




User-Centric System Fault Identification Using IO500 Benchmark

Radita Liem^{*}, Dmytro Povaliaiev^{*}, Julian Kunkel[‡], Jay Lofstead[†], Christian Terboven^{*}

^{*}Chair for High Performance Computing, IT Center, RWTH Aachen University

[‡]Göttingen University/GDWG

[†]Sandia National Laboratories



Introduction



I/O is a shared resource in a cluster. It creates performance variability on the running applications and makes predicting I/O performance difficult.

User often **relies on their own perception** that the application is slowing down¹ and the **current solutions mostly focused on system maintainer.**



We want to provide users with **easy-to-digest realistic expectations** about their application's I/O in the specific environment where their applications are running.

In this pursuit, **we propose to use the IO500 benchmark** as a way to manage user expectation for a HPC system.

IO⁵⁰⁰

¹J. Kunkel and E. Betke, "Tracking user-perceived i/o slowdown via probing," in High Performance Computing, M. Weiland, G. Juckeland, S. Alam, and H. Jagode, Eds. Cham: Springer International Publishing, 2019, pp. 169–182.

IO500 Benchmark Overview

- Developed in 2017 and quickly becoming the de facto benchmarking standard for HPC Storage².
- Designed to create balanced performance measurement
- The benchmark has 5 scenarios³ :
 - **IOR ‘easy’**: Free to tune IOR parameters. Typically file-per-process, large, aligned chunks. Creating **best possible bandwidth performance**.
 - **IOR ‘hard’**: Limited options to tune. Forced to use small unaligned I/O to a single shared file. Providing **worst possible bandwidth performance**.
 - **mdtest ‘easy’**: Free to tune mdtest parameters with zero size files in separate directory per process. Creating **best case scenario for metadata rate**
 - **mdtest ‘hard’**: Limited options to tune. Forced all processes to write into a single shared directory. Providing **worst case scenario for metadata rate**.
 - **Find**: Finding specific subset of files created by 4 other benchmarks.
- The numbers from the scenarios are combined using geometric mean and each scenario is running for 300s that eliminates the cache effect

The logo for the IO500 benchmark, featuring a large, stylized 'IO' in red and black, followed by the number '500' in black.

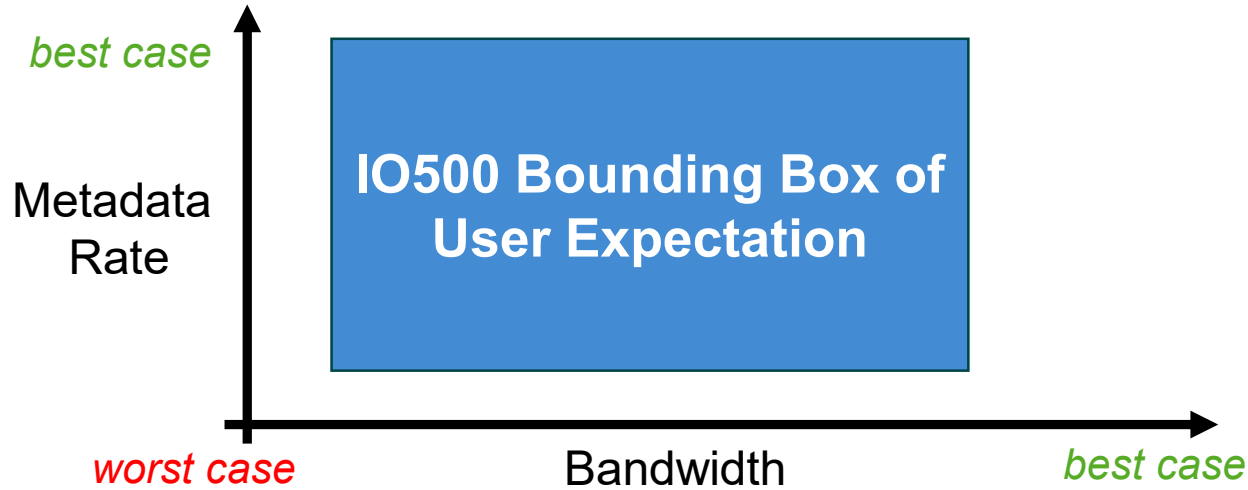
² J. Kunkel, “Virtual Institute for I/O,” *ISC’19: The IO-500 and the Virtual Institute of I/O*, 30-Oct-2020. [Online]. Available: <https://www.vi4io.org/io500/bofs/isc19/start>. [Accessed: 27-May-2021].

³ M. Rásó-Barnett, “Lustre and IO-500: Experiences with the Cambridge Data Accelerator”, 2019. [Online]. https://www.eofs.eu/_media/events/lad19/03_matt_raso-barnett-io500-cambridge.pdf. [Accessed: 02-Mar-2021]

IO500 Benchmark Usage

- IO500 benchmark's mdtest and IOR scenario can be used to form a bounding box of user expectations⁴ as illustrated by the figure below

IO⁵⁰⁰



- Worst case scenario** is coming from IOR and mdtest **'hard'** scenario
- Best case scenario** is coming from IOR and mdtest **'easy'** scenario
- 'Find'** is not used in this bounding box model since it is not as controlled as IOR and mdtest and will skew the IO500 numbers


⁴ A. Dilger, "IO500 | A storage Benchmark for HPC", 2019. [Online]. Available: https://wiki.lustre.org/images/9/92/LUG2019-IO500_Storage_Benchmark_for_HPC-Dilger.pdf. [Accessed: 02-Mar-2021]

IO500-based Workflow Proposed

1 Setting up the boundary of expectation

2 Mapping the application's performance


3 Tuning the application



Performance variability Application's I/O performance Tuning direction

Proof of Concept Experiment Setup

- The experimentation in this work covers **the first and second step of the workflow**. The second step of the work flow is still in our **exploratory stage** and the third step is for the future work




- Experiment environment:**

- CLAIX-2018 cluster at RWTH Aachen University (48 cores Intel Skylake, 384 GB memory), 40 Gb/s Ethernet.
- 4 nodes BeeGFS Filesystem, each with 480 GB SSD.
- IO500 benchmark - SC20 submission version.
- NAS Parallel Benchmark – BTIO “full” class A,B, and C on 4,9, and 16 processes

Results: Forming Bounding Box of User Expectation [1]


- Bounding box of **POSIX** API, each square represents individual run from the same IO configuration



This project is currently displayed in: <https://bit.ly/3BhhAFZ>

Results: Forming Bounding Box of User Expectation [2]

- Bounding box of **POSIX** API, read and write show different pattern



Bounding box from POSIX **write** result


Bounding box from POSIX **read** result

This project is currently displayed in: <https://bit.ly/3BhhAFZ>

Results: Anomalous Bounding Box

- Anomalous result in **MPI-IO** API: IOR 'Easy' score gets lower number than IOR 'hard'
- Broken node is most likely the reason behind these anomalous result

Bounding box skewed to the direction of IOR 'hard'




	IOR hard (GiB/s)	IOR easy (GiB/s)	mdtest hard (KIOPS)	mdtest easy (KIOPS)
■	0.90	1.89	10.74	135.79
■	0.94	0.47	10.50	106.45
■	1.14	0.47	12.86	114.73
■	0.83	1.89	11.12	124.06
■	1.46	0.47	13.78	130.29
■	0.88	0.47	13.50	103.47
■	0.99	0.47	13.24	122.10
■	0.91	0.47	13.04	135.73
■	0.95	0.46	13.10	140.41
■	0.85	0.47	13.02	142.20
■	0.96	1.90	11.00	141.28
■	0.86	1.85	10.81	146.24

This project is currently displayed in: <https://bit.ly/3BhhAFZ>

Results: Exploration on the I/O Performance Mapping [1]

- Exploration with BTIO shows the application's performance falls within the box for **MPI-IO API with cache effect**

Bounding boxes of IO500 running with default setup (300s)

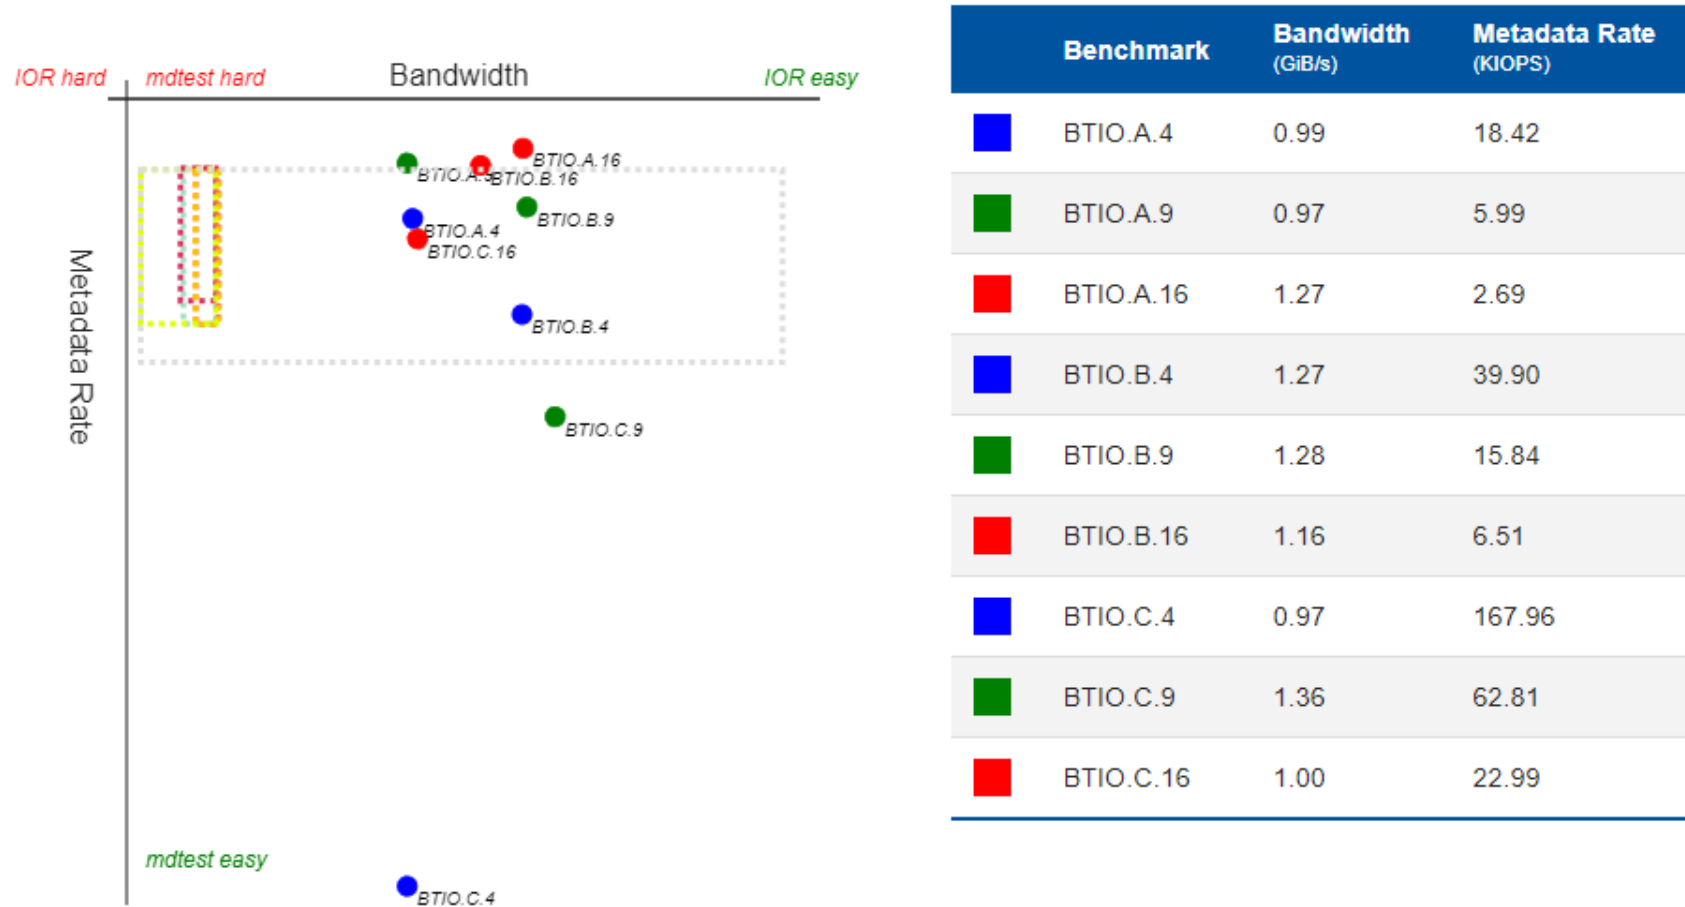


Benchmark	Bandwidth (GiB/s)	Metadata Rate (KIOPS)
BTIO.A.4	0.68	21.67
BTIO.A.9	0.64	14.18
BTIO.A.16	0.73	8.66
BTIO.B.4	0.85	17.08
BTIO.B.9	0.82	11.35
BTIO.B.16	0.75	9.58
BTIO.C.4	0.74	25.50
BTIO.C.9	0.91	11.71
BTIO.C.16	0.90	9.00

This project is currently displayed in: <https://bit.ly/3BhhAFZ>

Results: Exploration on the I/O Performance Mapping [2]


- However in the **POSIX API**, metadata rate calculation falls outside the cache box




This project is currently displayed in: <https://bit.ly/3BhhAFZ>

Results: Tail Latency Observations

- Tail latency is one of the issue in the IO500 benchmark community forums with discussions about the impact of the stonewall requirement.



Runtime variety of each IO 500 scenario.
Several tests have significant variability



Close up look on the runtime range for each scenario.

Discussion and Future Work

- **The workflow provides a way to identify system fault.**
 - It provides sanity check using the IO500 'hard' and 'easy' relationship.
 - Without users informed about this anomaly, they will just assume that the machine is operating slow but intended.
- The workflow shows a promising results and also **several tasks for future work:**
 - Refining the IO500 configuration
 - Testing on different filesystems, cluster, and APIs such as HDF5
 - Scaling up the experiment
- The exploratory work shows **the impact of cache effect** that needs further investigation and additional works on **exploring I/O tools and metrics**



Thank you!

Inquiry and question: **Radita Liem** (liem@itc.rwth-aachen.de)



UNIVERSITY OF GÖTTINGEN
GERMANY



Sandia
National
Laboratories



High
Performance
Computing



IT Center

RWTHAACHEN
UNIVERSITY