GNOSIS RESEARCH CENTER
ILLINOIS INSTITUTE OF TECHNOLOGY

Optimizing Workflow Performance by Elucidating Semantic Data Flow

Meng Tang [1] Nathan R. Tallent [2] Anthony Kougkas [1] Xian-He Sun [1]

[1]Illinois Institute of Technology [2]Pacific Northwest National Laboratory

Pacific Northwest NATIONAL LABORATORY

Introduction

- Distributed scientific workflows are data-intensive; bottleneck is usually data movement through storage systems. Therefore, it is critical to understand data flow.
- Many scientific datasets incorporate domain semantics with formats like HDF and NetCDF, enhancing the interpretability and context of the data for analysts.
- We shed new insight on workflow bottlenecks by understanding how semantic data sets flow through storage.

We unveil a fresh perspective with

1. careful runtime measurement,
2. recovering the mapping of domain semantics to low-level I/O operations, and
3. effective visualization and analysis of semantic flows.

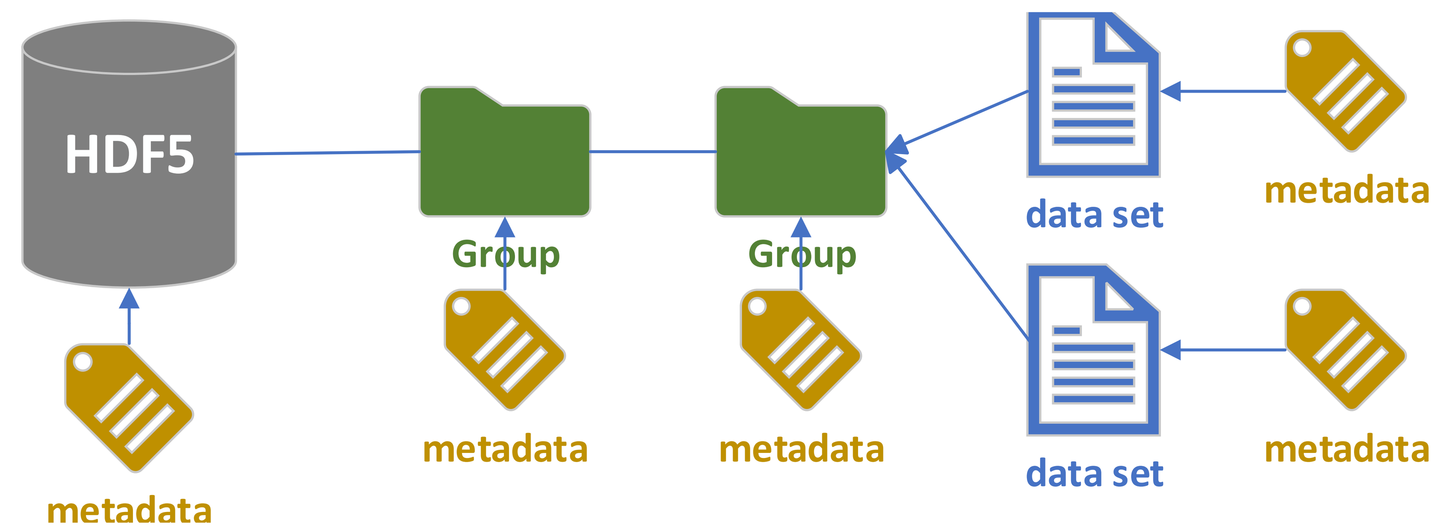

Figure 1: HDF5 Semantic Rich Structures

- HDF5's hierarchical organization of groups, datasets, and attributes enables the inclusion of context-rich metadata [1].
- This structural design effectively conveys data relationships, and annotations, enhancing data's meaning.

Challenges

- Mapping Data Semantics to I/O Access
- Tracking Data Flow Across Tasks
- Visualizing Coordination and Time

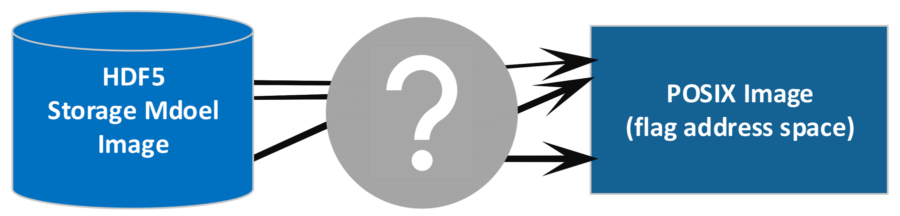

Approach

Step 1: Measures the key metrics to recover workflow task to data mapping.

Step 2: Uses HDF5 semantic object and file address mapping information to construct Semantic DAGs.

Step 3: Performance and data access related statistics are extracted. Interactive Sankey diagram is used to visualize lifecycle graphs.

Case Study I: DeepDriveMD

DeepDriveMD (DDMD) is a deep learning-driven molecular dynamics simulations workflow for protein folding [2].

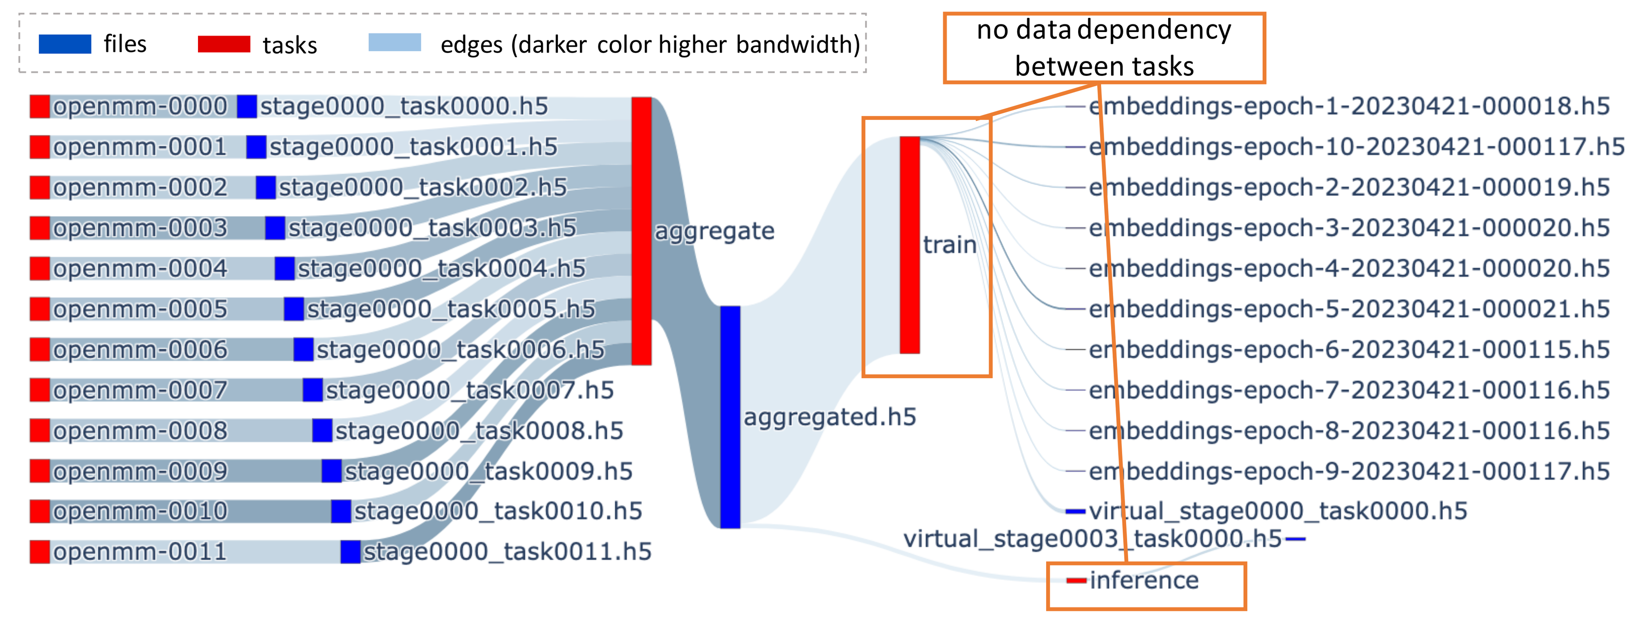

Figure 2: Four-Stages Pipeline Workflow (simulation, aggregate, train, and inference).

- Observation: no data dependencies between train and inference.
- Opportunity: inference and train tasks can be parallelized.

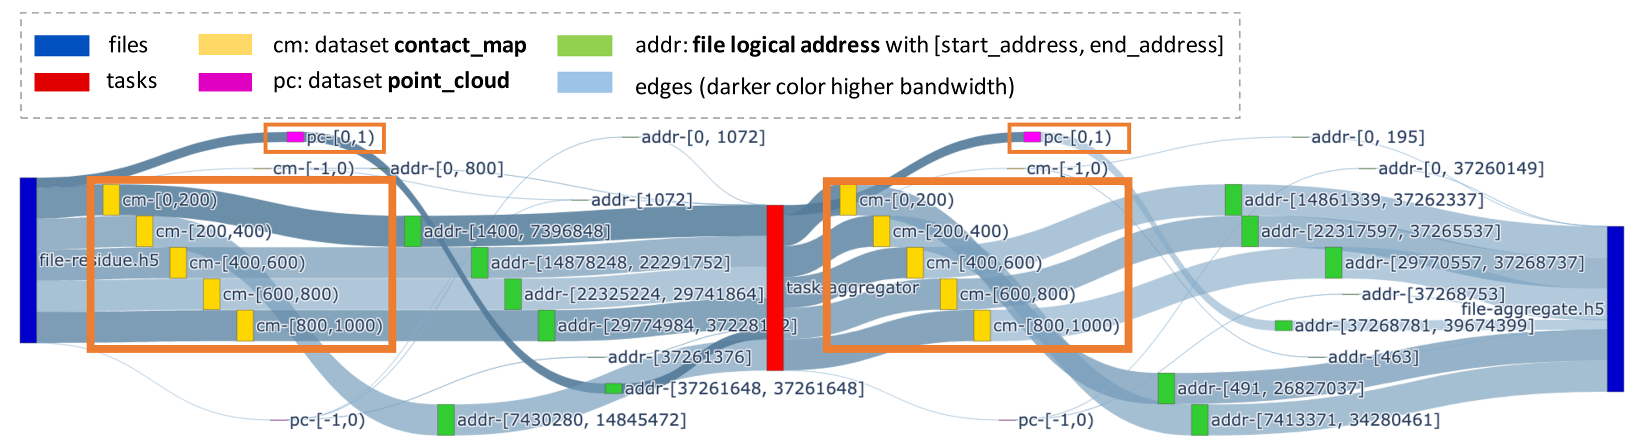

Figure 3: Aggregate Stage Close-Up TDD showing two datasets.

- Observation: aggregate task changes the data layout without content change.
- Opportunity: aggregate task can be removed.

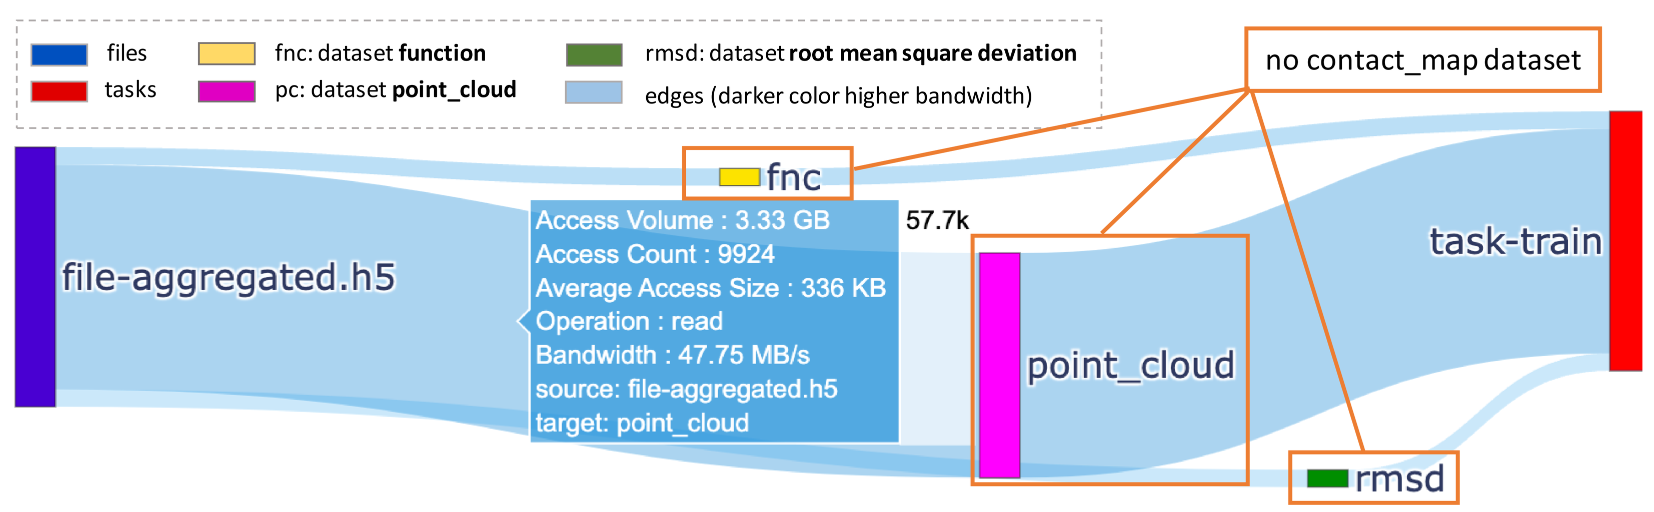

Figure 4: DDMD Train Stage Read File I/O Performance Detail.

Observation: train task not using all datasets from aggregated.h5.
Opportunities:

- aggregate rearranges data not used by a downstream task.
- when memory is limited, caching a subset of the aggregated.h5 does not violate task-data dependency.

Case Study II: Storm Tracking Workflow

Storm Tracking uses a flexible atmospheric feature tracking software package [3] for weather research and forecast datasets.

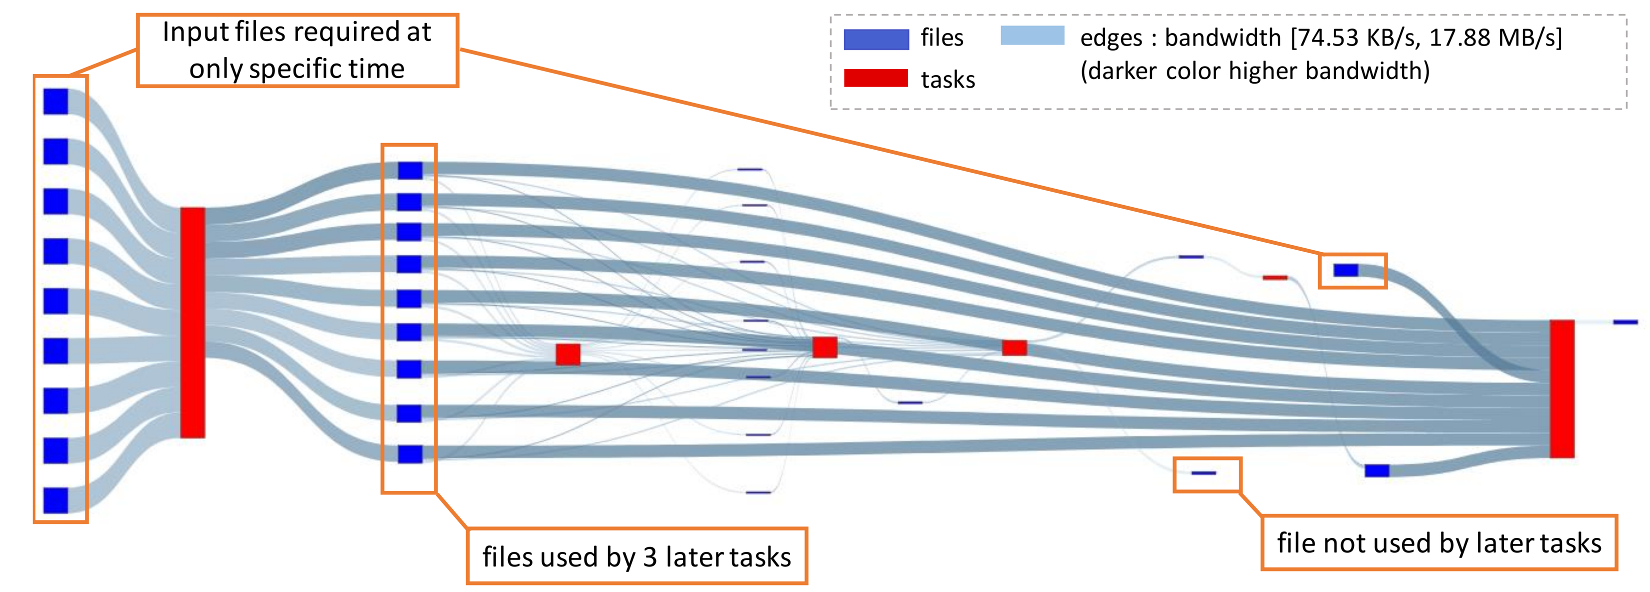

Figure 5: Six-Stages Pipeline Storm Tracking Workflow.

Observation:

- Inter-task Data Reuse: task 2, 4, and 6 uses files produced by the first task
- Time-dependents inputs: some input files are only required for specific tasks
- Data None-Used: file produced by task 4 is not used by any later task

Opportunity:

- Tasks that use common datasets can be scheduled on the same resource
- Input can be stage-in at different time points of the workflow
- Files not used by later tasks can be immediately offloaded to free up memory

Conclusions

- Lack automated method to comprehend data access within workflows.
- Semantic DAGs: enhancing traditional DAGs by incorporating tasks and filenames that trace the flow of semantic objects.
- Gathered statistics enable analysis of performance and data access patterns, enabling new insights on improving workflow.
- Future work: focus on creating an automated approach that leverages our enhanced understanding from this analysis to improve workflow performance.

Acknowledgement

This research is supported by the U.S. Department of Energy (DOE) through the Office of Advanced Scientific Computing Research's "Orchestration for Distributed & Data-Intensive Scientific Exploration.'' Also, the material is based upon work supported by the National Science Foundation under Grant no. NSF CSSI-2104013.

References

[1] M. Folk, G. Heber, Q. Koziol, E. Pourmal, and D. Robinson, "An overview of the hdf5 technology suite and its applications," in Proceedings of the EDBT/ICDT 2011 workshop on array databases, pp. 36–47, 2011.

[2] H. Lee, M. Turilli, S. Jha, D. Bhowmik, H. Ma, and A. Ramanathan, "Deepdrivemd: Deep-learning driven adaptive molecular simulations for protein folding," in 2019 IEEE/ACM Third Workshop on Deep Learning on Supercomputers (DLS), pp. 12–19, IEEE, 2019.

[3] Z. Feng, J. Hardin, H. C. Barnes, J. Li, L. R. Leung, A. Varble, and Z. Zhang, "Pyflextrkr: a flexible feature tracking python software for convective cloud analysis," EGUsphere, pp. 1–29, 2022.

mtang11@hawk.iit.edu