

# ExaFEL: Achieving Real-Time XFEL Data Analysis using Exascale Hardware



(in alphabetic order)

NERSC: Katie Antypas, Bill Arndt, Debbie Bard, Wahid Bhimji, Shane Canon, Bjoern Enders, Lisa Gerhardt, Laurie Stephey, Rollin Thomas, Felix Wittwer

LBNL: Aaron Brewster, Asmit Bhowmick, Anna Giannakou, Daniel Paley, Lavanya Ramakrishnan, Nick Sauter, Elyse Schriber

SLAC: Wilko Kroeger, Derek Mendez, Amedeo Perazzo, Murali Shankar, Elliott Slaughter, Monarin Uervirojnangkoorn, Jana Thayer, Chuck Yoon

+ All the great folks at OLCF and ALCF

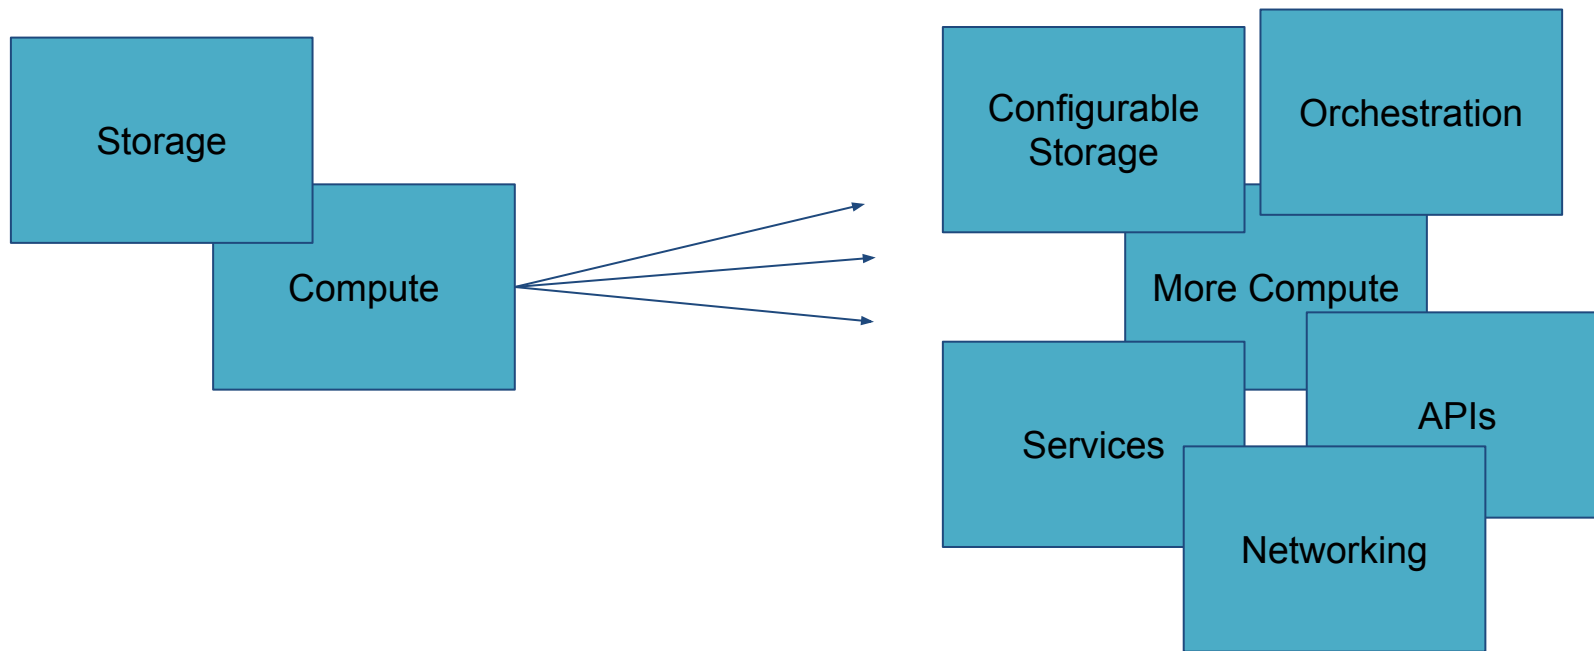
**Johannes Blaschke**

Data Science Engagement Group, NERSC  
IPAM, NMEW Workshop May 4, 2023

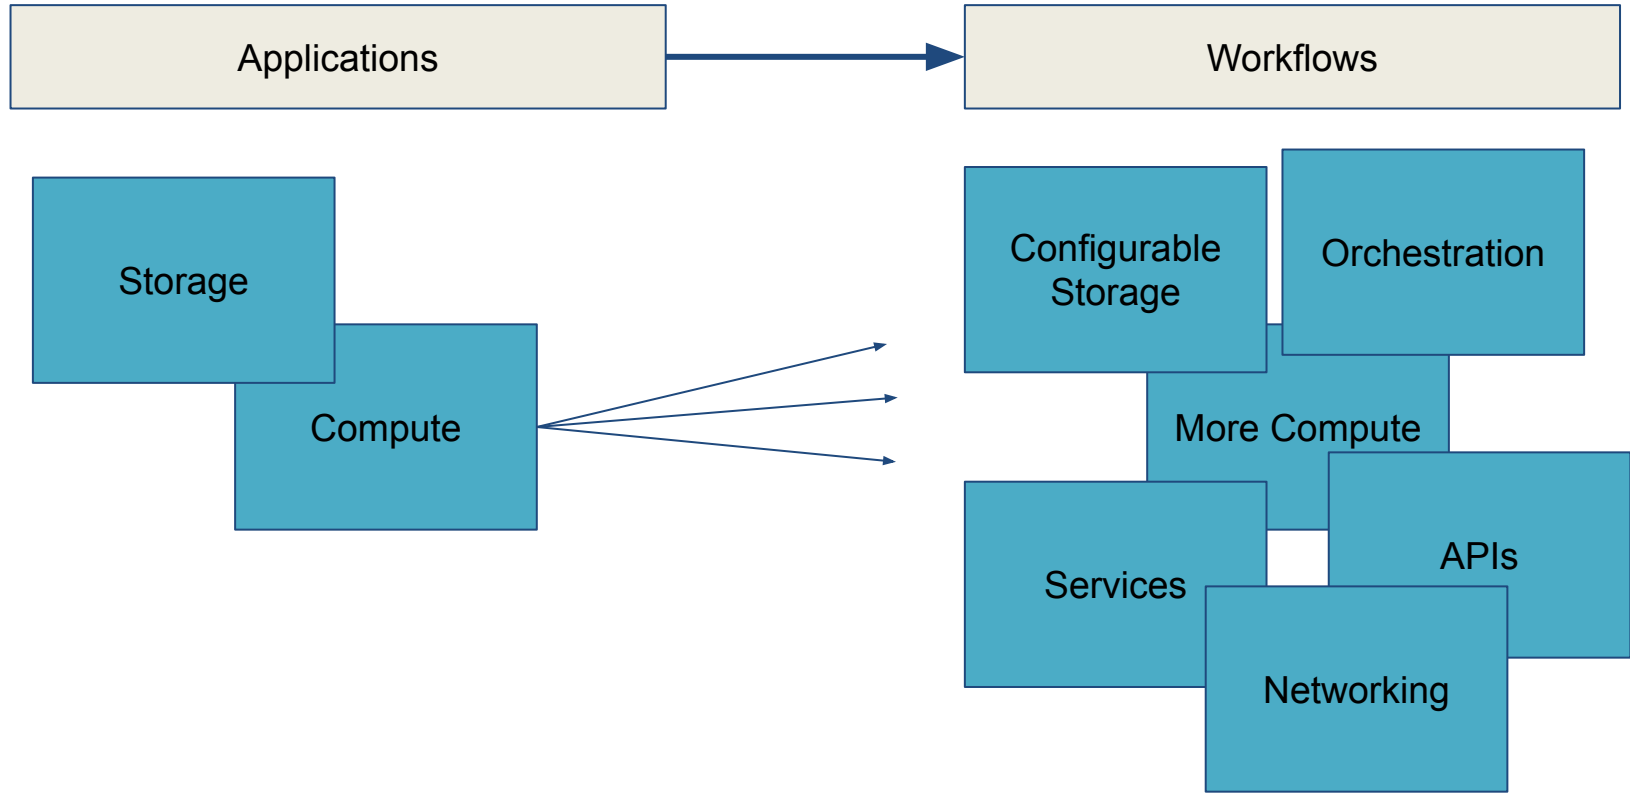
# HPC is Evolving



# HPC is Evolving

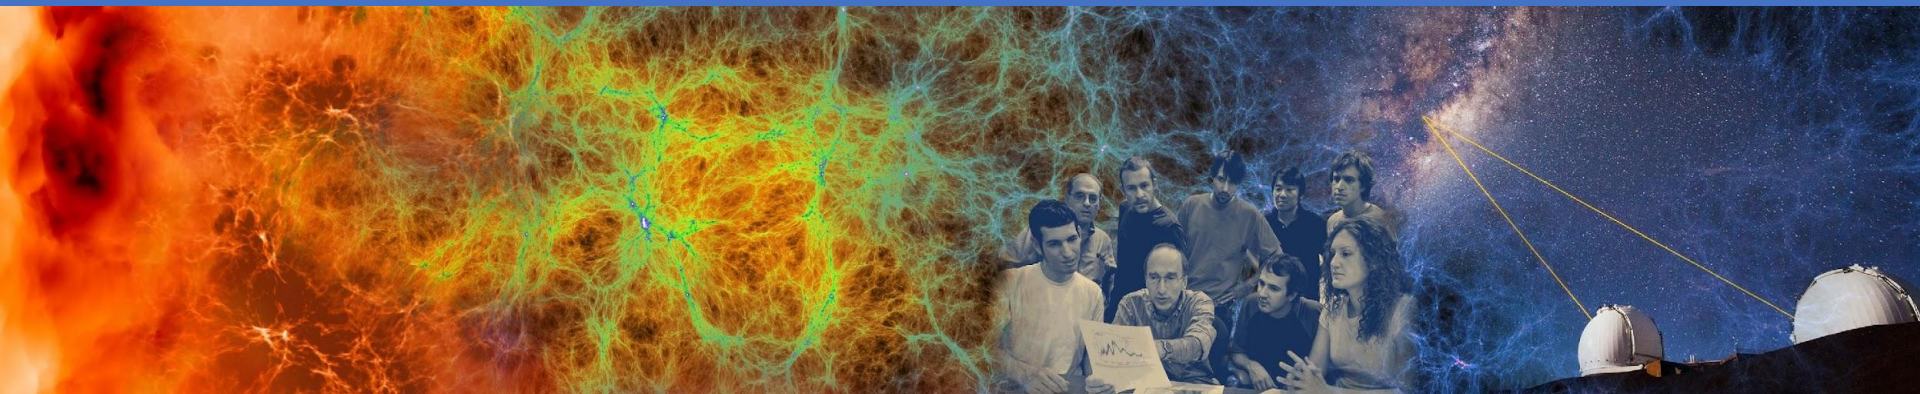


# HPC is Evolving



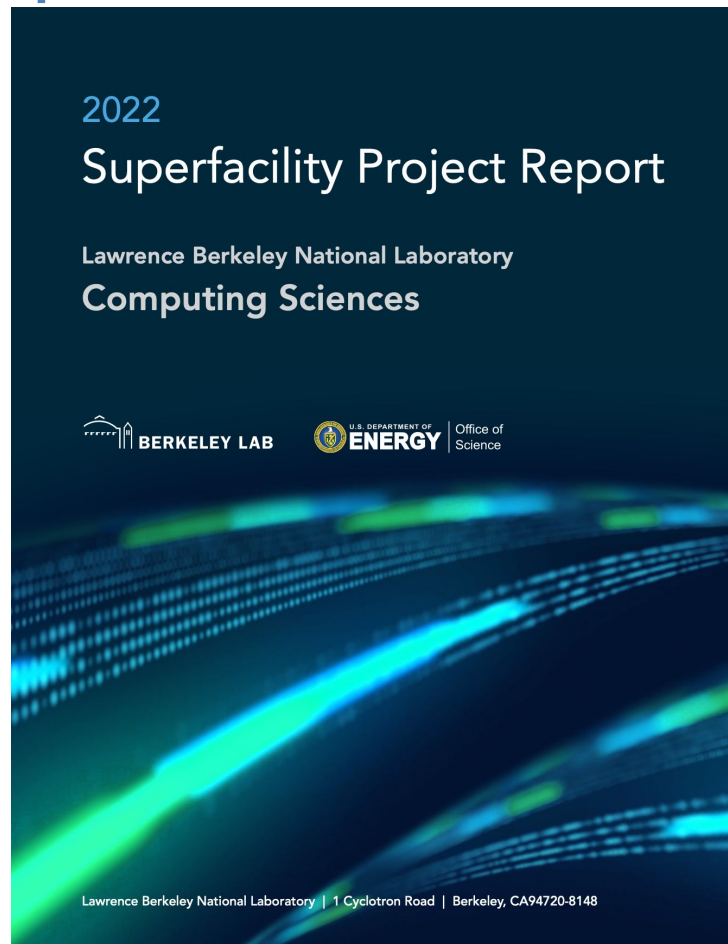


# Superfacility, and Integrated Research Infrastructure



# Bringing HPC “Closer” to Experiments

- Data collection rate at XFEL light sources expected to increase by 400x
- Expected to outpace local computing resources
- Large experiments require tight coupling between data collection and analysis  $\Rightarrow$  superfacility model



# New Generation of Workflows

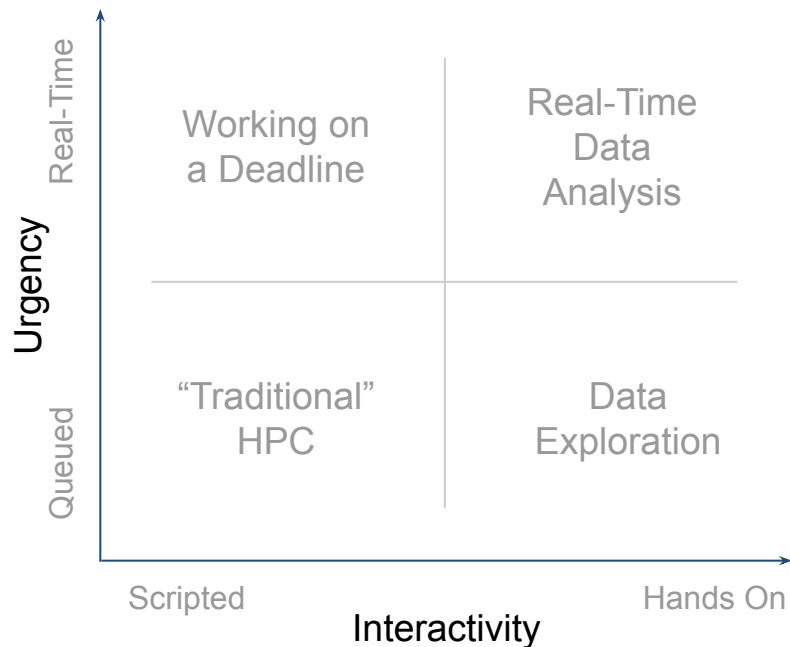
Similar need across many science domains! Eg.:

1. Astrophysics and Cosmology:  
*LSST DESC, CMB/TOAST, DESI, DUNE*
2. High-Energy and Particle Physics:  
*ATLAS, CMS, LZ*
3. Electron Microscopy and Nanoscience:  
*NCEM*
4. Nuclear Fusion Experiments:  
*KSTAR*

# New Generation of Workflows

Similar need across many science domains! Eg.:

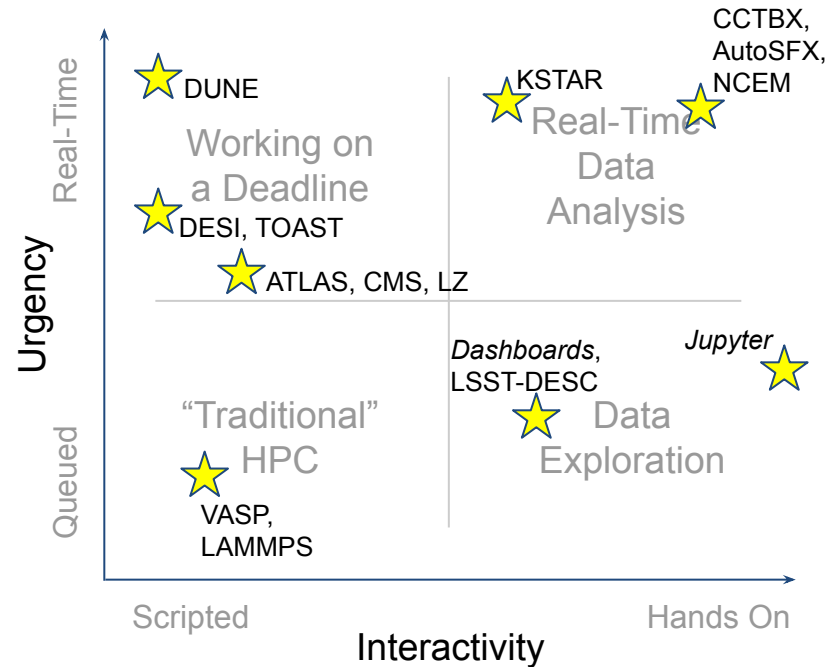
1. Astrophysics and Cosmology:  
*LSST DESC, CMB/TOAST, DESI, DUNE*
2. High-Energy and Particle Physics:  
*ATLAS, CMS, LZ*
3. Electron Microscopy and Nanoscience:  
*NCEM*
4. Nuclear Fusion Experiments:  
*KSTAR*

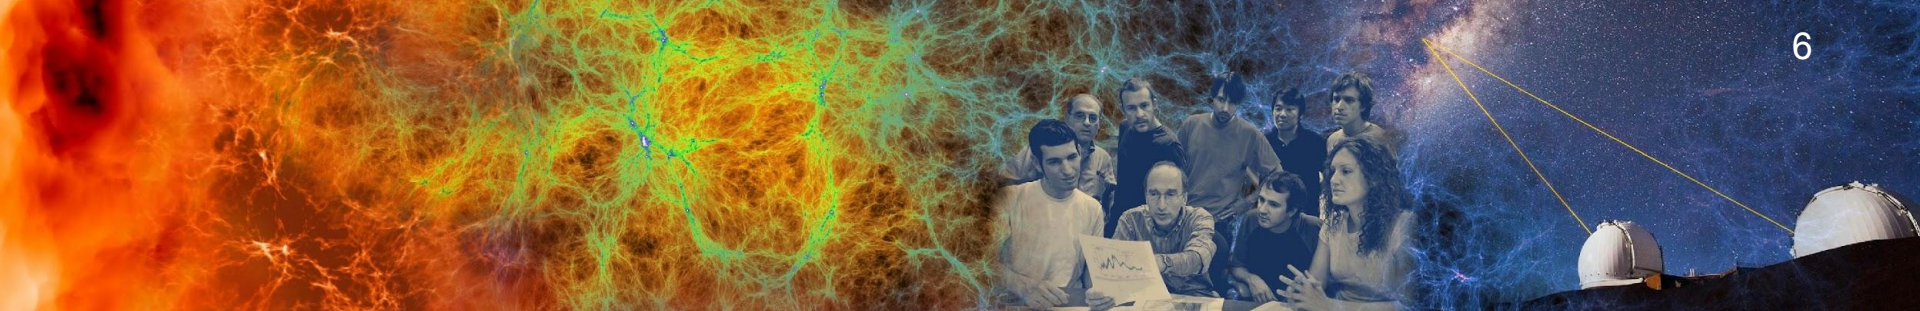


# New Generation of Workflows

Similar need across many science domains! Eg.:

1. Astrophysics and Cosmology:  
*LSST DESC, CMB/TOAST, DESI, DUNE*
2. High-Energy and Particle Physics:  
*ATLAS, CMS, LZ*
3. Electron Microscopy and Nanoscience:  
*NCEM*
4. Nuclear Fusion Experiments:  
*KSTAR*



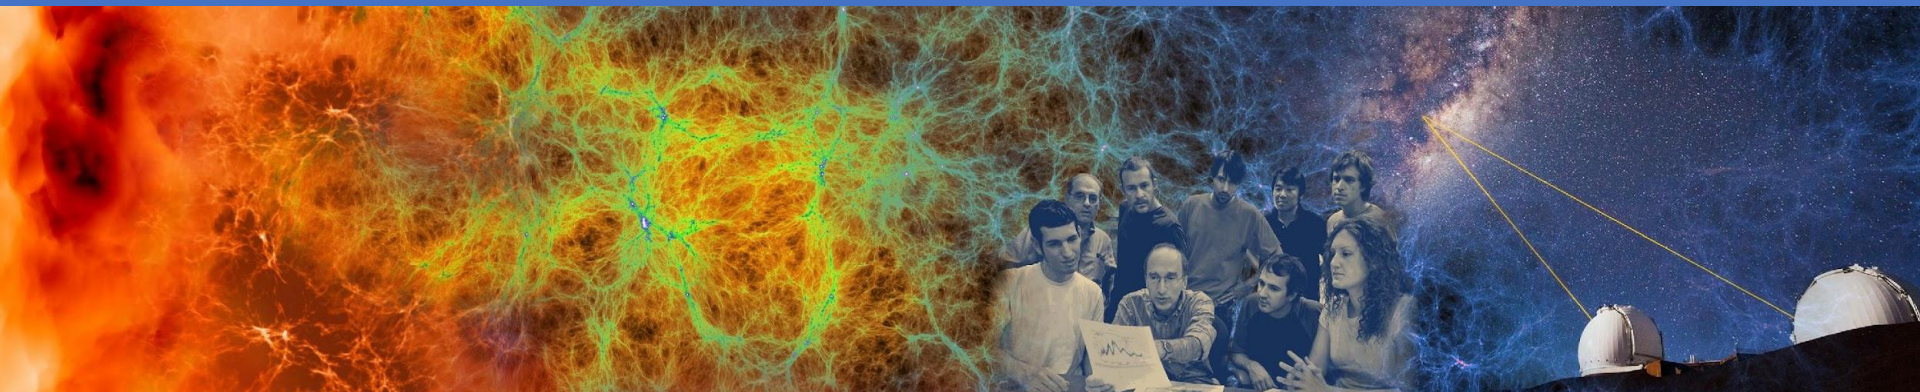


# Outline

1. The Human Perspective: *XFEL + Superfacility*
2. The Engineering Nuts and Bolts: *Performance Optimization Towards the Exascale*
3. Lessons Learned

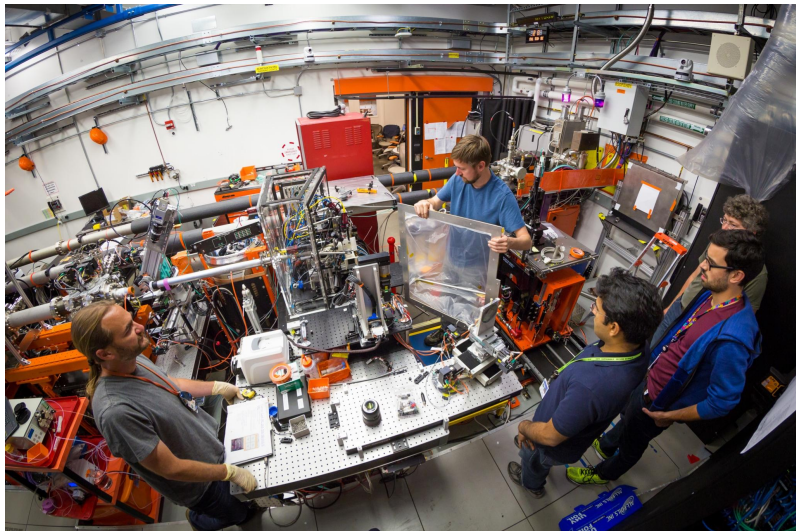


# The Human Perspective: XFEL + Superfacility

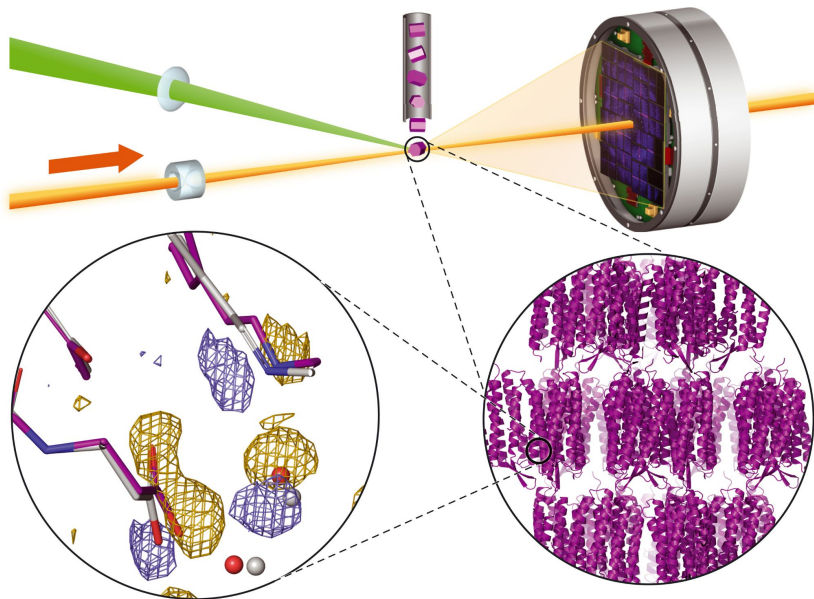
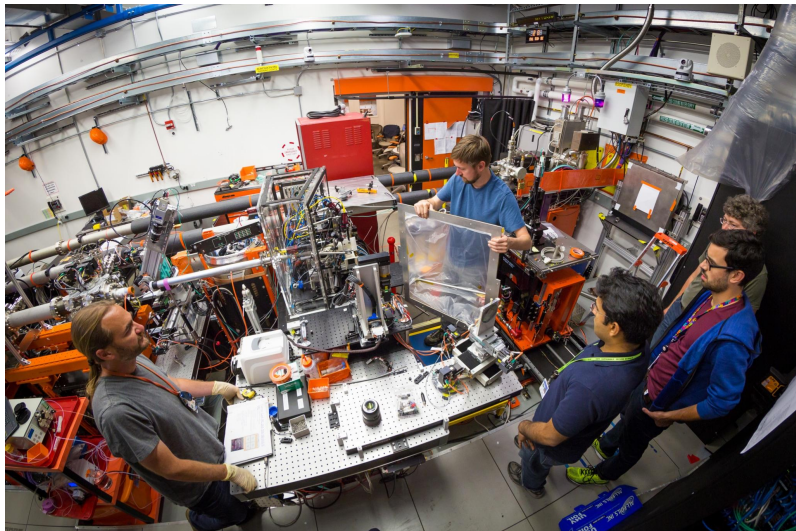




# XFEL: Serial Crystallography

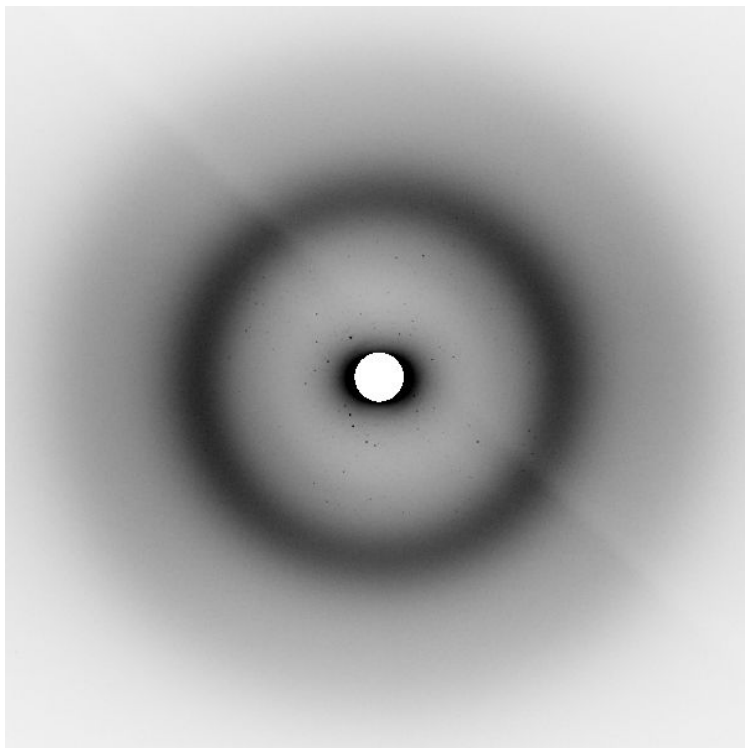


# XFEL: Serial Crystallography

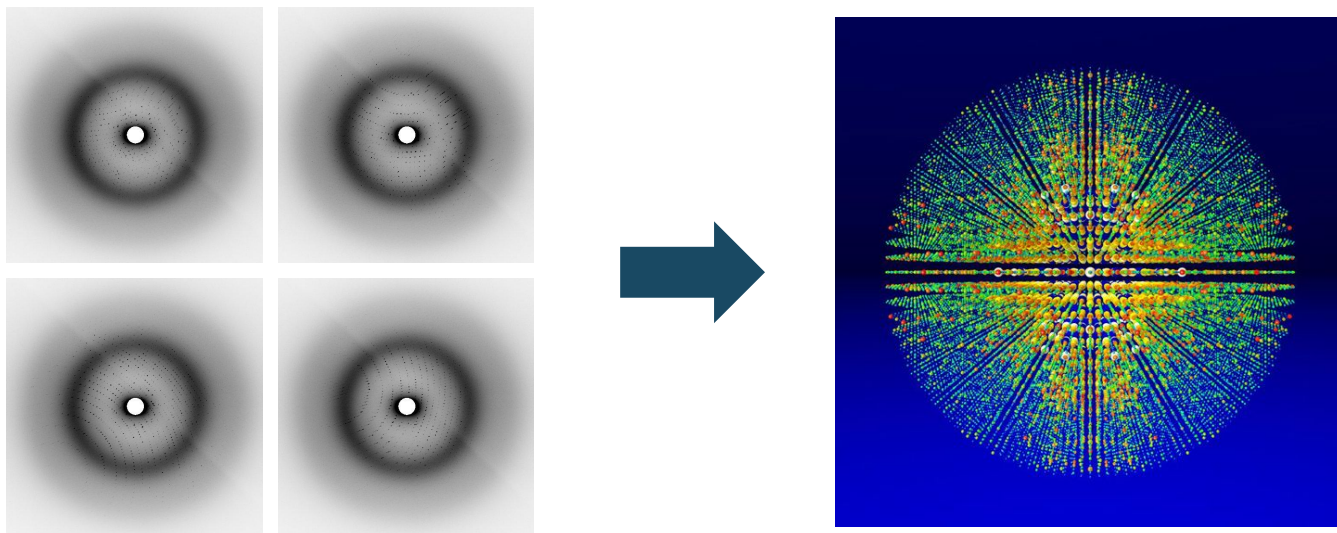


(Brändén, Science 2021)

# XFEL: Serial Crystallography



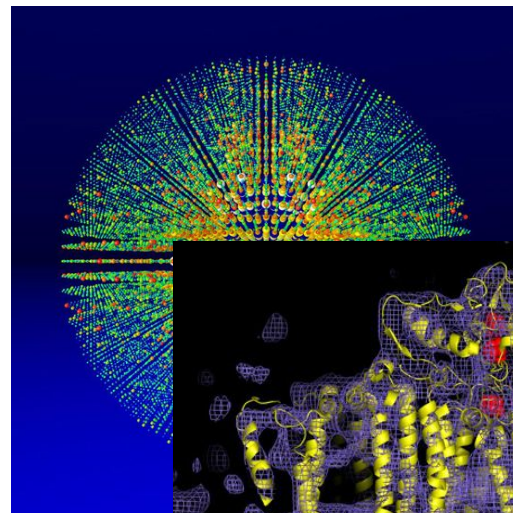
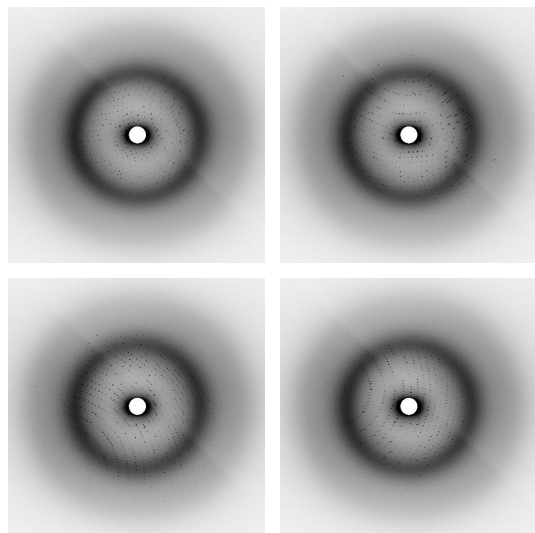
# XFEL: Serial Crystallography



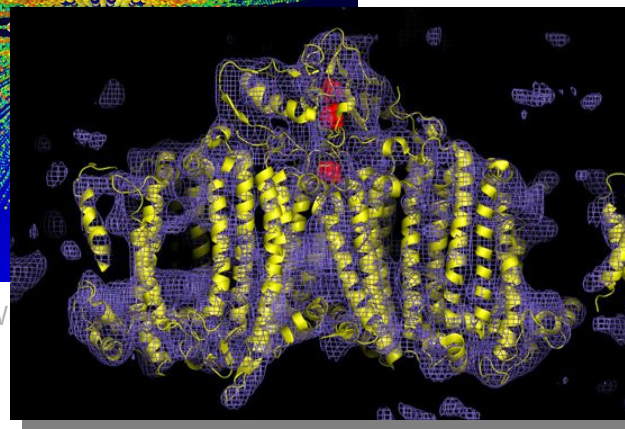
(Thomas White, CFEL)



# XFEL: Serial Crystallography



(Thomas W

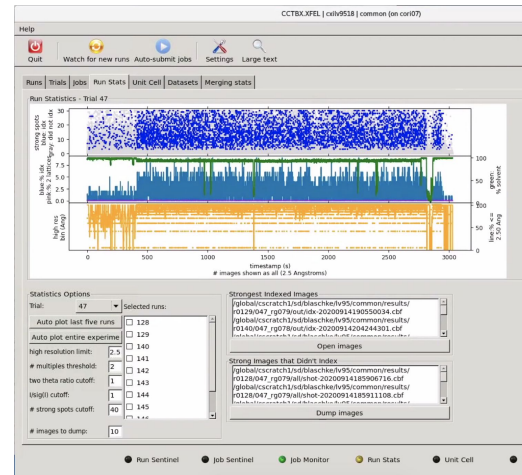
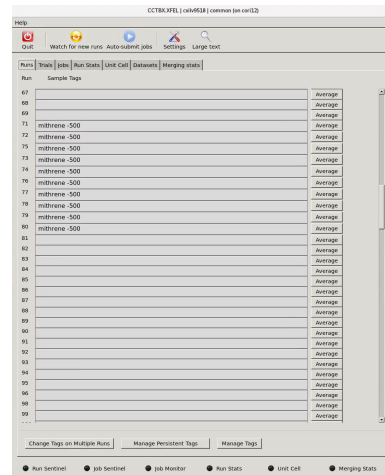


(Chapman, Nature 2011)



# When should I move onto the Next Sample?

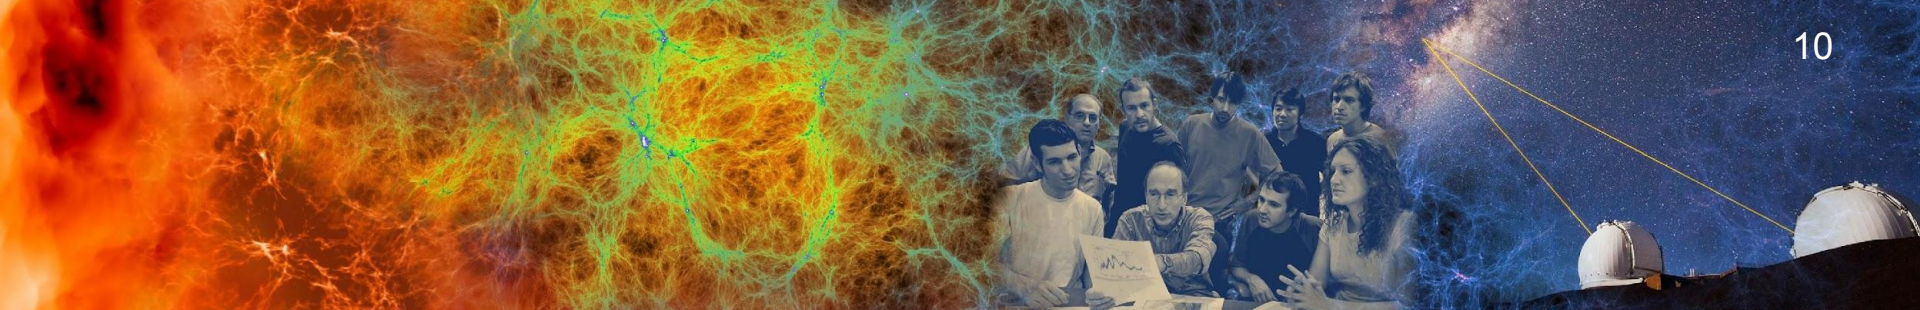
- Beamtime is scarce!
- Critical live feedback:
  - Does the beam hit the sample?
  - Do we see crystals?
  - Does the data make sense?
  - What is the quality of the data?
- Can I move on to the next sample?



Incoming data

Build jobs

Monitor Analysis



# Experimentalists Are In The Driver's Seat

Live Data Analysis for Experiments using the NERSC+LCLS Superfacility



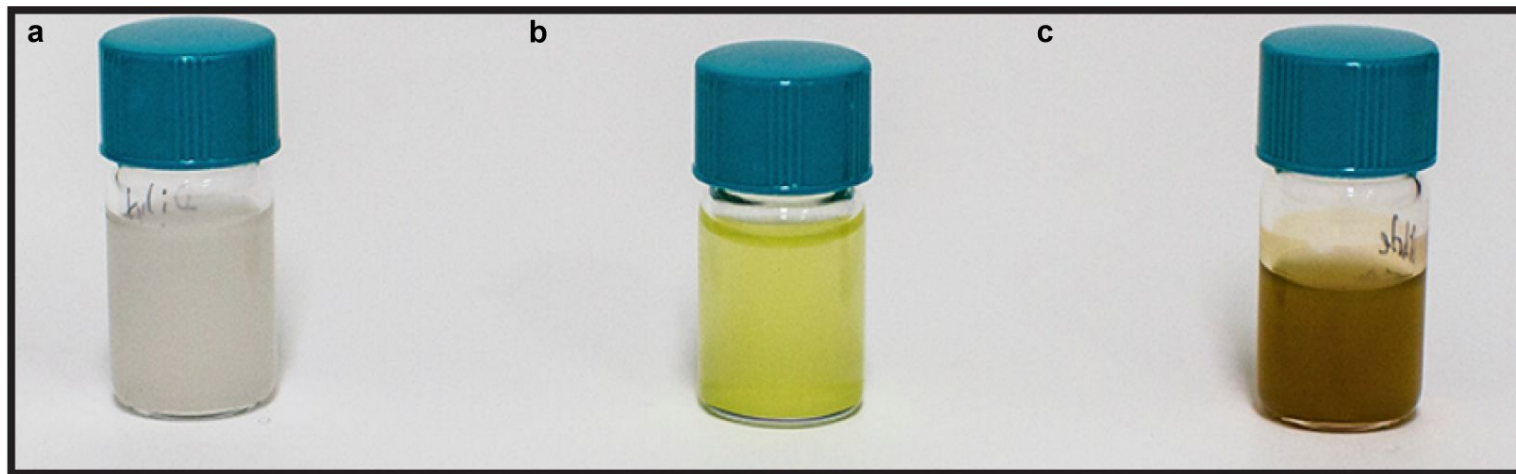


# Example: Mithrene derivatives

Thiorene  
(sulfur)

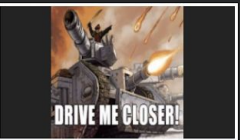
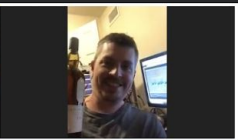
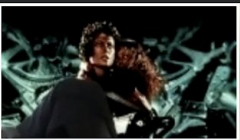
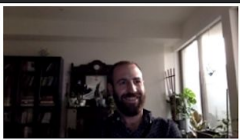
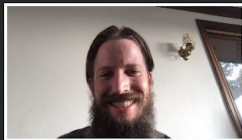
Mithrene  
(selenium)

Tethrene  
(tellurium)



Inactive

Photo-active



number of threads should not exceed about 65500. -c sets the reflection buffer size. This depends on the CPU cache size but will rarely need changing.

-g sets the number of reflection groups used for calculating R<sub>int</sub>. This must be greater than 1 but not greater than the total number of reflections for refinement. It sets the current reflection group number. This may not be less than 1 nor greater than the number set by -g, these command line flags override other ways of defining the-R reflection. The -m value is also used as a seed for the WIGL pseudo-random shifts.

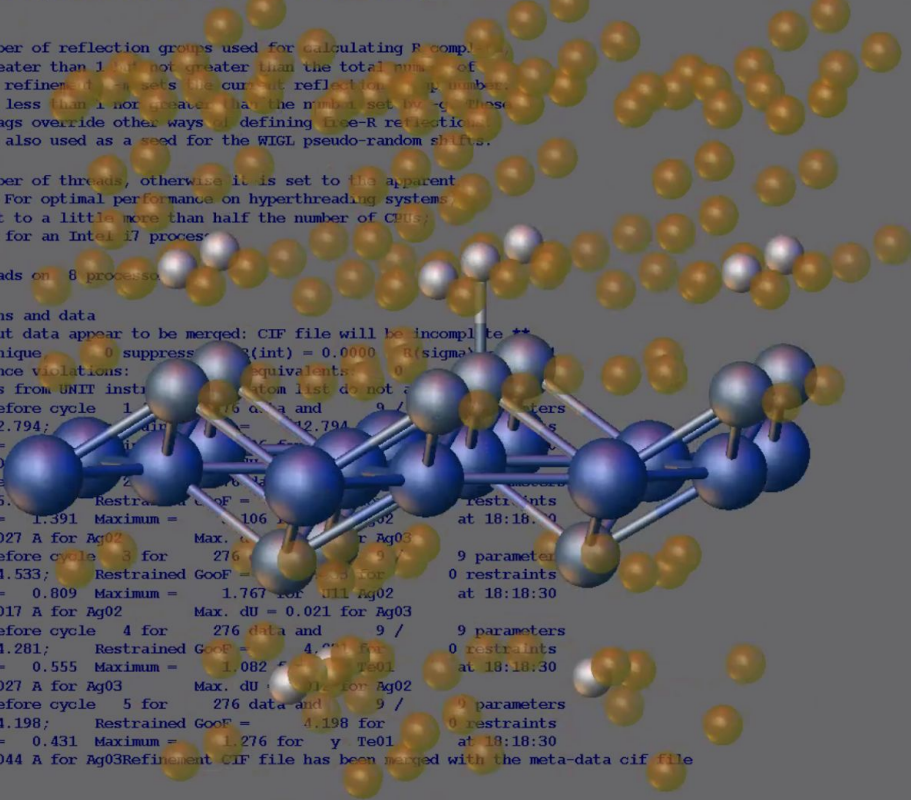
-t sets the number of threads, otherwise it is set to the apparent number of CPUs. For optimal performance on hyperthreading systems, -t should be set to a little more than half the number of CPUs, e.g. -t4 or -t5 for an Intel i7 processor.

Running 8 threads on 8 processors

Read instructions and data

\*\* WARNING: Input data appear to be merged; CIF file will be incomplete \*\*

```
Data: 276 unique reflections, 0 suppressions, R(int) = 0.0000, R(sigma) = 0.0000
Systematic absence violations: 0
** Cell contents from UNIT instructions do not match atom list do not match
wR2 = 0.7663 before cycle 1 for 276 data and 9 / 9 parameters
Goof = S = 12.794; Restrained Goof = 4.281 for 0 restraints
Mean shift/esd = 0.0000 Maximum = 0.0000 at 18:18:30
Max. shift = 0.0000 A for Ag02 Max. du = 0.0000 for Ag02
wR2 = 0.6673 before cycle 2 for 276 data and 9 / 9 parameters
Goof = S = 5.198; Restrained Goof = 4.198 for 0 restraints
Mean shift/esd = 1.391 Maximum = 1.062 for y Te01 at 18:18:30
Max. shift = 0.027 A for Ag02 Max. du = 0.021 for Ag03
wR2 = 0.6561 before cycle 3 for 276 data and 9 / 9 parameters
Goof = S = 4.533; Restrained Goof = 4.533 for 0 restraints
Mean shift/esd = 0.809 Maximum = 1.767 for Ag02 at 18:18:30
Max. shift = 0.017 A for Ag02 Max. du = 0.021 for Ag03
wR2 = 0.6514 before cycle 4 for 276 data and 9 / 9 parameters
Goof = S = 4.281; Restrained Goof = 4.281 for 0 restraints
Mean shift/esd = 0.555 Maximum = 1.082 for y Te01 at 18:18:30
Max. shift = 0.027 A for Ag03 Max. du = 0.021 for Ag02
wR2 = 0.6480 before cycle 5 for 276 data and 9 / 9 parameters
Goof = S = 4.198; Restrained Goof = 4.198 for 0 restraints
Mean shift/esd = 0.431 Maximum = 1.276 for y Te01 at 18:18:30
Max. shift = 0.044 A for Ag03Refinement CIF file has been merged with the meta-data cif file
```



tethrene tri...rk0 v0 all a C2/c

CHAGTe

a = 5.900(3) Å α = 90° Z = 16  
 b = 7.424(4) Å β = 97.686(15)° Z' = 2  
 c = 30.258(17) Å γ = 90° V = 1313.3(12) Å<sup>3</sup> wR<sub>2</sub> = 0.119 R<sub>1</sub> = 0.031

Goof = 4.198; Restrained Goof = 4.198 for 0 restraints

Home Work View Tools Info

Solve Refine Draw Report

Program: ShelXL Least Squares Cycles: 10 Peaks: 20

hkl file: hkl: Sun Sep 13 14:00:00 2020

Weight: 1001.200 1.001.000 EXTI

Use Solvent Mask: This is the Olex2 implementation of BYPASS (a.k.a. SQUEEZE)

Refinement Settings Extra

Toolbox Work

Labels: Labels OFF/ON

CHAGTe ... Add H

Select atom(s) and then: mFit mSplit Split SAM Split

MAP: Show Map Map Settings

Peak & Uiso Sliders

Growing

Finishing

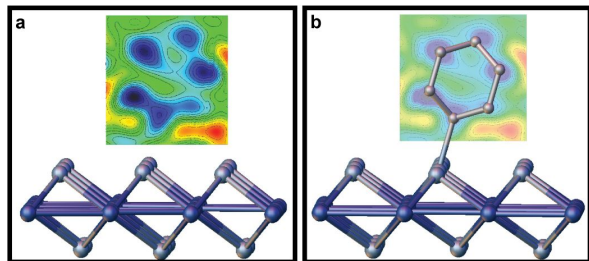
History

Select

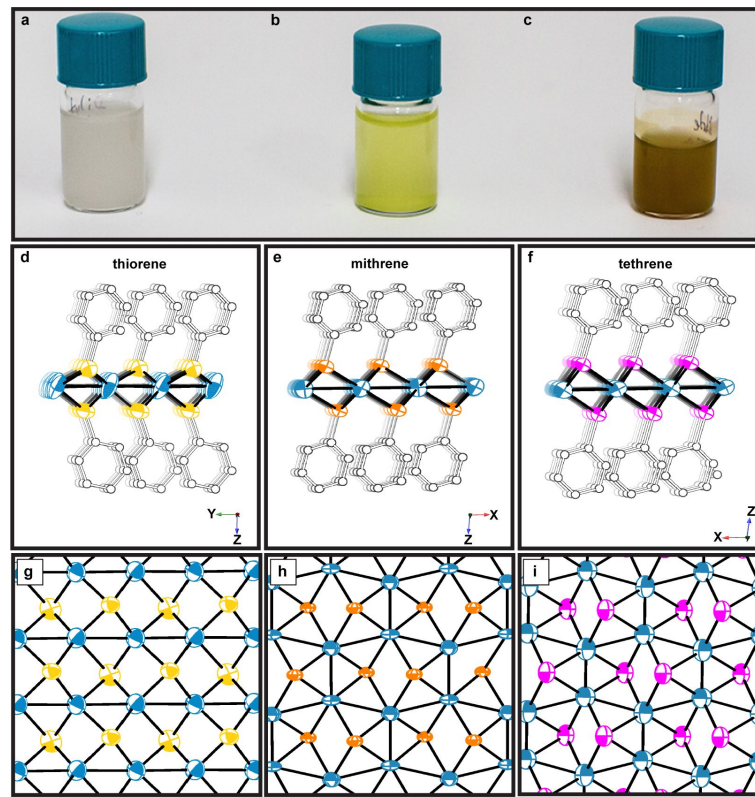
Naming

Sorting

# XFEL structures reveal why thiorene is not photo-active



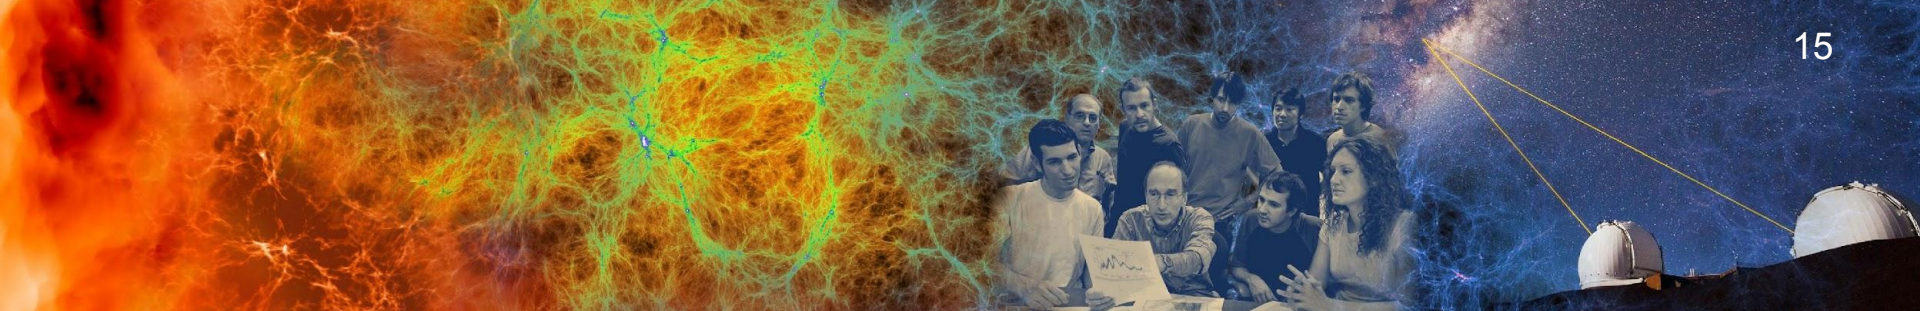
Schriber EA, Paley DW *et. al.* (2021) “Chemical Crystallography by Serial Femtosecond X-ray Diffraction.” *Nature* **601**, 360-365.



Inactive

Photo-active



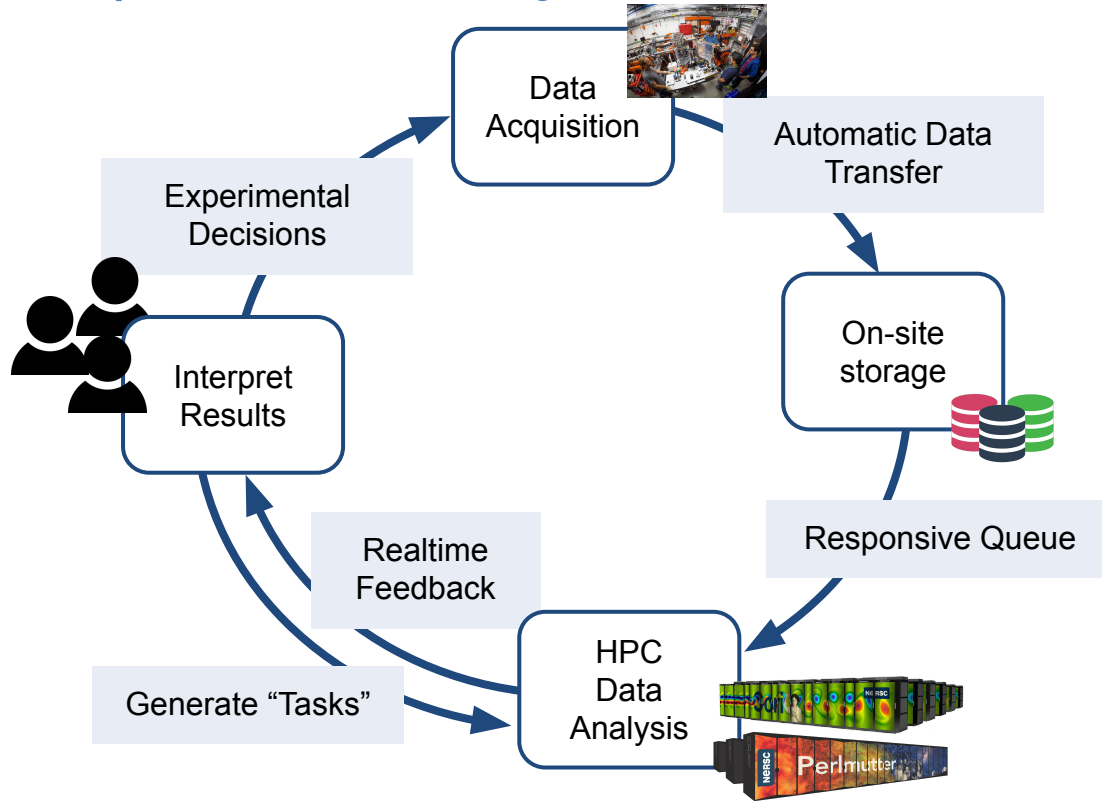


# How is the NERSC + LCLS Superfacility Used?

# How is the NERSC + LCLS Superfacility Used?

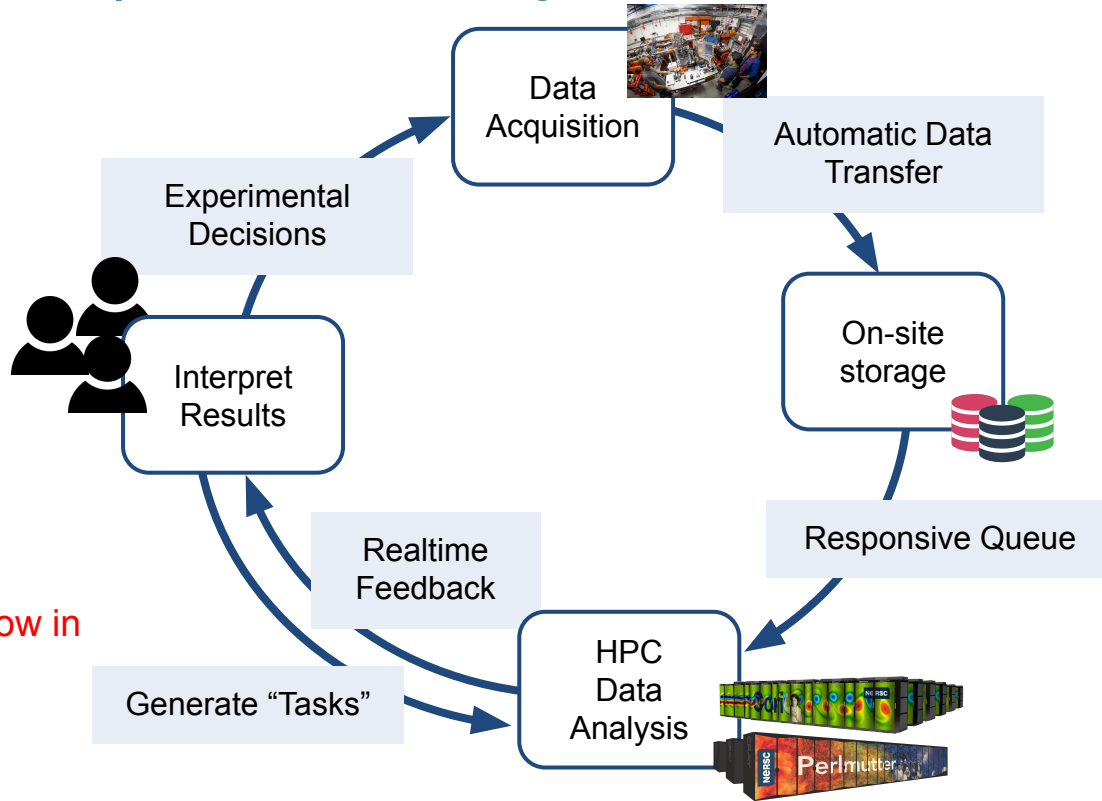
- Difficult to define what “interactivity” means
  - Eg: Interactive can mean exploration, vs human-in-the-loop
  - Each use HPC differently
- ExaFEL: Case study of an interactive HPC workflow
- NERSC staff deploy code mine log files, and instrument workflows during runtime in order to monitor the whole workflow

# (Realtime) Data Analysis Workflow





# (Realtime) Data Analysis Workflow



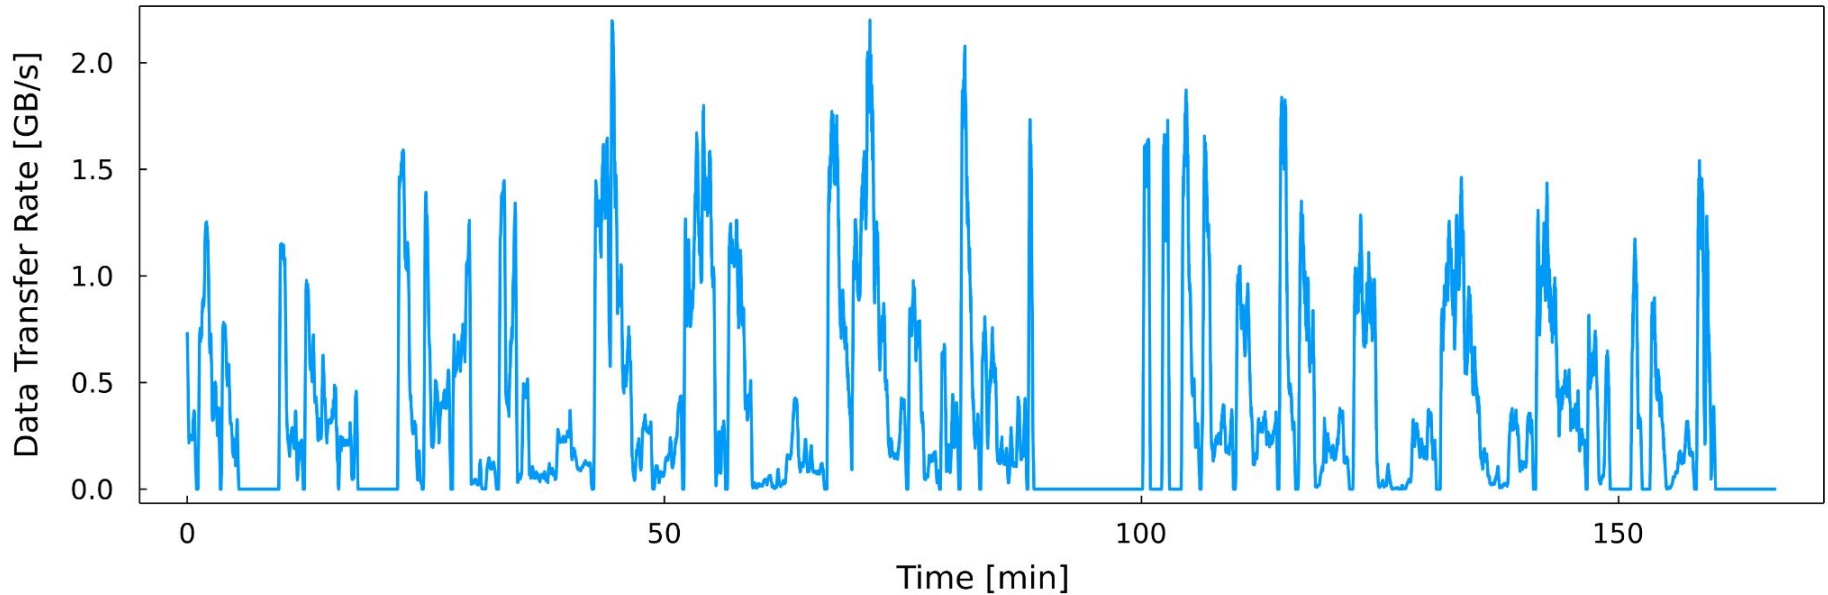
1. How is data getting to NERSC?

3. How responsive is the job queue?

2. How do users interact with workflow in real time?

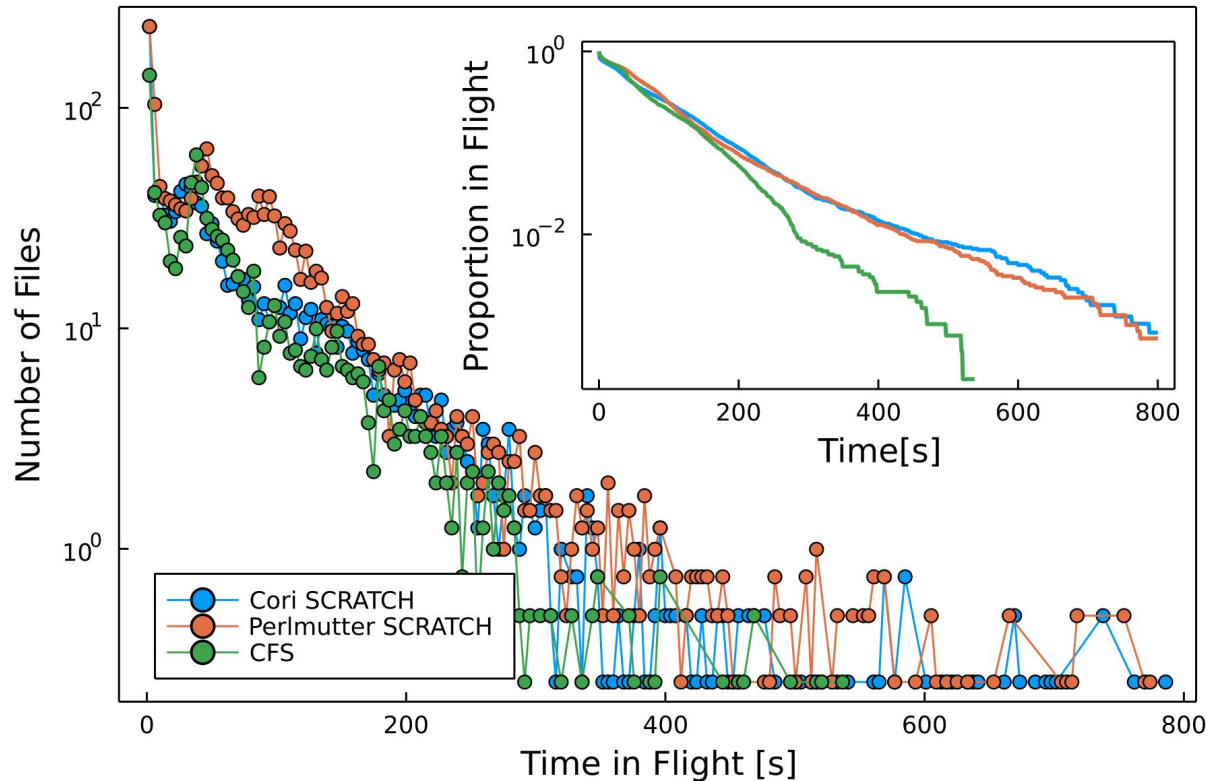
# 1. How is Data Getting to NERSC?

- Files are automatically transferred using XRootD immediately after data is collected 15 TB/day

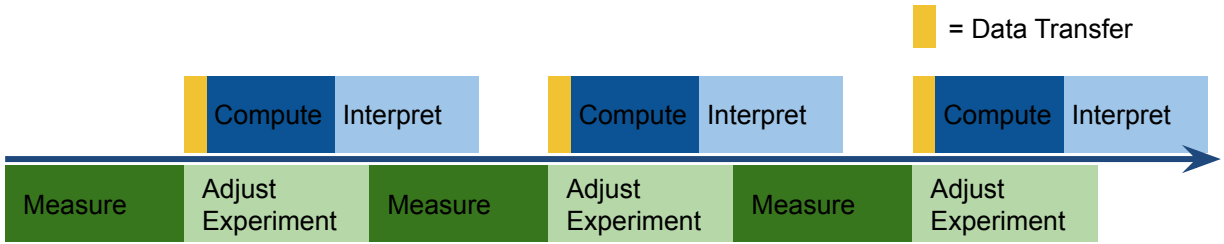
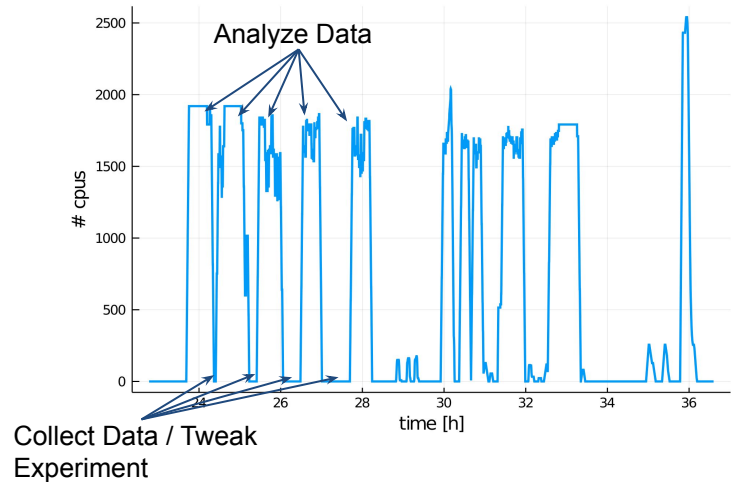
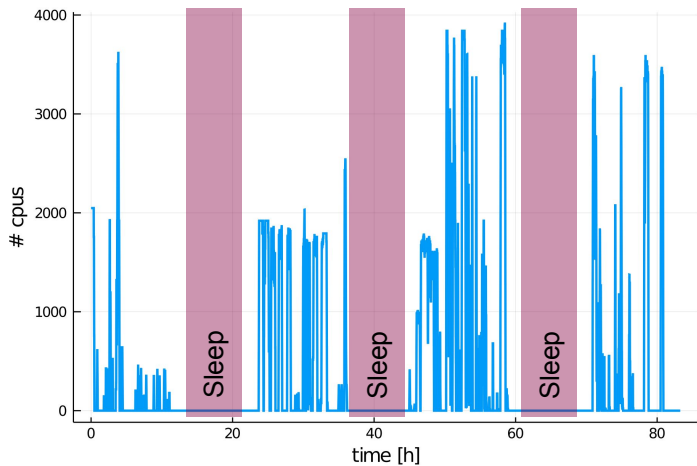


# 1. How is Data Getting to NERSC?

- Tail latency can disrupt the entire workflow!
- Choice of file system matters



# 2. How do Users Interact with the Workflow?

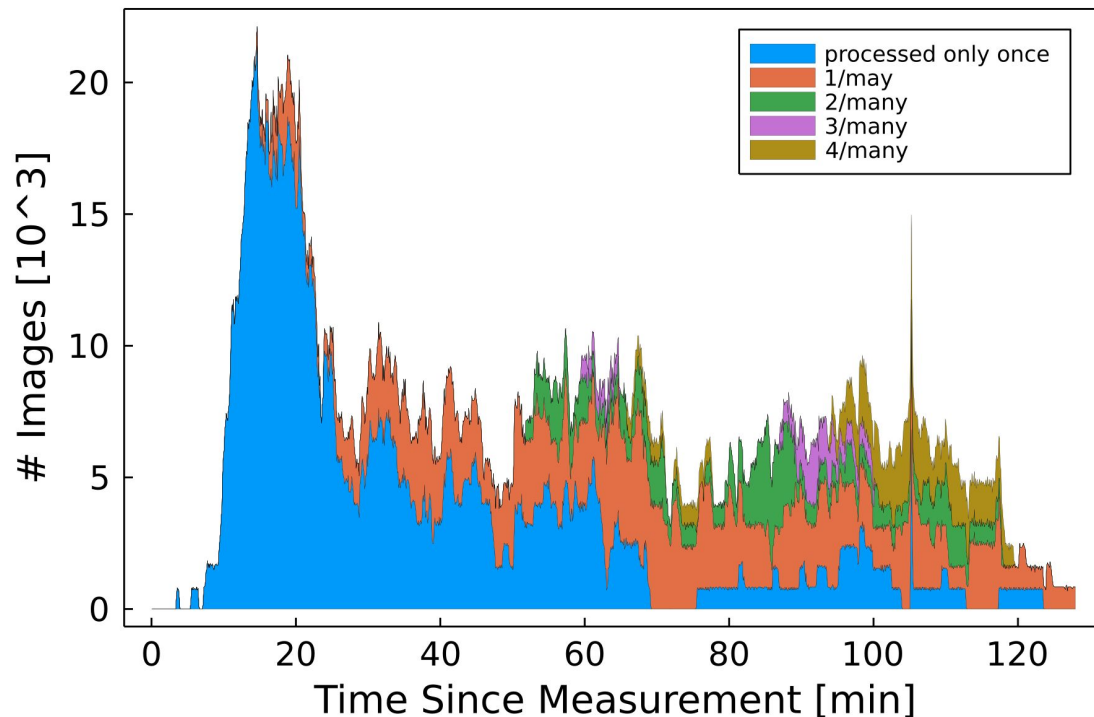


## 2. How do Users Interact with the Workflow?

- An HPC center bridges 3 time scales:
  - Hours => Experiment, Job Scheduler
  - Minutes => Humans
  - Milliseconds => Compute Nodes
  
- Objective: measure the human timescale
  - How much of the workflow requires “constant” human interaction?

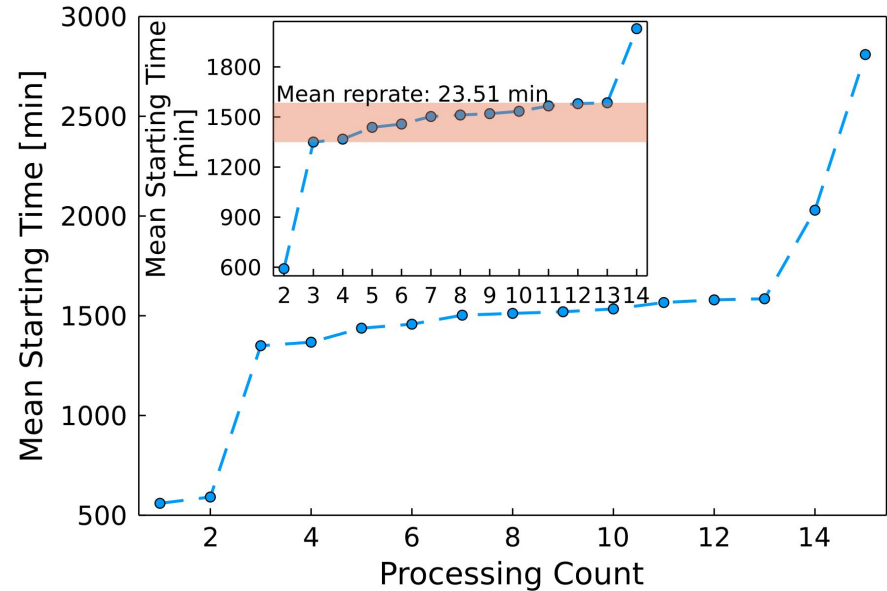
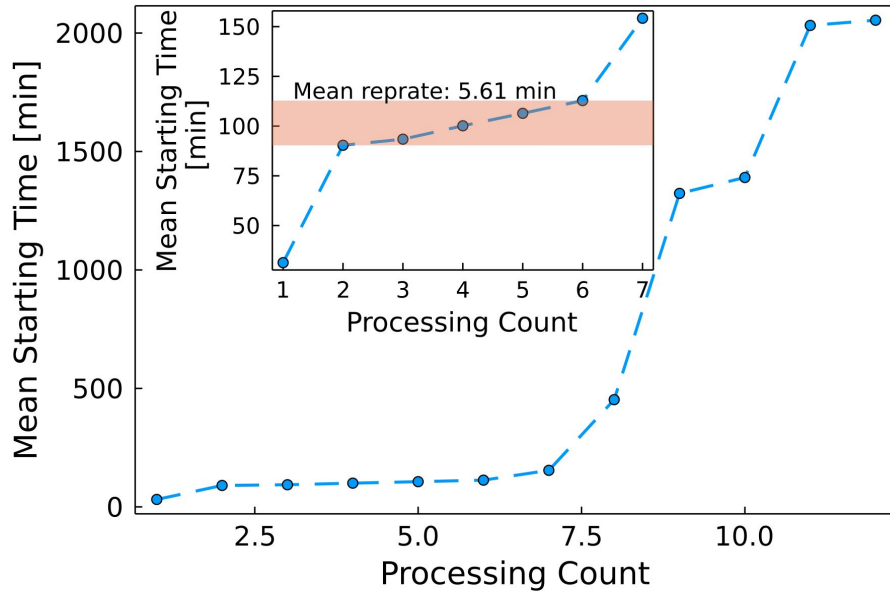
## 2. How do Users Interact with the Workflow?

- Most images are processed only once
  - This step can be automated once tuned
- Reprocessing as a “fingerprint” of interactivity



# Interactive Engagement with Data Sets

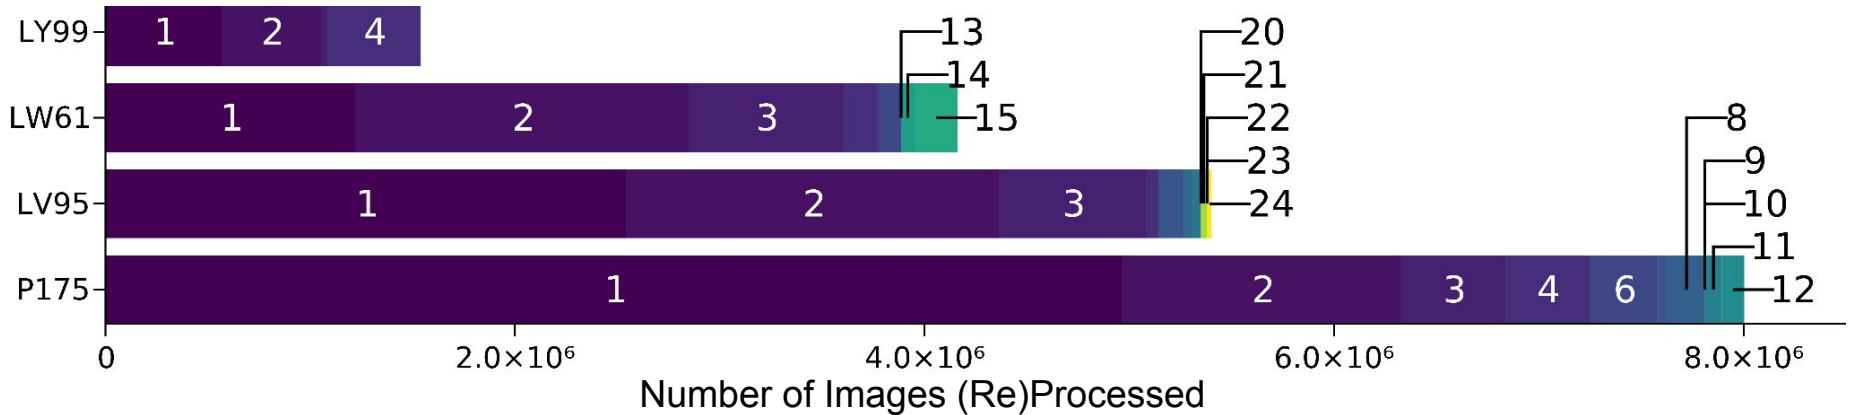
- Repeated reprocessing indicates user is exploring data (troubleshooting?)
- Details different from experiment to experiment





# Interactive Engagement with Data Sets

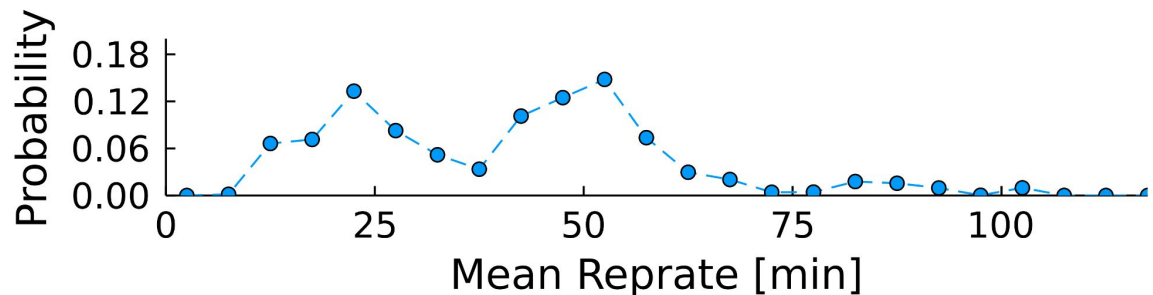
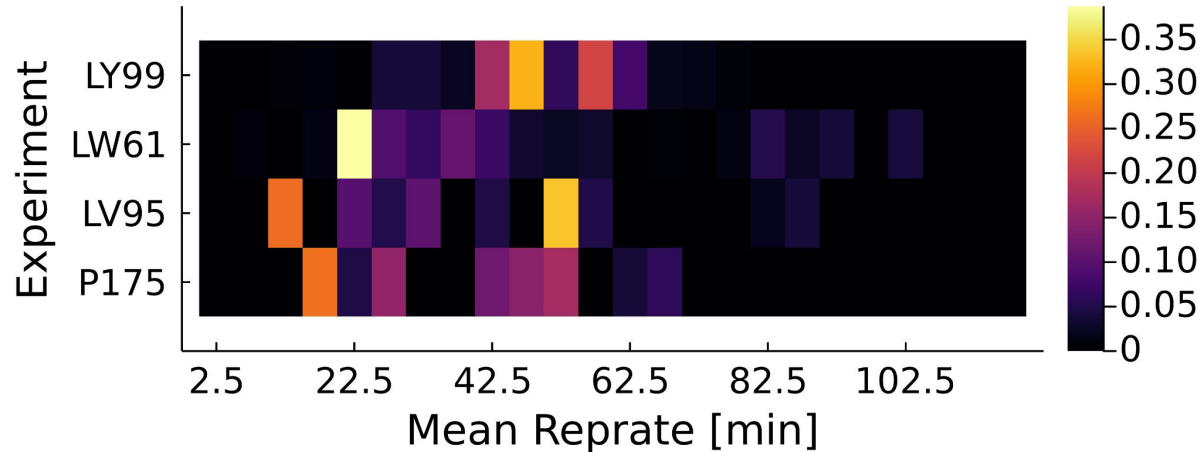
- Every beamtime is unique!



- Only ~5 % of images reprocessed > 10x

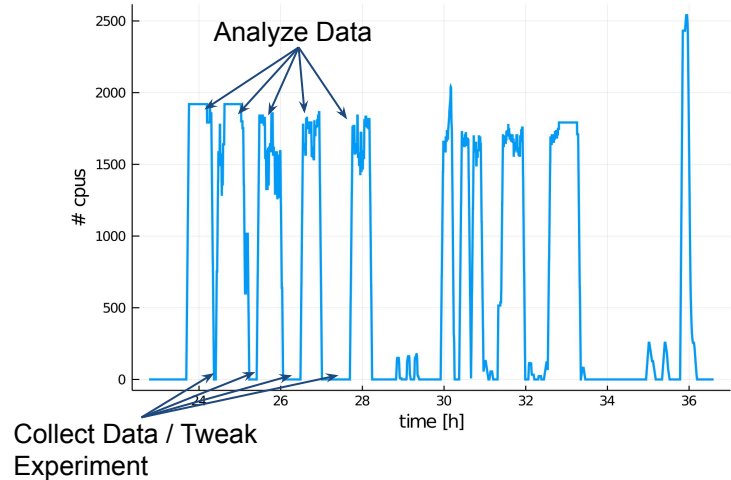
# Interactive Engagement with Data Sets

- Most reprocessing takes place on a 25-50 min cadence
- In line with the cadence of measurement runs at LCLS



# 3. How Responsive is the Job Queue?

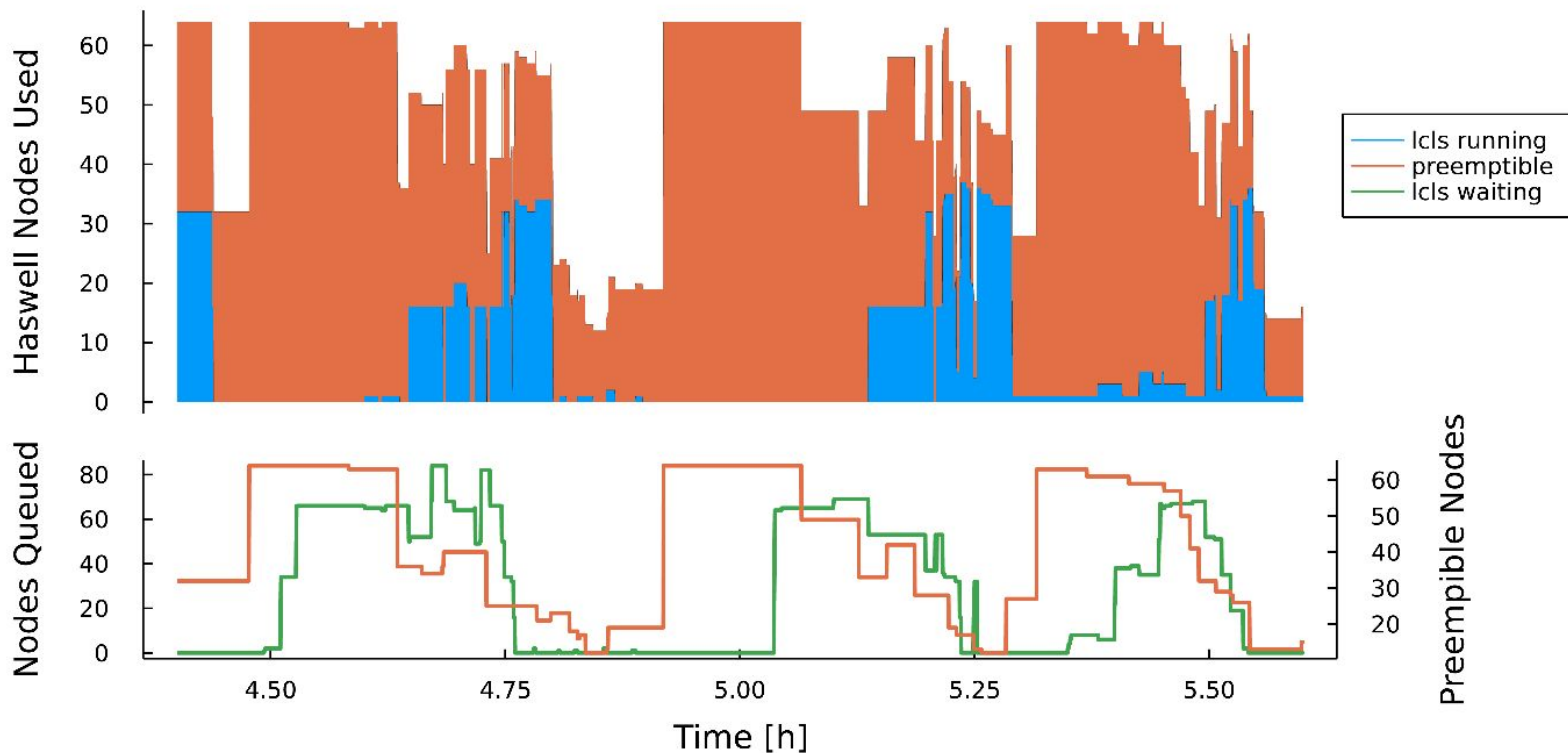
- We don't know when data will need to be analyzed
- Once data is available, queue needs to respond ASAP
- Use reservations
  - 60-80 Nodes
- Idle reservations “waste” compute => try sharing with preemptible jobs



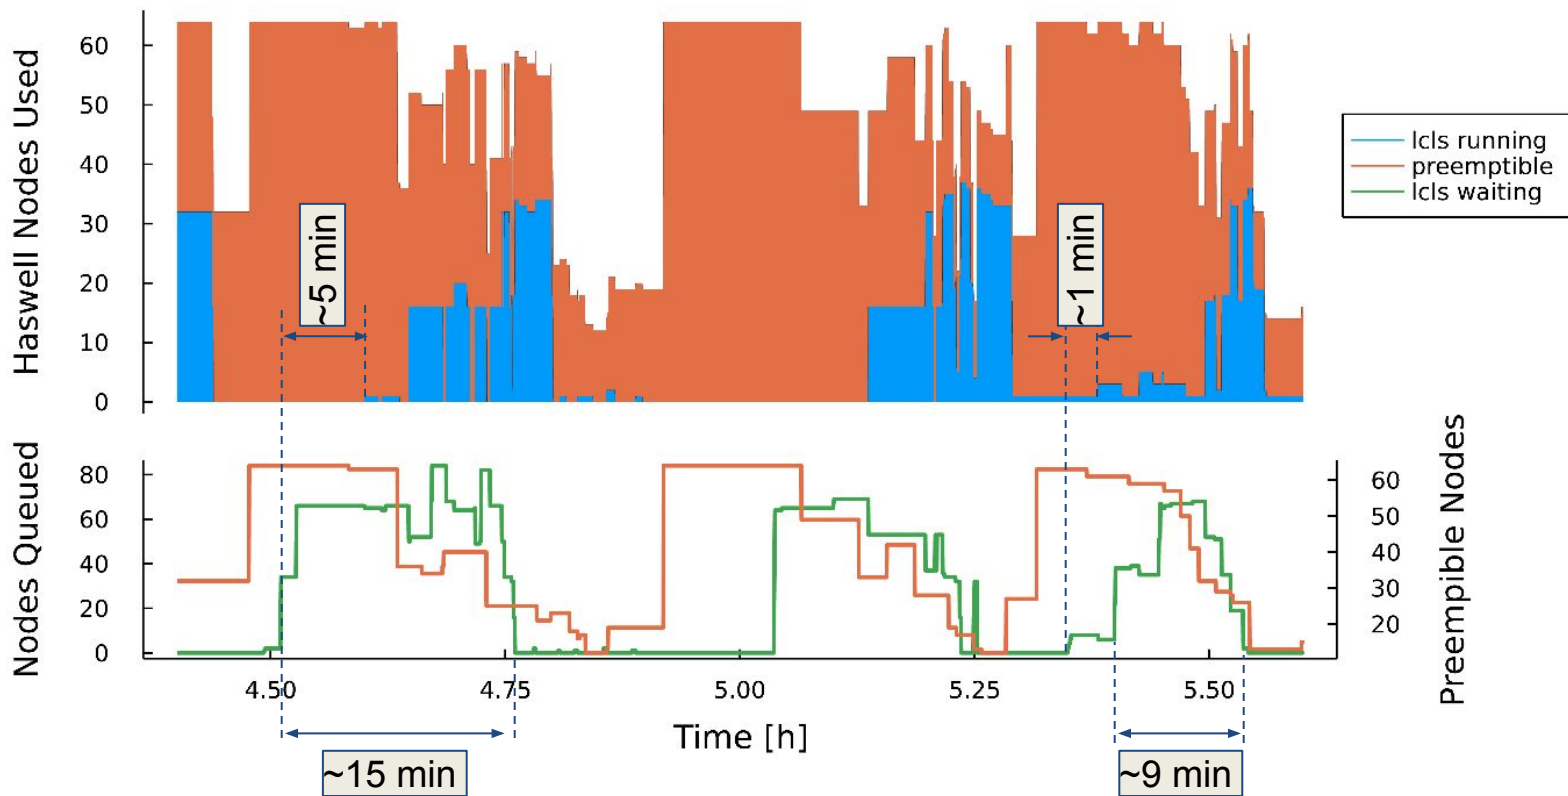
# Preemptible Reservations

- Allow preemptible jobs to use reservation:
  - `MaxStartDelay=${MINUTES}`  
Eg: `MaxStartDelay=5`
- Make a job preemptible
  - `#SBATCH --signal=R:INT@${SECONDS}`  
Eg. `#SBATCH --signal=R:INT@300`
- Jobs submitted to the regular QOS where:  
`$MINUTES * 60 >= $SECONDS`  
are allowed to run in the reservation
  - A job being preempted will receive `SIGINT` and then has `$SECONDS` many seconds to shut down gracefully (before being killed)

# Testing Preemptible Reservations



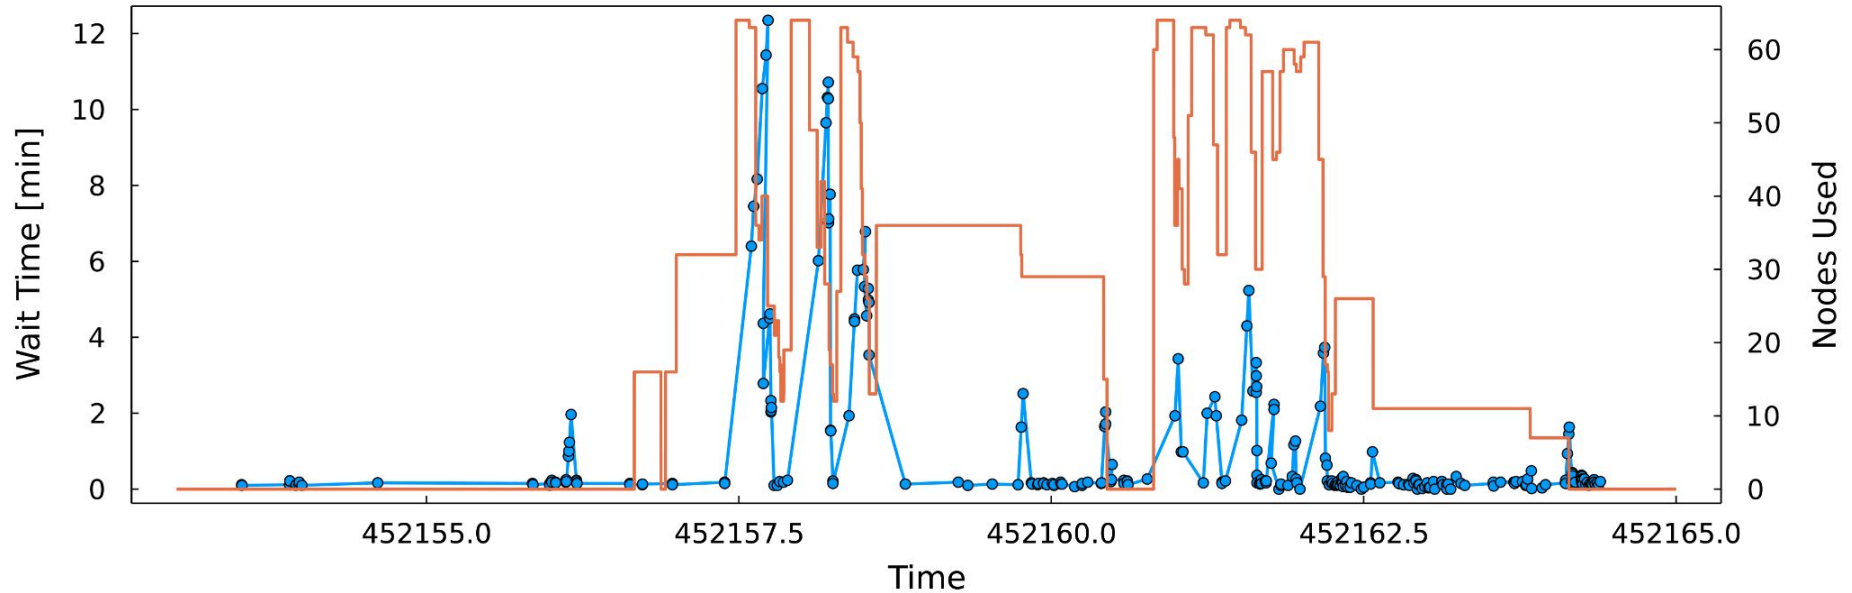
# Testing Preemptible Reservations



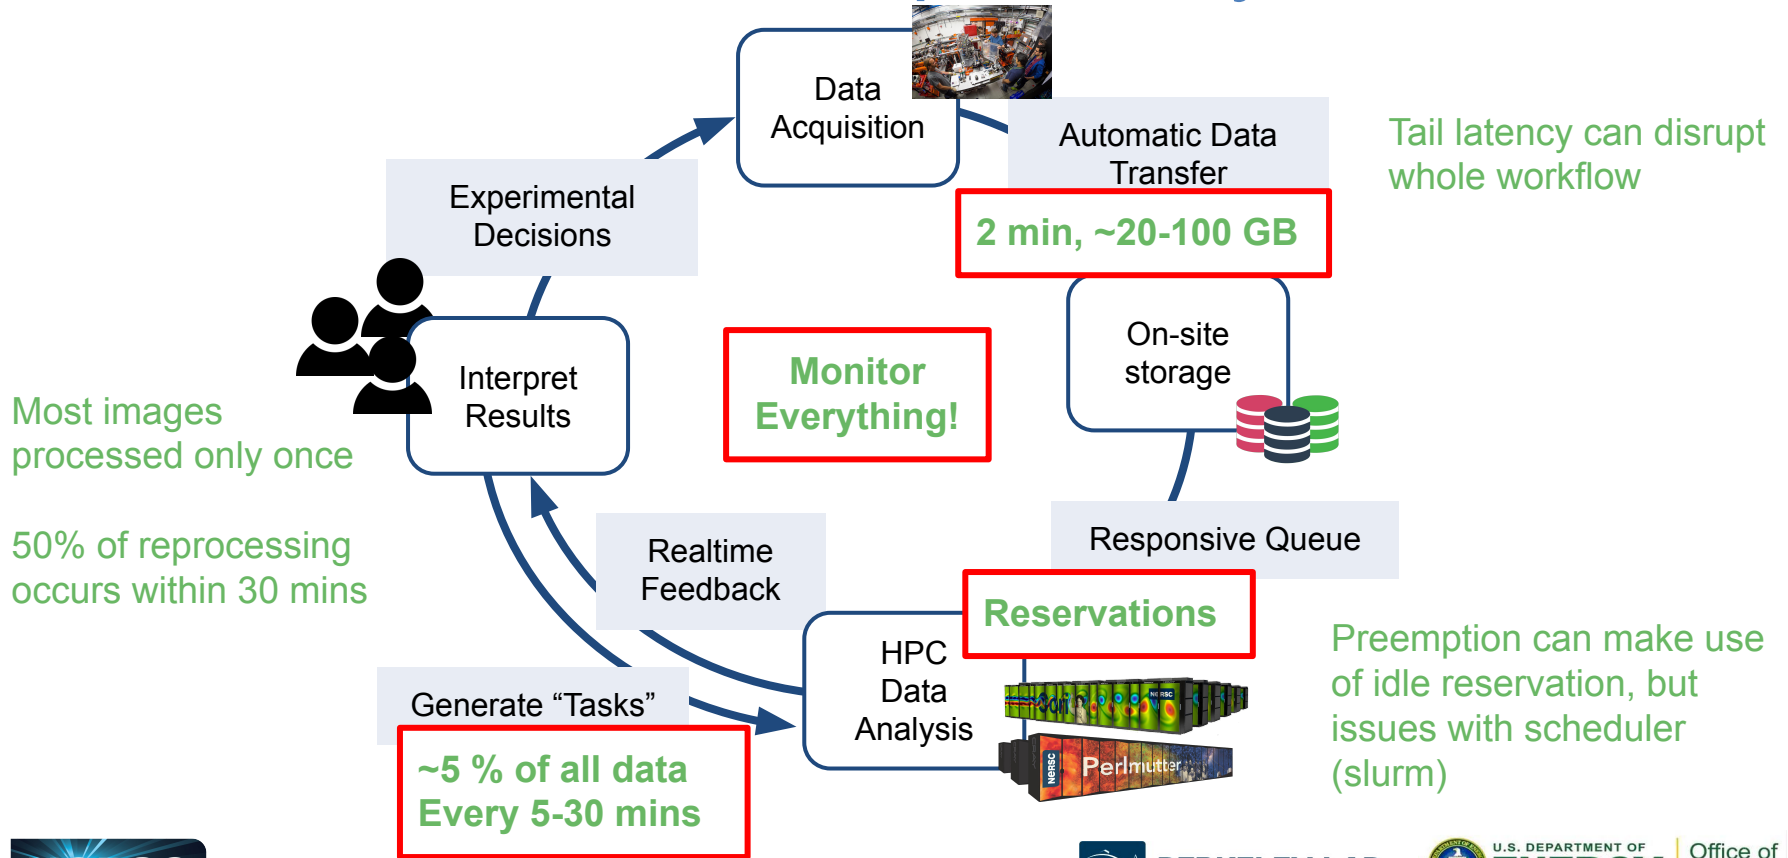


# Testing Preemptible Reservations

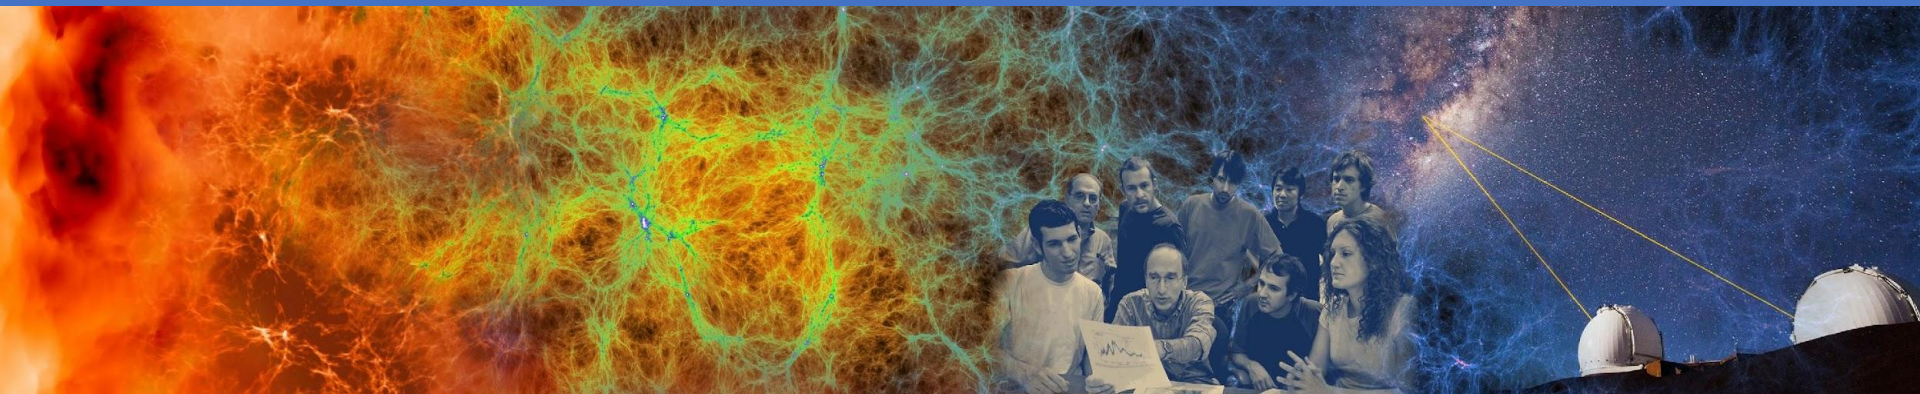
- Preemptible jobs don't leave reservation quickly enough



# Conclusion: XFEL + Superfacility

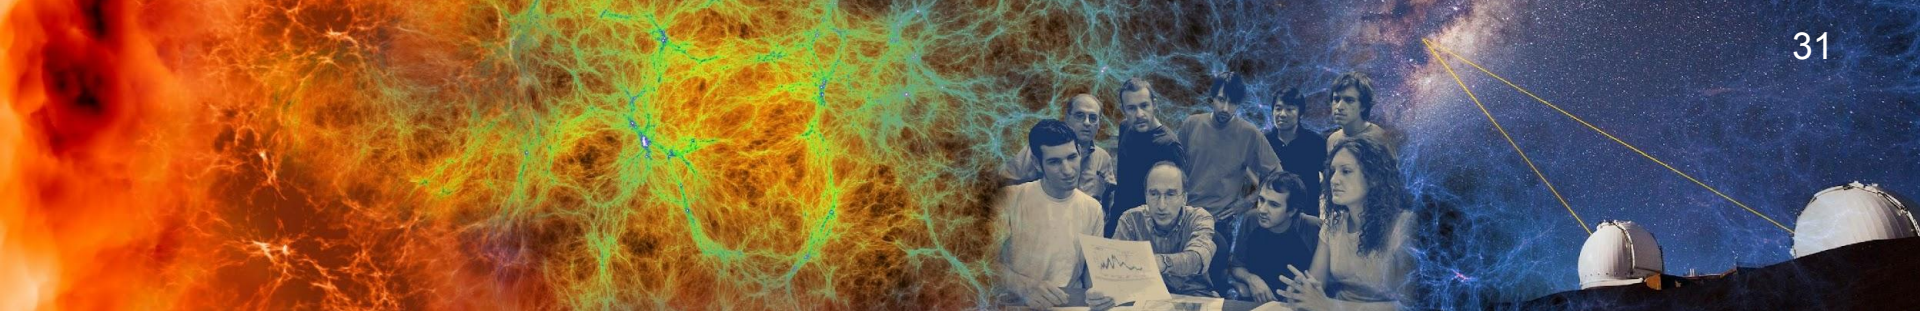


# Engineering Nuts and Bolts: Performance Optimization Towards the Exascale

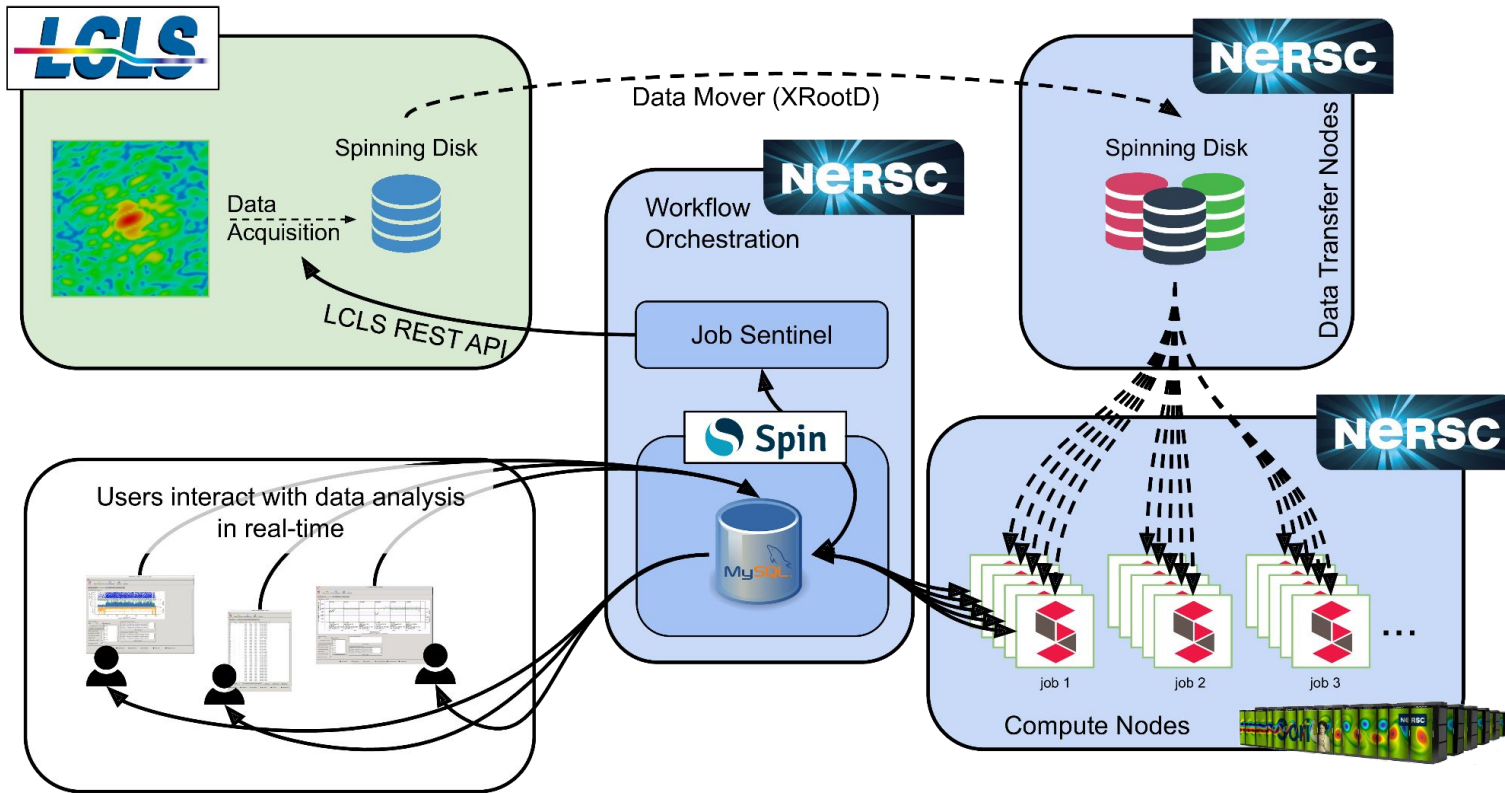




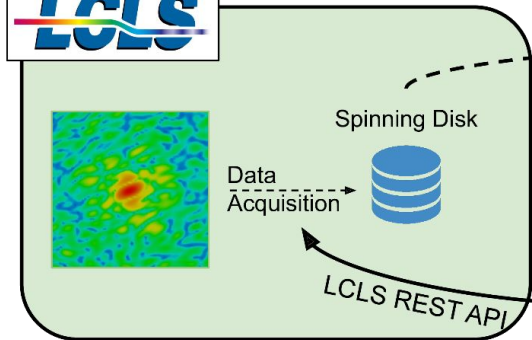




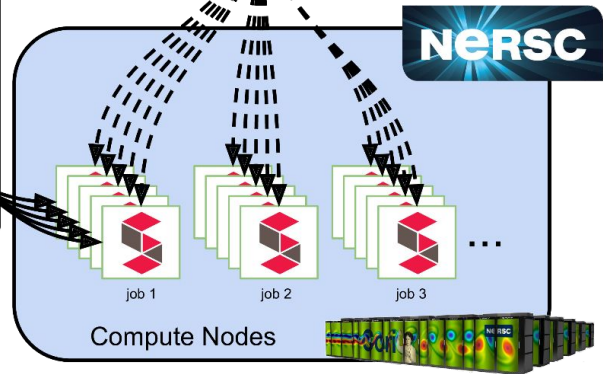
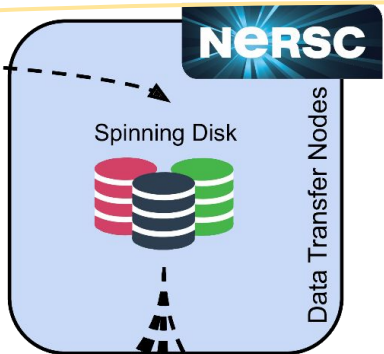
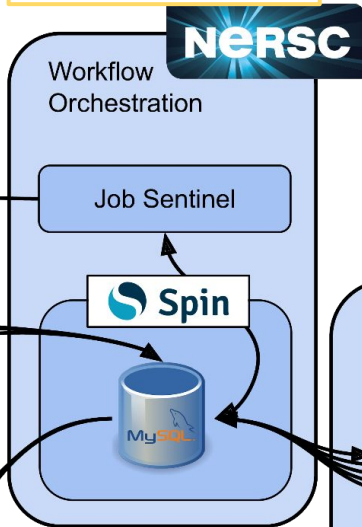
# Deploying CCTBX at NERSC





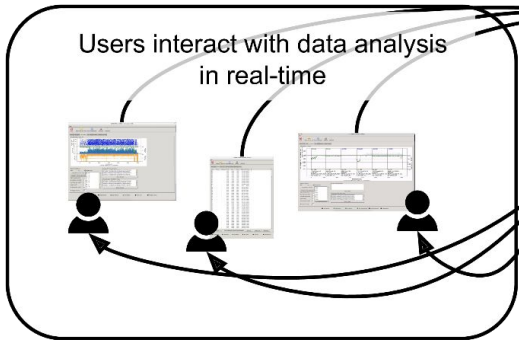


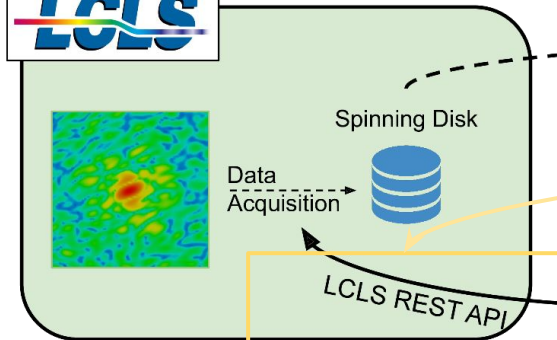
Data Mover (XRootD)



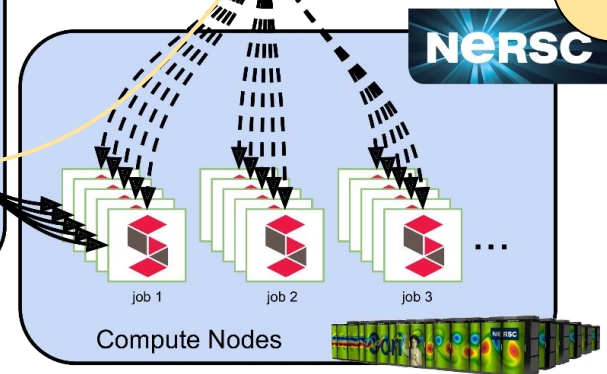
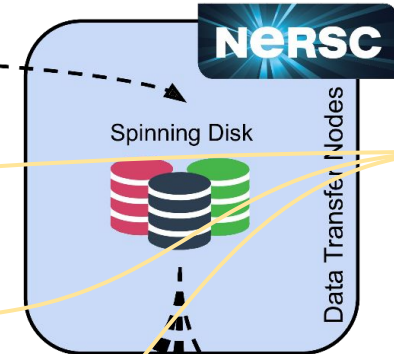
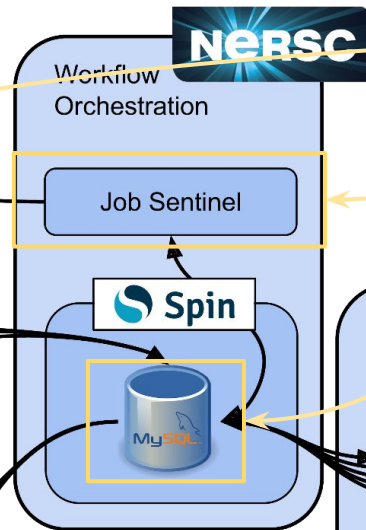
XRootD automatically transfers data to NERSC DTNs

Users interact with data analysis in real-time

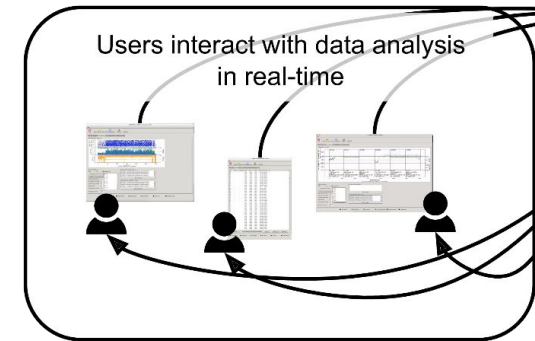


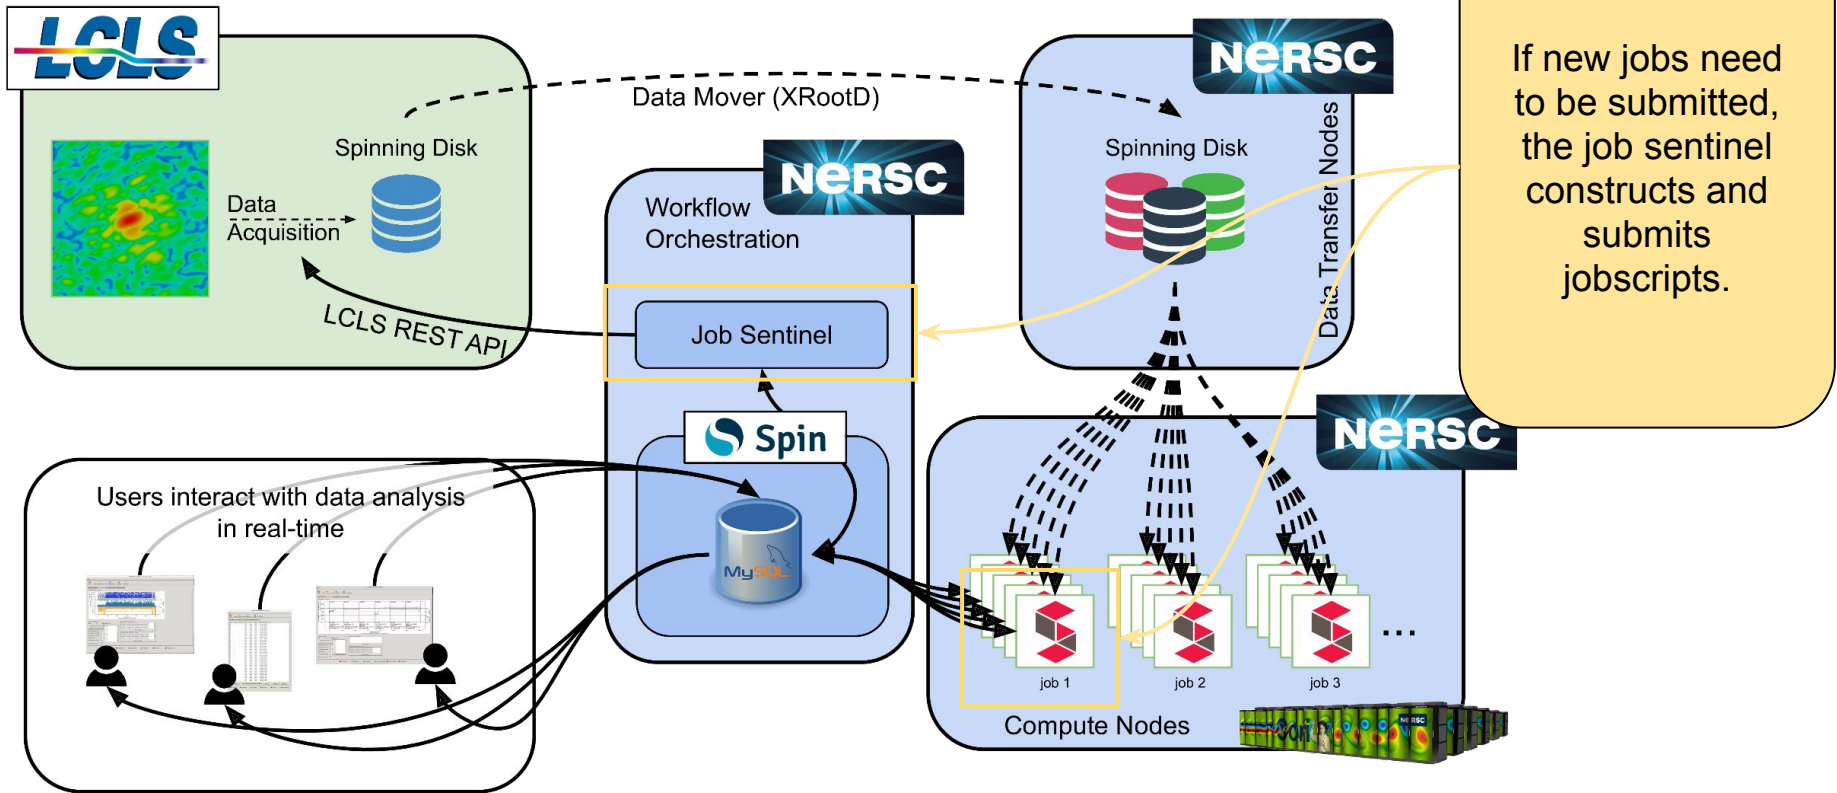


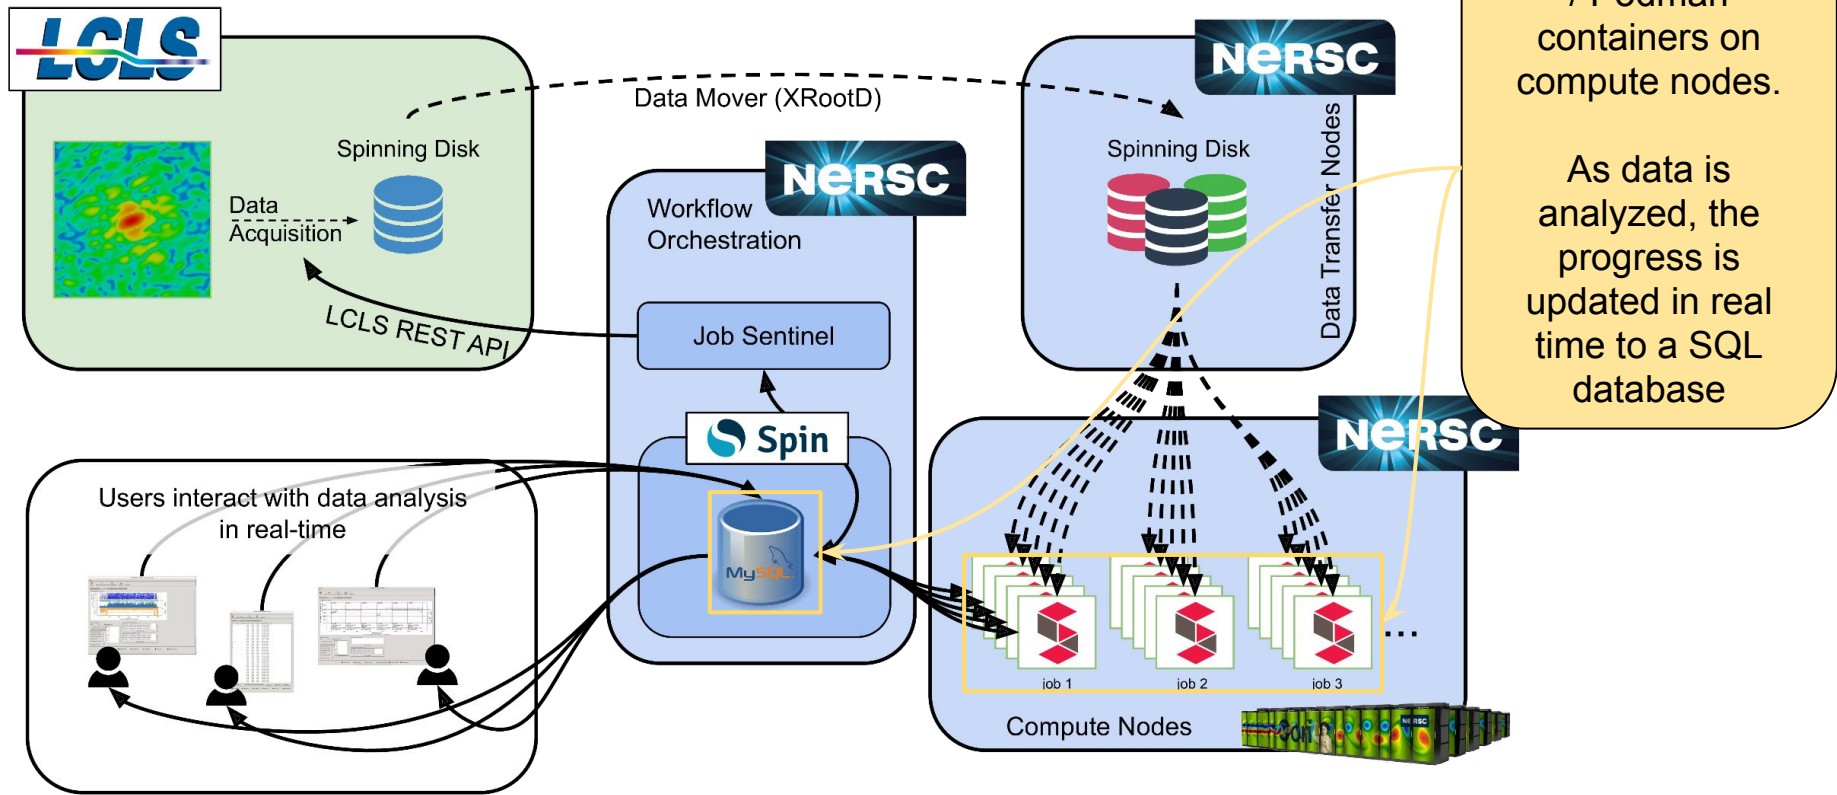
Data Mover (XRootD)

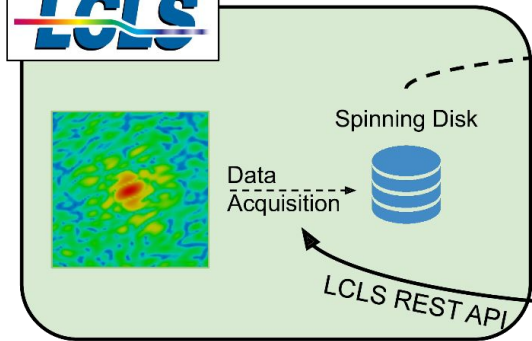


Job sentinel regularly queries the LCLS REST API for new experiments. It then compares the experiment list with the record of completed data analyses.

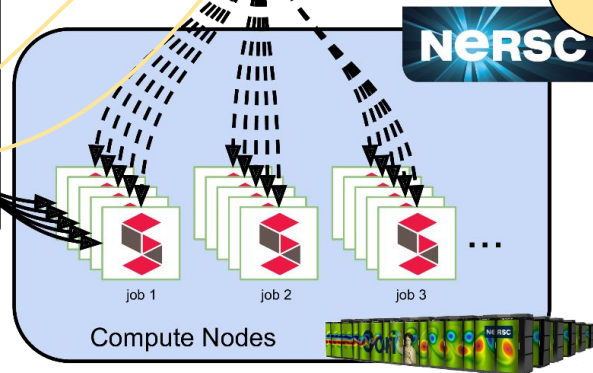
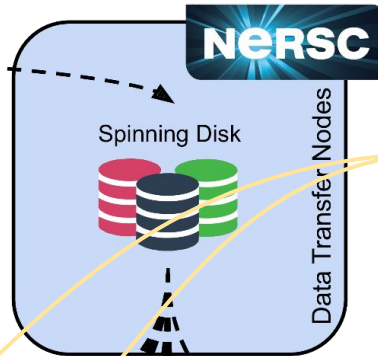
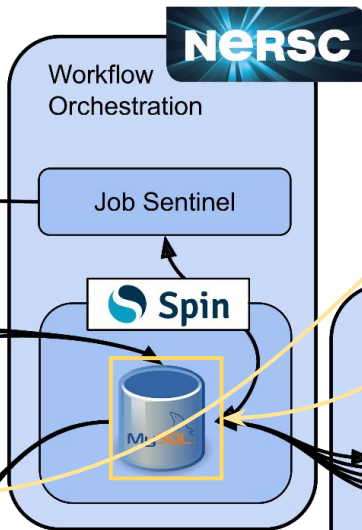








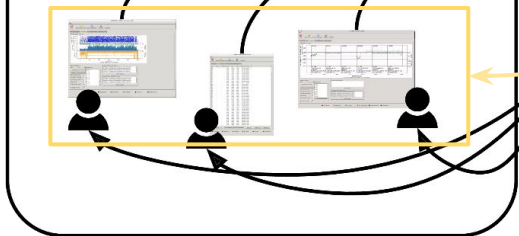
Data Mover (XRootD)



Users (AoE) can connect to the SQL database, to observe the progress of the data analysis.

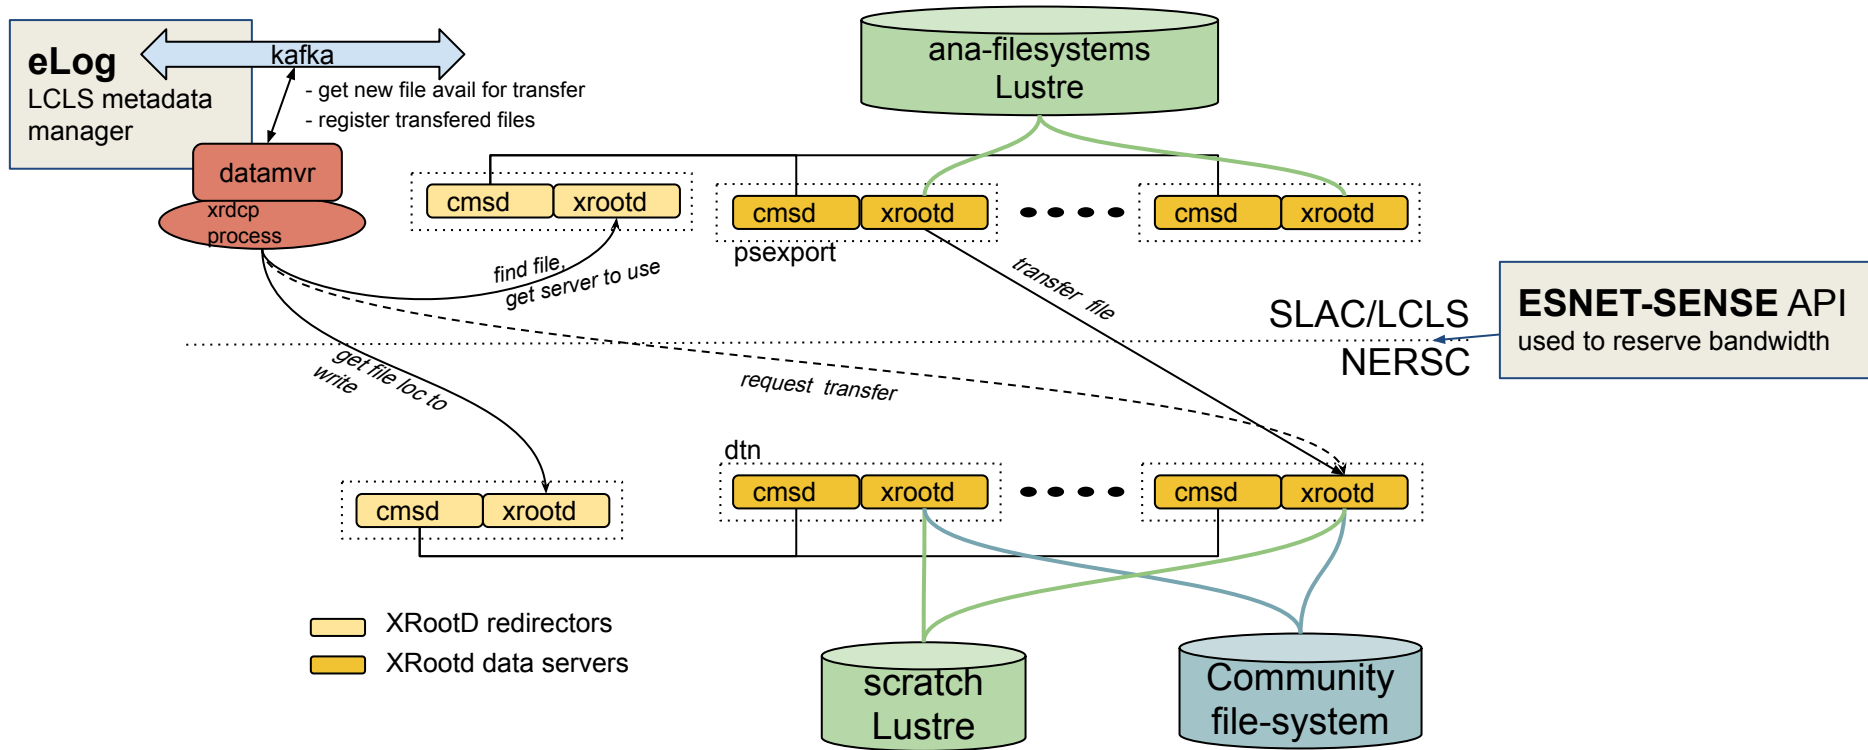
Users also commit new parameters to be analyzed

Users interact with data analysis in real-time



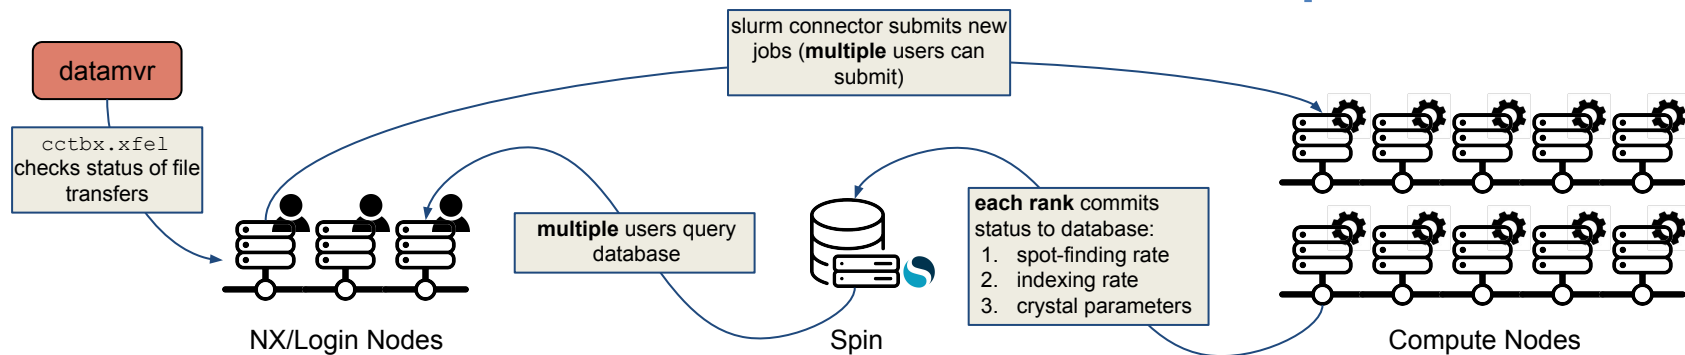


# Data Movement XRootD clusters

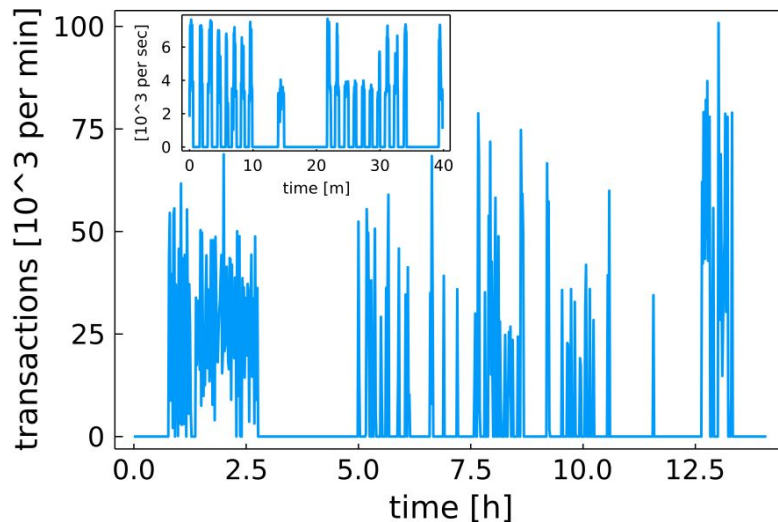




# Workflow Coordination hosted on Spin



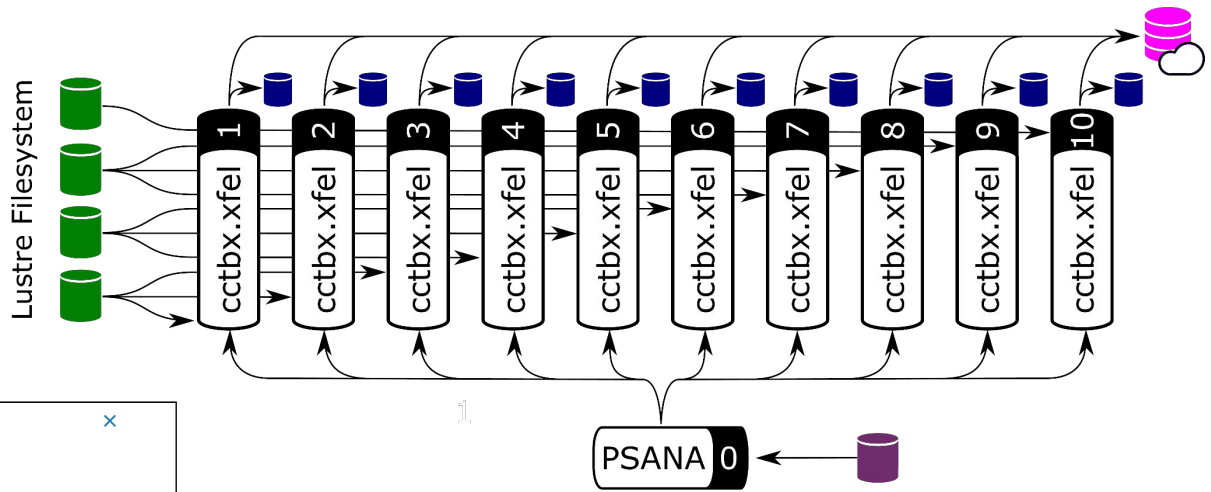
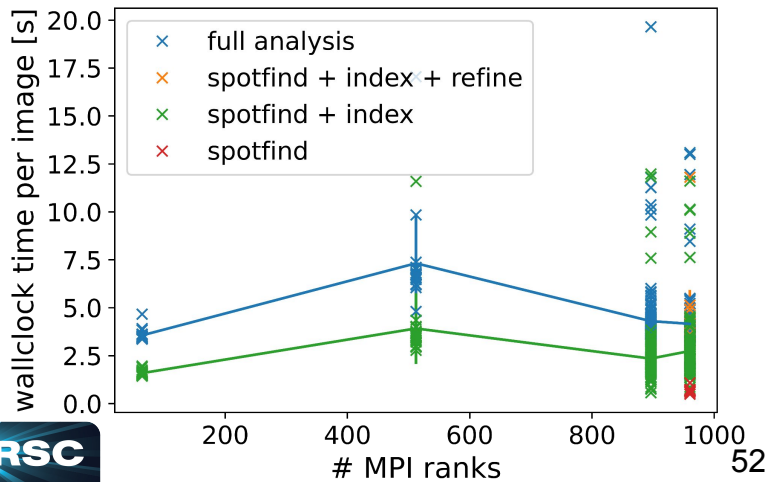
- Spin was capable of handling over 7k transactions/sec
- Transactions and connections need to be pooled when exceeding 4000 connections (ranks)



# Data Analysis

- Data analysis follows sequential stages:
  - spotfinding
  - indexing Bragg spots
  - model refinement
  - integrating Bragg spots

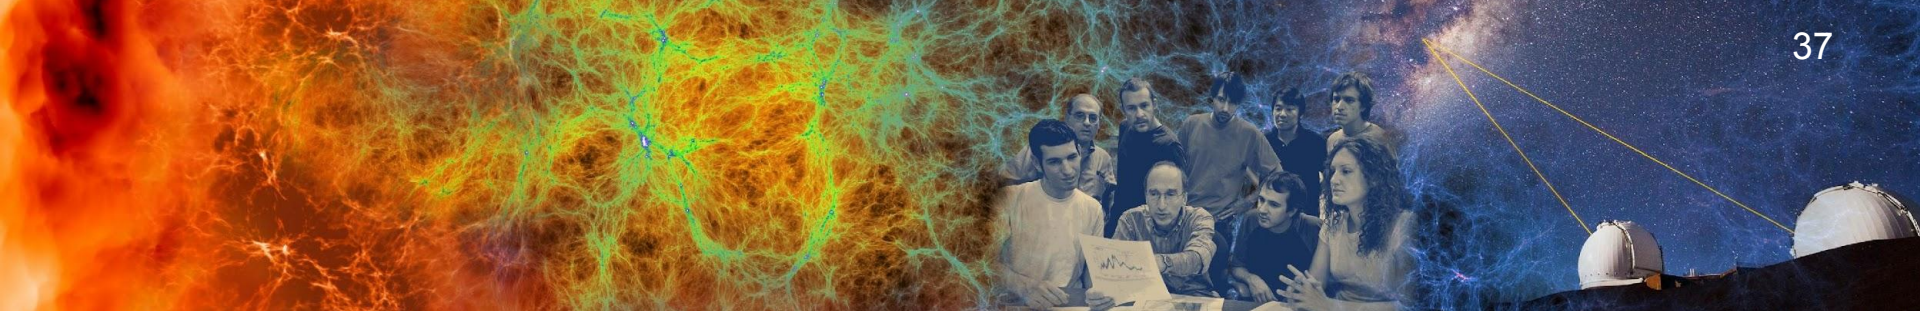
Reprocessing



Process Name MPI Rank microservices platform

MySQL database xtc index file xtc data file

results file



# Performance Optimization For The Exascale

# How's the Computation Weather Today?

- Computational Weatherplot:
  - Each line shows work done by one MPI rank
  - There is no “*the* cctbx.xfel workload”

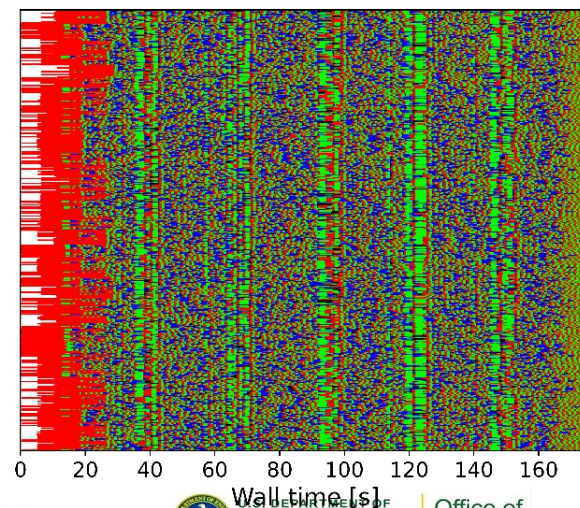
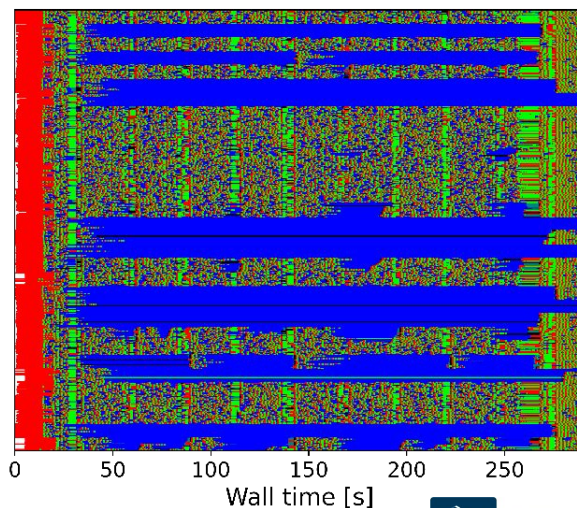
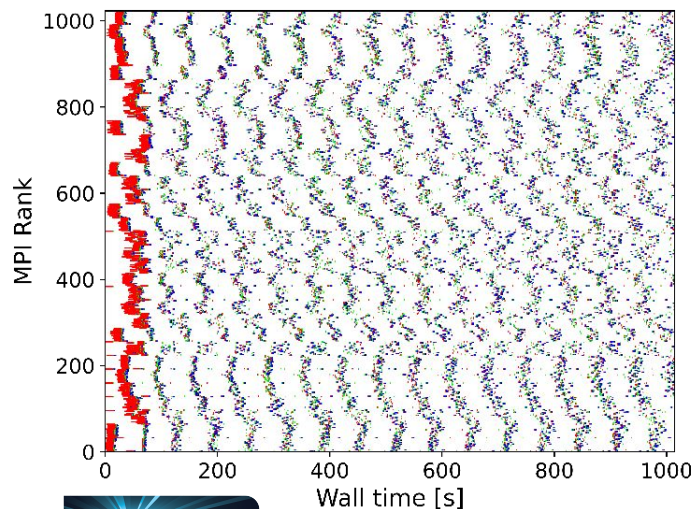
Start-Up and I/O (PSANA)

Spot Detection (DIALS)

Indexing (DIALS)

Refinement (DIALS)

Integrating (DIALS)





# How's the Computation Weather Today?

- Computational Weatherplot:
  - Each line shows work done by one MPI rank
  - There is no “*the* cctbx.xfel workload”

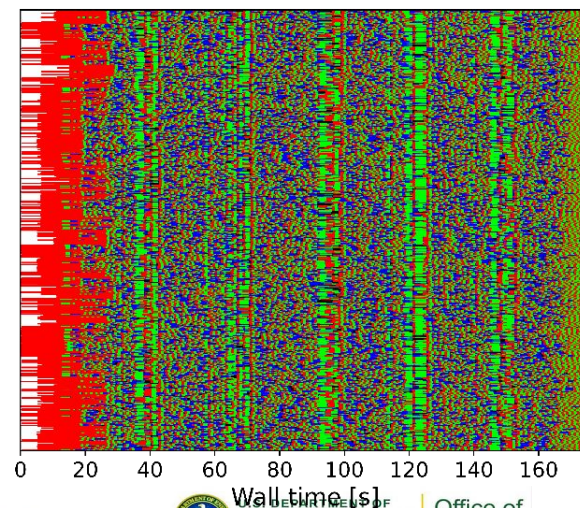
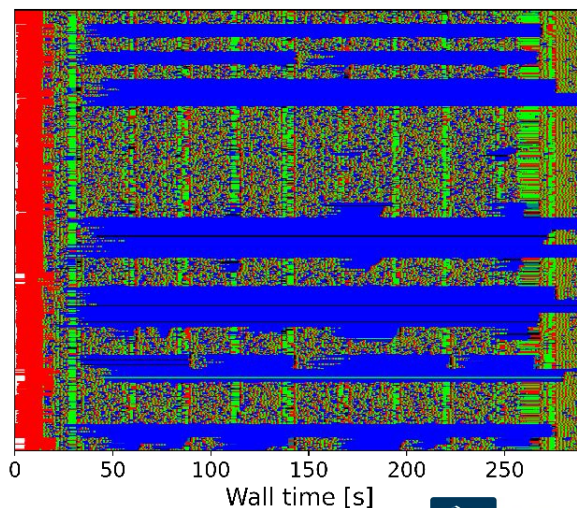
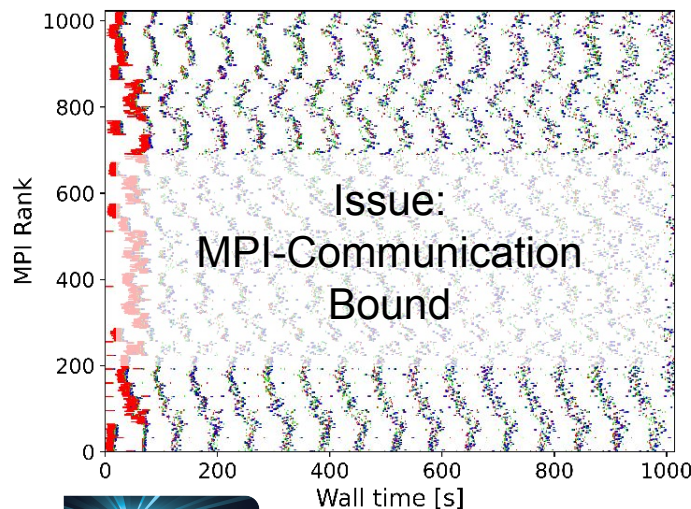
Start-Up and I/O (PSANA)

Spot Detection (DIALS)

Indexing (DIALS)

Refinement (DIALS)

Integrating (DIALS)



# How's the Computation Weather Today?

- Computational Weatherplot:
  - Each line shows work done by one MPI rank
  - There is no “*the* cctbx.xfel workload”

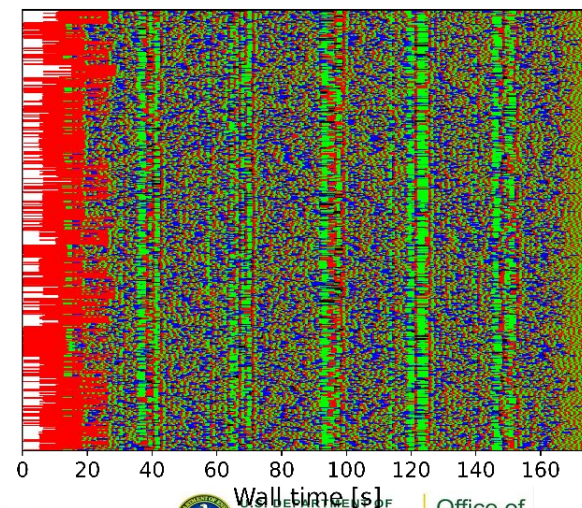
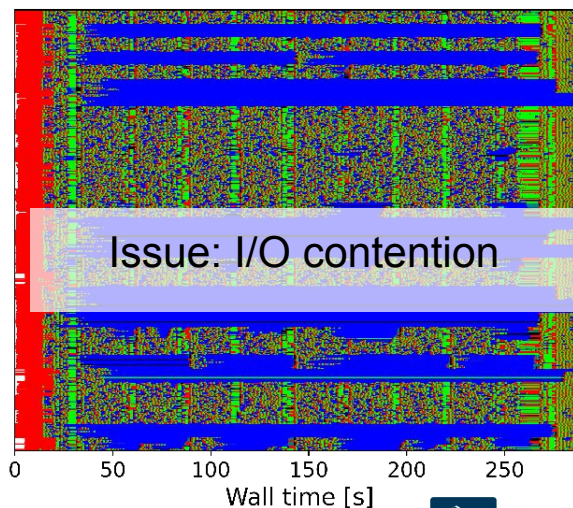
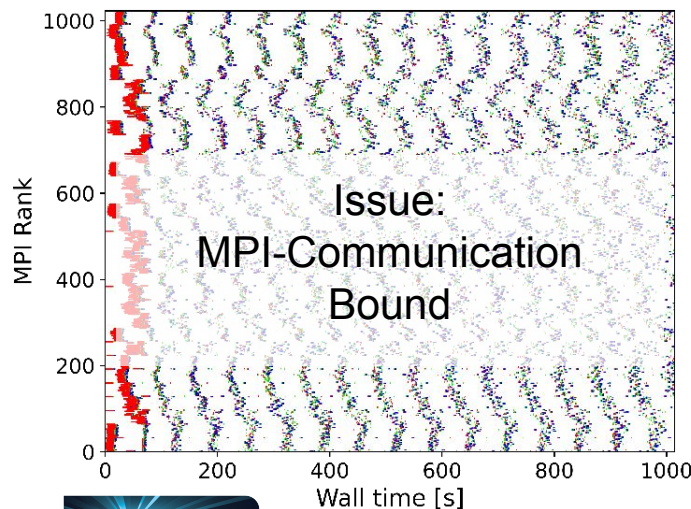
Start-Up and I/O (PSANA)

Spot Detection (DIALS)

Indexing (DIALS)

Refinement (DIALS)

Integrating (DIALS)

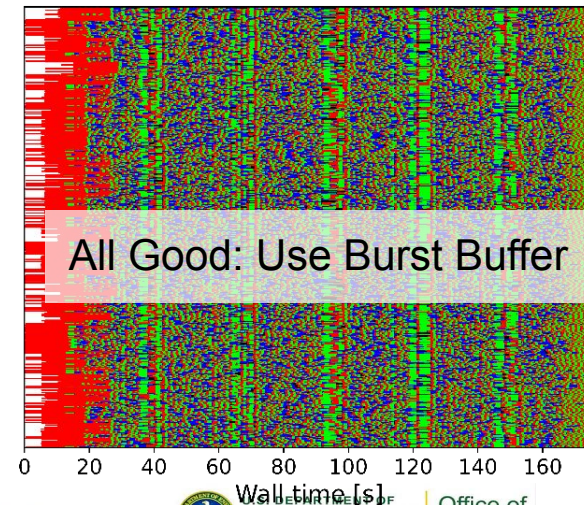
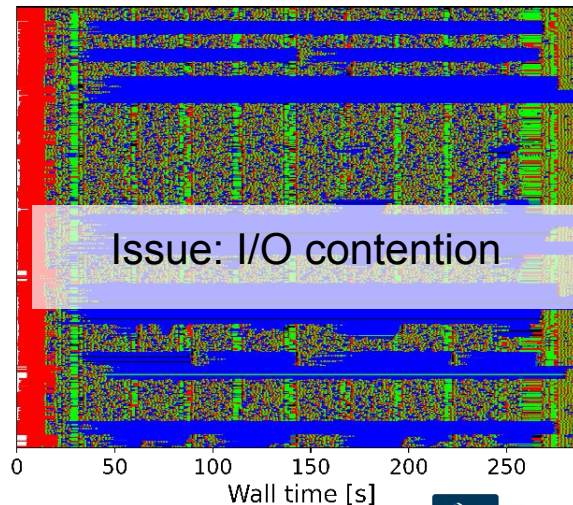
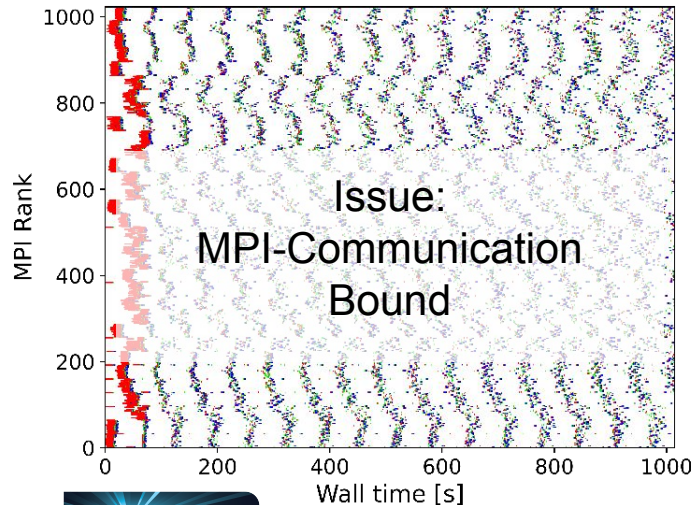




# How's the Computation Weather Today?

- Computational Weatherplot:
  - Each line shows work done by one MPI rank
  - There is no “*the* cctbx.xfel workload”

Start-Up and I/O (PSANA)  
Spot Detection (DIALS)  
Indexing (DIALS)  
Refinement (DIALS)  
Integrating (DIALS)



# Performance Portability

Work by Felix Wittwer



- Use Kokkos (<https://github.com/kokkos>) to generate code for each architecture
- Kokkos kernel abstractions are hardware independent

```
add_array.cu

__global__
void addArray(double* lhs,
              float* rhs,
              int total_pixels) {
    int j = blockDim.x * blockIdx.x + threadIdx.x;
    if (j < total_pixels) {
        lhs[j] = lhs[j] + (double) rhs[j];
    }
}
```

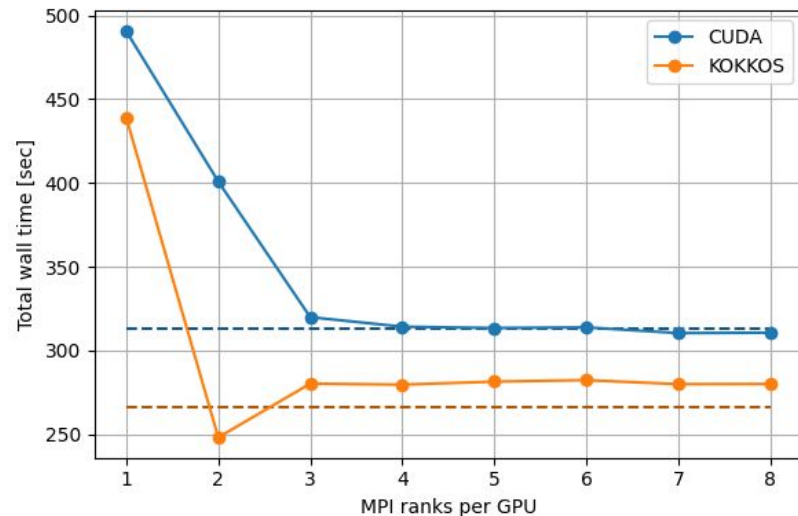
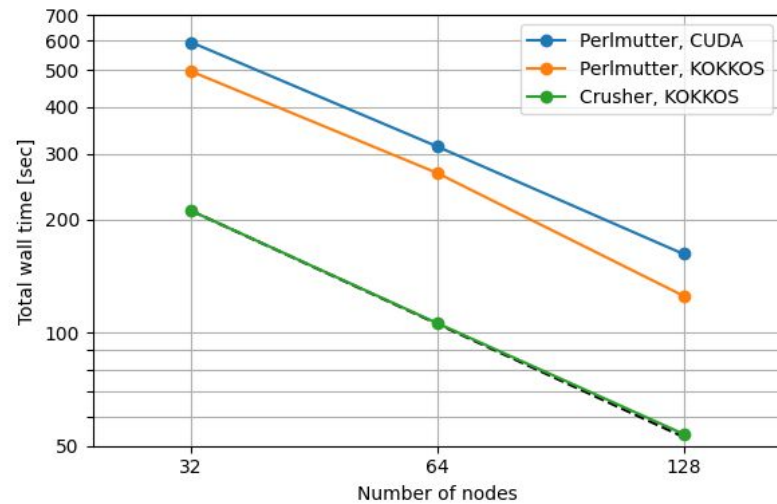


```
add_array.cpp

void addArray(Kokkos::View<double*> lhs,
              Kokkos::View<float*> rhs,
              int total_pixels) {
    Kokkos::parallel_for(
        "addArray", total_pixels,
        KOKKOS_LAMBDA (const int& j) {
            lhs(j) = lhs(j) + (double) rhs(j);
        }
    );
}
```

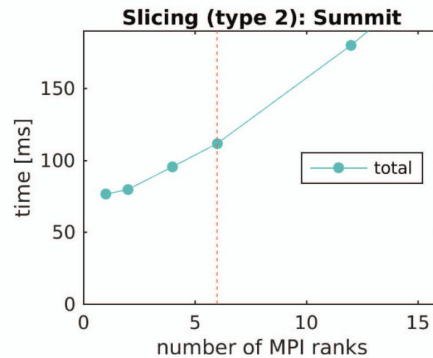
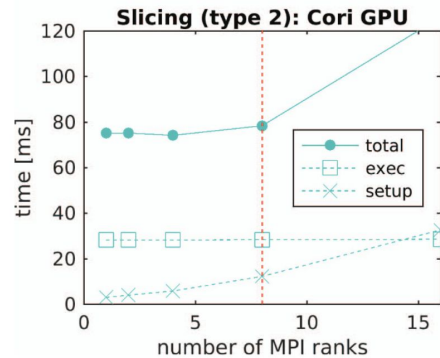
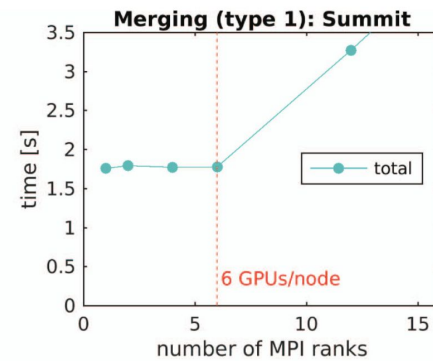
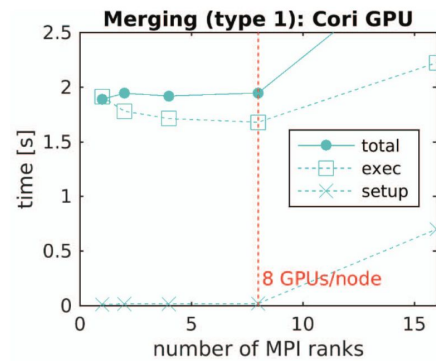
# Performance Portability

- Kokkos-generated code outperformed original CUDA kernels
  - Kokkos generated kernels better at filling available resources on GPU
- Hide workflow and I/O latency by sharing GPU among several MPI ranks
  - Ideal number of ranks/GPU depend on how much of the GPU each kernel can use



# Vendors + Libraries

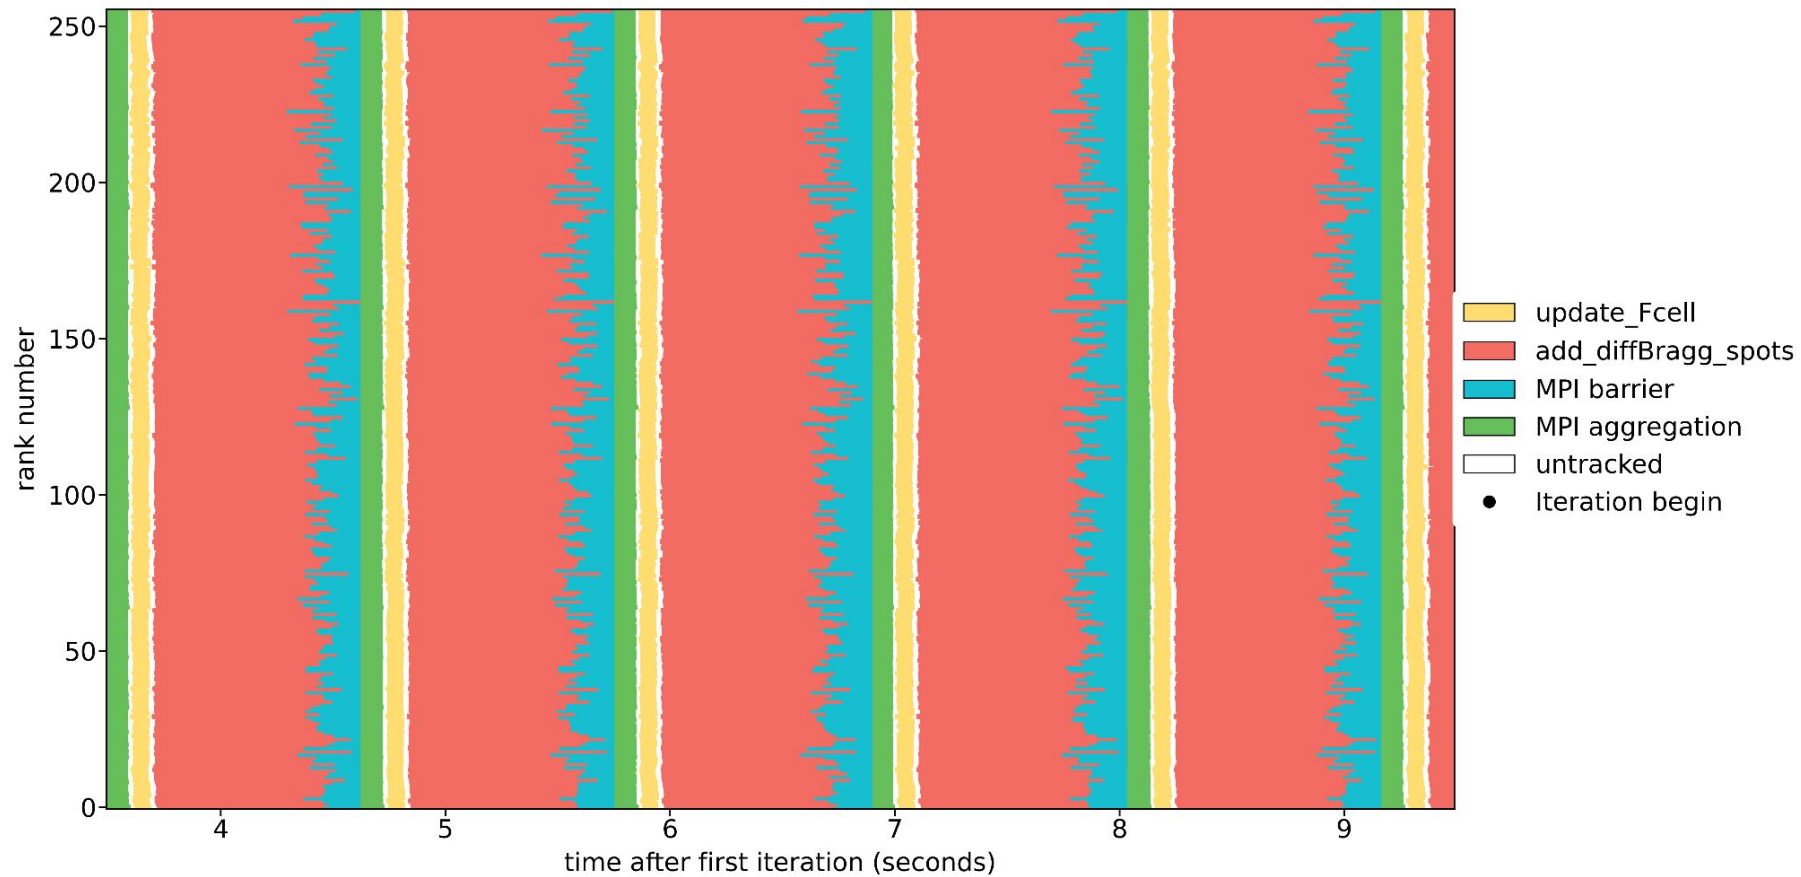
- Kokkos does not provide portable interfaces to vendor libraries (eg. cuBLAS, cuFFT/rockBLAS, rocFFT)
- Requires active **engagements with vendors and developer community** to design portable and performant libraries
- <https://github.com/elliottslaughter/cufinufft>
- <https://github.com/JBlaschke/PybindGPU>



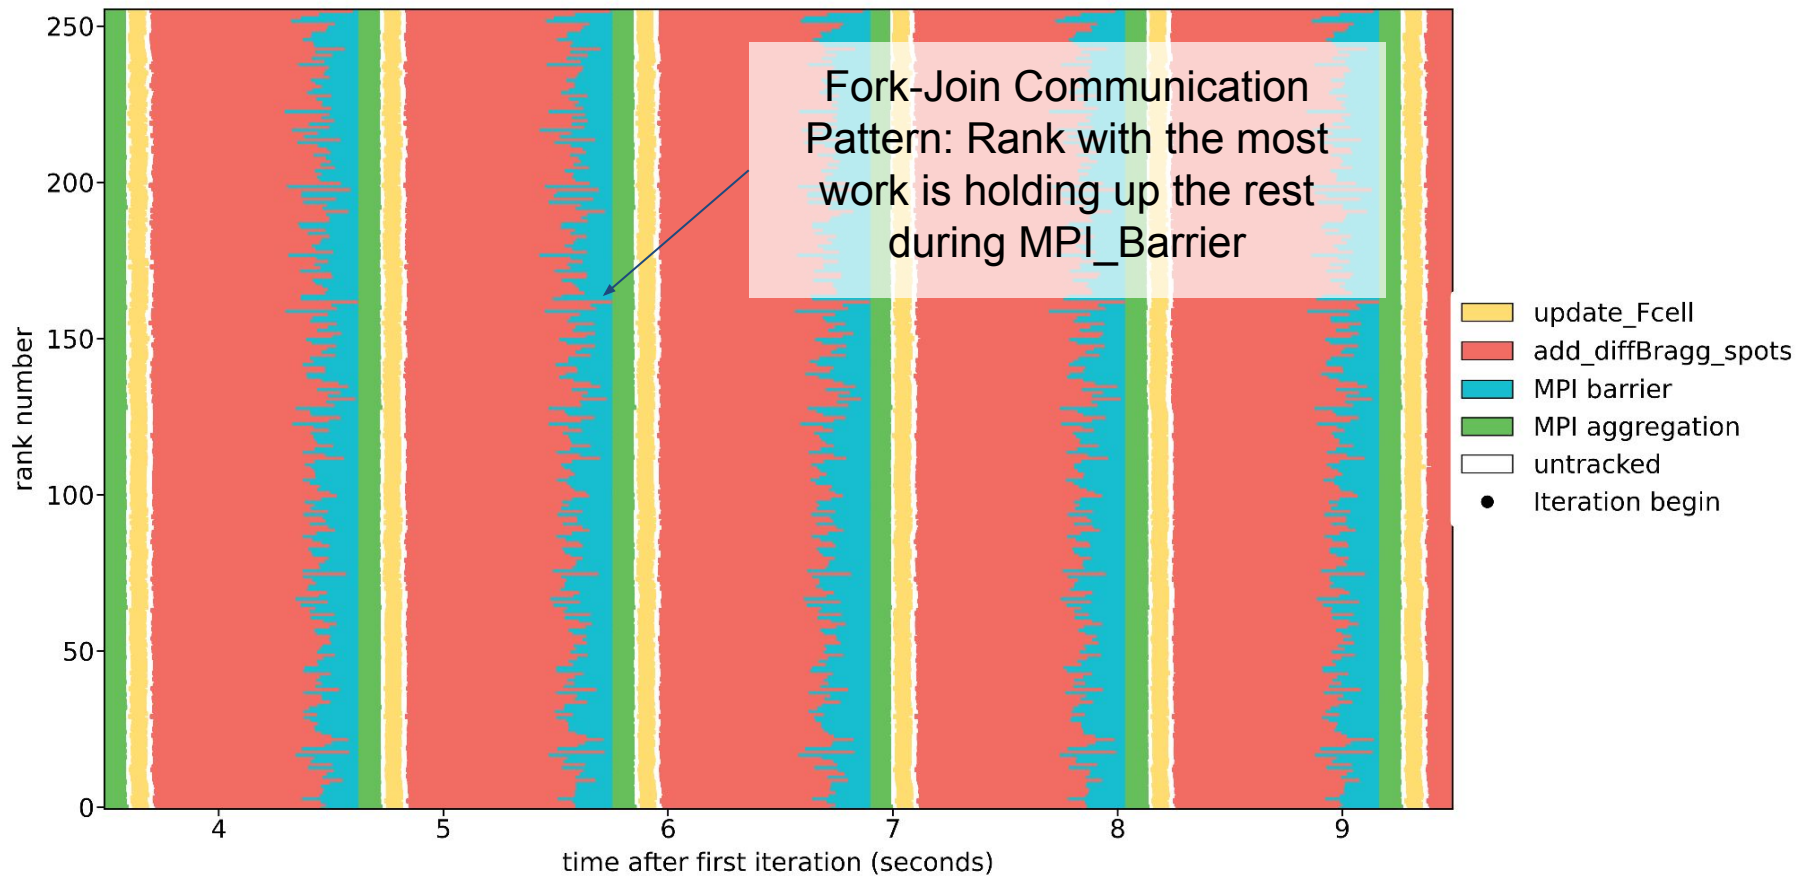
<https://arxiv.org/abs/2102.08463>



# Load Balancing

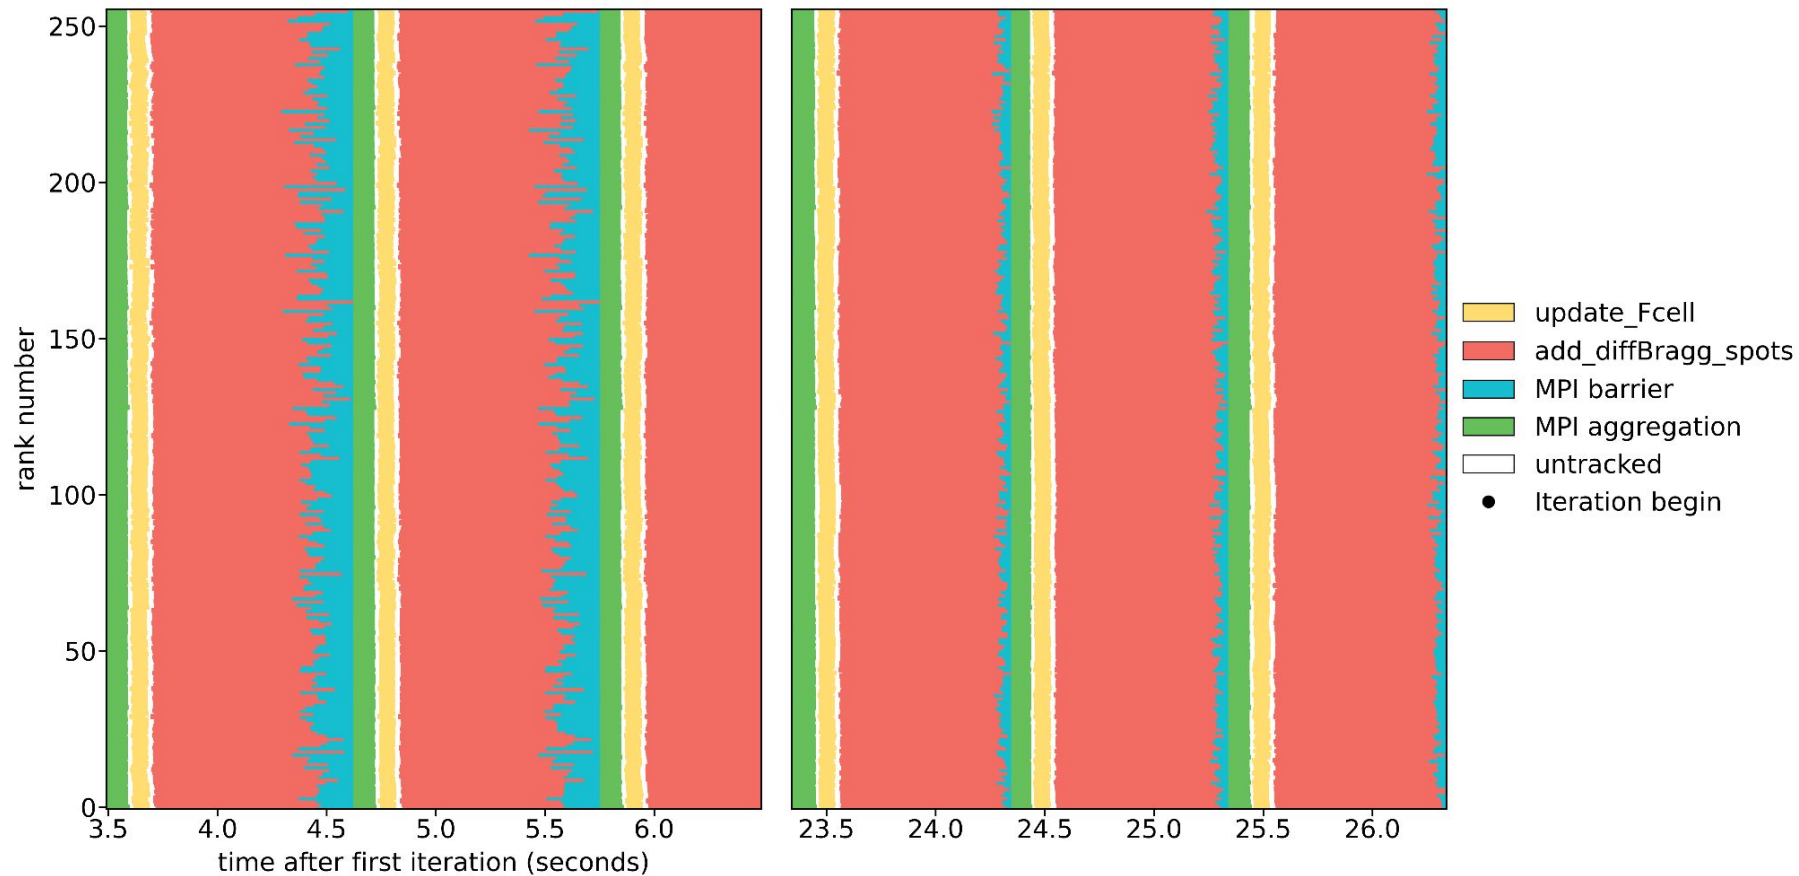


# Load Balancing



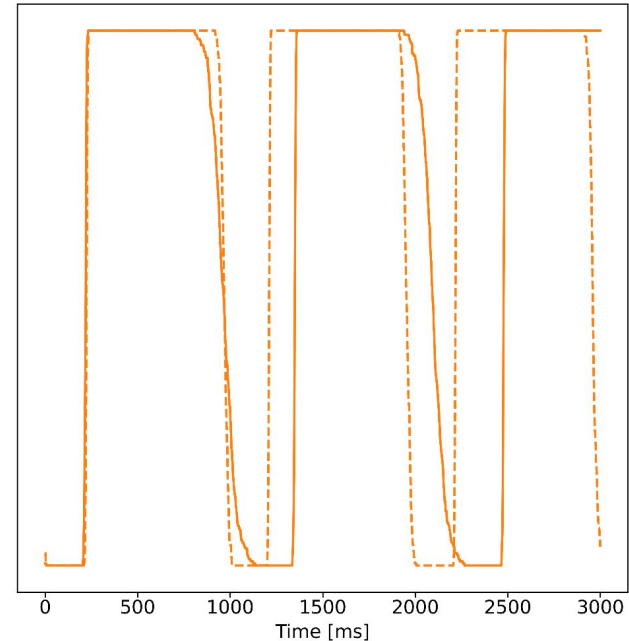
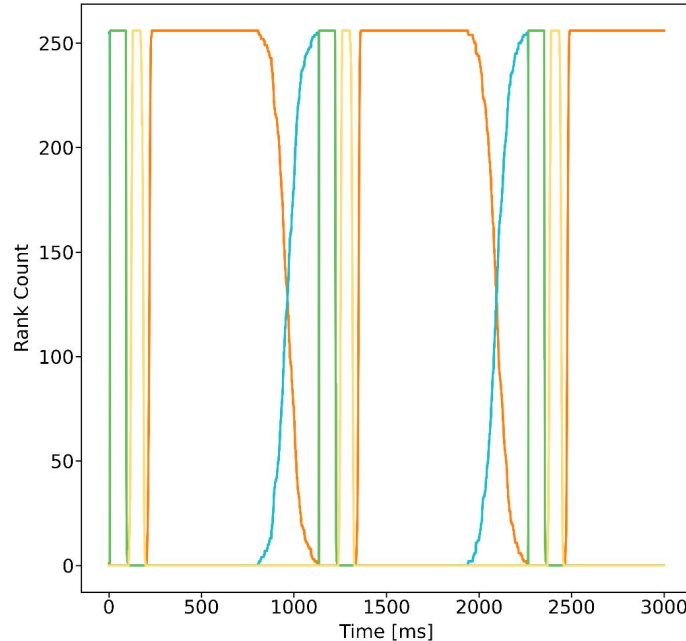


# Load Balancing



# Load Balancing

- Simple heuristic load-balancing (based on number of active pixels) results in very good load balancing
- Quick change results in 20% speedup





# Lesson Learned:

Performance + Portability

Planning

Packaging

Priority

Packets

# Performance + Portability

- When developing code of HW accelerators:
  - Keep portability a high priority: <https://github.com/kokkos/kokkos>
  - Kernel abstractions are a useful tool for getting good performance out of unfamiliar hardware (on a budget)
- Measure Everything! (profiling, debugging, tail latencies)
  - Instrumentation libraries (e.g. <https://github.com/LLNL/GOTCHA>, <https://github.com/NERSC/timemory> ) do this with minimal effort
  - Workflow managers come with some monitoring tools out of the box
- Fast Feedback!
  - Measured performance data needs to be communicated to users in a timely manner
- Porting ExaFEL is still a high-touch effort

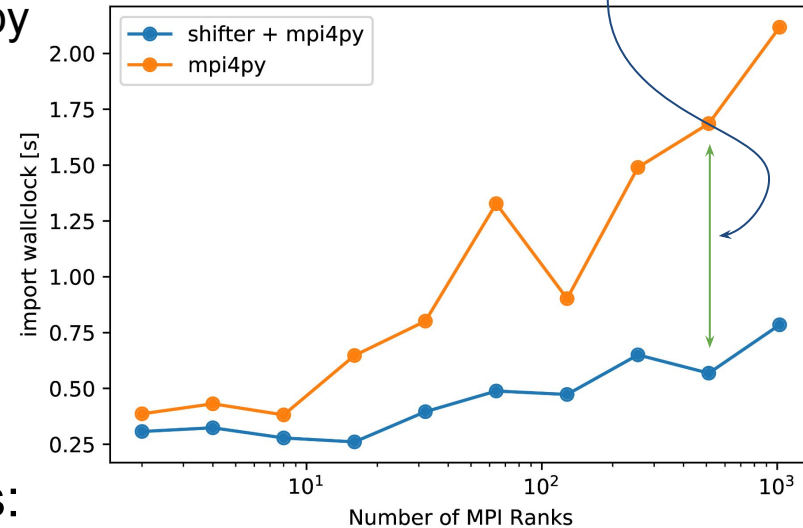
# Planning: Workflow Coordination

- New class of workflows include multiple data processing steps - Central **scalable** management and persistent state becomes necessary
  - Often hosted on login nodes  $\Rightarrow$  scaling becomes an issue
  - CCTBX: MySQL database for logging process state (completed, running, waiting)
  - Others: MongoDB (fireworks), or sometimes just writing to file
- Collaborative workspaces to facilitate dataset access and analysis
  - NERSC features **collaborative** accounts (users from a particular group can login with the collaborative account credentials)
- Use a workflow manager! (or a **data center's API**)
  - Data centers can customize these to account for their system needs

# Packaging: Containers

- Design images to be portable:
  - Can Streamline cross-facility workflows by standardizing image build process
  - Optimize dependency handling, data analysis workflows tend to have complex dependencies
  - Shifter / Podman mounts NERSC-specific libraries into image
  - **Isolate Anaconda's conflicts with OS**
- Significantly lower import/loading times:
  - Image cached to node-local storage
  - Improved scalability even for few MPI Ranks

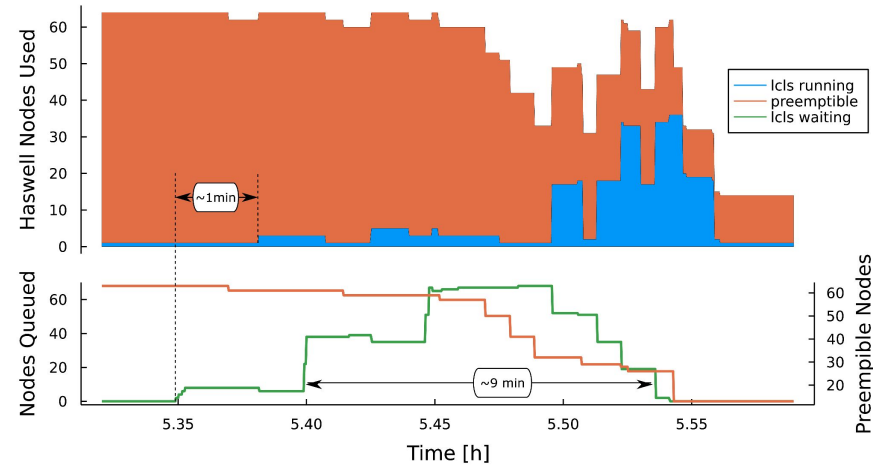
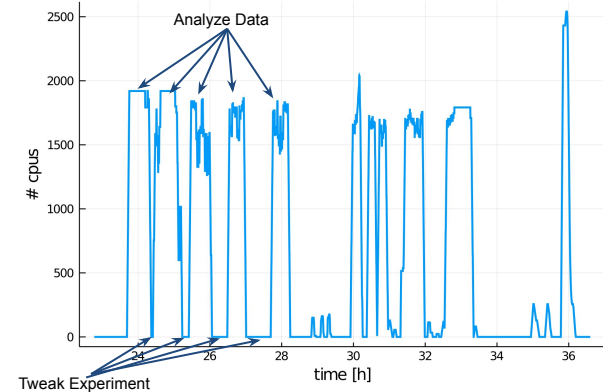
Dynamic linker's symbol resolution uses node-local storage  $\Rightarrow$  many ranks do not flood file-system.





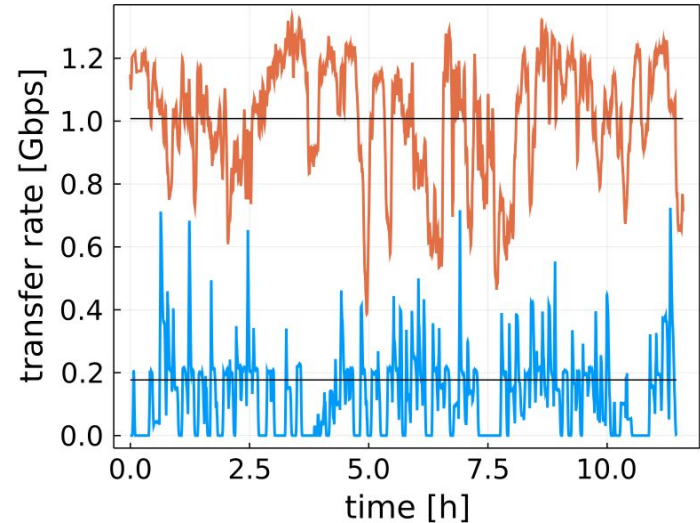
# Priority: Job Queue and Urgent Computing

- Real-time data analysis workloads are bursty:
  - Experimental operators require analysis to make decisions about future runs
  - Regular queue turnaround time insufficient for real-time workflow
- Reserving nodes for bursty use is wasteful:
  - Nodes would sit idle
  - Solution: share reserved nodes with preemptible jobs
- Future: Explore flexible resource allocation (eg. malleable jobs)



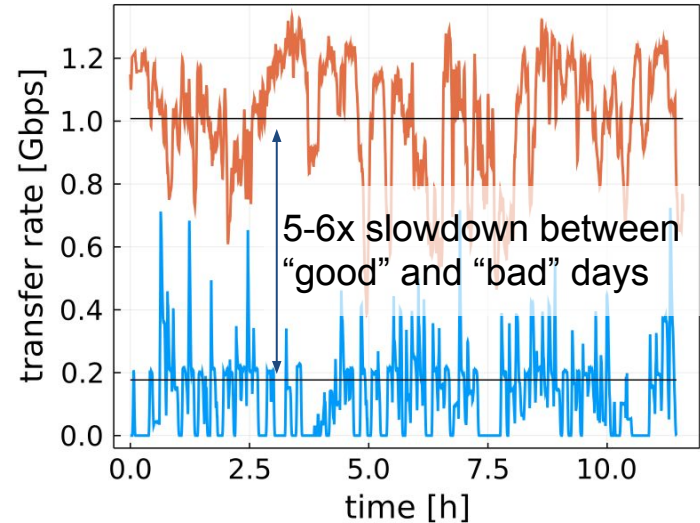
# Packets: Network and I/O

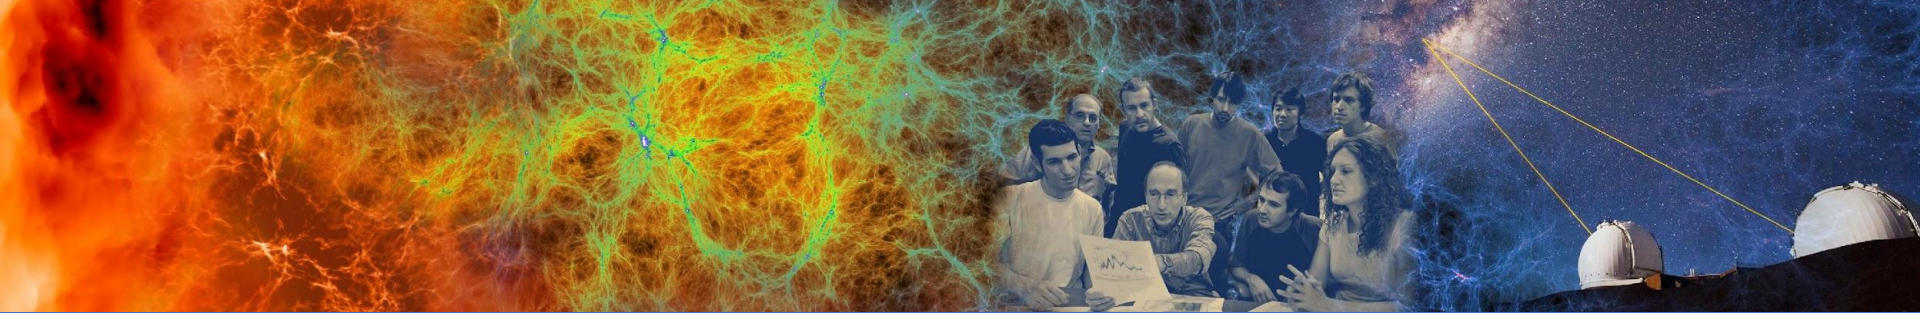
- In data analysis workflows, file systems and network can become bottlenecks
- I/O Optimization is Crucial:
  - Optimize software (eg. python logger) for high-frequency parallel I/O
  - Write logs to Burst Buffer
- Experience during Beamtimes:
  - Transfers ran smoothly (most of the time), can switch redirect destination in emergencies
  - **FS performance can limit the transfer rate**
- Future Work:
  - Explore in-memory burst buffers, eg <https://github.com/LLNL/UnifyFS>
  - Offer more fine-grained control over which data is transferred and when
  - Better monitoring and alerting



# Packets: Network and I/O

- In data analysis workflows, file systems and network can become bottlenecks
- I/O Optimization is Crucial:
  - Optimize software (eg. python logger) for high-frequency parallel I/O
  - Write logs to Burst Buffer
- Experience during Beamtimes:
  - Transfers ran smoothly (most of the time), can switch redirect destination in emergencies
  - **FS performance can limit the transfer rate**
- Future Work:
  - Explore in-memory burst buffers, eg <https://github.com/LLNL/UnifyFS>
  - Offer more fine-grained control over which data is transferred and when
  - Better monitoring and alerting





# Extra Bits



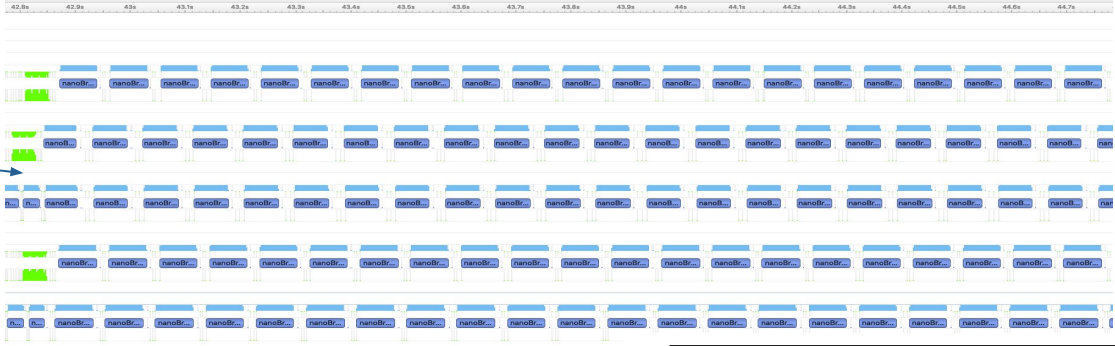
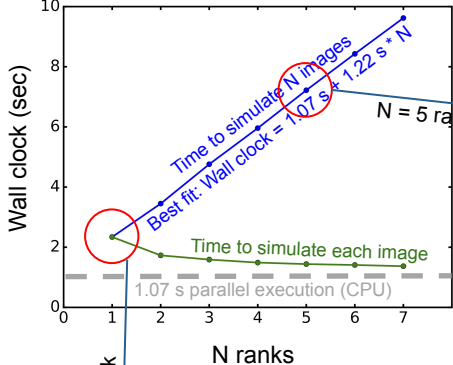
BERKELEY LAB



U.S. DEPARTMENT OF  
**ENERGY**

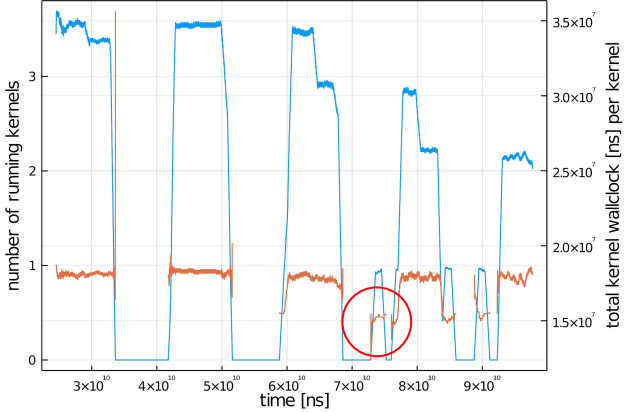
Office of  
Science

# Sharing GPUs Between MPI Ranks



■ Host to device memcopy  
■ Kernel execution

— 100 ms



# Example Issue: LCLS Slow File Transfer

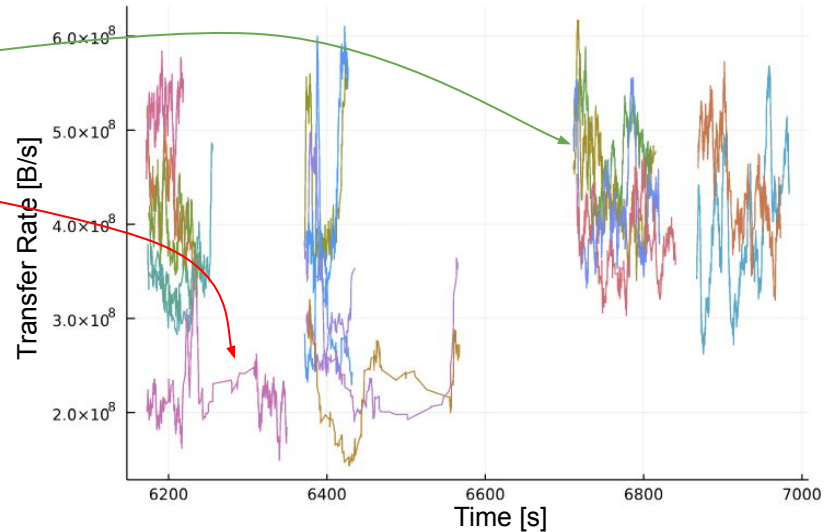
1. Files transfer from LCLS via LCLS SCRATCH => ESNet => NERSC Scratch
2. Most files transfer quickly
3. Some files come in slowly

4. Open Work Items:

- a. Why?
  - i. LDMS on DTN (identify slow FS)?
  - ii. Work with ESNet?

- b. Accept slow files

=> Design workflow so that slow files don't hold up other work





# Fireworks Workflow Manager

```
workflow.py

qadapter_1 = CommonAdapter(
    q_type = "SLURM",
    rocket_launch=f"rlaunch --json -l '{lp_str}' -w '{fw_str_1}' singleshot",
    constraint="gpu",
    account="nstaff",
    walltime="'00:02:00'",
    qos="regular",
    nodes="2"
)

fw_1 = Firework(ScriptTask.from_str(
    "srun -n 2 cctbx.python -m mpi4py.bench helloworld"
), spec={"_category": "n2", "_fworker": "mpi_2_fworker"}, name="mpi_2")

workflow = Workflow(
    [fw_1, fw_2],
    {fw_1: fw_2},
    name="multi_queue"
)

launchpad.add_wf(workflow)
rapidfire(launchpad, fworker_1, qadapter_1, sleep_time=60, reserve=True)
```

# This is really the end

