KV-CSD: Accelerating Scientific Data Analytics Using Ordered, Hardware-Accelerated Key-Value Stores

Qing Zheng

Scientist, Los Alamos National Laboratory (LANL)

LA-UR-23-28796

Overview

Problem

Scientific data analytics often read more data than is necessary

Trend

Smart devices increasingly popular

Approach

Ordered, hardware-accelerated KV stores for efficient data indexing and retrieval

How Scientific Simulations Run

Time based bulk-synchronous parallel programs

Iterate between compute & I/O phases

Analytics occur after simulation

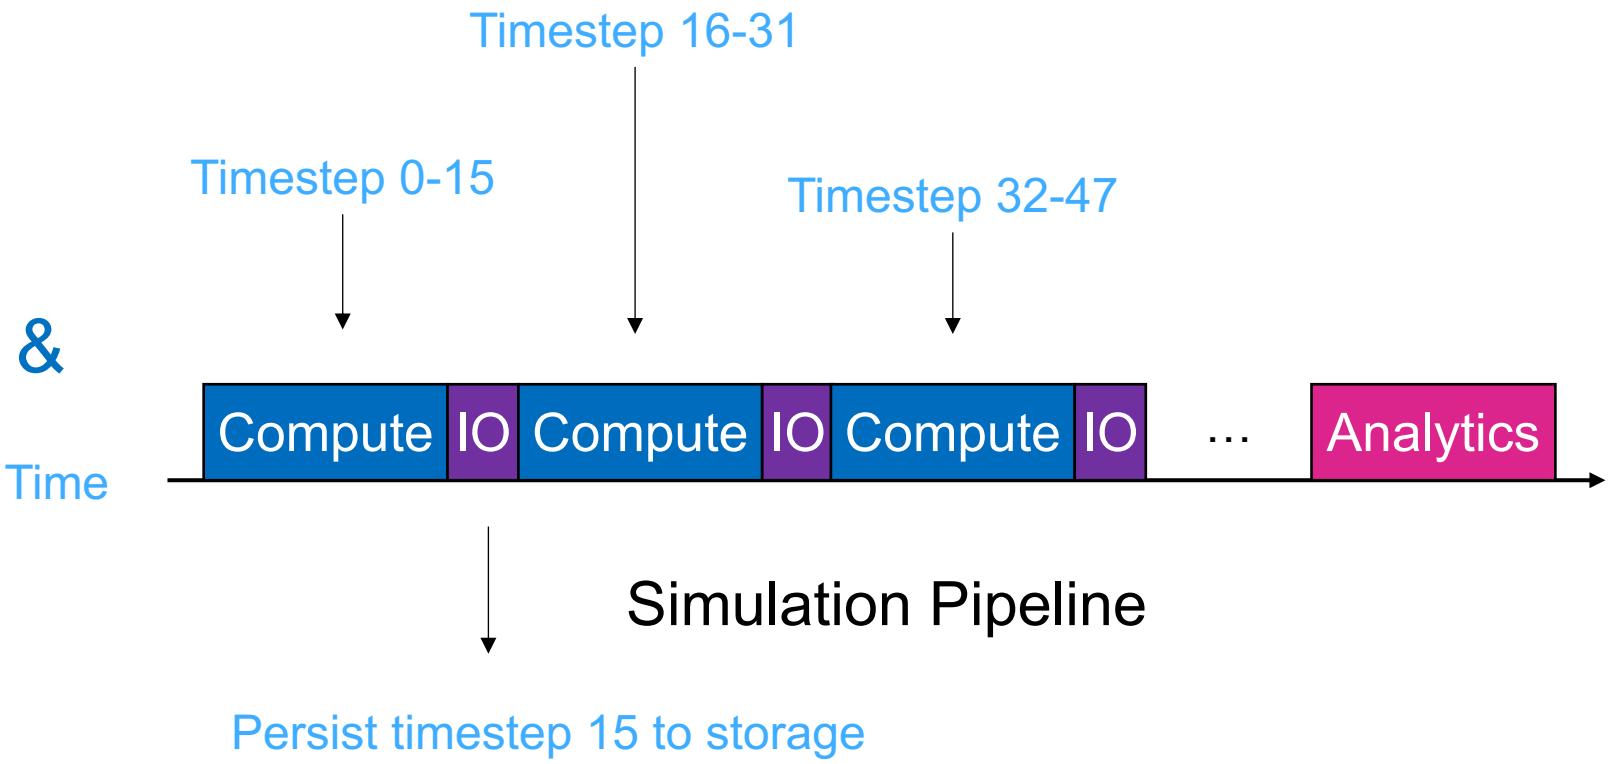

How Data is Stored Today

Flash Memory Summit

Through filesystems

Data stored as one big or many small files per timestep

Data typically accompanied by metadata that describes the data (type, dimension, …)

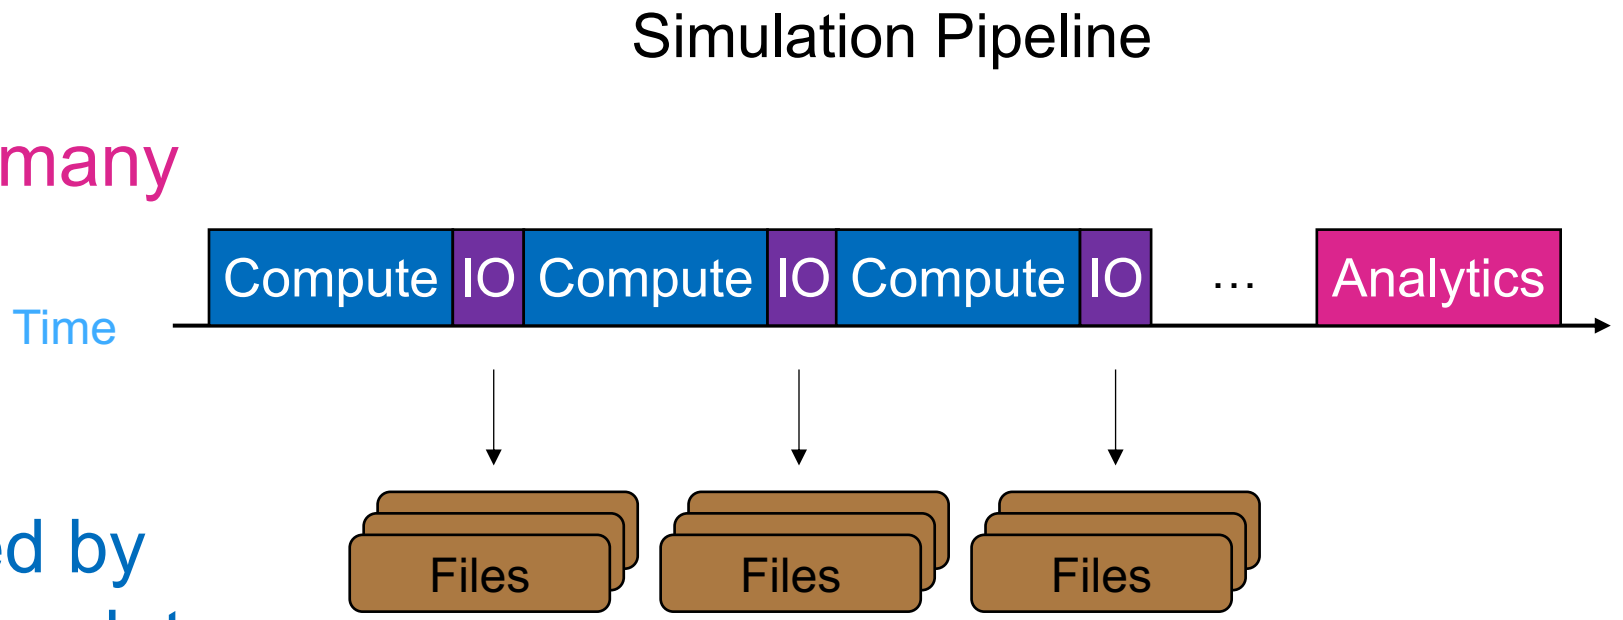

Why Reading Back Just Interesting Data is Difficult

- Data may not be persisted in the same order as queries
- Also, re-sorting data using app compute nodes is becoming increasingly inefficient
- Computational storage offers new ways of acceleration

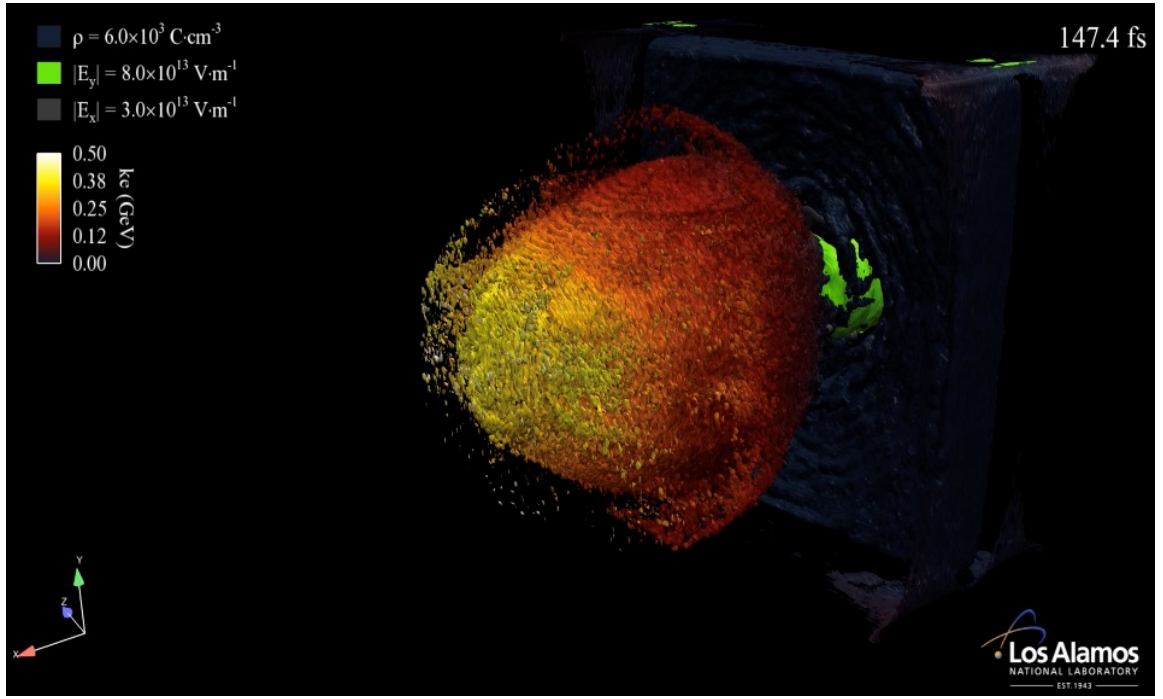

Image from LANL VPIC simulation done by L. Yin, et al at SC10

Example: a simulation may store its particles in particle ID order but queries may target particles' energy levels

Towards Ordered, KV-Based Computational Storage

App converts data to KV pairs and bulk inserts them into storage

One KV namespace per app process per timestep

Storage sorts data by key asynchronously and builds secondary indexes per app query needs

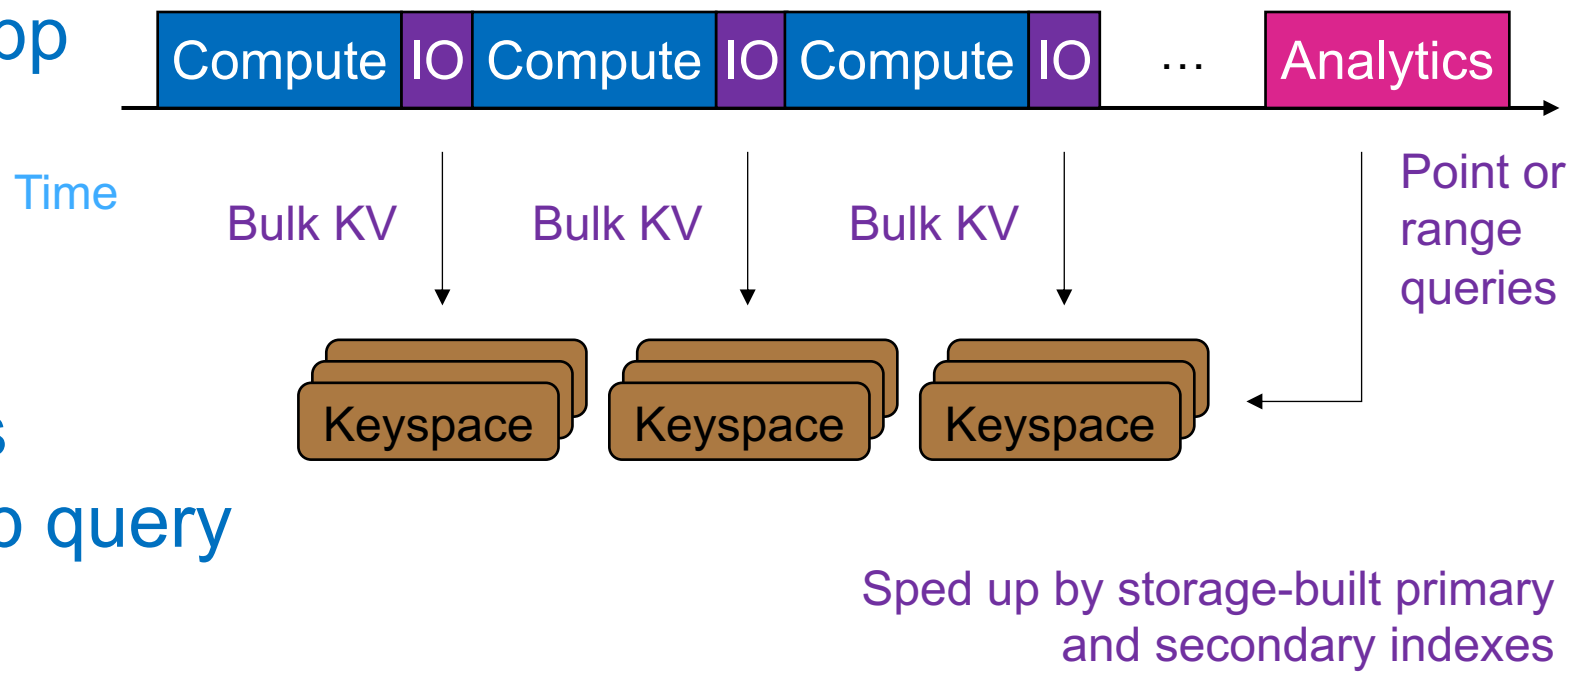

KV-CSD: A LANL/SK hynix Collaboration

Prototype consists of an ARM SoC and an E1.L ZNS SSD

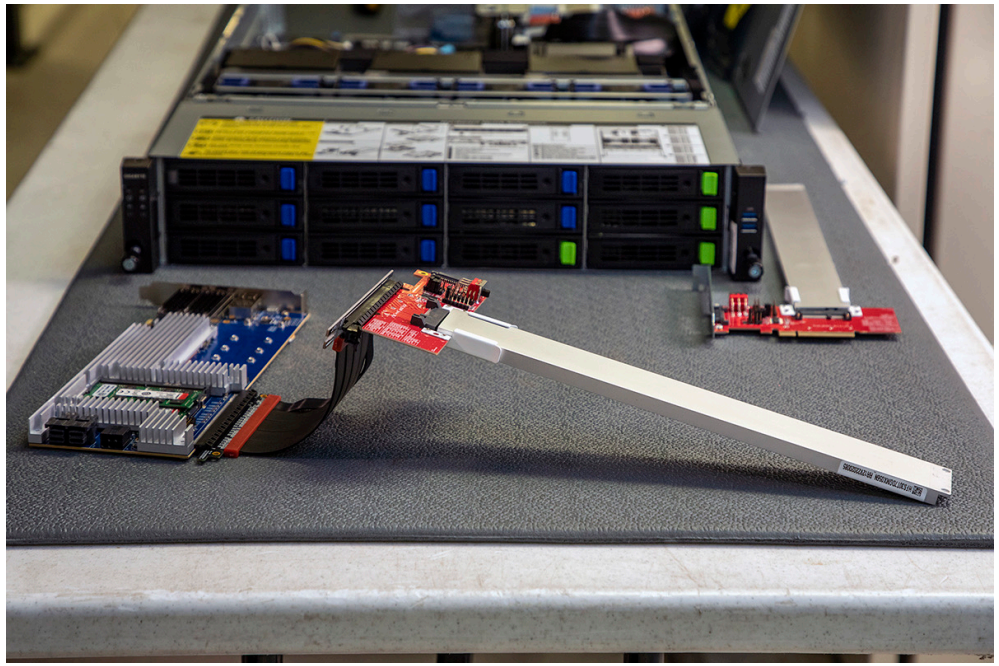

Custom NVMe KV commands for bulk data insertion and index construction

Local PCIe for now (NVMeOF in future)

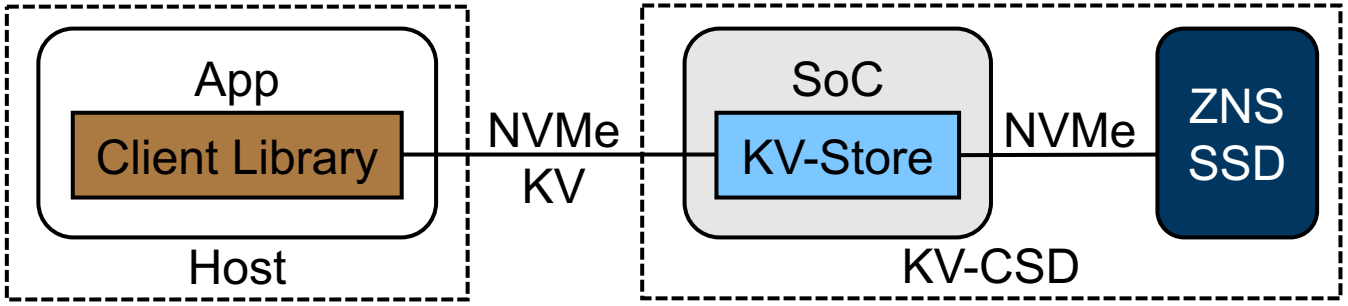

KV-CSD Drastically Reduces Time-to-Insight

Compared to files:

KV-CSD avoids full data scans and returns only interesting data

Compared to re-sorting data before queries:

KV-CSD allows overlapping data operations with simulation computation to minimize overall data processing time

Compared to software KV stores such as RocksDB:

KV-CSD speeds up processing and minimizes host-storage data movement by performing data operations very close to it

More on KV-CSD

- FMS presentation by SK hynix:
  - Architecture of a Query Accelerating KV-CSD in an HPC System
- SDC presentation by LANL:
  - KV-CSD: An Ordered, Hardware-Accelerated Key-Value Store For Rapid Data Insertion and Queries
- KV-CSD Paper (IEEE Cluster Computing Conference):
  - KV-CSD: A Hardware-Accelerated Key-Value Store for Data-Intensive Applications

Conclusion

Flash Memory Summit

Efficient data retrieval performance is key to scientific analytics

Computational storage provides new ways of accelerating data-intensive analytics workloads

Preliminary results are very promising

More work/collaboration/innovation is needed for production deployment

