

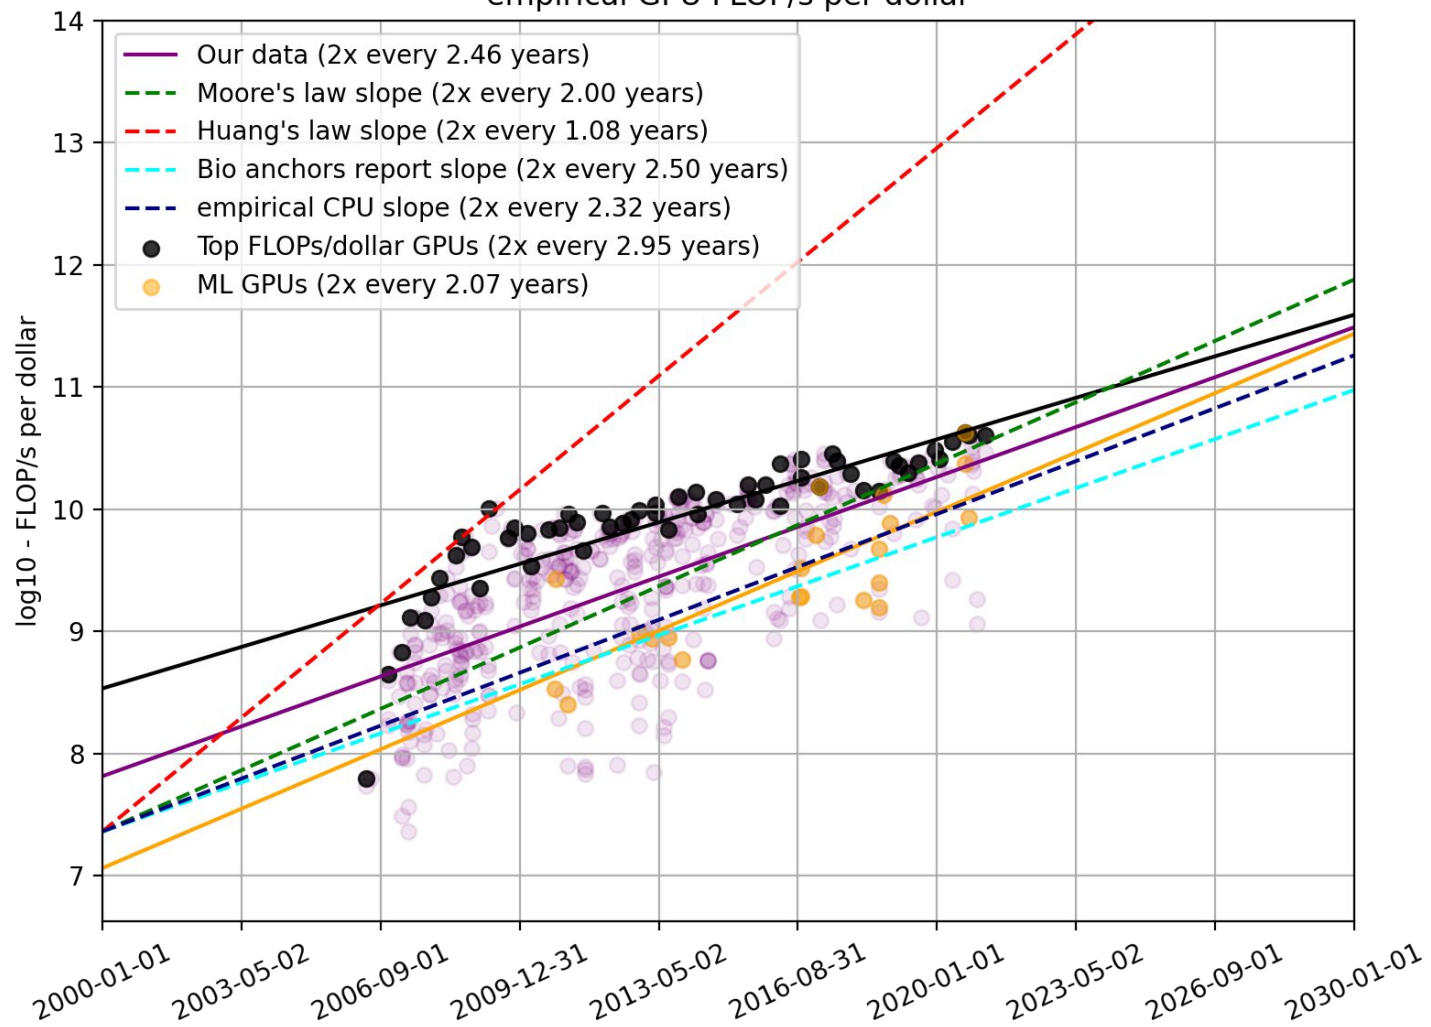
Composable Hardware Accelerated Executors

Felipe Aramburu
Voltron Data Inc.

Why


- Multiple Hardware Providers
 - GPU, FPGA's, Multicore CPUs
- Multiple Execution Frameworks
 - GPU
 - Rapids, OpenCL, Cub
 - CPU
 - Velox, Arrow Compute, Xsimd
- Rapidly evolving interconnects
 - Nvidia NVLink, PCI-E 5.0, CXL
- Changes to Memory / solid state cost and performance

empirical GPU FLOP/s per dollar



Historical cost of computer memory and storage

This data is expressed in US dollars per terabyte (TB). It is not adjusted for inflation.



Source: John C. McCallum (2022)

OurWorldInData.org/technological-change • CC BY

Note: For each year, the time series shows the cheapest historical price recorded until that year.

Some Components of Composable Executors

- Memory Management
- Logical and Physical Plans
- Task Management
- Data Flow

Memory Management

- Executors Allocators must be customizable
 - Tracking allocations
 - Managing resource constraints
- Need APIs for reserving memory
- Memory needs to be managed in centralized place

Logical and Physical Plans

- Logical Plans are necessary to integrate with multiple Client types
- Physical Plan are necessary to leverage multiple Backends
 - A physical plan per backend is often necessary
 - Heterogenous plans are also possible

Task Management

- Tasks Need to be able to recover from resource exhaustion
- Tasks Should be specific to an execution backend
- Tasks should be pipelined to executors in order improve performance

Data Flow

