

# Using Multiple Composable, Hardware-Accelerated Executors (Lightning Talk)

Felipe Aramburu<sup>1</sup>

<sup>1</sup>Voltron Data Inc., 650 Castro Street Suite 120, PMB 96571, San Francisco, CA 94041

## Abstract

By design, a composable data system architecture provides the flexibility to use many execution paradigms. At Voltron Data, to prove this out, we are developing an execution system that uses the same logical plans on both CPU and GPU executors: (1) GPU executor is RAPIDS cuDF from NVIDIA, and (2) the CPU executor is Velox from Meta.

The number of FLOPs is steadily increasing and the costs for memory is decreasing quickly [1], [2]. There is a diverse lineup of accelerated hardware options, from FPGAs, to GPUs, Vector Processors, and increasingly larger multicore CPUs. This combined with the evolution of high speed interconnects like NVIDIA NVlink, CXL, PCI-E 5.0 mean that teams building data analysis systems that need to be able to quickly adapt to new hardware and execution capacities.

This talk will briefly introduce and discuss using multiple composable, open-source executors using hardware accelerators in data systems. In order to be able to exchange executors easily there are a few things that the developers building these systems must keep in mind:

Resource management cannot be the domain of the executor. Executors will frequently have to share resources, so resource management must be handled outside of them. This also means that the execution backends must have a way of requesting resources from your resource managers.

Separate logical plans and tasks from physical ones. Different backends will have implemented different primitives. It is up to the developer of these systems to find a way to convert a logical plan into a physical one that can be executed on the backend. Open Standards like Substrait (<https://substrait.io/>) can be used to provide a common framework for expressing these logical and physical plans.

Executors need to be able to reschedule tasks that failed to recoverable errors. With multiple execution backends that know nothing of each other, situations exist where a particular a resource gets exhausted, or a plan that should be able to be processed fails on one of your execution backends. When this occurs logical plans can be rescheduled on the system so that they can be retried later by the same executor or by another.

## Keywords

Composable Data Management Systems, Hardware Accelerated Executors

## References

- [1] Marius Hobbhahn and Tamay Besiroglu (2022), "Trends in GPU price-performance". Retrieved from: '<https://epochai.org/blog/trends-in-gpu-price-performance>'
- [2] John C. McCallum (2022), "Historical Cost of Compute and Storage". Retrieved from: '<https://ourworldindata.org/grapher/historical-cost-of-computer-memory-and-storage>'

*Joint Workshops at 49th International Conference on Very Large Data Bases (VLDBW'23) – Second International Workshop on Composable Data Management Systems (CDMS'23), August 28 - September 1, 2023, Vancouver, Canada*

✉ [felipe@voltrondata.com](mailto:felipe@voltrondata.com) (F. Aramburu)

🌐 <https://www.voltrondata.com/> (F. Aramburu)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)