increase in the checkpoint and computation size. Unfortunately, popular scientific self-describing file formats such as netCDF [8] and HDF5 [3] are designed with a focus on portability and flexibility. Extra care and careful crafting of the output structure and API calls is required to optimize for write performance using these APIs. To provide sufficient write bandwidth to continue to support the demands of scientific applications, the HPC community has developed a number of I/O middleware layers, that structure output into write-optimized file formats. However, the obvious concern with any write optimized file format would be a corresponding penalty on reads. In the log-structured filesystem [13], for example, a file generated by random writes could be written efficiently, but reading the file back sequentially later would result in very poor performance. Simulation results require efficient read-back for visualization and analytics, and though most checkpoint files are never used, the efficiency of a restart is very important in the face of inevitable failures. The utility of write speed improving middleware would be greatly diminished if it sacrificed acceptable read performance. In this paper we examine the read performance of two write-optimized middleware layers on large parallel machines and compare it to reading data natively in popular file formats.

The two I/O middleware layer examined in this paper are the Adaptable IO System (ADIOS) [12, 9], a library-based approach developed at Oak Ridge National Laboratory to provide a high-level IO API that can be used in place of netCDF or HDF5 to do much more aggressive write-behind and efficient reordering of data locations within the file; and

the Parallel Log-structured Filesystem (PLFS) [5], a stackable FUSE [1] filesystem approach developed at Los Alamos National Laboratory that decouples concurrent writes to improve the speed of checkpoints. Since ADIOS is an I/O componentization that affords selection of different I/O methods at or during runtime, through a single API, users can have access to MPI-IO, Posix-IO, HDF5, netCDF, and staging methods [4]. The ADIOS BP file format [10] is a new log-file format that has a superset of the features of both HDF5 and netCDF, but is designed to be portable and flexible while being optimized for writing. PLFS, on the other hand, is mounted as a stackable filesystem on top of an existing parallel filesystem. Reads or writes to the PLFS filesystem are transparently translated into operations on per-process log files stored in the underlying parallel filesystem. Since PLFS performs this translation without application modification, users can write in HDF5, netCDF, or app-specific file formats and PLFS will store the writes in a set of efficiently written log-formatted files, while presenting the user with a logical ‘flat’ file on reads. Despite their different approaches, the commonality behind both of these middleware systems is that they both write to a log file format. As shown in previous publications [11, 5], writes are fully optimized in both systems, sometimes resulting in 100x improvements over writing data in popular file formats. But as mentioned above, writing is only one part of the story. In this paper we examine the read performance of our middleware layers on large parallel machines and compare these to reading data either natively or from other popular file formats. We compare the reading performance in two different scenarios: 1) Reading back restarts from the same number of processors as wrote the data and 2) Reading back restart data from a different number of processors.

We observe that not only can write-optimized I/O middleware be built to not penalize read speeds, but for important workloads, techniques that improve write performance can, perhaps counterintuitively, improve read speeds over reading from a contiguously organized file format. In the remainder of this paper, we investigate this further through case studies of PLFS and ADIOS on simulation checkpoint restarts.

2. PLFS

As described above, the Parallel Log-Structured Filesystem (PLFS) is a FUSE-based stackable filesystem that accelerates the write performance of N-1 checkpoints by decoupling the concurrent access of multiple processes writing to a single, shared file into many non-concurrent writes appending to per-process log files.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Supercomputing PDSW'09 Nov. 15, 2009. Portland, OR, USA
Copyright 2009 ACM 978-1-60558-883-4/09/11 ...\$10.00.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE NOV 2009	2. REPORT TYPE	3. DATES COVERED 00-00-2009 to 00-00-2009			
4. TITLE AND SUBTITLE ...And eat it too: High read performance in write-optimized HPC I/O middleware file formats		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University, Parallel Data Laboratory, Pittsburgh, PA, 15213		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	5	

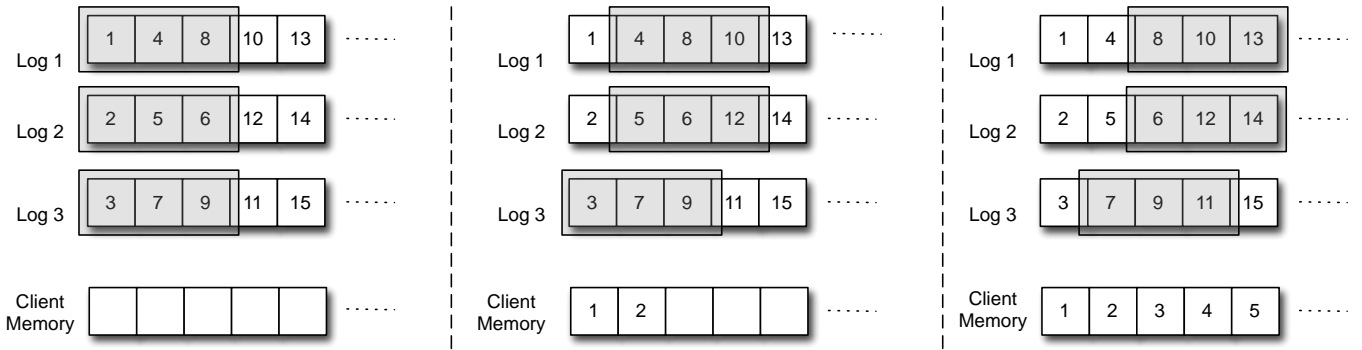


Figure 1: Progress of sequential reading of a checkpoint log file

Filesystems perform well on sequential access to single-writer files compared to strided, concurrent writes to a shared file. This is true even of parallel filesystems; although they spread I/O across multiple servers and spindles, parallel filesystems still must deal with contention within in a single file object or stripe. So it is unsurprising that PLFS’s log-structured writing has achieved large write performance improvements on real applications and checkpoint benchmarks [5]. However, for pathological combinations of write and read patterns (for example, a random write pattern read back sequentially), a log file format is expected to exhibit poor read performance due to additional seeks. Surprisingly, this expected poor read performance was not borne out during PLFS experiments.

To understand why PLFS performs so well on its read path, it helps to understand that while reading sequentially from a log file format consisting of random writes can be slow, the checkpointing and scientific applications that write to PLFS do so in a structured manner. An analysis of the trace repository of checkpoints run through PLFS available on the web [2] reveals two important observations. First, the multiple processes writing a checkpoint do interleave their writes with one another; this is the behavior that allows middleware layers such as PLFS and ADIOS to improve write bandwidth. The second observation is that all writes from each individual process are strictly increasing in their logical offsets.

Figure 1 is a graphical representation of three time steps of a single client reading sequentially from a PLFS checkpoint file created by three writing processes and stored as log structured files on an underlying parallel filesystem. Squares represent segments of client memory or entries in a log file, and the numbers inside correspond to their logical offset. Grey boxes represent the parallel filesystem’s read-ahead buffers for each log. Note that each log file is monotonically increasing in logical offset, and the next request from the client is always at the front of one of the log files. Due to this property, as the client continues reading from the checkpoint file, rather than the log format resulting in expensive seeks, the read ahead buffers will slide forward through the log files, reading them sequentially in parallel. ADIOS’s BP format is somewhat different, but by storing variables together in log formatted files, it too allows for efficient read back in an analogous manner.

Below we discuss how this property allows for efficient read-back of PLFS checkpoint files on uniform and non-uniform restart, by examining one checkpoint benchmark:

The `mpi_io_test` benchmark from LANL [7]. This benchmark is designed to represent a simple checkpoint I/O workload. In the examples below, it is configured to write a single 20 GB file in 47KB strided writes from a varying number of processes on the LANL Roadrunner system. One implication of this workload is that as the number of writers grows, the amount written by each writer (and the corresponding size of the PLFS log file) shrinks. We write both to PLFS and to the underlying parallel filesystem directly and compare the results.

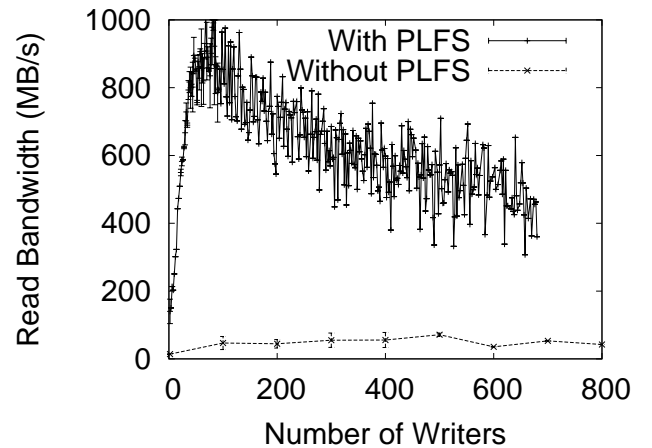


Figure 2: Uniform Reads for PLFS

2.1 Uniform Restart in PLFS

Reading back a checkpoint file with the same number of processes is highly efficient in PLFS. During a checkpoint restart, each reading process will read back one writing process’s data in the checkpoint in the same pattern in which it was written. In the case of PLFS, since the checkpoint was written in per-process monotonically increasing offsets, each reader will access a single log structured file in sequential order. The scenario is similar to that shown in Figure 1, except with multiple clients each reading individual log files allowing the underlying parallel filesystem to make efficient use of read ahead buffers and disks. Results are shown in Figure 2. Performance quickly ramps up, but gradually falls off as the numbers of writers for the 20 GB file increases, the size of each individual log file shrinks, and the benefit of

prefetching is reduced.

By contrast, all processes reading directly from a single file stored on the parallel filesystem elicits poor performance as multiple readers are at any given time all issuing a series of small reads to a relatively narrow window that gradually moves through the file. The result is possible contention within the file, more seeks, and poor utilization of spindles and servers (since a narrow region of the file is generally not spread as widely across disks and servers as multiple log structured files). The result is that the read back speed of readers without PLFS does not scale noticeably with more processes.

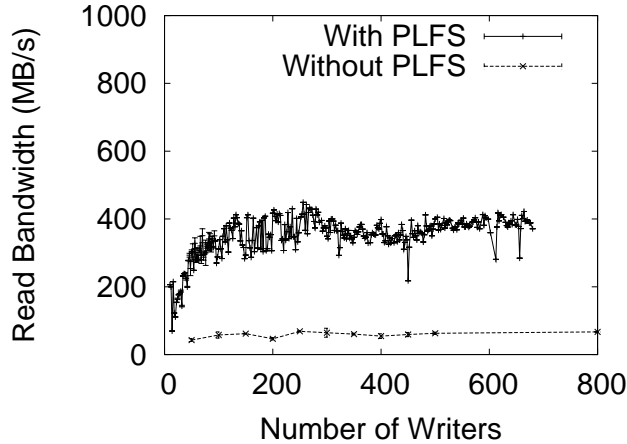


Figure 3: Non-Uniform Reads for PLFS

2.2 Non-Uniform Restart in PLFS

Reading back a checkpoint file with fewer processes is somewhat less efficient in PLFS, as shown in Figure 3. Here a file is written by a varying number of processes and read back by eight fewer reader processes. The situation is similar to the previous section, however, now each log structured file is accessed by multiple processes instead of just one. The result is some additional read contention within each PLFS log file.

Reading directly from a single, flat file stored on the underlying parallel filesystem is analogous to the uniform case, however, and once again we see it exhibit a non-scaling read performance.

3. ADIOS

The ADIOS architecture is designed for portable, scalable performance in two key ways. First, the componentization of the IO implementation affords selecting the highest performance output mechanism for a particular platform without requiring any source code changes. Second, the default BP file format achieves excellent write performance by allowing delaying consistency checks [10]. This enables aggressive buffering and decouples processes during the output process. The output format is organized around the concept of a *Process Group*, a collection of data typically written by a single process, similar in spirit to PLFS’s per-writer logs. Each process is assigned a section of the file to write local data with sufficient annotations to identify the items written and place the portions of global items properly in the

global space. For example, global arrays are written as a local array in the assigned process group with annotation of the global array it is part of, the global dimensions, the local dimensions, and the offsets into the global space this piece represents. Scalars, other local values, and data attributes are stored similarly. This format decouples processes and data reorganization when writing by avoiding constructing a contiguous global array on disk when writing and selecting a slice appropriate for the local process on read, assuming the same number of processes and decomposition are used when reading. The perceived potential penalty for the gain in write performance for this decision is the inability to read the data back in efficiently. For restarts of a p process output read back in on p processes, this is an optimal format. Each process reads the data from exactly one process group independently and requires no data reorganization. For other process configurations, the performance story was unknown. This evaluation resolves the restart performance question.

3.1 Evaluation

Four different experiments are performed on one of the most difficult data decompositions, a 3-D domain decomposition. For this application, the simulation domain is a 3-D space divided into boxes along the X, Y, and Z dimensions with each process responsible for one rectangular area. For netCDF and HDF5 formatted data, this output is reorganized into a contiguous format for the entire 3-D domain with selective reading on restarts. The BP format stores each rectangular area as a single block in a process group with each area in a non-contiguous space in the file. In order to read the entire 3-D domain from a single process, in order, the system would need to skip around the file or read in chunks reorganizing in memory. An IO kernel for the Pixie3D MHD code [6] is tested.

Pixie3D has three modes of operation tested for this study. The Small configuration consists of 256 KB of data per process, Medium has 16 MB per process, and 128 MB per process for Large. The data is divided up evenly among 8 variables and also contains many small scalar values. To test the restart read performance, various process counts from 128 to 2048 are employed to write the data with the same or one-half the number of process used to write to read the data back in. netCDF formatted data is compared.

For all tests, the best results for a series of four runs for each data point is selected for the BP and netCDF performance. The horizontal axis represents the number of writers employed to generate the restart data no matter the number of processes used to read the data back in.

The experiments are performed on Jaguar, the Cray XT4 at ORNL, striping across all 144 Lustre storage targets and using a stripe size of 1 MB.

3.2 Pixie3D Uniform Restarts

To establish a baseline, a uniform restart is tested. This is using the same number of processes to read the restart as wrote it originally. Figure 4 shows the performance results. For large data, the performance approaches the IOR benchmark performance [15] for the machine.

3.3 Pixie3D Small

For the small data, good read performance just cannot be attained no matter the number of processes used to read or write the data. The overall data size is too small to overcome

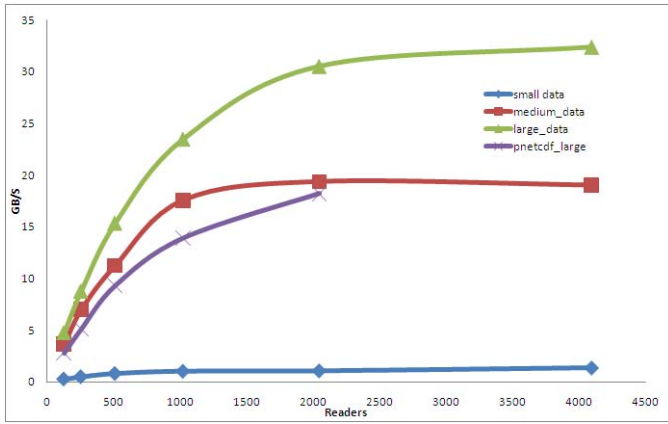


Figure 4: Uniform Reads for the Pixie3D data

the inherent overhead of the parallel file system. Figure 5 shows the performance for reading the restart output on half as many processes as wrote the restart. The horizontal axis represents the number of processes that wrote the data originally.

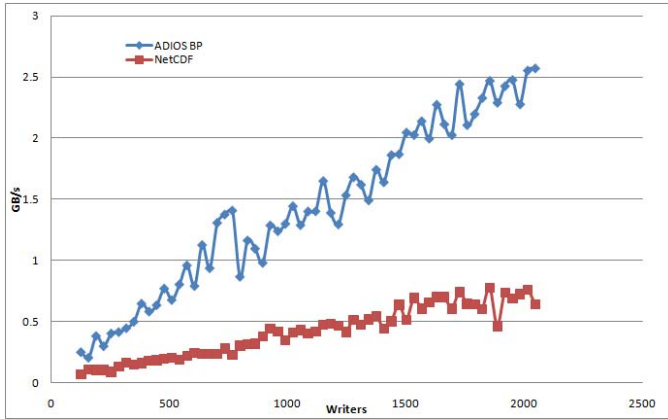


Figure 5: Small Model, Half Process Count

The BP formatted data was able to be read faster than the netCDF formatted data. The trendlines for the performance clearly show the performance gap should continue to widen as the process count increases.

3.4 Pixie3D Medium

Restarts using the medium data model can achieve much better performance, but still cannot achieve more than a fraction of the theoretical maximum performance for the system. Figure 6 shows the performance for restarting using half as many processes.

For the ‘half’ case, the BP performance consistently outperforms netCDF with the performance gap narrowing slightly as the process count increases.

3.5 Pixie3D Large

The large data cases finally reach the maximum general performance seen for applications in production use. Figure 7 shows the performance for using half as many processes to restart.

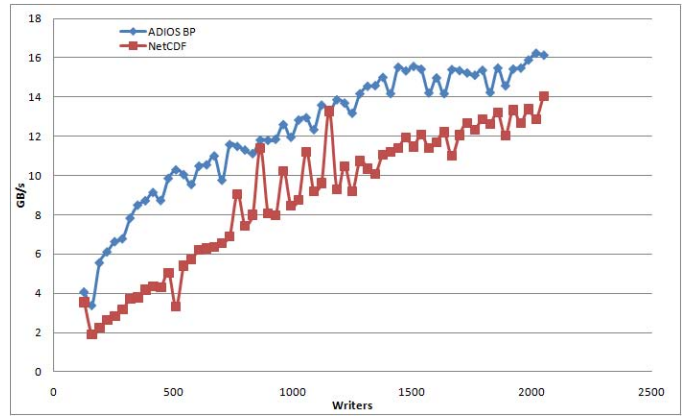


Figure 6: Medium Model, Half Process Count

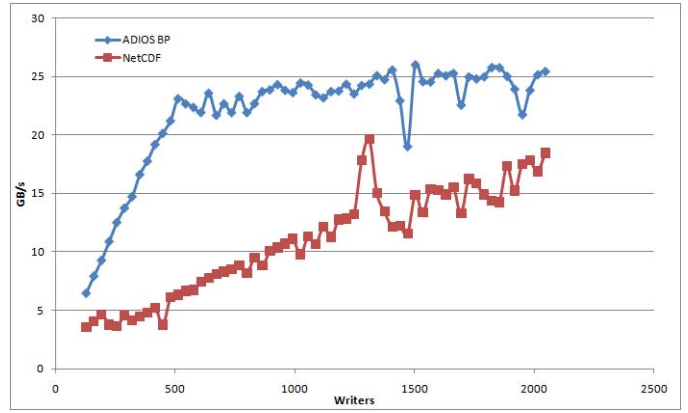


Figure 7: Large Model, Half Process Count

3.6 Discussion

Overall, the performance for all configurations of BP data is either absolutely better or about the same as a contiguous format like netCDF. Comparing the half-processes restarts with the uniform restarts, the performance for large and medium data is about 80% of the uniform read rate. Small data is about the same performance. Though we do not have room to show it here, we also tested ADIOS read performance on additional domain decompositions, reading back with different numbers of readers (including restarts involving more reading processes than writers). In most cases we were 2x faster compared to reading directly from netCDF, and read speed was always superior on non-uniform restarts.

An additional set of tests using a non-integer factor difference between the writers and readers yields similar results. For example, using 64 writers and 80 readers represents moving from a 4x4x4 setup to 4x4x5. For small data, BP is consistently 2x faster compared with netCDF, for medium, it is 20% faster, and for large data, the performance is essentially identical. The detailed results were omitted for space reasons.

4. CONCLUSION AND FUTURE WORK

Previously, both ADIOS’s BP format and PLFS’s log-based format have shown excellent performance for writing data due to reorganizing the data when writing. This

generated two main open questions. First, what is the performance for reading all of the data, such as for a restart. Second, for reading a subset of the data, such as for a typical analysis task, how much does the log-based format penalize the user. As has been shown, both of these log-based formats do not suffer from the expected penalties for storing the data in the log-based formats when reading restarts. This holds not just for reading on the same number of processes as wrote the data, but also when fewer read the data back in.

When reading on the same number of processes as generated the data, no reorganization effort is required for any process to read the desired data and the data is stored in larger, single chunks on the various storage targets. Each process can read large, contiguous blocks of data from the file reducing the likelihood of interfering with other processes reading other data. Restarting on a different number of processes is more challenging and requires a deeper analysis.

The belief that a canonical storage format is most efficient for reading is based on the notion that most often, contiguous chunks of data will be read. The closer that data is stored together, the better the overall performance. While this may have been true on a single spindle with a single reader, with parallel processes and file systems, a different consideration must be made. Parallel file systems rely on large, contiguous reads or writes to all storage targets in parallel for optimal performance. Although neither BP nor the PLFS format is necessarily directly knowledgeable about the underlying storage organization, the formats are designed understanding that the file system employed will use the strategy of striping the data across many storage targets while performing fewer, larger operations to attain high performance on these systems.

The second question of the performance for data analysis tasks is still open. These results demonstrate that a log-based format has viable performance for restarts of various data sizes and process counts but not necessarily superior for read workloads overall. Particularly missing is evidence that the log-based format works well for analysis patterns. For example, though it still increases write speeds dramatically, without domain knowledge of the variables being written PLFS is slightly penalized when reading back Pixie3D data, unlike ADIOS. Both systems may need to be enhanced to handle analytic read patterns that differ drastically from write and restart patterns. Work is ongoing to evaluate how the restart performance relates to the analysis task performance.

5. ACKNOWLEDGEMENTS

The work in this paper is based on research supported in part by the Department of Energy under award number DE-FC02-06ER25767, the Department of Energy under Contract No. DEAC05-00OR22725 at Oak Ridge National Laboratory, the resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, by the National Science Foundation, under award number OCI-0852543, by the Los Alamos National Laboratory and the Institute for Reliable High Performance Information Technology, under contract number 54515-001-07, and by Sandia National Laboratories under contract DE-AC04-94AL85000, a grant from NSF as part of the HECURA program, a grant from the Department of Defense, a grant from the Office of Science through the SciDAC program, and the

SDM center in the OSCR office, and CPES in the OFES and OASCR office. We also thank the member companies of the PDL Consortium (including APC, DataDomain, EMC, Facebook, Google, Hewlett-Packard, Hitachi, IBM, Intel, LSI, Microsoft, NEC, NetApp, Oracle, Seagate, Sun, Symantec, and VMware) for their interest, insights, feedback, and support.

6. REFERENCES

- [1] FUSE: Filesystem in Userspace. <http://fuse.sourceforge.net/>.
- [2] PLFS I/O Map Repository. <http://institute.lanl.gov/plfs/maps/>.
- [3] The HDF Group. <http://www.hdfgroup.org/>.
- [4] H. Abbasi, J. Lofstead, F. Zheng, S. Klasky, K. Schwan, and M. Wolf. Extending i/o through high performance data services. In *Cluster Computing*, Austin, TX, September 2007. IEEE International.
- [5] J. Bent, G. Gibson, G. Grider, B. McClelland, P. Nowoczynski, J. Nunez, M. Polte, and M. Wingate. Plfs: A checkpoint filesystem for parallel applications. *SC Conference*, 0, 2009.
- [6] L. Chacón. A non-staggered, conservative, $\nabla s B \Rightarrow 0$, finite-volume scheme for 3D implicit extended magnetohydrodynamics in curvilinear geometries. *Computer Physics Communications*, 163:143–171, Nov. 2004.
- [7] Gary Grider and James Nunez and John Bent. LANL MPI-IO Test. <http://institutes.lanl.gov/data/software/>, jul 2008.
- [8] H. L. Jenter and R. P. Signell. NetCDF: A Public-Domain-Software Solution to Data-Access Problems for Numerical Modelers, 1992.
- [9] O. R. N. Laboratory. <http://adiosapi.org/>.
- [10] J. Lofstead, F. Zheng, S. Klasky, and K. Schwan. Input/output apis and data organization for high performance scientific computing. In *In Proceedings of Petascale Data Storage Workshop 2008 at Supercomputing 2008*, 2008.
- [11] J. Lofstead, F. Zheng, S. Klasky, and K. Schwan. Adaptable, metadata rich io methods for portable high performance io. In *In Proceedings of IPDPS'09, May 25-29, Rome, Italy, 2009*.
- [12] 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 *CLADE '08: Proceedings of the 6th international workshop on Challenges of large applications in distributed environments*, pages 15–24, New York, NY, USA, 2008. ACM.
- [13] M. Rosenblum and J. K. Ousterhout. The design and implementation of a log-structured file system. *ACM Transactions on Computer Systems*, 10:1–15, 1991.
- [14] B. Schroeder and G. A. Gibson. Understanding failures in petascale computers. *Journal of Physics: Conference Series*, 78:012022 (11pp), 2007.
- [15] H. Shan and J. Shalf. Using ior to analyze the i/o performance for hpc platforms. Cray User's Group, 2007.