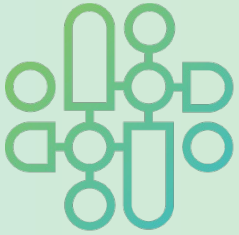

SC23
Denver, CO | i am hpc.

The I/O Trace Initiative: Building a Collaborative I/O Archive to Advance HPC

Nafiseh Moti[1], André Brinkmann[1], Marc-André Vef[1], Philippe Deniel[2], Jesus Carretero[3], Philip Carns[4], Jean-Thomas Acquaviva[5], Reza Salkhordeh[1]

[1]Johannes Gutenberg University Mainz, Germany
[2]CEA, France
[3]Universidad Carlos III Madrid, Spain
[4]Argonne National Laboratory, USA
[5]DDN, France

IO-Tracing: Motivation

- I/O bottlenecks are common in HPC
- Applications suffer from cross-application I/O interference
- Yet, details of I/O behavior of HPC applications largely unknown
- Analysis of small set of applications with limited set of inputs not sufficient

- HPC IO trace archive required by application and system developers

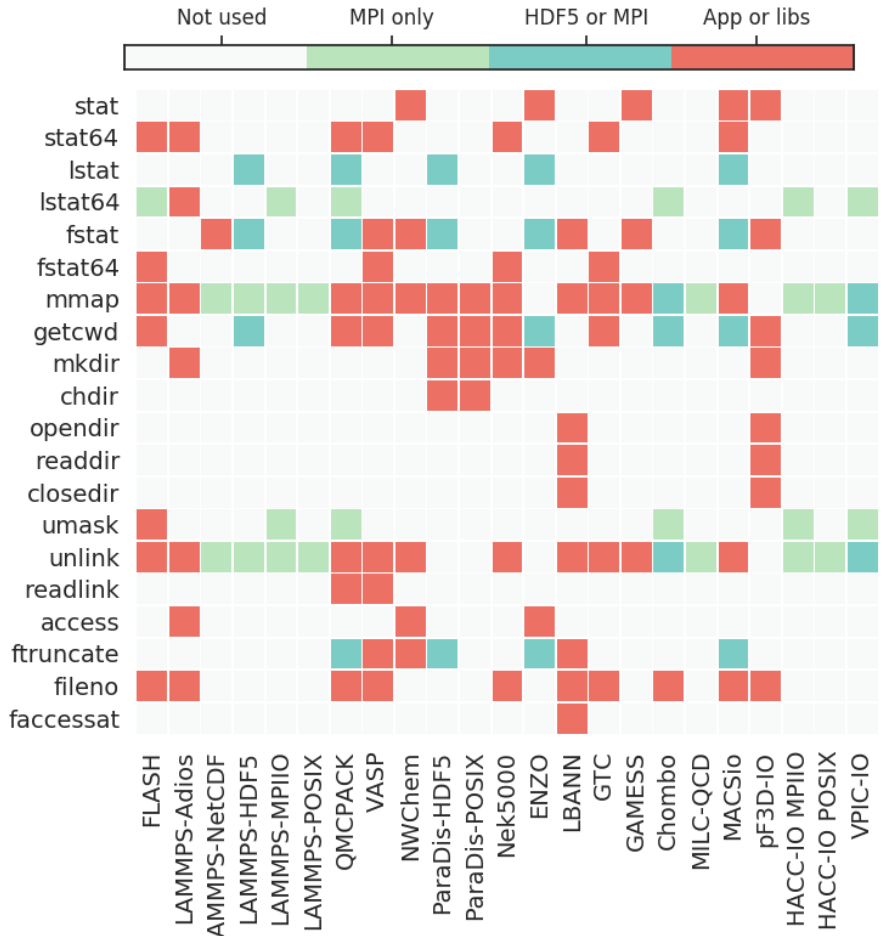

Goals and Requirements

- Allow app developers to understand behavior of their applications
- Storage system developers must receive an overview about IO behavior of a huge fraction of I/O heavy and important HPC applications

- Traces therefore must cover different input parameters, scales, and even systems for many investigated application
- Trace output has to be standardized to enable automated analysis
- Collecting traces has to be as easy as possible
- Archive will follow FAIR principles and provide DOIs using Zenodo

Darshan: Trace and profiling environment

- Profiling, tracing, and analysis built on Darshan
- Darshan is lightweight and scalable I/O characterization tool for HPC
- Comprehensive and widely available
- Long-term commitment by Argonne National Lab
- Easy to use (no code changes, minimal overhead)

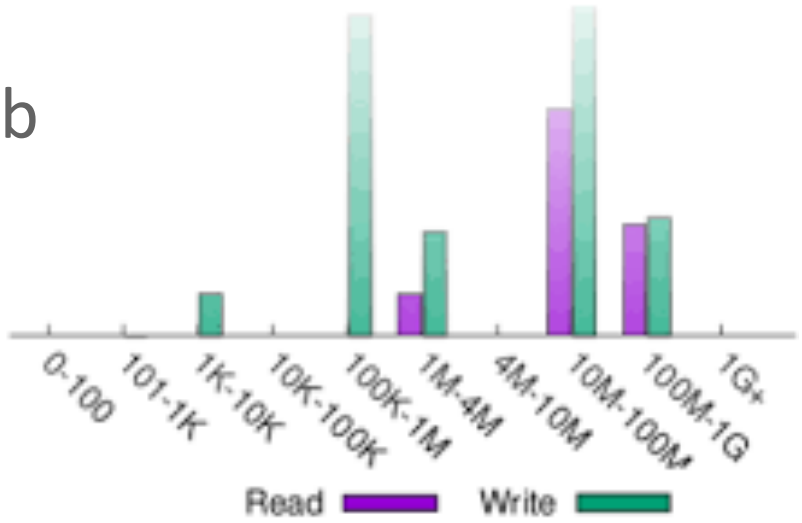

Quantum Espresso: Compare usage patterns[1]

Molecular dynamics:

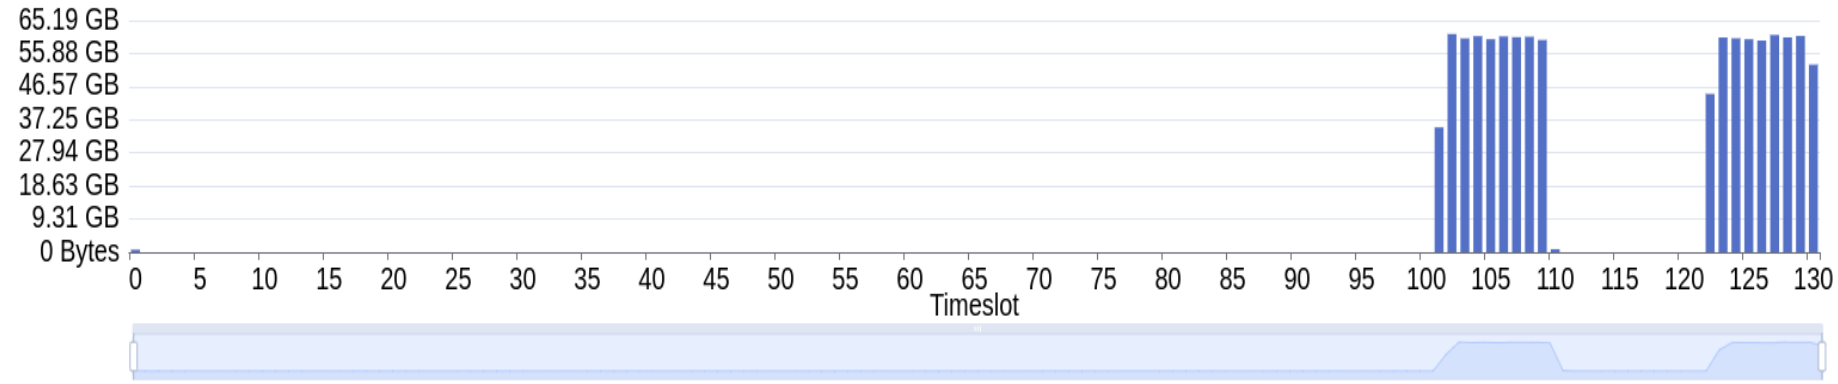

Phonon frequencies

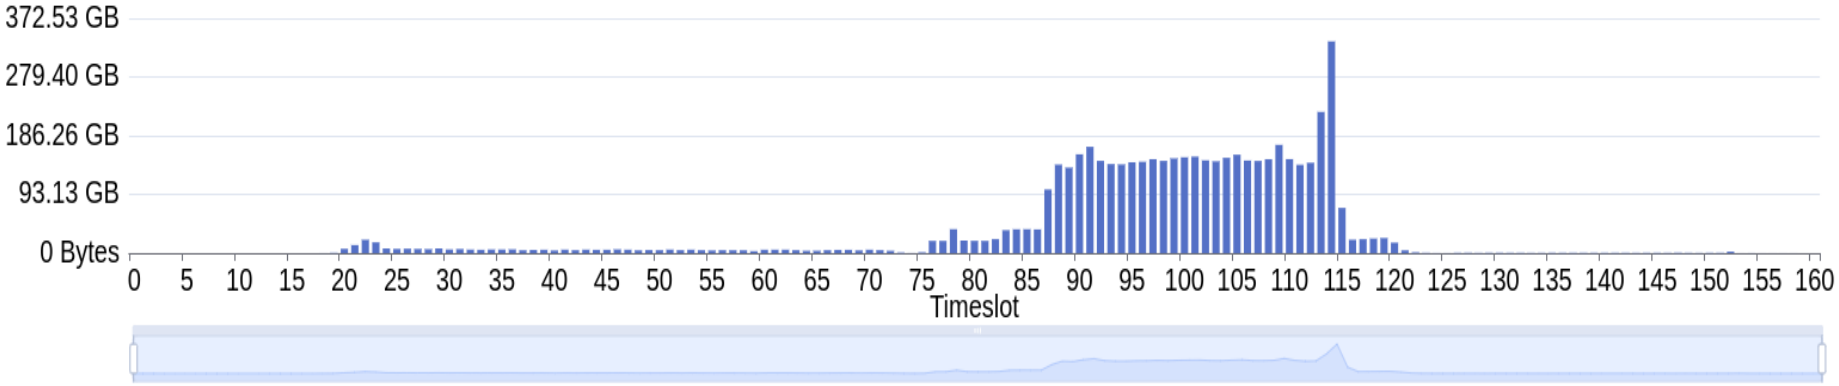

[1]One example from the online demo at BoF-session

Conclusion

- The I/O Trace Initiative: Pioneering I/O management in HPC and AI/ML
- Hub for I/O trace collection and analysis: https://hpcioanalysis.zdv.uni-mainz.de
- Community engagement: Jointly share, find, and use I/O traces
- Advanced features: Submission, archiving, searching, and visualization
- Commitment to FAIR Principles: Enhancing I/O trace accessibility
- On-going dedication: Platform evolution to support HPC and AI/ML communities