

# Advancing Automated I/O Analysis with Multi-Perspective Views

Izzet Yildirim<sup>1</sup>, Hariharan Devarajan<sup>2</sup>, Anthony Kougkas<sup>1</sup>, Xian-He Sun<sup>1</sup>, Kathryn Mohror<sup>2</sup>

[iyildirim@hawk.iit.edu](mailto:iyildirim@hawk.iit.edu), [hariharandev1@llnl.gov](mailto:hariharandev1@llnl.gov), [akougkas@iit.edu](mailto:akougkas@iit.edu), [sun@iit.edu](mailto:sun@iit.edu), [kathryn@llnl.gov](mailto:kathryn@llnl.gov)

<sup>1</sup>Illinois Institute of Technology, <sup>2</sup>Lawrence Livermore National Laboratory

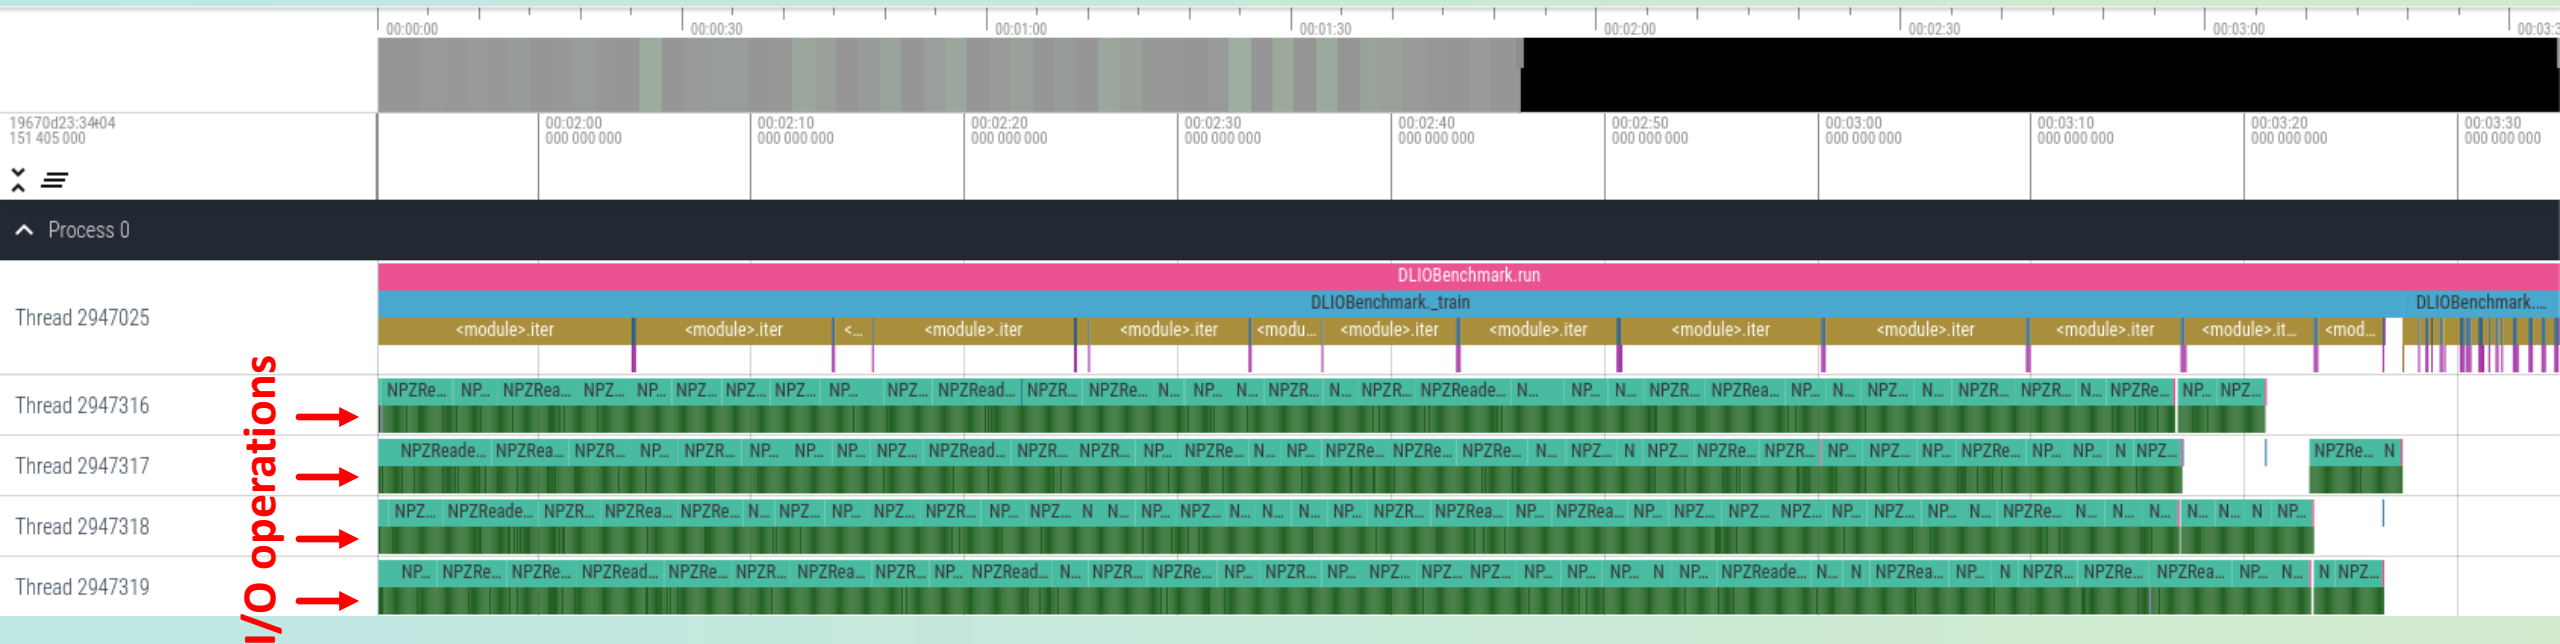
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under the DOE Early Career Research Program (LLNL-PRES-857392). Also, the material is based upon work supported by the National Science Foundation under Grant no. NSF OAC-2104013, OCI-1835764, and CSR-1814872.

# Problem

- I/O analysis often rely on heuristics and I/O analysis tools are built on top of those heuristics
- I/O analysis tools can benefit from transitioning from heuristics-based approaches towards trace-data-driven decision-making
- According to our findings, data-driven decision-making can be helpful in two areas:
  - Reducing the search space
    - DXT and Recorder offer highly detailed traces which comes at the cost of significantly larger trace sizes
  - Data-driven identification of problematic I/O accesses
    - I/O analysis involves querying I/O traces to identify issues by applying heuristics
    - While heuristics accumulated through extensive studies are highly valuable, the evolving nature of workloads necessitates new studies to generate new heuristics
      - e.g., deep learning applications, async checkpointing

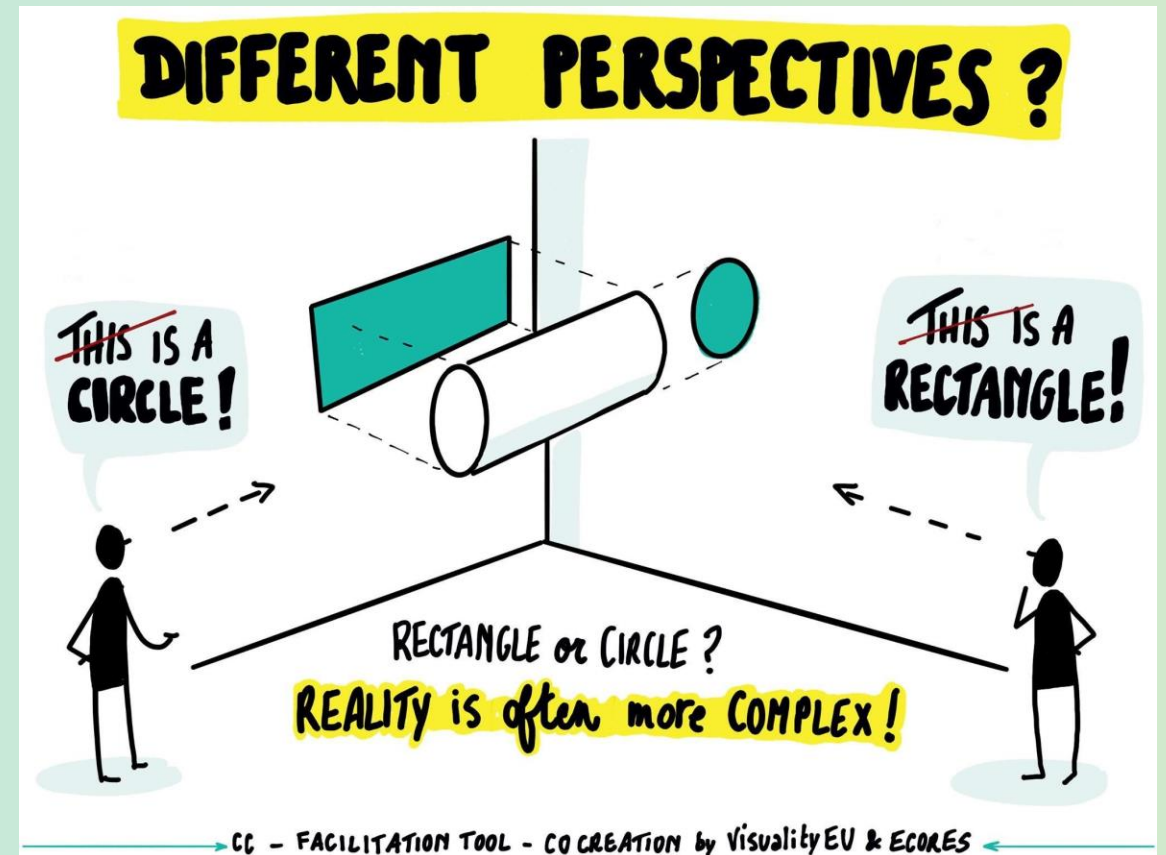
# Motivation

- The following is from DLIO Benchmark for a TensorFlow workload with NumPy arrays
  - 5TB in size, accesses 5376 files of 132MB in size, 4 read threads, 1000 epoch = multimillion I/O accesses
- Analyzing this data from a time-based perspective would not reveal that the first thread doesn't do any I/O, unlike the other threads



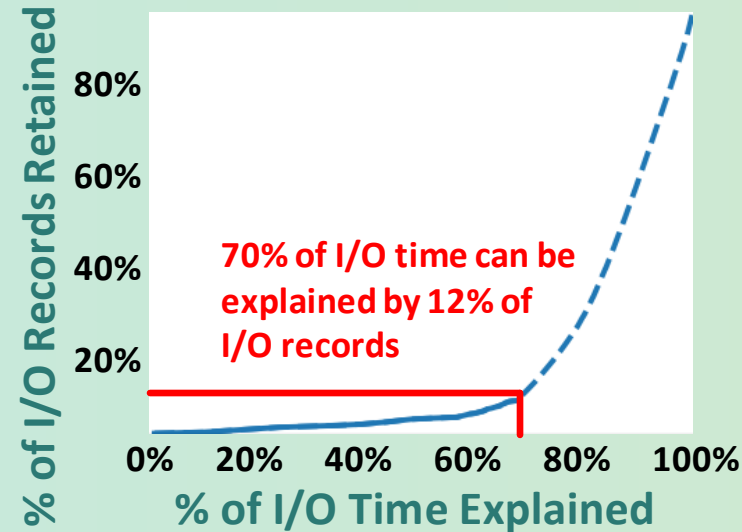
# Multi-Perspective Views

- So, reality is often more complex!
- Multi-perspective views enable us to look at the same trace data from multiple viewpoints
- This helps us not only to reduce the amount of data analyzed per view, but also to identify problematic I/O accesses more accurately



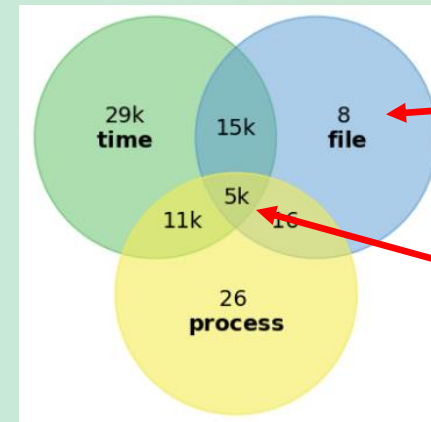
# Multi-Perspective Views

- As demonstrated, large-scale workflows may accumulate well over 1 billion I/O accesses
  - Our convergence analysis reduces the search space significantly while retaining the most critical I/O accesses
- Multi-perspective views include time-, file- and process-based views
  - The figure shows that certain problematic I/O accesses can only be identified using specific views
- This demonstrates that the trace data itself can be used to pinpoint the most critical issues



View (-based)	# of I/O access
time	62,779 (0.5%)
file	43,710 (0.3%)
process	60,726 (0.5%)
time > file	43,702 (0.3%)
time > process	60,724 (0.5%)
file > time	43,709 (0.3%)
file > process	41,750 (0.3%)
process > time	60,716 (0.5%)
process > file	41,729 (0.3%)

**% of I/O records retained for each view**



**8 problematic I/O accesses can only be found via file-based view**

**5k problematic I/O accesses are detectable across all views**

# Thank you!

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under the DOE Early Career Research Program (LLNL-PRES-857392). Also, the material is based upon work supported by the National Science Foundation under Grant no. NSF OAC-2104013, OCI-1835764, and CSR-1814872.

