

Pangeo Benchmarking Analysis: Object Storage vs. POSIX File System

Haiying Xu

National Center for Atmospheric Research
Boulder, USA
haiyingx@ucar.edu

Kevin Paul

National Center for Atmospheric Research
Boulder, USA
kpaul@ucar.edu

Anderson Banihirwe

National Center for Atmospheric Research
Boulder, USA
abanihi@ucar.edu

Abstract—Pangeo is a community of scientists and software developers collaborating to enable Big Data Geoscience analysis interactively in the public cloud and on high-performance computing (HPC) systems. At the core of the Pangeo software stack is (1) Xarray, which adds labels to metadata such as dimensions, coordinates and attributes for raw array-oriented data, (2) Dask, which provides parallel computation and out-of-core memory capabilities, and (3) Jupyter Lab which offers the web-based interactive environment to the Pangeo platform. Geoscientists now have a strong candidate software stack to analyze large datasets, and they are very curious about performance differences between the Zarr and NetCDF4 data formats on both traditional file storage systems and object storage. We have written a benchmarking suite for the Pangeo stack that can measure scalability and performance information of both input/output (I/O) throughput and computation. We will describe how we performed these benchmarks, analyzed our results, and we will discuss the pros and cons of the Pangeo software stack in terms of I/O scalability on both cloud and HPC storage systems.

Index Terms—Pangeo, benchmark, cloud, HPC, object store, throughput, I/O

I. INTRODUCTION

Geoscientists confront many challenges when attempting to process Big Data. For example, climate scientists have generated petabytes of data as part of the Coupled Model Intercomparison Project (CMIP6) [1]. Such datasets are so large they become difficult to deal with when the volume of data exceeds computer memory or storage capacity. Under such circumstances, scientists tend to access and analyze only a small percentage of the total data, instead of analyzing the entire dataset. How to make this Big Data accessible and easy to analyze by scientists is the goal of the Pangeo community [2]–[5].

The core components of the Pangeo software stack [6] include Xarray [7], Dask [8], and Jupyter [9]. Xarray is an open source Python package for computing with labeled multi-dimensional arrays. Xarray employs the same Common Data Model that is used by NetCDF [10], which is a self-describing scientific data format and is widely used in the geosciences.

Xarray’s labeled array structure is used by many other Pangeo software packages and many computation and vi-

sualization tools. Dask provides a flexible mechanism for parallel data analytics. Dask is used internally within Xarray to provide scalability in the Pangeo platform. Dask’s lazy execution of code allows for delaying the loading of extremely large datasets into memory until the final necessary step, and its distributed schedulers allow supercomputers and cloud computing clusters to easily perform parallel computations in batch mode and interactively with Jupyter notebooks. Xarray’s data structures implement a layer around NumPy [28] (in-memory) or Dask (out-of-memory) arrays. Jupyter provides the interactive development environment to the Pangeo platform, allowing scientists to interactively run HPC jobs, read and write large files and plot results.

Our Pangeo benchmarking suite uses Xarray to read and write Dask arrays from/to NetCDF files or Zarr stores and uses Dask to schedule and run benchmarking jobs on high performance computing clusters (HPC) and cloud compute platforms, such as AWS [12]. The benchmarks are run interactively via Jupyter Notebook so as to provide a great user experience during debugging and visualization of results. Our benchmarking suite reads and writes both NetCDF [10] and Zarr [11] data to measure I/O throughput performance. NetCDF is both a data format and a collection of data access libraries for self-describing, array-like data, commonly used by geoscientists to create, read/write, and share scientific data. NetCDF is developed and maintained by Unidata [10] and is widely used in areas such as climatology, meteorology, oceanography and GIS, for weather forecasting and climate change research. Zarr is both a data format and an open source Python package delivering data access functions for chunked, compressed, multi-dimensional arrays. Zarr arrays can be saved as a key-value store on object storage systems and as a hierarchical directory store on POSIX [18] file systems.

With Dask, the Pangeo stack can load small chunks of data from 100GB+ datasets, providing a low-memory footprint that enables scientists to potentially analyze large datasets on their laptop. Similarly, Dask provides an easy mechanism for parallelizing I/O, regardless of the data format (Zarr or NetCDF). The question we have asked in motivating this

benchmarking study is, “When does it make sense to use NetCDF or Zarr?” Is Zarr only useful on object storage systems? Or is it also performant on POSIX-based file storage systems? And which data format is best on which storage system?

II. BENCHMARKING VARIABLES

We set up several varying parameters for our tests: the storage system (POSIX or object), the data format (NetCDF or Zarr), Dask chunk size S_{chunk} , and Dask cluster size (i.e., number of HPC nodes) N_{node} .

A. Storage system: object store vs. POSIX-compliant file systems

The HPC system we used in our benchmarks was the Cheyenne supercomputer at the National Center for Atmospheric Research’s (NCAR’s) [16] Wyoming Supercomputing Center (NWSC) [17]. It features 145,152 Intel Xeon processor cores in 4,032 dual-socket nodes (36 cores/node) and 313 TB of total memory. The traditional POSIX-compliant [18] data storage system was provided by DataDirect Networks (DDN) [19] and transfers data at the rate of 200 GBps and has a total usable capacity of 38PB. It uses a high-performance GPFS [29] shared file system.

Additionally, an ActiveScale object storage system from Quantum [20] has also recently been deployed at the NWSC. This object storage [14] system has 5PB of usable capacity and can transfer single-stream data at 160 MBps (8 GBps aggregate total) and is AWS S3 compliant. S3 object storage is essentially a key-value storage system. The keys are unique strings, and the values are bytes of data. Data is read and written using HTTPS calls. Object storage does not follow a hierarchical model or use a directory tree. Instead, every discrete unit of data is collected in a single flat repository. This unique design makes object storage more scalable, meaning that we can make thousands of concurrent reads and writes easily. For instance, AWS S3 provides 3500 PUT and 5500 GET requests per second for each prefix of a bucket [22].

We employed the `fsspec` Python package [21] in our benchmarks to abstract the storage layer. For example, `fsspec` uses the same API to open a NetCDF file on POSIX or S3. The `fsspec` package is a Python library for working with different storage technologies and to manipulate remote file systems such as `s3fs` [24] (an interface to S3) or `gcsfs` [23] (an interface to Google Cloud Storage). Thus, our benchmarking suite is easy to extend to different cloud storage systems. Since all of our test data is local on the same HPC system, we need not download or transfer data. Instead, we can directly measure the system’s read/write performance from/to the storage systems. See Figure 1 for these two system configurations.

B. Data format: NetCDF vs. Zarr

With these two different storage systems, we tested two data formats: NetCDF and Zarr. Figure 2 illustrates the structure of a Zarr store persisted on a POSIX filesystem. With the chunked data design, Zarr is an ideal format for object storage.

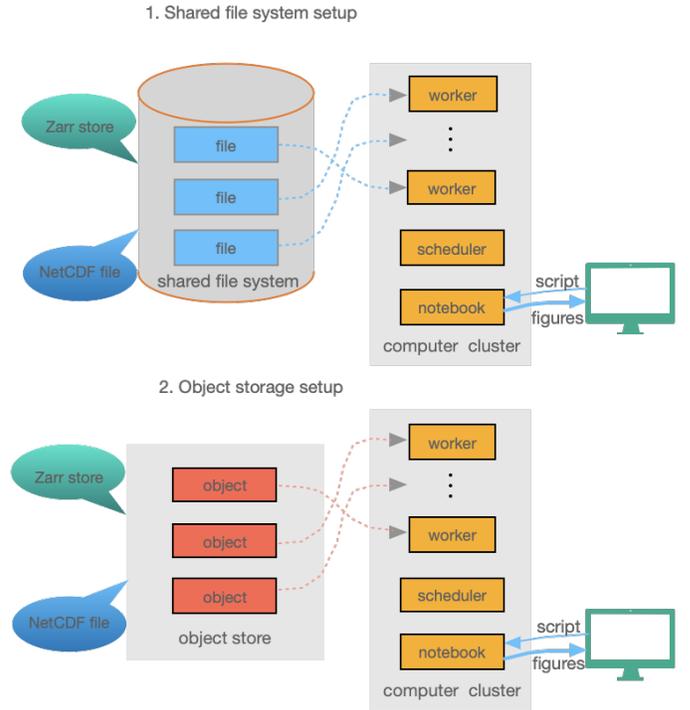


Fig. 1: Storage system configurations

Xarray (with Dask) can read and write both formats on a POSIX file system. On S3, Xarray (with Dask) can read and write Zarr, and it can read (not write) NetCDF. Since writing NetCDF data directly to S3 is currently not supported, we generate the NetCDF dataset on the POSIX system and then upload the data to S3 storage for testing. Also, since it takes a very long time to write a NetCDF file exceeding 5GB, we used Xarray’s `save_mfdataset` function to write multiple NetCDF files in parallel [15], improving the NetCDF write throughput significantly. In this study, we always generated eleven NetCDF files for this purpose.

C. Dask chunk size

We wrote this benchmarking suite not to depend on existing data. Instead, we generated data on the fly using Xarray to create a Dataset with 1 data variable `sst` having 3 dimensions: `time`, `lon` (longitude), and `lat` (latitude). We then randomly generated the values of the data array (`sst`). We chunked the `sst` data array as shown in Figure 3. Each chunk size S_{chunk} ranges from 64MB to 768MB, and the number of chunks N_{chunk} is computed according to:

$$N_{chunk} = \lceil S_{chunk} / (lon * lat * typesize) \rceil, \quad (1)$$

where `typesize` is the size (in bytes) of the data variable datatype. Thus, the chunks displayed in Figure 3 have the following shape:

$$chunks = (N_{chunk}, lon, lat). \quad (2)$$

By varying the chunk size, we can find out which in-memory chunk size provides the best read/write performance.

D. The number of HPC nodes

We vary the number of HPC nodes used in these benchmarks to evaluate how the I/O throughput scales. We create a Dask cluster with Dask workers across multiple processors. To do this, we use Dask Jobqueue [26] to create a Dask cluster on our HPC system using the PBS [27] job scheduler. Our benchmarking suite varies the size of the Dask cluster from 1 to 12 whole Cheyenne compute nodes (i.e., exclusive access to all 36 cores on a single node) to ensure no other jobs could run on the node. While each Cheyenne compute node has 36 cores, only 10 cores are used for Dask workers, one thread per worker, allowing approximately 11GB of memory per worker (109GB of memory per node).

III. BENCHMARKING STANDARD AND SETUP

We have chosen the read throughput and write throughput as our benchmarking standards. We developed our benchmarking suite to perform both weak and strong scaling studies of the Pangeo stack to obtain them.

A. Weak scaling study

Our weak scaling study measured the read/write throughput as the number of nodes varied for a fixed *dataset size per processor*. The scaled read/write throughput is calculated as the total dataset size divided by the total run time. Ideally, the throughput is expected to scale linearly with respect to the number of cores with a positive slope smaller than one.

In this study, we fixed the chunk size S_{chunk} , using 10 chunks per worker and 10 workers per node, and varied the total dataset size S_{total} with the number of nodes N_{node} as shown here:

$$S_{total} = S_{chunk} * N_{node} * 10 * 10 . \quad (3)$$

We performed four different weak scaling studies using 4 different fixed chunk sizes S_{chunk} of 64, 192, 384 and 786 MB. For each study, the number of nodes N_{node} was changed over the values 1, 2, 3, 6, and 12. In total, 120 different configurations of alternative operation (read or write), data format, storage system, node count, and chunk size were performed as part of the weak scaling study. We ran each configuration 3 or 4 times to obtain a total of 391 runs on the Cheyenne supercomputer.

B. Strong scaling study

Our strong scaling study measured the read throughput (no write) as the number of nodes varied for a fixed *total dataset size*. From this study, we can ascertain an upper limit on the scaling of object storage for the two different data formats by measuring the total read time for each configuration of the study. In this analysis, we fixed the total dataset size S_{total} at 460 GB. The number of nodes N_{node} was changed over 1, 2, 3, 6, and 12, and the chunk size S_{chunk} was varied across 64, 192, 384 and 768 MB, with the number of chunks N_{chunk} computed from Equation (1). The total number of

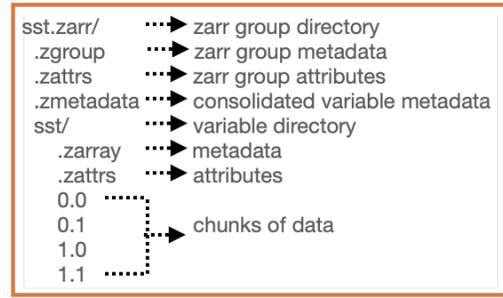


Fig. 2: A Zarr Store on POSIX

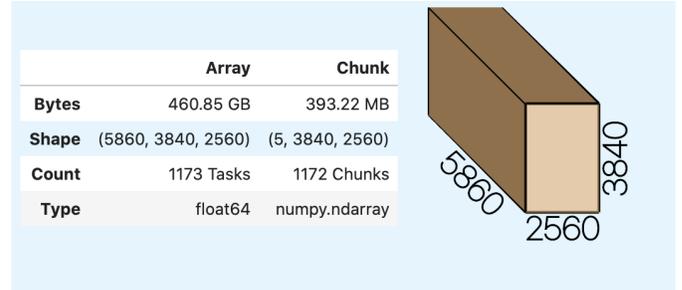


Fig. 3: Dask chunks

chunks N_{total} was computed according to (4), such that the total dataset size stayed constant.

$$N_{total} = S_{total} / (N_{chunk} * lon * lat * typesize) . \quad (4)$$

In total, 80 configurations of varying data format, storage system, node count, and chunk size were performed as part of the strong scaling study. We ran each configuration 4 to 7 times to obtain a total of 489 runs on the Cheyenne supercomputer.

To fine tune Zarr reading, we set `use_listings_cache=True` in `fsspec.filesystem`, which reduced read times. Also, we set `skip_instance_cache=True` to make sure no file data caching was done during benchmarking.

IV. PERFORMANCE ANALYSIS AND DISCUSSION

A. Weak scaling analysis

The weak scaling study results for the read and write operations on Cheyenne are shown in Figures 4 and 5, respectively.

The read study shows that NetCDF on POSIX is roughly 20% faster than Zarr on POSIX, and both are 10 times better than Zarr on S3, which is almost certainly caused by the S3 store having a maximum 8 GBps aggregate stream while the POSIX system has a maximum 200 GBps transfer rate. The best read scaling for Zarr on the POSIX system was found using a 768MB chunk size and is shown in the red lines on Figures 4(b). These lines show roughly linear scaling with the number of nodes, as we expect.

The write study shows that Zarr on the POSIX system has the best performance and scales linearly with increasing node count (Figure 5(b)). Due to Zarr's one-file-per-chunk format, Zarr on the POSIX system is 10% faster than NetCDF on the POSIX system with 1 node, and 9 times faster with 12

nodes (Figures 5(b) and 5(a) red lines). Zarr on S3 is 50% slower on 1 node than NetCDF, and 45% faster on 12 nodes (Figures 5(c) and 5(a)). Our study shows that Zarr write performance scales better than NetCDF write performance on either storage system, which makes Zarr look like an appealing option for actual model output and not simply data analysis.

Note that the optimal performance with Zarr is obtained with a chunk size between 384MB and 768MB, and scaling with these chunk sizes is better than with other smaller chunk sizes.

Large chunk sizes reduce the total number of chunks leading to less communication overhead but larger memory usage.

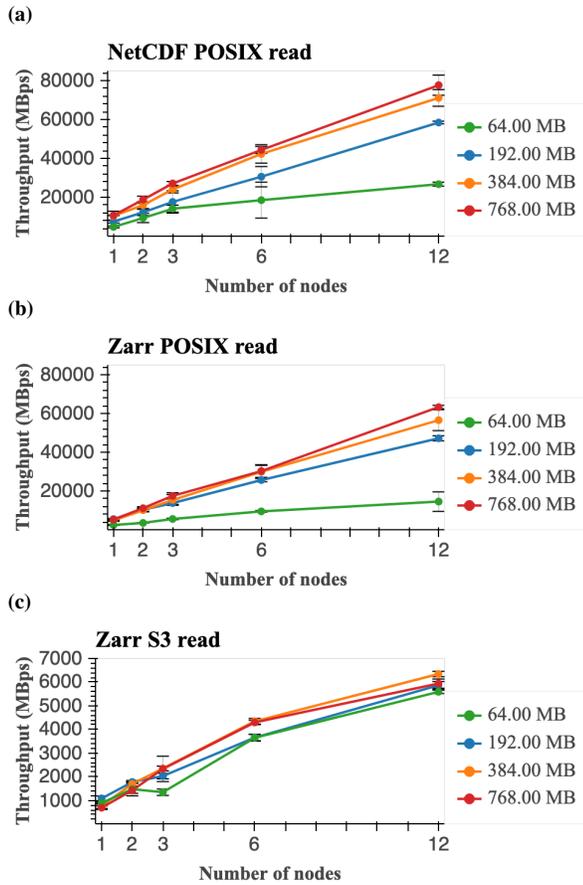


Fig. 4: Weak scaling analysis results for *read throughput* using the Cheyenne supercomputer. The *x*-axis displays increasing number of nodes for each test. The *y*-axis shows the corresponding values of *read throughput*. The chunk sizes of 64, 192, 384, and 786 MB are shown with green, blue, orange and red colors, respectively. The solid lines connect the median *read throughput* on varying number of nodes, and the error bars indicate a single standard deviation from the mean value. From the top, Figures (a), (b) and (c) show NetCDF on POSIX, Zarr on POSIX and Zarr on S3, accordingly. Note: Figure (c) ranges in 0 to 7000 MBps which is different from Figures (a) and (b).

B. Strong scaling analysis

The strong scaling study results for the read operation on Cheyenne with the POSIX storage system and S3 object storage system are shown in Figures 6 and 7, respectively.

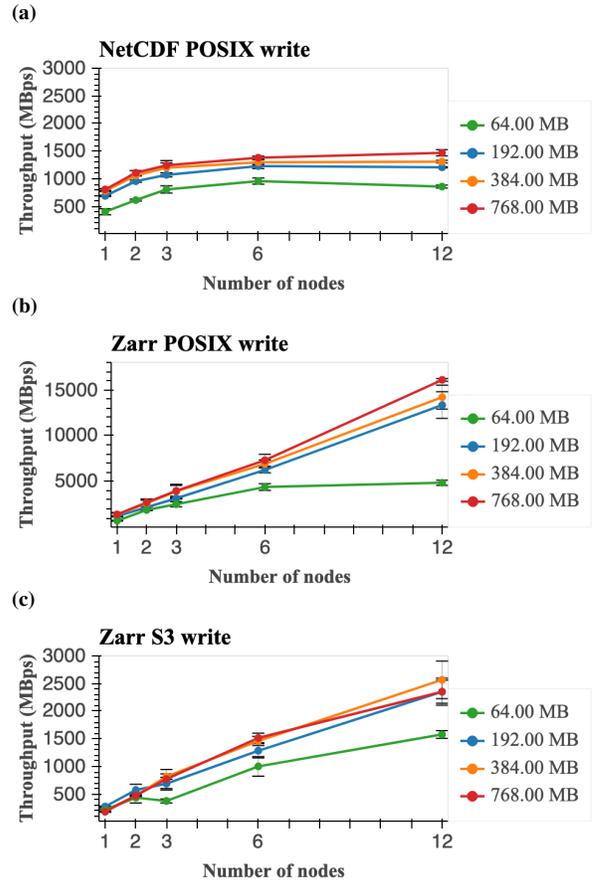


Fig. 5: Weak scaling analysis results for the *write throughput* using the Cheyenne supercomputer. The *x*-axis displays increasing number of nodes for each test. The *y*-axis shows the corresponding values of *write throughput*. The chunk sizes of 64, 192, 384, and 786 MB are shown in green, blue, orange and red colors, respectively. The lines connect the medians of *write throughput* on varying number of nodes, and the error bars indicate a single standard deviation from the mean value. From the top, Figures (a), (b) and (c) show NetCDF on POSIX, Zarr on POSIX, Zarr on S3 accordingly. Note: Figure (b) ranges in 0 to 16000 MBps which is different from Figures (a) and (c).

Reads of NetCDF data from POSIX are 0–45% faster than reads of Zarr on POSIX. The comparison between the two data formats with chunk size 64 MB on 1 nodes gives the 45% difference, 0% occurred at chunk size 64 MB on 12 nodes, and average 30% on 6 and 15% on 12 nodes for any chunk size. This indicates NetCDF has better performance on POSIX systems, but Zarr scales better as shown in Figures 6(a) and 6(b). Another observation is that Zarr test results have larger standard deviations, tests on the daytime and nighttime can have quite different outcomes because the available network bandwidth varies largely in a day. If we want to use the Zarr format on POSIX systems, we will likely need greater network bandwidth.

Reads of NetCDF data from S3 are 3% slower to 13% faster than reads of Zarr data on S3 with 1 node, roughly 6% faster on 6 nodes, and then 8% faster with 2 and 12 nodes as shown in Figures 7(a) and 7(b). Overall, there is no significant

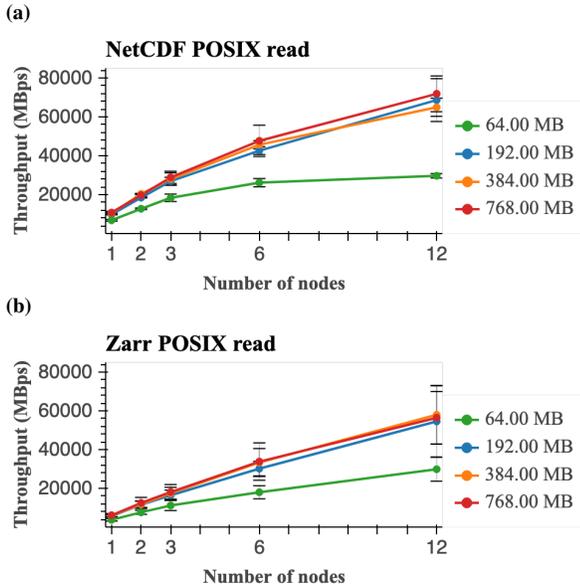


Fig. 6: Strong scaling analysis results for *POSIX read throughput* using Cheyenne supercomputer. The x axis displays increasing number of nodes for each test. The y axis shows the corresponding values of *read throughput*. The chunk sizes of 64, 192, 384, and 786 MB are in the green, blue, orange and red colors respectively. The lines connect the medians of *read throughput* on varying number of nodes, and the error bars indicate a single standard deviation from the mean value. From the top, figures (a) and (b) show the NetCDF on POSIX, Zarr on POSIX, accordingly.

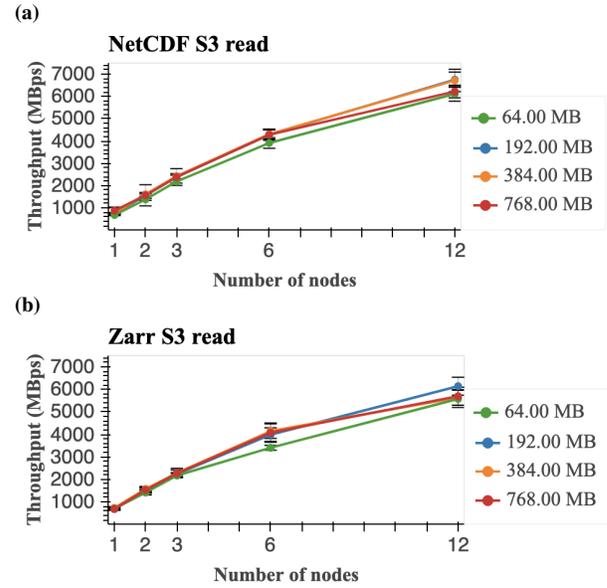


Fig. 7: Strong scaling analysis results for *S3 read throughput* using Cheyenne supercomputer. The x axis displays increasing number of nodes for each test. The y axis shows the corresponding values of *read throughput*. The chunk sizes of 64, 192, 384, and 786 MB are in the green, blue, orange and red colors respectively. The lines connect the medians of *read throughput* on varying number of nodes, and the error bars indicate a single standard deviation from the mean value. From the top, figures (a) and (b) show the NetCDF on S3, Zarr on S3, accordingly.

difference between these two test sets. Again, this could be due to S3 storage having 8 GBps aggregate transfer rate at most. Dask's parallelism scales well with both data formats until using 6 nodes, starting to saturate when using 12 nodes, and NetCDF appears to saturate sooner than Zarr. A chunk size of 192 MB was found to be the optimal chunk size during this study.

V. CONCLUSIONS AND FUTURE WORK

The Pangeo community is working to provide geoscientists with a powerful software stack to generate and analyze data in parallel. However, scientists are concerned about whether this software stack and I/O performance with the Zarr format is comparable to traditional analysis methods with NetCDF. Additionally, scientists are concerned about what to expect from object storage systems in comparison with traditional POSIX filesystems, especially as cloud-based analysis becomes increasingly popular. We have made a first attempt at addressing these concerns with this Pangeo I/O benchmarking study.

Our weak scaling analysis shows that write rates with Zarr are 9 times better than those for NetCDF on POSIX systems, since Zarr's chunked data format with Dask improves parallelism and obtains the optimal scaled write speedup. While we could have parallelized over more NetCDF files, the case considered is more common, aggregating multiple data chunks in a single NetCDF file. When using fewer processor counts, read rates with Zarr on POSIX are slower than those

for NetCDF on POSIX. With increasing processor counts, Zarr's read rates are almost identical to those of NetCDF. Read and write rates of Zarr on S3 object store show good scaling as well, but are much slower than those for Zarr on our POSIX system. This is likely limited due to the slower data transfer rate from NCAR's S3 object store. The strong scaling analysis further indicates that read throughput of NetCDF and Zarr on S3 are identical. This can be seen as reassurance for geoscientists that the read performance will remain intact, if not better, if they switch from NetCDF to Zarr. Numerous benefits come with such a switch, including parallel output with both lossy and lossless compression and flexible storage APIs. The Zarr format within the Pangeo software stack will be definitely beneficial to the geoscience community.

In the future, we intend to run this benchmarking suite on available cloud platforms with high throughput scalable object storage, such as on Amazon AWS and Google Cloud Platform. We will also study asynchronous I/O with NetCDF and Zarr to check if it can improve the I/O performance.

Currently, the code for these benchmarks [25] is available at GitHub and can be executed with the appropriate Conda environment from a clone of the git repository. In the future, the suite will use a Docker container for easy use.

ACKNOWLEDGMENT

Dr. Paul and Mr. Banihirwe were both supported by NSF Earthcube award 1740633.

REFERENCES

- [1] World Climate Research Programme, "CMIP Phase 6", World Climate Research Programme, <https://www.wcrp-climate.org/wgcm-cmip/wgcm-cmip6> (accessed on Aug. 10, 2020).
- [2] Pangeo Team, "Pangeo: A community platform for Big Data geo-science," Pangeo Team, <https://pangeo.io> (accessed Jul. 30, 2020).
- [3] R. Abernathey, K. Paul, J. Hamman, M. Rocklin, C. Lepore, M. Tippet, et al., "Pangeo NSF Earthcube Proposal," Aug. 31, 2017. Accessed on: Aug. 15, 2020. [Online]. Available: <https://doi.org/10.6084/m9.figshare.5361094.v1>
- [4] N.H. Robinson, J. Hamman, R. Abernathey, "Science needs to rethink how it interacts with big data: Five principles for effective scientific big data systems," *ArXiv*, vol. abs/1908.03356, 2019.
- [5] G.E. Bontemps, J. Hamman, "The PANGEO big data ecosystem and its use at CNES," in *Proc. of 2019 Big Data from Space*, Munich, Germany, 2019, pp. 49–52. [Online]. Available: <https://doi.org/10.2760/848593>
- [6] Pangeo Team, "Packages - Pangeo documentation," Pangeo Team, <http://pangeo.io/packages.html> (accessed on Jul. 15, 2020).
- [7] Xarray Developers, "xarray: N-D labeled arrays and datasets in Python," xarray.pydata.org, <http://xarray.pydata.org/en/stable/> (accessed on Jul. 20, 2020).
- [8] Dask core Developers, "Dask: Why Dask?," [dask.org https://docs.dask.org/en/latest/why.html](https://docs.dask.org/en/latest/why.html) (accessed On Jul. 20, 2020).
- [9] Project Jupyter, "Project Jupyter — Home," [jupyter.org, https://jupyter.org/](https://jupyter.org/) (accessed on Aug. 17, 2020).
- [10] Unidata, "Network Common Data Form (netCDF) version 4.5.3," Boulder, CO: UCAR/Unidata, Jun. 2020. [Online]. Available: <http://doi.org/10.5065/D6H70CW6>
- [11] A. Miles, J. Kirkham, M. Durant, J. Bourbeau, T. Onalan, J. Hamman, etc., "zarr-developers/zarr-python: v2.4.0," Jan. 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.3773450>
- [12] J. Polzehl and V. Spokoyny, "Propagation-separation approach for local likelihood estimation," *Probab. Theory Related Fields*, vol. 135, no. 3, pp. 335-362, 2006.
- [13] M. P. Forum, "MPI: A Message-Passing Interface Standard," University of Tennessee, USA, Tech. Rep. 10.5555/898758, 1994.
- [14] M. Factor, K. Meth, D. Naor, O. Rodeh, J. Satran, "Object storage: The future building block for storage systems," in *Proc. Local to Global Data Interoperability Conf. (LGD12005)*, Jun. 20-24, 2005, pp. 119 - 123.
- [15] A. Banihirwe, "Writing multiple netCDF files in parallel with xarray and dask," [ncar.github, https://ncar.github.io/xdev/posts/writing-multiple-netcdf-files-in-parallel-with-xarray-and-dask/](https://ncar.github.io/xdev/posts/writing-multiple-netcdf-files-in-parallel-with-xarray-and-dask/) (accessed Jun. 30, 2020).
- [16] UCAR, "National Center for Atmospheric Research," <https://ncar.ucar.edu>
- [17] Computational and Information Systems Laboratory. *Cheyenne: HPE/SGI ICE XA System (Climate Simulation Laboratory)*. Boulder, CO: National Center for Atmospheric Research. doi:10.5065/D6RX99HX. 2019.
- [18] IEEE, Institute of Electrical and Electronics Engineers, Inc., *IEEE Standard for Information Technology - Portable Operating System Interface (POSIX): System Application Program Interface (API), Amendment 1: Realtime Extension (C Language), IEEE Std 1003.1b-1993*. USA: IEEE Standards Office, 1994
- [19] DataDirect Networks, "DataDirect Networks (DDN) Storage," DataDirect Networks, <https://www.ddn.com/> (accessed Aug. 15, 2020).
- [20] Quantum Corp. "ActiveScale Object Storage — Quantum," [quantum.com, https://www.quantum.com/en/products/object-storage/](https://www.quantum.com/en/products/object-storage/) (accessed Aug. 23, 2020).
- [21] M. Durant, "FSSPEC: Filesystem interfaces for Python," [filesystem-spec.readthedocs.io, https://filesystem-spec.readthedocs.io/en/latest/](https://filesystem-spec.readthedocs.io/en/latest/) (accessed Jan. 23, 2020).
- [22] Amazon Web Services, INC., "Best Practices Design Patterns: Optimizing Amazon S3 Performance — Amazon Simple Storage Service," [aws.amazon.com, https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html](https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html) (accessed Aug. 26, 2020).
- [23] Continuum Analytics, "GCSFS — GCSFS documentation," [gcsfs.readthedocs.io, https://gcsfs.readthedocs.io/en/latest/](https://gcsfs.readthedocs.io/en/latest/) (accessed Aug. 21, 2020).
- [24] Continuum Analytics, "S3FS — S3FS documentation," [s3fs.readthedocs.io, https://s3fs.readthedocs.io/en/latest/](https://s3fs.readthedocs.io/en/latest/) (accessed Jun. 06, 2020).
- [25] H. Xu, "NCAR/benchmarking: Benchmarking & Scaling Studies of the Pangeo Platform," [github.com/NCAR, https://github.com/NCAR/benchmarking](https://github.com/NCAR/benchmarking) (accessed Aug. 29, 2020).
- [26] J. Hamman, "Dask-jobqueue," <https://github.com/dask/dask-jobqueue/> (accessed Jul. 20, 2020).
- [27] B. Nitzberg, J.M. Schopf, J.P. Jones, "PBS Pro: Grid Computing and Scheduling Attributes," in *Grid Resource Management: State of the Art and Future Trends*, J. Nabrzyski, J.M. Schopf, J. Weglarz, Eds., Boston, MA, USA: International Series in Operations Research & Management Science, 2004, pp. 183–190. [Online]. Available: https://doi.org/10.1007/978-1-4615-0509-9_13
- [28] T. E. Oliphant, *A guide to NumPy*, USA: Trelgol Publishing, 2006, vol. 1.
- [29] H. Roger, "Tiger Shark - a scalable file system for multimedia," in *IBM Journal of Research and Development*, Vol. 42, No. 2, Mar. 1998.