UC SANTA CRUZ | Baskin Engineering

iris hep

Skyhook: Managing Columnar Data Within Storage

Jayjeet Chakraborty, Carlos Maltzahn

University of California, Santa Cruz

CROSS

cross.ucsc.edu

Problem

- Exploratory data analysis requires a dataset to be viewed in different ways
  - Derived columns
  - Different subsets of columns
- Current practice: datasets are copied to create new views
  - Increases overall analysis time
  - Uses up unnecessary space
  - Requires manual work to relate copies to each other
  - Difficult to keep track of how the data evolved over time
  - Makes workflows more complex

Our Vision

- Dataset “repositories” that support different views over the same data without extra copying
  - Process data directly over data lakes without ingesting into databases
- Support version control of data through time travel, roll back, schema evolution using transactions
- Create and track views and their provenance
- Ability to join (e.g. different compression levels of) datasets
- Scalable and distributed data processing
- And, reduced data movement over the entire system

Challenges with High Energy Physics Data

- Huge sized data, large number of columns, most with embedded values
- Complex nested schema does not fit into traditional RDBMS systems
- Systems need to fit well into the Python ecosystem
- Leverage latest data management and processing technologies

Our vision seems to fit well in solving data management challenges of HEP data !

Solution

- Zero-Copy In-Memory Data Format
  - Eliminate serialization costs while moving data between different processes
- Distributed Active Storage Layer
  - Reduce data movement by filtering data within the storage layer
- Distributed Compute Layer
  - Distributed compute operations such as Joins, GroupBy using MapReduce/BSP over MPI/UCX/RDMA
- Transactions over Datasets: *Lakehousing*
  - Features of warehousing such as **transactions**, **views**, **time travel**, **schema evolution** but directly over data lakes
- Expressive Query interface and Query compiler
  - Different query interfaces generating standard query plans acceptable by popular data processing systems

Implementation Plan

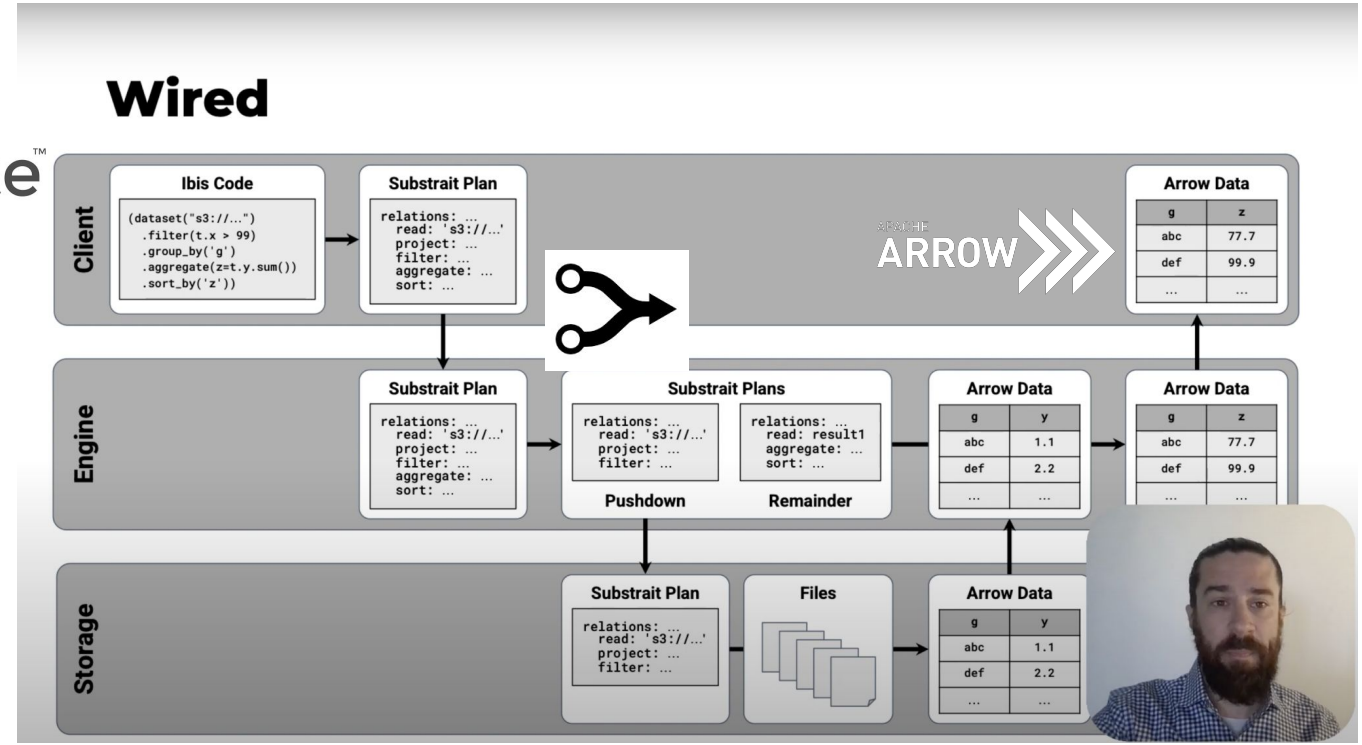

Ian Cook, "Arrow and Substrait: Better Together," The Data Thread, 6/23/22, https://youtu.be/5JjaB7p3Sjk

Implementation Plan

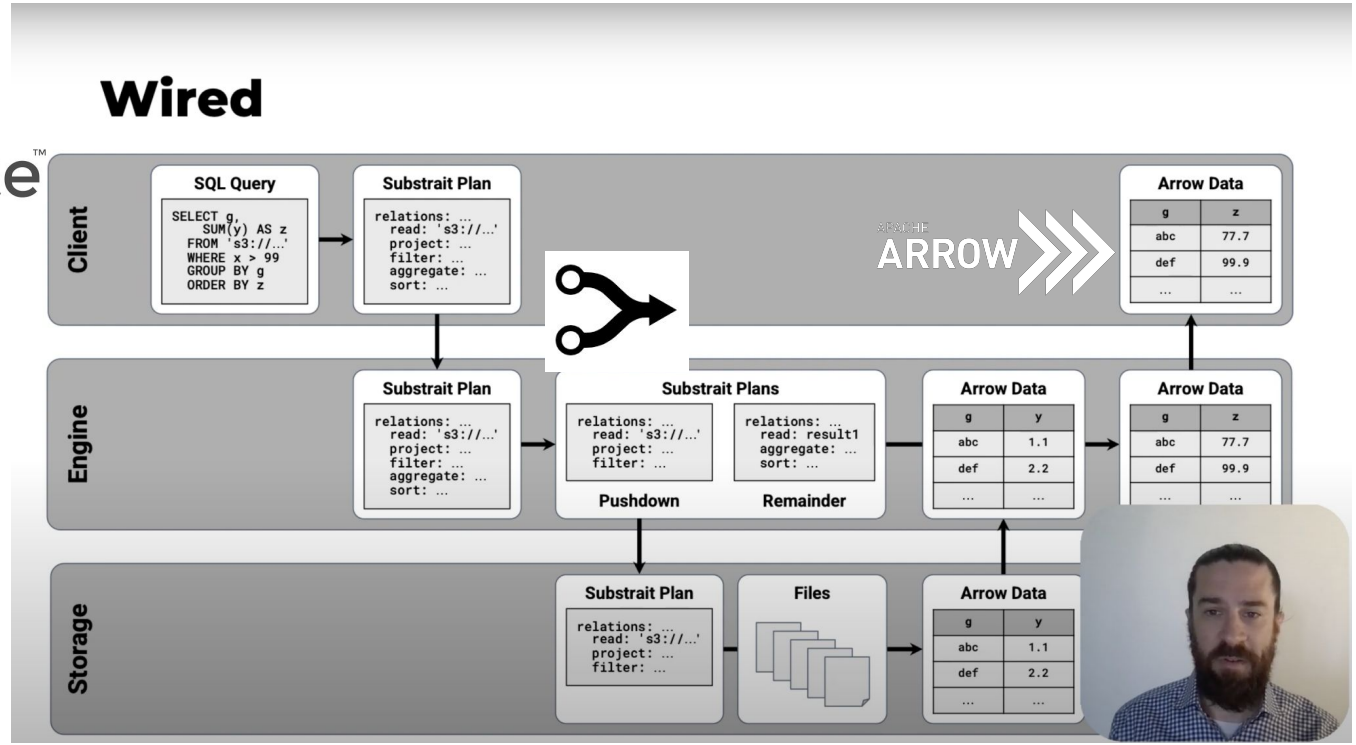

Ian Cook, “Arrow and Substrait: Better Together,” The Data Thread, 6/23/22, https://youtu.be/5JjaB7p3Sjk

Skyhook Architecture

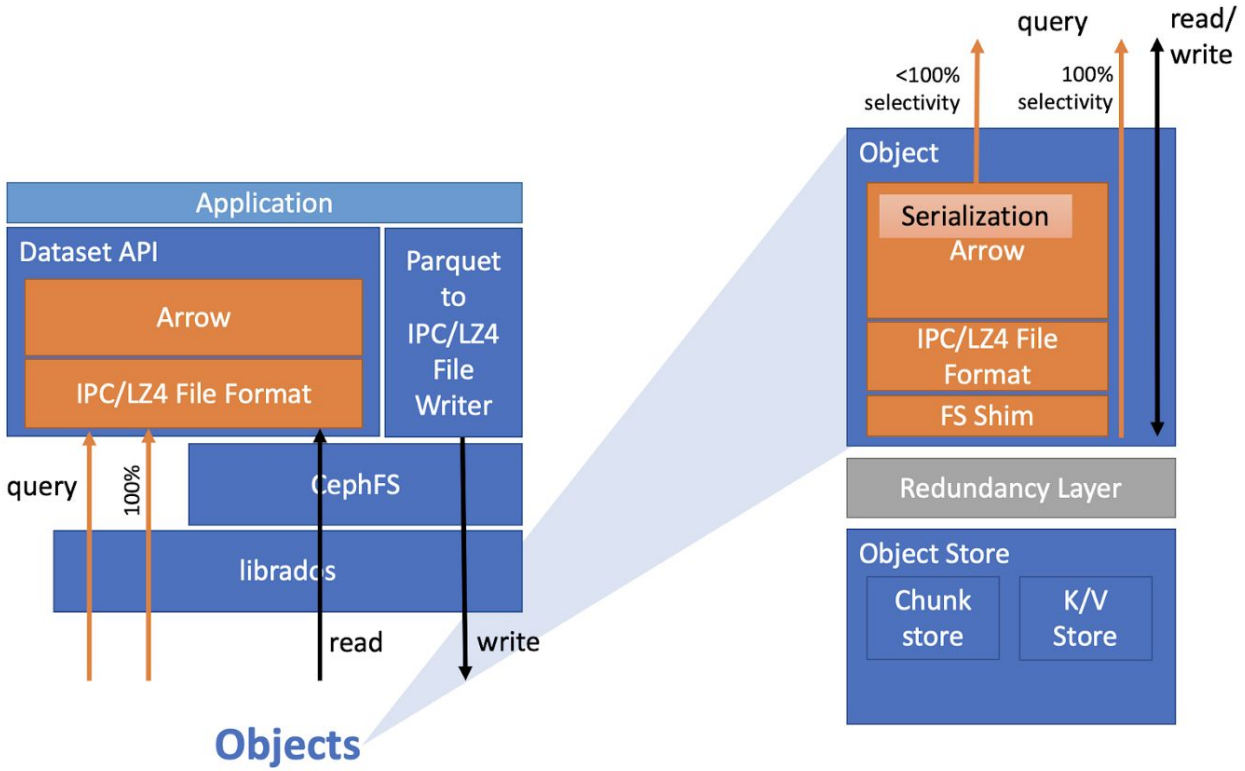

Current Status: Skyhook upstreamed in Apache Arrow !

Skyhook: Bringing Computation to Storage with Apache Arrow

PUBLISHED 31 Jan 2022
BY Jayjeet Chakraborty, Carlos Maltzahn, David Li, Tom Drabas

CPUs, memory, storage, and network bandwidth get better every year, but increasingly, they're improving in different dimensions. Processors are faster, but their memory bandwidth hasn't kept up; meanwhile, cloud computing has led to storage being separated from applications across a network link. This divergent evolution means we need to rethink where and when we perform computation to best make use of the resources available to us.

SkyhookDM is now a part of Apache Arrow!

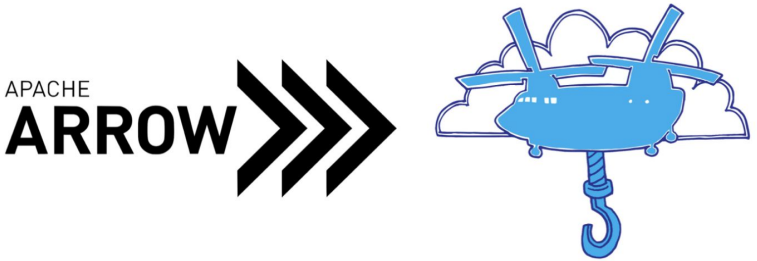

We are happy to announce that Skyhook Data Management is now officially a part of the Apache Arrow project mainline and is planned to be included in release 7.0.0. SkyhookDM is a plugin for offloading computations involving data processing operations into the storage layer of distributed and programmable object storage systems. It is be[illegible] aintained by researchers at

Acknowledgements

This work was supported by the National Science Foundation under Cooperative Agreement OAC-1836650 (IRIS-HEP.org), and the Center for Research in Open Source Software, UC Santa Cruz (cross.ucsc.edu)

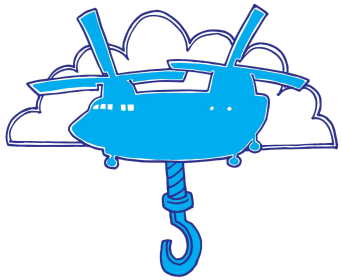

Thank You

jayjeetc@ucsc.edu

https://iris-hep.org/projects/skyhookdm.html