I/O in WRF: Evaluating Modern Parallel I/O Techniques

RAPIDS

Scientific Achievement

This work presents a comparative study of several parallel I/O implementations in the Weather Research and Forecasting model (WRF). The I/O libraries under study include PnetCDF, HDF5 via NetCDF4, and ADIOS. Our evaluation and performance analysis can guide I/O strategies for modern parallel codes.

Significance and Impact

- When data is stored in the **canonical layout**, I/O request aggregation, a crucial optimization, can achieve 2x performance improvement.
- When data is stored in the **time-log layout**, the HDF5 Log VOL can achieve a 4x faster write performance than ADIOS.
- Storing data in log layouts needs an off-line conversion to a canonical layout for post-simulation analysis. Such conversion costs are not negligible.

Technical Approach

- An I/O request aggregation method has been integrated into PnetCDF.
- A HDF5 Log VOL has been developed that works seamlessly with existing HDF5 applications without code changes.

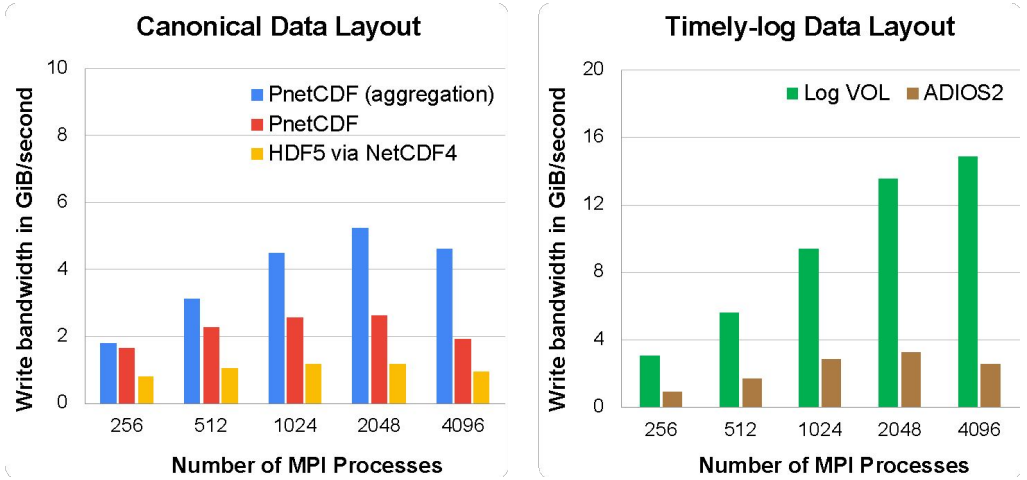

Write bandwidths of WRF evaluated on Cori supercomputer at NERSC. When output data is stored in the canonical layout, the PnetCDF aggregation strategy effectively improved the I/O performance. When output data is stored in the log layout, the HDF5 Log VOL outperformed ADIOS.

PI(s)/Facility Lead(s): Rob Ross (ANL) Local Lab POC: Wei-Keng Liao (Northwestern University)
Collaborating Institutions: Northwestern University, Argonne National Laboratory
ASCR Program: RAPIDS
ASCR PM: Kalyan Perumalla
Publication for this work: : Huang, et al. ”I/O in WRF: A Case Study in Modern Parallel I/O Techniques,” In The International Conference for High Performance Computing, Networking, Storage and Analysis November 12–17, 2023. doi:10.1145/3581784.3613216
Code artifact: https://doi.org/10.5281/zenodo.8081498