Towards A Practical Provenance Framework for Scientific Data on HPC Systems

Runzhou Han, Suren Byna* and Mai Zheng

Dept of Electrical and Computer Engineering, Iowa State University

** Lawrence Berkeley National Labs*

FAST'22
20th USENIX Conference on File and Storage Technologies
FEBRUARY 22–24, 2022
SANTA CLARA, CA, USA
www.usenix.org/fast22

Motivation

- A need for data provenance in HPC scientific workflows
  - Workflow & scientific data reusability, scientific data lineage tracking, workflow optimization
- Challenges
  - The scientists' diverse needs is unclear
  - The variety of the workflow characteristics (e.g., different I/O interfaces and file formats)

Example 1: DASSA, an acoustic sensing workflow [1]

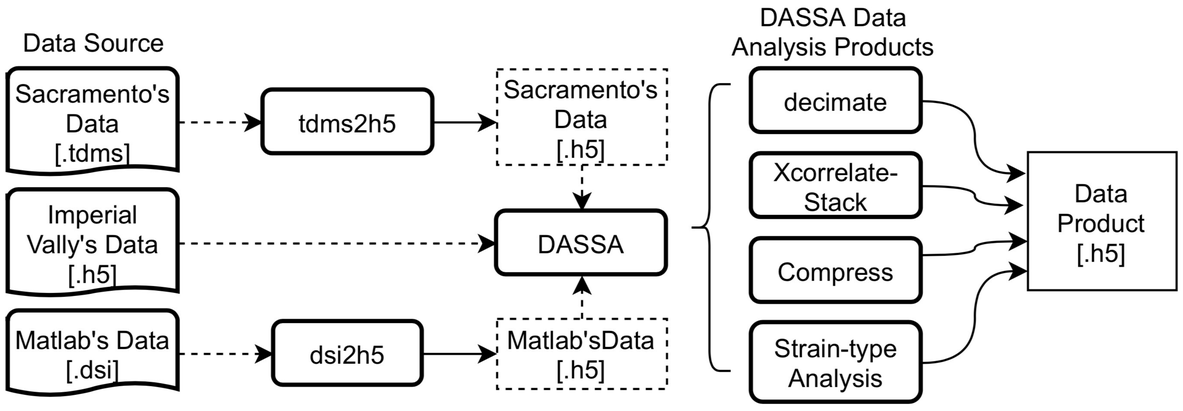

Example 2: Top Reco, an ML workflow for top quark reconstruction [2]

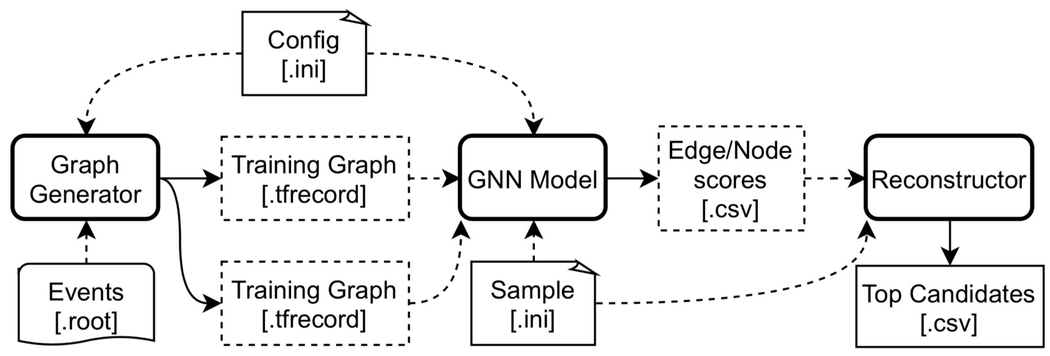

- A gap between the HPC data provenance and existing solutions
  - Vague provenance model
  - Incompatible provenance software on HPC environment
  - Limited transparency

Software	Type	Base Model	Environment	Language	Transparent?
Karma	Middleware	OPM	Cloud	Java	Yes
Komadu	Provenance system	W3C PROV-DM	Cloud	Java	No
Tanerna	Workflow system	W3C PROV-DM	Desktop	Java	Hybrid
ProvLake	Provenance system	W3C PROV-DM	Cloud	Python	No

Table 1: Existing Provenance Software for Scientific Workflows

Methodology

- Study variant workflows and their providers' provenance needs
- Propose a provenance model for HPC data provenance
- Build a provenance framework
  - Provenance Tracking for capturing diverse I/O operations
  - Provenance Storage for persisting the captured provenance information as standard RDF triples [3]
- User Engine for querying and visualizing provenance

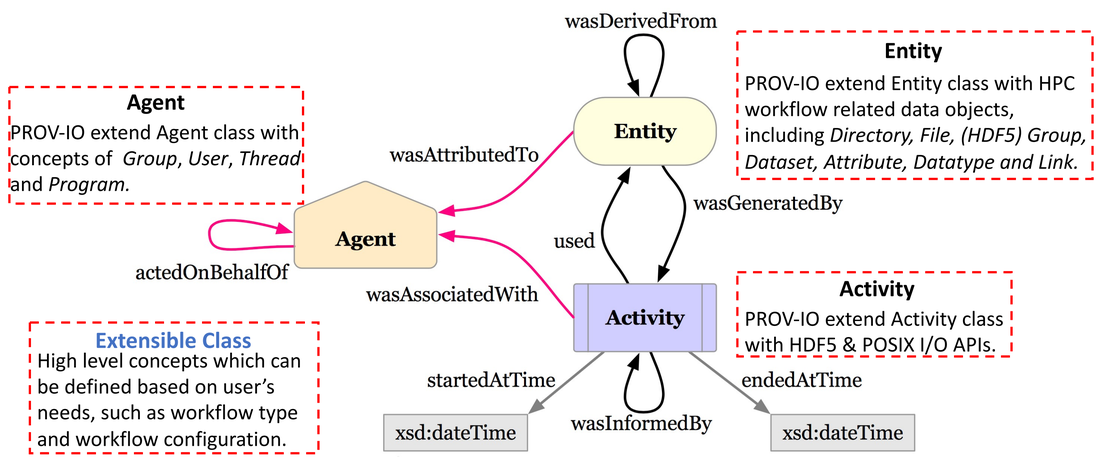

Figure 1: PROV-IO Model. An extension of W3C PROV-DM

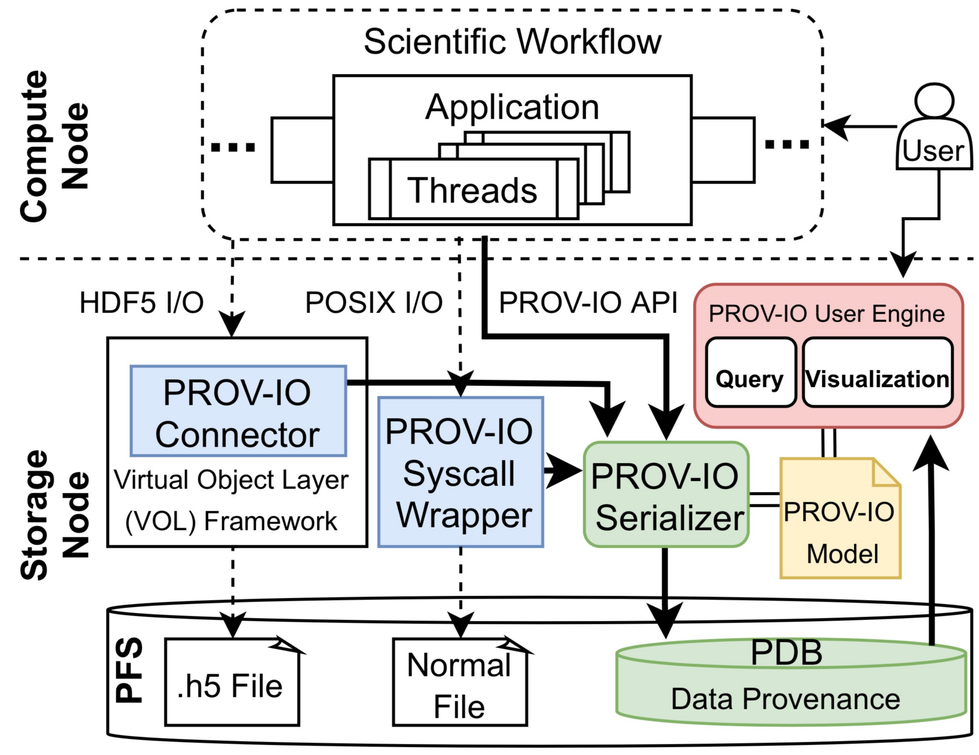

Figure 2: PROV-IO framework

Preliminary Results

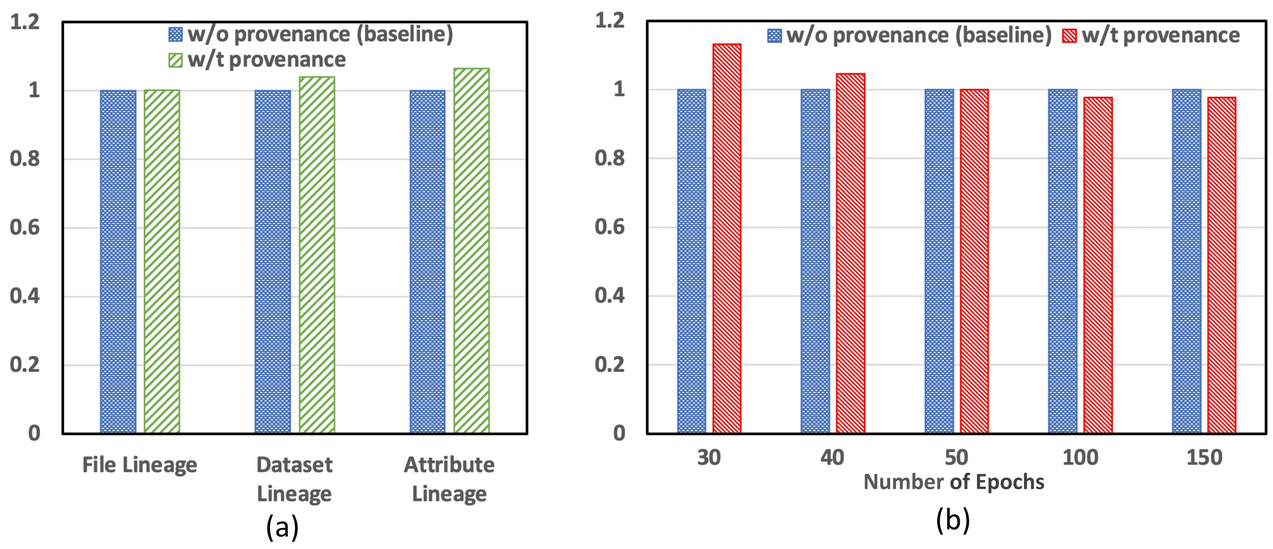

Figure 3: Performance of Provenance Tracking. (a) Tracking performance of DASSA. (c) Tracking performance of Top Reco. In the two figures, tracking performance is measured with normalized workflow completion time with a baseline of 1

- In DASSA workflow, the maximum provenance tracking overhead is about 5.2% when tracking the provenance of the most fine-grained data object in HDF5 (Attribute lineage)
- In Top Reco workflow, the maximum provenance tracking overhead is 13.2% when the number of epochs is 30. With more training epochs, provenance tracking overhead becomes negligible

Future Work

- Fully implement PROV-IO Syscall Wrapper
- Improve PDB & User engine
  - Optimization existing database solutions and PROV-IO user APIs

Acknowledgements

- This work was supported in part by NSF under grant CCF-1910747. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of NSF

References

[1] Bin Dong, Verónica R. Tribaldos, Xin Xing, Suren Byna, Jonathan Ajo-Franklin, Kesheng Wu. DASSA: Parallel DAS Data Storage and Analysis for Subsurface Event Detection. In Proceedings of IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2020.

[2] A graph neural network-based top quark reconstruction package. https://indico.cern.ch/event/932415/contributions/3918265/attachments/2086561/3505362/GNN_Top_Reco_-_Allison_Xu.pdf

[3] Resource description framework. https://www.w3.org/RDF/