COLORADO SCHOOL OF MINES

HPC Budgeting and Systems

How to get your HPC workload done?

September 9, 2024

Learning Objectives

- Identify types HPC Workloads
- Explain Simulation Metrics used to benchmark HPC Workflows
- Perform benchmarking and identify efficient job allocation
- Determine allocation allotment to complete a HPC project
- Define the two HPC options from Mines and two off-campus resources

Life Cycle of Research Computing Projects

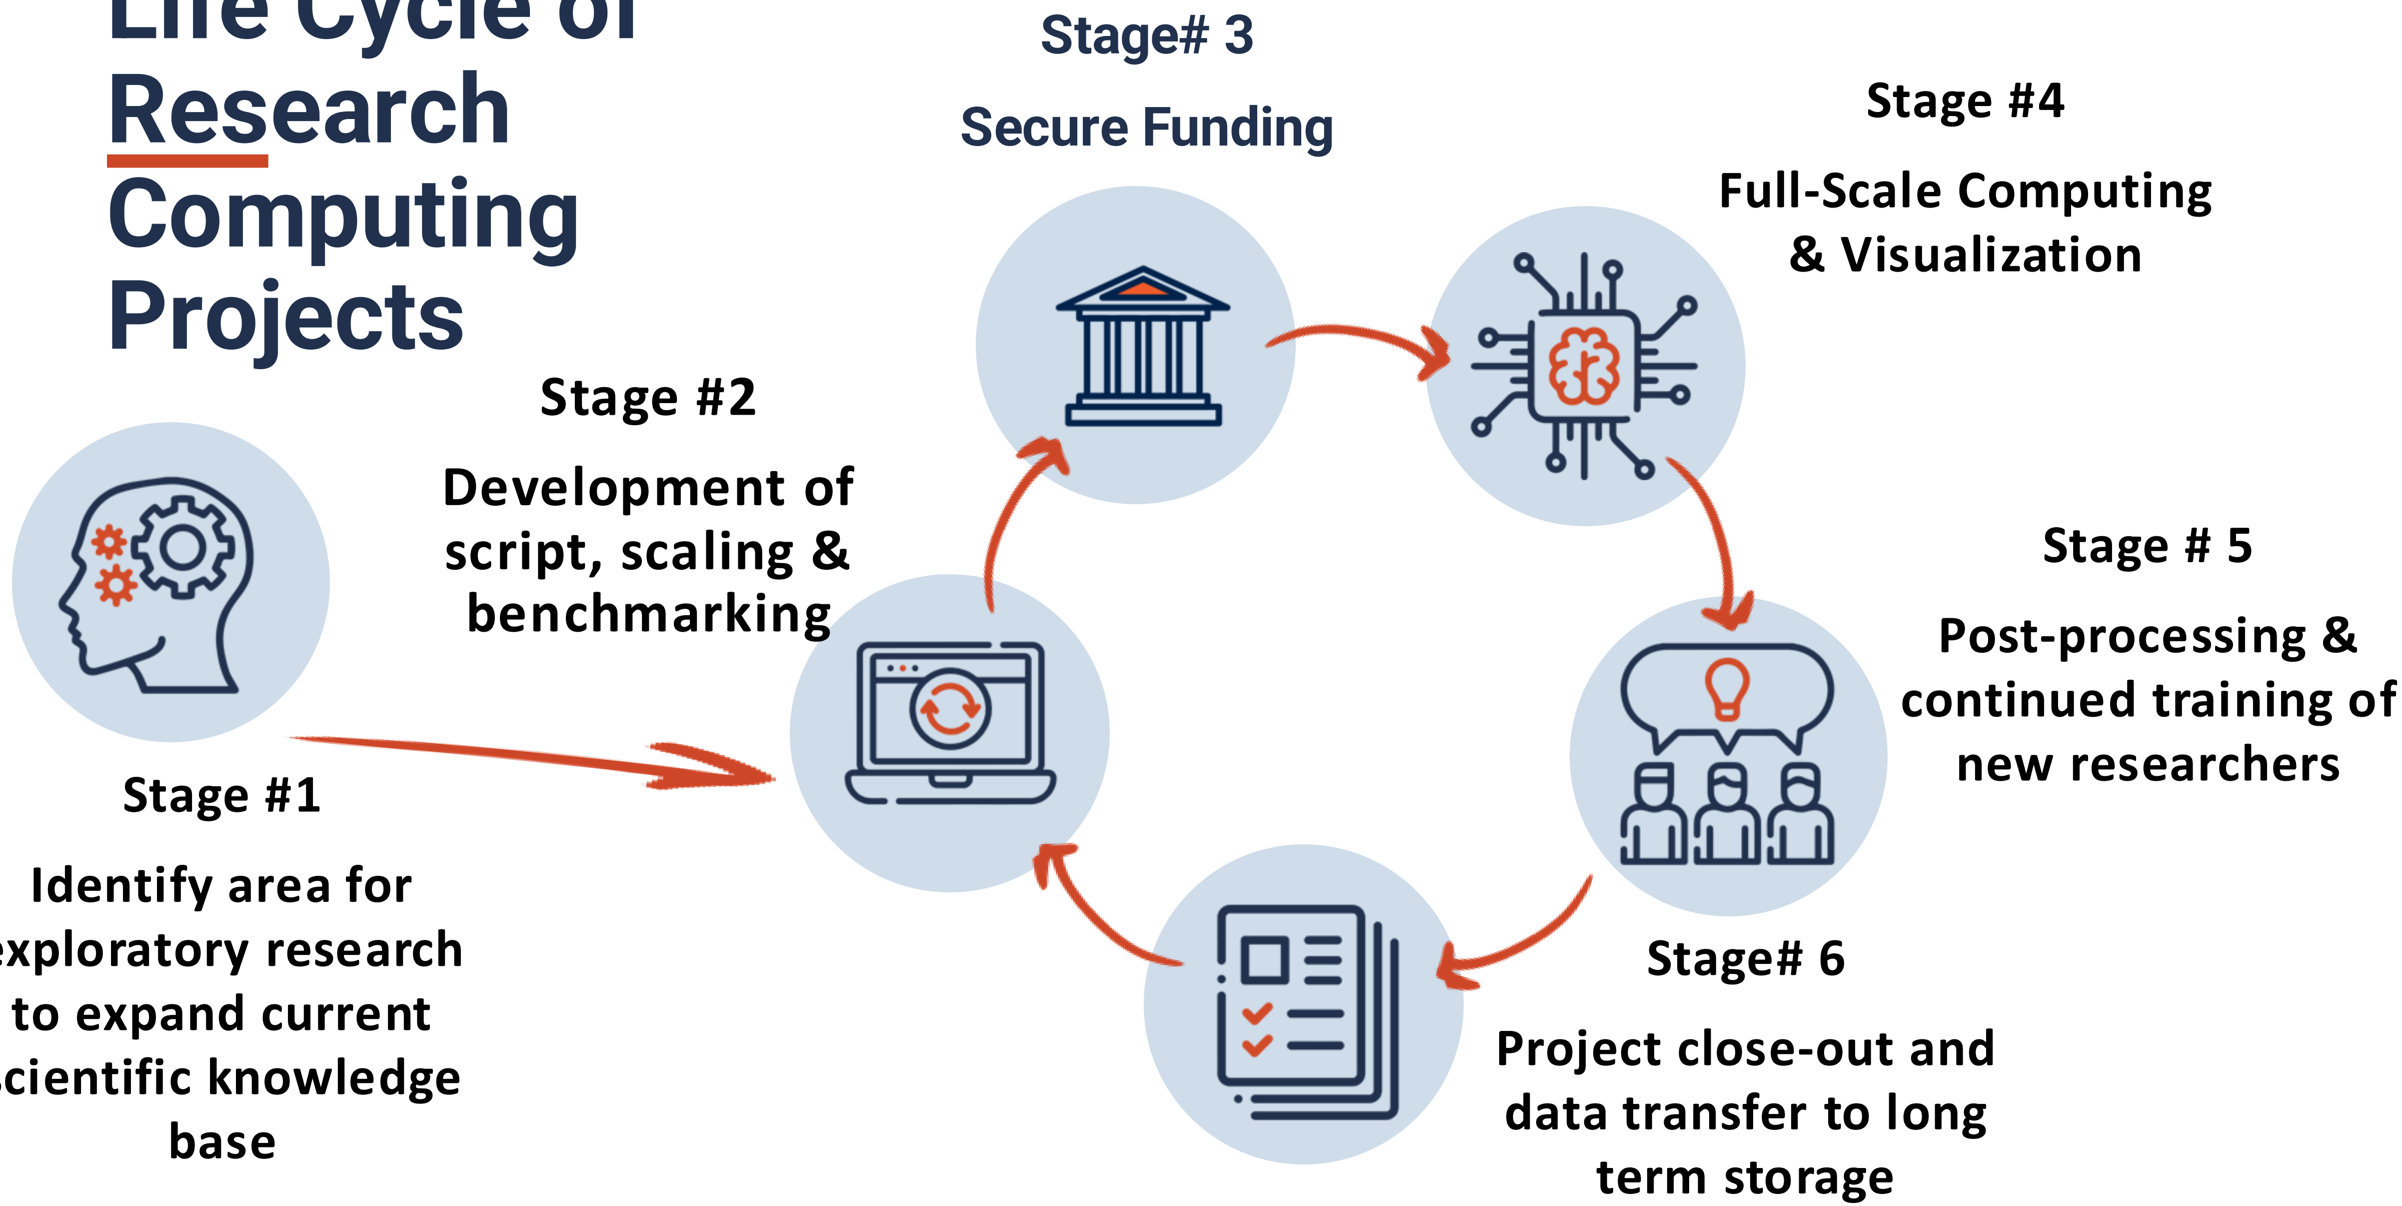

COLORADO SCHOOL OF MINES

Research Computing (rc.mines.edu)

Type of High-Performance Computing Workloads

From Serial to Parallel and Beyond

September 5, 2024

Computational Workloads and Workflows

Serial Workloads

A separate parameter solved on each core/processor

- One protein compared to a binding location
- One mesh-node optimized
- One scenario with a set of parameters
- Temporal/Transient Simulations

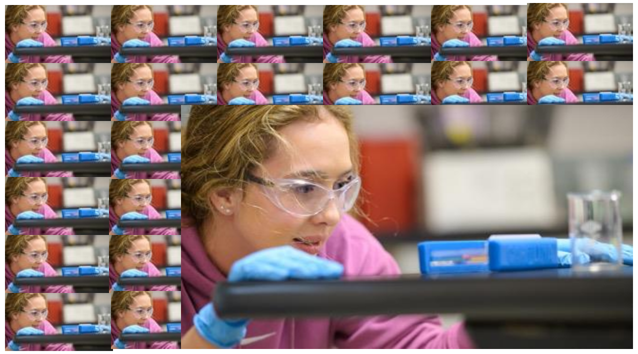

Parallel Workflows

Multi-node or Multi-threaded

Single or Multi GPU

Perfectly Parallelization of Serial Workflows

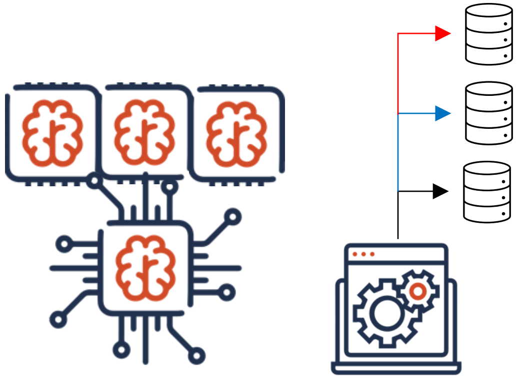

Disk or Memory Resource Requirements

High Memory per CPU

High Data Input/Output

Serial Workloads: One Scenario or Parameter

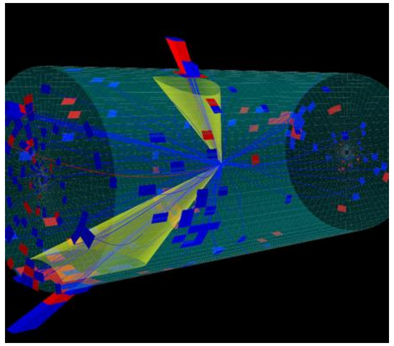

Monte Carlo Simulations

Argonne National Laboratory – Particle Collision Simulations

Soak Time	Pressure Level H	Pressure Level I	Pressure Level J
D	A,B,C	A,B,C	A,B,C
E	A,B,C	A,B,C	A,B,C
F	A,B,C	A,B,C	A,B,C
G	A,B,C	A,B,C	A,B,C

A	Hard Maple
B	Black Walnut
C	Redwood
D	20 min.
E	40 min.
F	80 min.
G	120 min.
H	Atm. Pressure
I	100 psig
J	800 psig

Parameter Sweeps

Design of Experiment problems – Single parameter change on fixed scenario

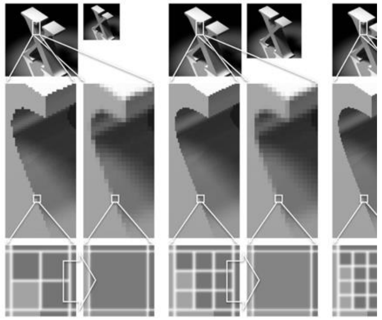

Pixel Rendering and Anti-Aliasing

Single pixel rendering reduction in image size

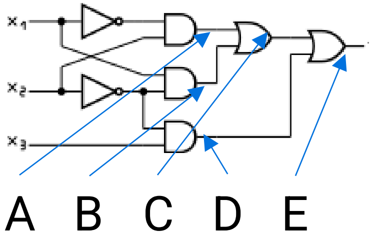

Circuit Value Problems

To solve E you must solve C and D…

To Solve C you must solve A and B

To solve D you must solve A, B, and C

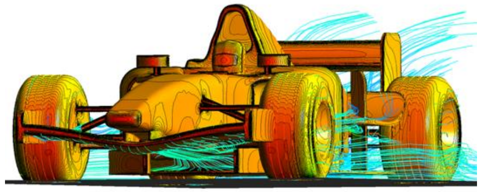

Large-Scale Iterative Simulations

Unsteady flow, transient, highly iterative simulations

Perfectly Parallel
(Embarrassingly parallel)

Linear Solutions

Parallel Workflows: Multiple Scenarios or Parameters

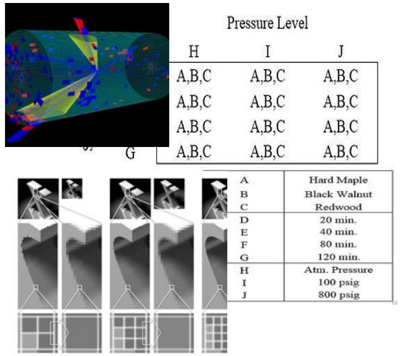

Multi- Monte Carlo, Parameter Sweeps, Pixel Rendering

Using Multiple jobs to execute the code

Scaling up has no limit

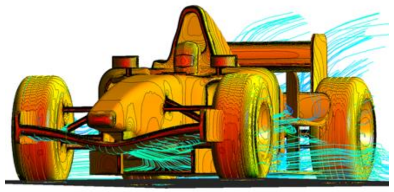

Large-Scale Iterative Simulations

Fluid flow simulations with Ansys Fluent, OpenFoam

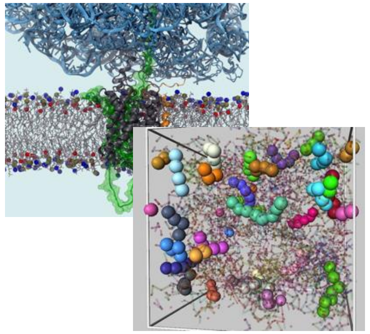

Molecular Dynamics

Proteins or Material Science using Gromacs, VASP, LAMMPS

Geophysics Simulations

Earthquakes, Inverse-wave problem, etc

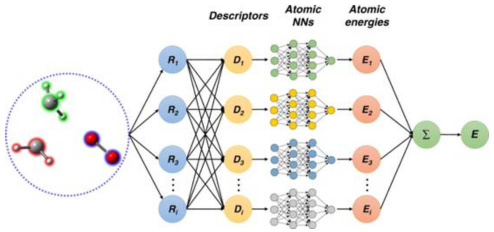

Machine Learning

Computer Vision, Large-Language, Neural Networks, Training, Inference, etc.

Perfectly Parallel (Embarrassingly parallel)	Optimization Need to Problem Size

High Memory or Data I/O Workflows

High Memory usage to store large matrixes

Computing a large matrix for data reduction

Mesh optimization problems

Visualization of data sets

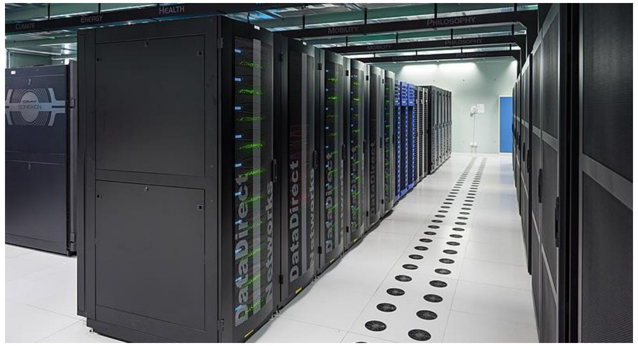

High Data Input and Output

Image processing of neural networks

Movie frame rendering

File server and database search

Code Development

Un-optimized coding leads to limited problem sizes that often must be simplified

Research Computing (rc.mines.edu)

High-Performance Computing Simulation Metrics

Metrics are everywhere

September 5, 2024

Simulation Reporting Performance

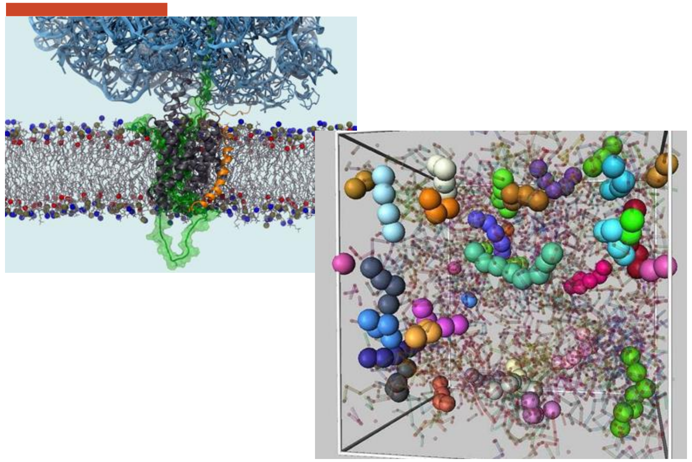

Molecular Dynamics

Nano-seconds of Simulated Time
Time in Actual Day

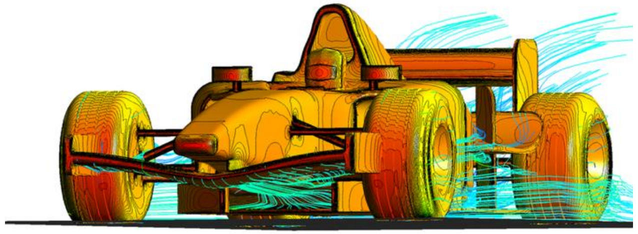

Computational Fluids Dynamics (CFD)

Iteration
Second

Simulation Time
Second

Jobs Completed
Day

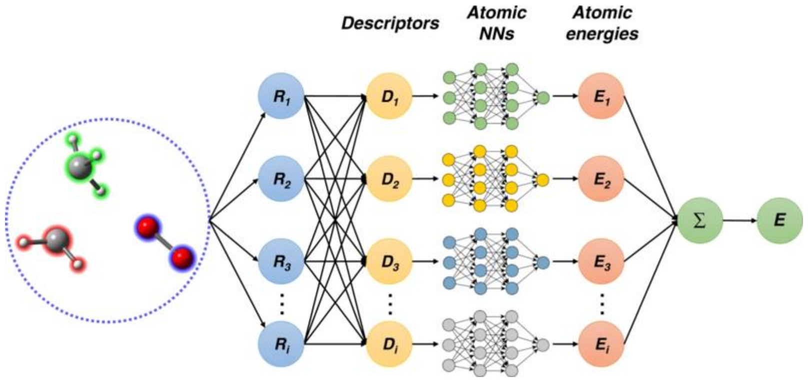

Machine Learning

Epochs
Hour

Models
Day

Parameters to Determine

Nodes x Cores

Total number of Cores used for simulation

Total Jobs

Jobs needed to complete research project

Core Run Time (Cores * Hours)

Work done by each CPU (CPUh)

Storage of Input and Output

Data sets inputs (multi-use or single-use)

Cost / (Core-Hour)

Wendian = $0.02/(Core-Hour)
AWS (apex.mines.edu) $0.04 and up and GPUs even more $96/hour

Storage during & after Simulation

S3-Bucket, Orebits, fast NVMe

Job Parametric Study

Number of Simulations to Assess Science Domain

GPU-hour (GPUh)

Increased costs for specialize resource

Wendian GPU V100 = 6 core-hour or $0.12

COLORADO SCHOOL OF MINES

Research Computing (rc.mines.edu)

Benchmarking HPC and Efficient Jobs

How many cores is enough?

September 5, 2024

XDMoD

open.xdmod.org

- Management capabilities
  - Monitoring standard metrics: utilization
  - Metrics designed to identify underperforming systems hardware and software
  - Reporting job level performance data for every job running on the HPC

1. A tool to effectively and efficiently use an allocations and optimize HPC resources
2. Ability to monitor, diagnose, and tune system performance and measure the performance of all applications running
3. Easily obtain detailed analysis of application performance to aid in optimizing code performance
4. A diagnostic tool to facilitate HPC planning and analysis
5. Metrics to help measure scientific impact.
6. Analyses of the operational characteristics of the HPC environment can be carried out at different levels of granularity
   - job, user, or on a system-wide basis.

xdmod.mines.edu

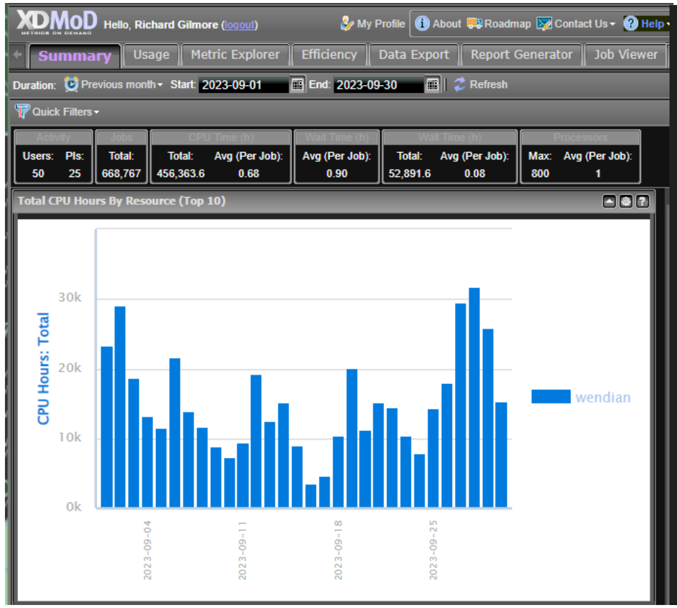

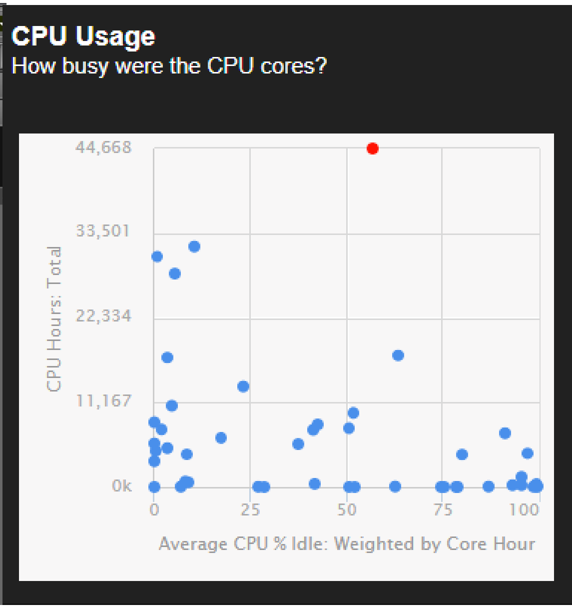

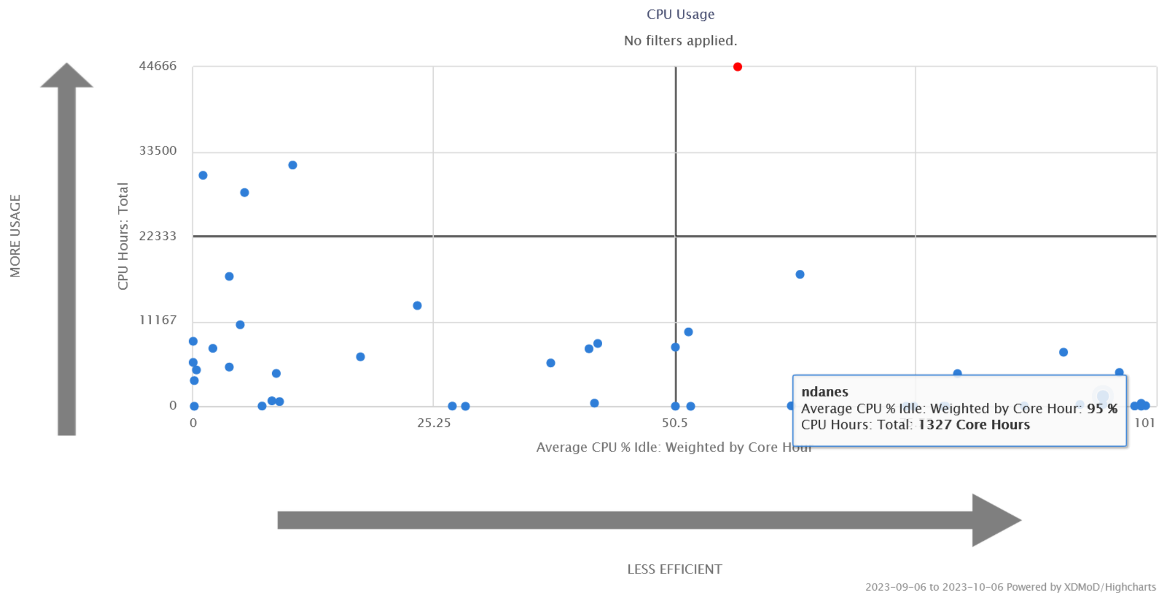

seff <job-id>

[Seff SchedMD/SLURM](#)

- Command Line Utility (after job completion)

```
[rgilmore@wendian001 ~]$ seff 23942161
Job ID: 23942161
Array Job ID: 23942161_0
Cluster: wendian
User/Group: ndanes/ndanes
State: FAILED (exit code 1)
Nodes: 1
Cores per node: 4
CPU Utilized: 00:00:01
CPU Efficiency: 25.00% of 00:00:04 core-walltime
Job Wall-clock time: 00:00:01
Memory Utilized: 1.52 MB
Memory Efficiency: 0.01% of 20.37 GB
```

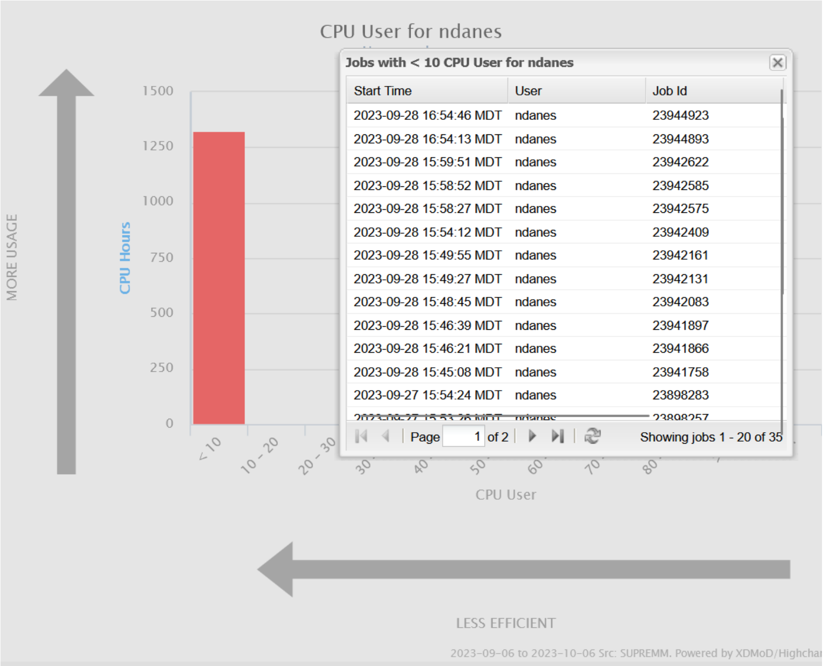

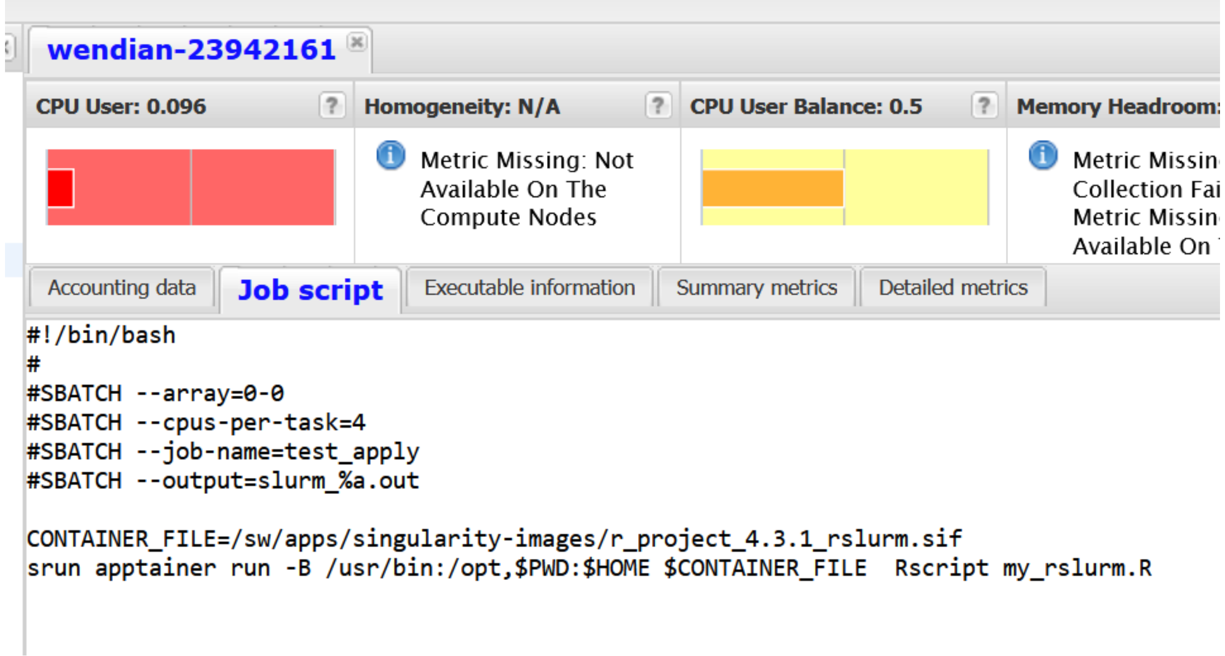

Rscript – 3 nodes

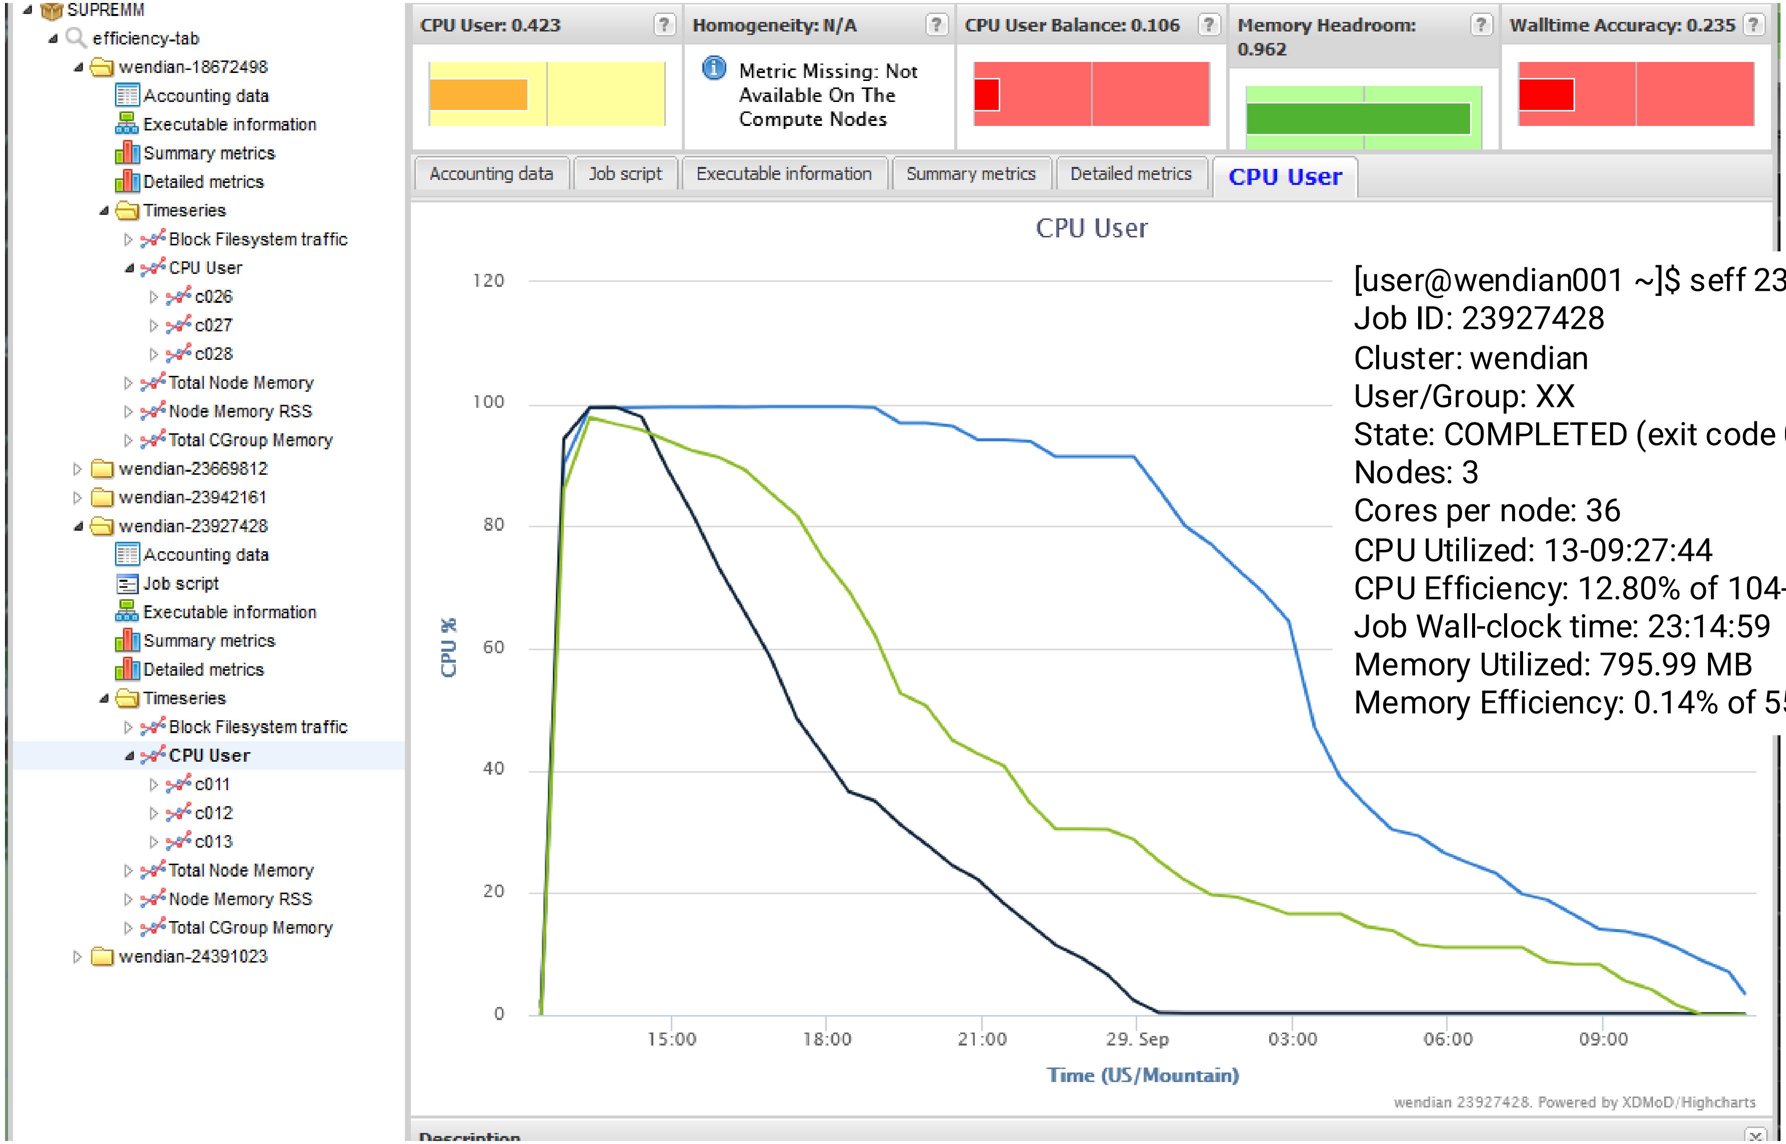


```
[user@wendian001 ~]$ seff 23927428
Job ID: 23927428
Cluster: wendian
User/Group: XX
State: COMPLETED (exit code 0)
Nodes: 3
Cores per node: 36
CPU Utilized: 13-09:27:44
CPU Efficiency: 12.80% of 104-14:58:12 core-walltime
Job Wall-clock time: 23:14:59
Memory Utilized: 795.99 MB
Memory Efficiency: 0.14% of 550.02 GB
```

COLORADO SCHOOL OF MINES

Research Computing (rc.mines.edu)

XDMoD lab of your GROMACS Jobs (day after completion)

What was the most efficient? What was the fastest?

September 5, 2024

GROMACS Benchmarking

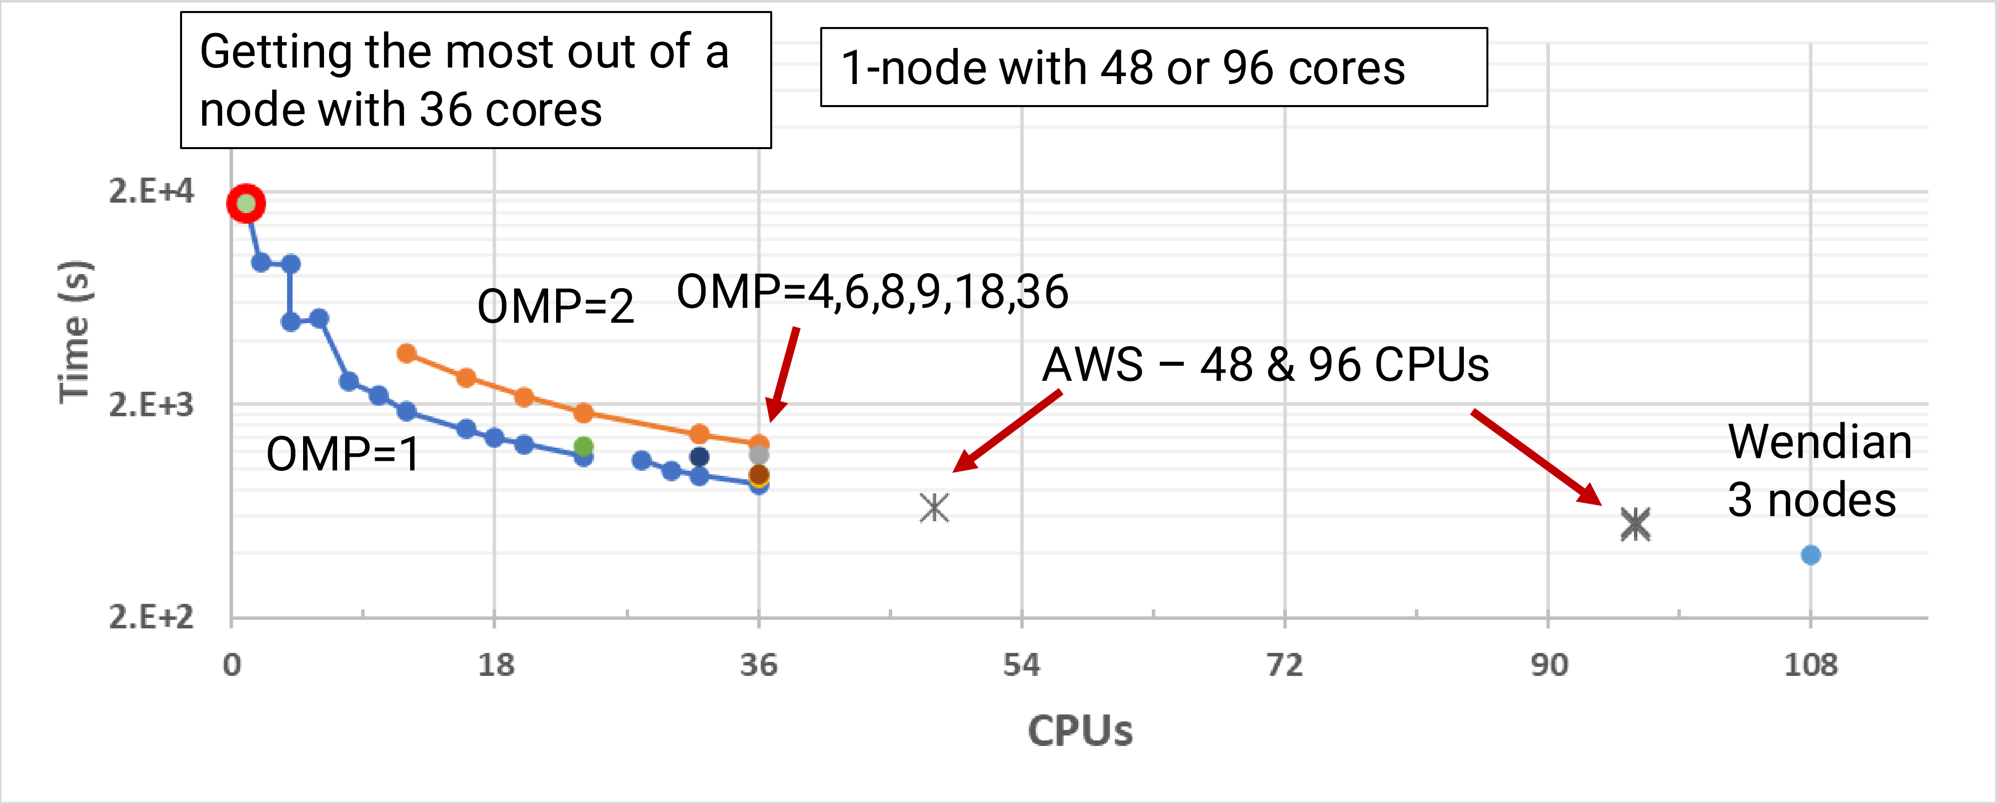

Where is the processing bottle neck?

COLORADO SCHOOL OF MINES

Cyber-Infrastructure and Advanced Research Computing (CIARC.mines.edu)

The Cost of HPC

All the simulations that you need to get done

October 12, 2023

Equation for HPC Project Cost

$$\left[\frac{\text{Nodes} \times \text{Cores}}{\text{Total Cores}}\right]\left[\text{Hours}\right]\left[\frac{\$}{\text{Core-Hour}}\right]\left[\text{Jobs}\right] =$$

- [Nodes X Cores / Total Cores] [Hours]: Optimize with benchmarks & Design of Experiment
- [$ / Core-Hour]: HPC Cluster Specific
- [Jobs]: Learning, Production, Post-Processing

GROMACS Benchmarking

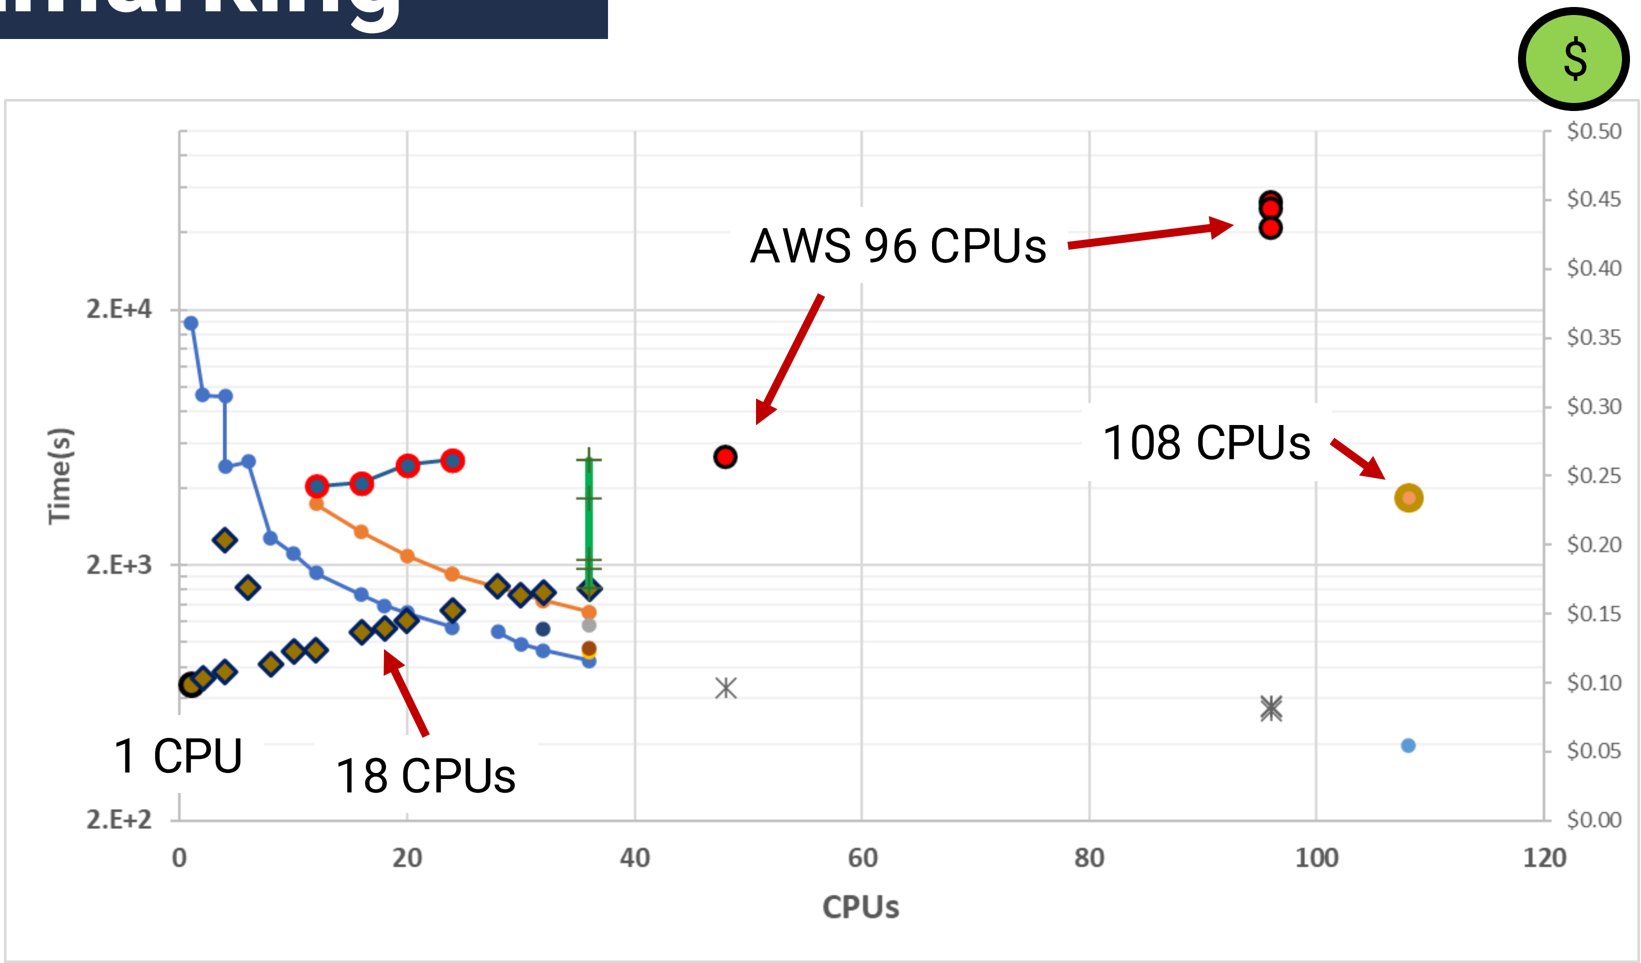

GROMACS Benchmarking

What is the simulation metric we need to optimize?

- Jobs/day
- Iteration/sec
- Nano-sec/day

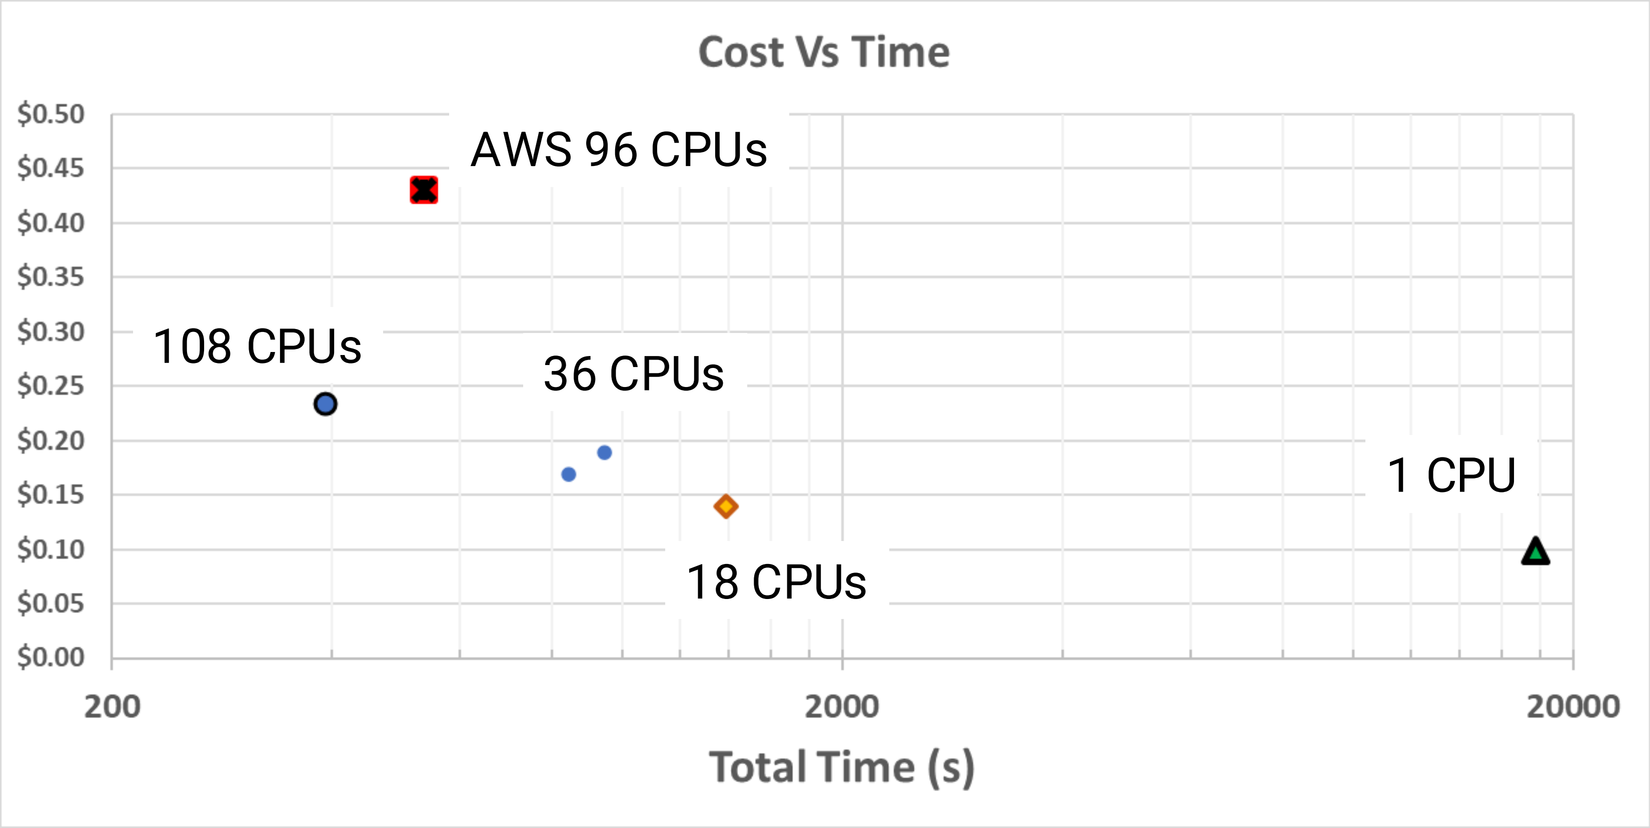

GROMACS Estimate Example

$$\left[\frac{\text{Nodes X Cores}}{18}\right]\left[\text{0.39-Hours}\right]\left[\frac{\$0.02}{\text{Core-Hour}}\right]\left[\frac{\text{300-jobs, 10,000-jobs, 1,000-jobs}}{11{,}300}\right] = \$1600$$

Wendian pricing

Optimize with benchmarks & Design of Experiment

HPC Cluster Specific

Learning, Production, Post-Processing

GROMACS Estimate Example

$$\left[\frac{\text{Nodes} \times \text{Cores}}{108}\right]\left[0.11\text{-Hours}\right]\left[\frac{\$0.02}{\text{Core-Hour}}\right]_{\text{Wendian pricing}}\left[\frac{\text{300-jobs, 10,000-jobs, 1,000-jobs}}{11{,}300}\right] = \$2700$$

Wendian pricing

Optimize with benchmarks & Design of Experiment

HPC Cluster Specific

Learning, Production, Post-Processing

GROMACS Estimate Example

[Nodes X Cores / 96] [0.15-Hours] [AWS pricing $0.03 / Core-Hour] [300-jobs, 10,000-jobs 1,000-jobs / 11,300] = $4900

- Optimize with benchmarks & Design of Experiment
- HPC Cluster Specific
- Learning, Production, Post-Processing

Life of a HPC Project (1-researcher)

Gromacs

- Learning
  - 6-9 months
  - Big and small models
  - No optimization
  - Cost $1000s/m
- Production
  - 12-18 months
  - Right sized models
  - Optimized
  - Cost $1000s/m
- Post-processing (papers, etc.)
  - 6 months
  - Rebuild models
  - No optimization
  - Cost $1000s/m

$25-40k

Fluent

- Learning
  - 12 months
  - Big models
  - Cost $1000s/m
- Production
  - 12 months
  - 1-2 models
  - Optimized
  - Cost $1000s/m
- Post-processing (paper, etc)
  - 6 months
  - Rebuild models
  - Cost $1000s/m
- Mesh independence
  - 3 months
  - 4-8 Really big models
  - Cost $1000s/m

$70k

Research Computing (rc.mines.edu)

HPC Resources available from many sources

Can you get free HPC jobs?

September 5, 2024

HPC Systems Options for Researchers

- FREE
  - ACCESS (access-ci.org)
  - RMACC (Colorado.edu/rc/articles/rmaccsummit)
- PAID
  - Mines HPC: Wendian (rc.mines.edu)
  - Cloud: AWS

Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS)

- NSF funded Supercomputer hosted around the country
  - Georgia Tech, Indiana, Johns Hopkings, Open Science Grid, Pittsburgh, Purdue, San Diego, Texas A&M, etc.
- System connection available through Allocations Process (allocations.access-ci.org)
  - Authentication through CILOGIN

Create an account through provided CILOGIN.ORG

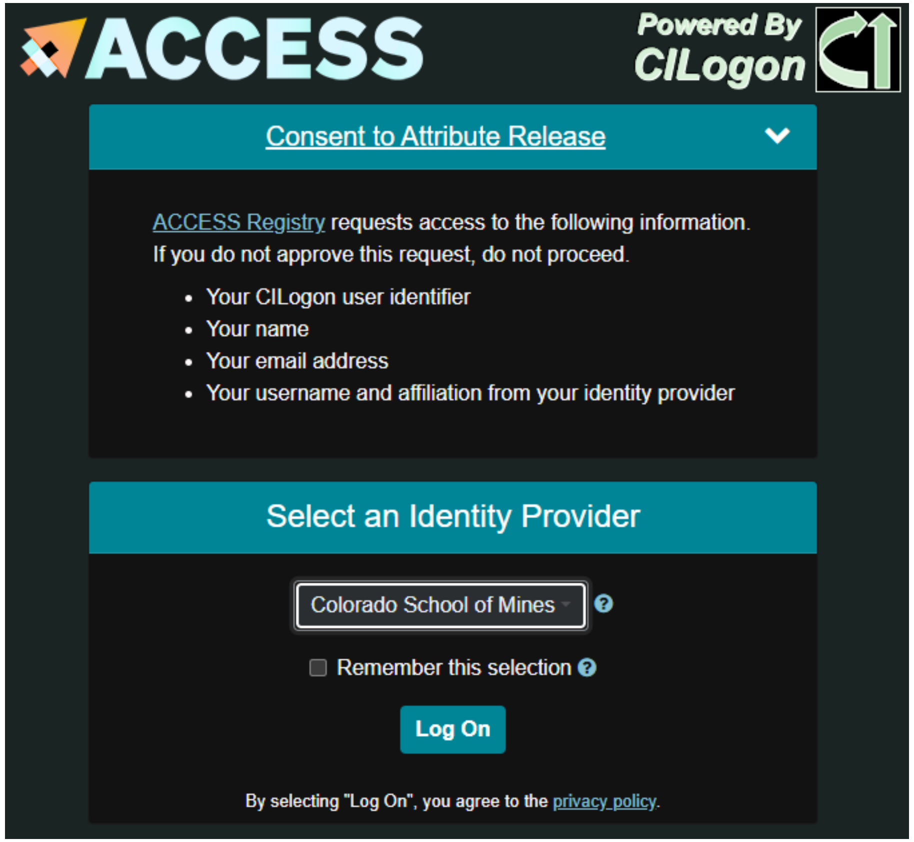

ACCESS Allocation Options

Allocation Type	Credit Threshold
Explore ACCESS	400,000
Discover ACCESS	1,500,000
Accelerate ACCESS	3,000,000
Maximize ACCESS	Not awarded in credits

- Exchange for resource-specific allocations
  - Often unit is 1-to-1 for CPU/hours
  - Higher exchange ratio for GPU/hours and newer systems
  - All rates are set by institution providers

  Exchange Calculator
  - Examples: Johns Hopkins (GPU cluster) 1000:19, Johns Hopkins (CPU cluster) 1000:1000 Purdue (CPU) 1000:833, TACC (Knights Landing Phi nodes) 1000:56

ACCESS Support and Training

All Provided by the MATCH Research Support team

- Create a Ticket
- Documentation
  - User guides, resource provides information, Support Knowledge Base
- Community Slack Channel
- Computational Science and Support Network (CSSN) Travel Grants and Rewards

MATCHPlus for 3-6 month projects

MATCHPremier for 12-18 month projects