I/O Traces of HPC Applications

Chen Wang[1] chenw5@illinois.edu

Kathryn Mohror[2] kathryn@llnl.gov

Marc Snir[1] snir@illinois.edu

1. University of Illinois at Urbana-Champaign
2. Lawrence Livermore National Laboratory

Introduction

- A better understanding of I/O behaviors of an application can:
  - Help users optimize I/O performance and identify potential I/O bottlenecks.
  - Help file system and I/O library developers make better optimization decisions.

- The first release of the dataset contains I/O traces from 14 HPC applications.

Introduction

- All traces were generated using Recorder[1], a multi-level tracing tool.
  - Trace record includes enter/exit timestamp of the function call, function name and all function parameters except the data buffer.
  - Currently includes POSIX, MPI-IO and HDF5.
  - I/O traces with more details take more time to collect. But once generated can be used repeatedly by researchers to perform different analysis tasks.
- Traces are in Recorder-specific binary format.
  - Tools are provided to read, decompress and decode the trace. These tools are written in C with python bindings provided for easy analysis.
  - More information on reading, visualizing and analyzing the traces can be found here: https://github.com/uiuc-hpc/Recorder

[1] Wang, Chen, Jinghan Sun, Marc Snir, Kathryn Mohror, and Elsa Gonsiorowski. "Recorder 2.0: Efficient parallel I/O tracing and analysis." In *2020 IPDPS Workshops (IPDPSW)*, pp. 1-8. IEEE, 2020.

I/O Traces of 14 HPC applications

- 14 HPC applications and I/O emulators.
- Applications perform I/O using POSIX, MPI-I/O and high level libraries include HDF5, NetCDF, Silo, and ADIOS.
- 8 nodes x 8 MPI ranks/node @ LLNL Quartz system.
- Traces and visualizations are available at https://doi.org/10.6075/J0Z899X4

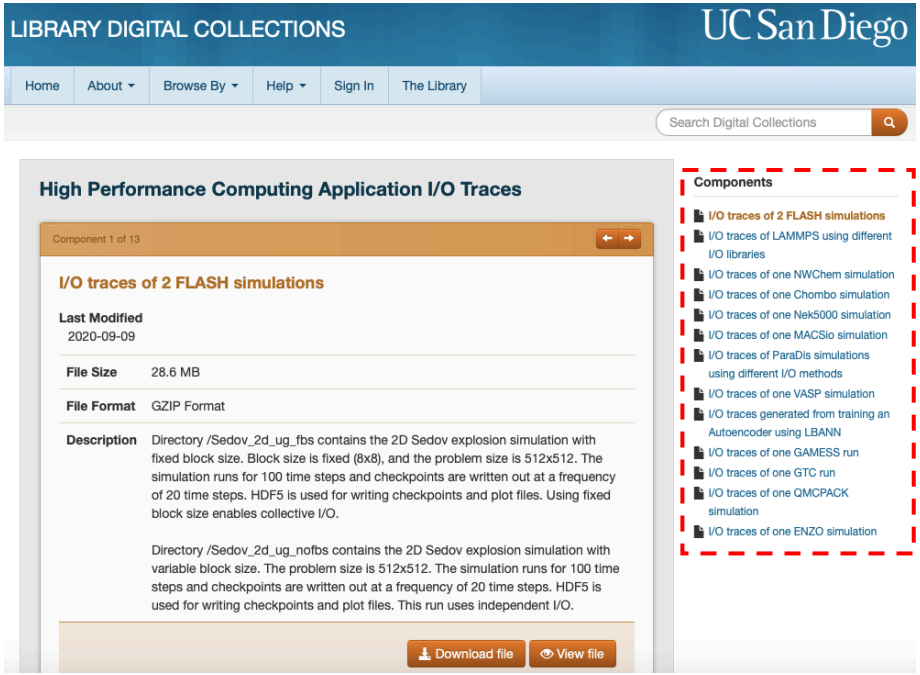

I/O Traces of 14 HPC applications

Following files are included for each application run:

- 64 trace files (*.itf) and 64 (*.mt) metadata files, one for each MPI rank.
- recorder-report.html: visualization report generated by Recorder.
- Input files if used.
- Other files like logs written by the application.

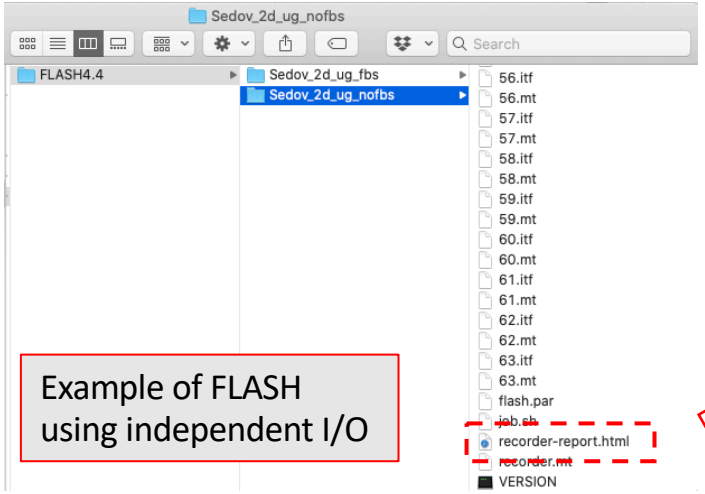

Example of FLASH using independent I/O

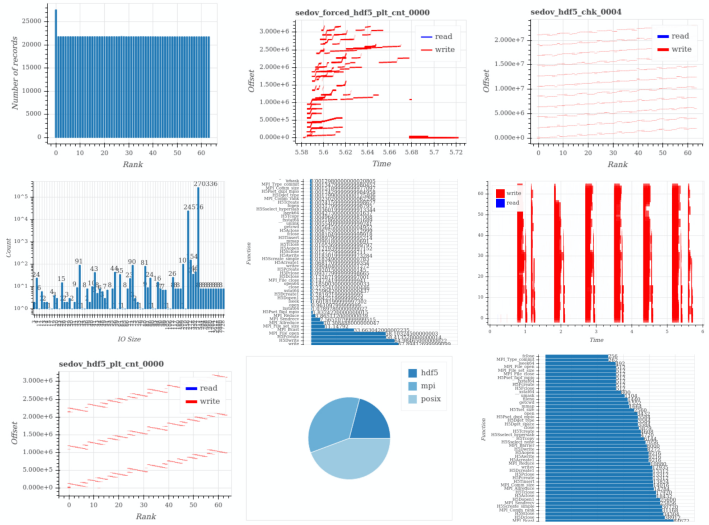

Thank You

Contact:

Chen Wang (chenw5@illinois.edu)

Kathryn Mohror (kathryn@llnl.gov)

Marc Snir (snir@illinois.edu)