Oct. 21, 2025, SLAC National Accelerator Laboratory
Productive, Performant Software for Large Scale Scientific Data Analysis

The Grand Bargain of Structural Biology at the kHz Frontier

Nicholas Sauter & David Mittan-Moreau

Lawrence Berkeley National Laboratory

BERKELEY LAB
Bringing Science Solutions to the World

What is serial crystallography?

Serial crystallography

Serial vs. traditional crystallography
- Room temperature
- Free from radiation damage
- Pump probe

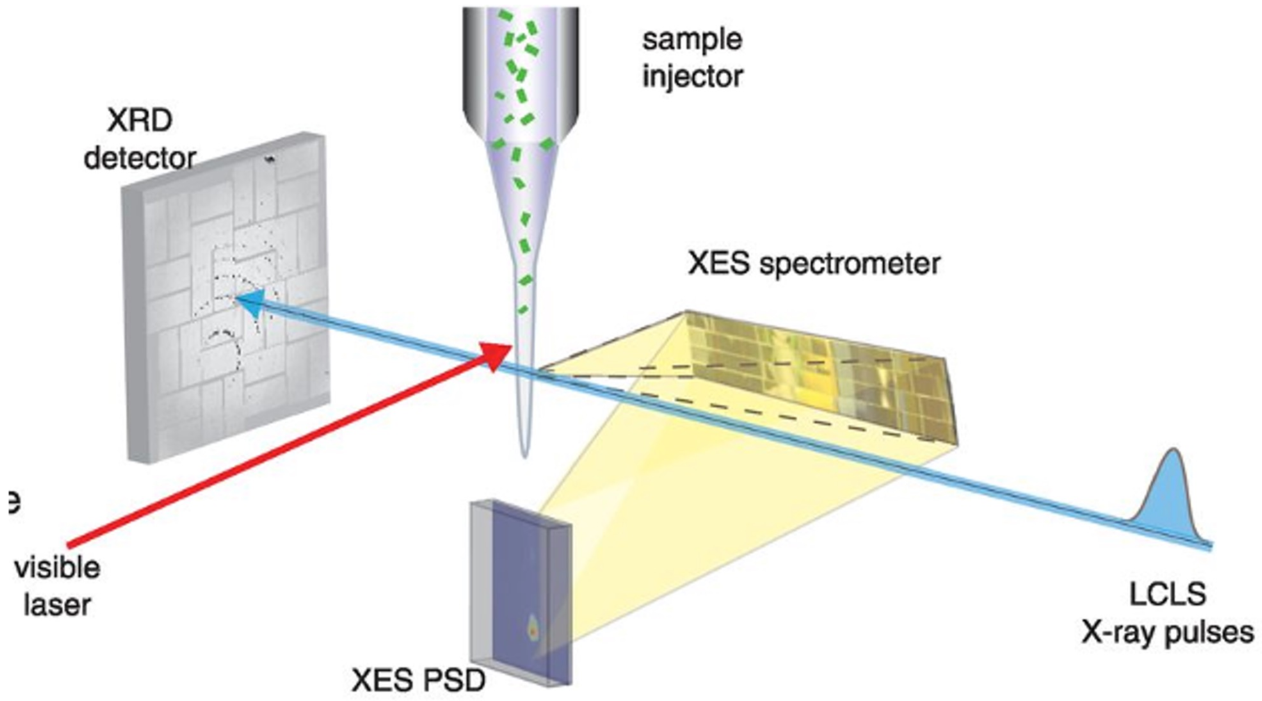

A 15-year track record of XFEL serial crystallography

Light activation

Substrate mixing / gas activation

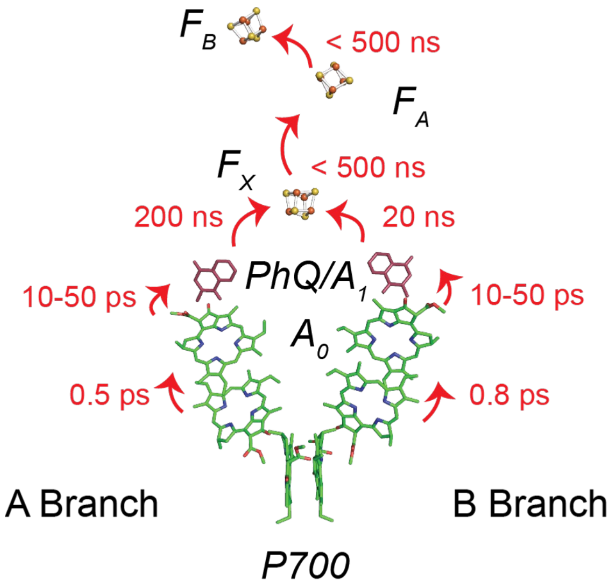

Photosystem I:
Captures sunlight energy

$NADP^+ + 2e^- + H^+ + 2\gamma \rightarrow NADPH$

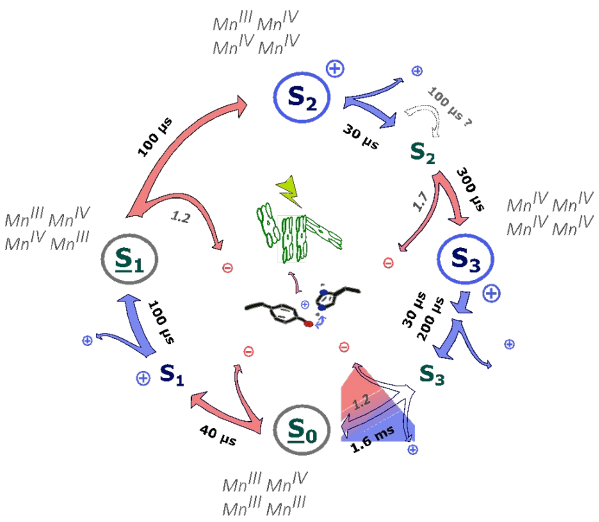

Photosystem II:
Generates atmospheric oxygen

$2H_2O + 4\gamma \rightarrow O_2 + 4e^- + 4H^+$

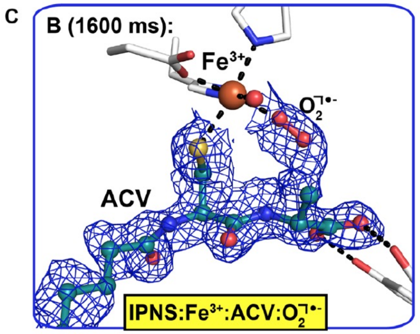

Isopenicillin N synthase:
Production of β-lactam antibiotics

$ACV + O_2 \rightarrow IPN + 2H_2O$

fs / ps **ns** **µs** **ms** **s**

A 15-year track record of XFEL serial crystallography

Light activation

Substrate mixing / gas activation

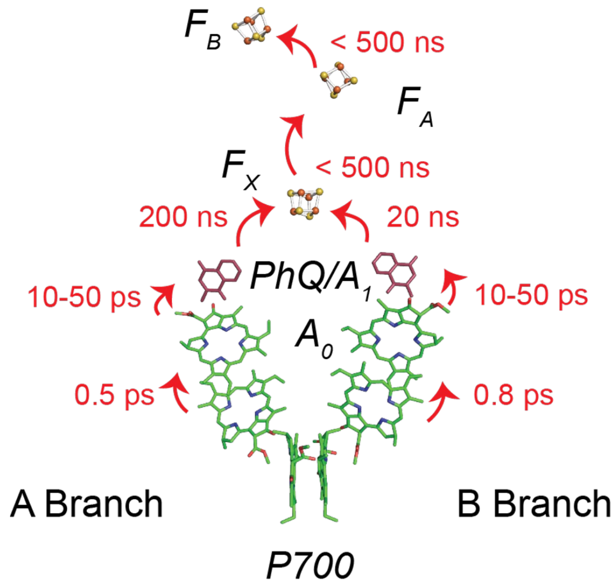

Photosystem I:
Captures sunlight energy

$NADP^+ + 2e^- + H^+ + 2\gamma \rightarrow NADPH$

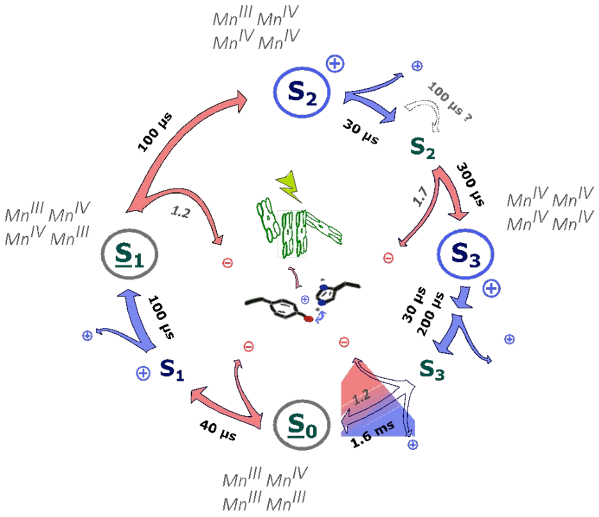

Photosystem II:
Generates atmospheric oxygen

$2H_2O + 4\gamma \rightarrow O_2 + 4e^- + 4H^+$

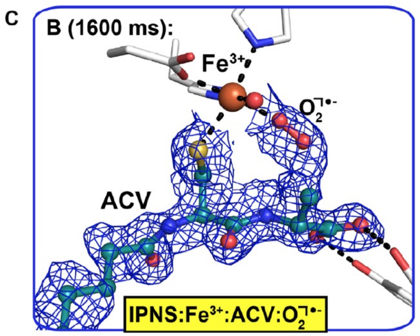

Isopenicillin N synthase:
Production of β-lactam antibiotics

$ACV + O_2 \rightarrow IPN + 2H_2O$

fs / ps — **ns** — **µs** — **ms** — **s**

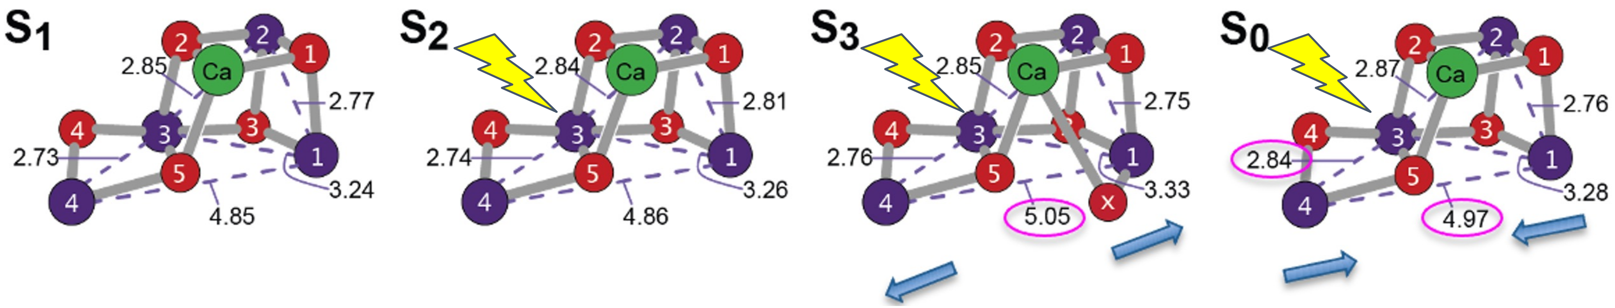

A 15-year track record of XFEL serial crystallography

Light activation

Substrate mixing / gas activation

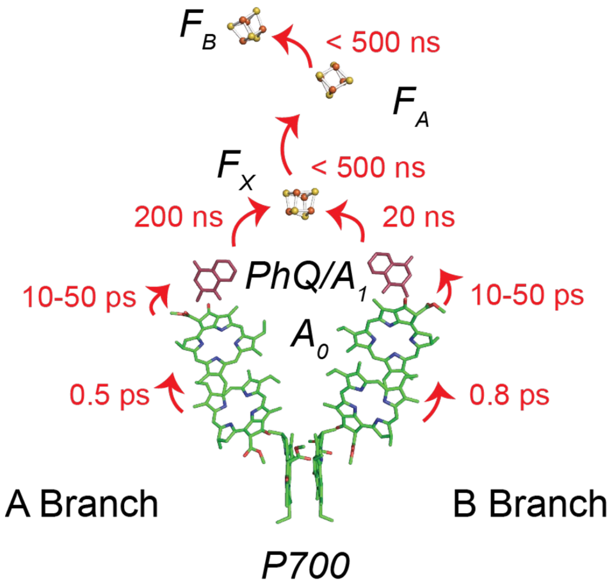

Photosystem I:
Captures sunlight energy

$NADP^+ + 2e^- + H^+ + 2\gamma \rightarrow NADPH$

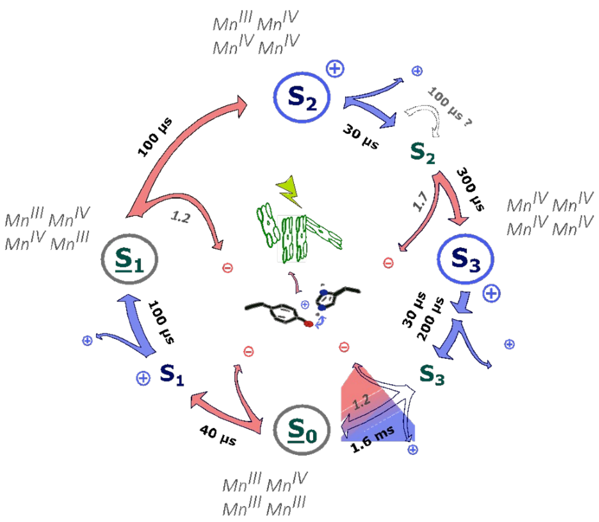

Photosystem II:
Generates atmospheric oxygen

$2H_2O + 4\gamma \rightarrow O_2 + 4e^- + 4H^+$

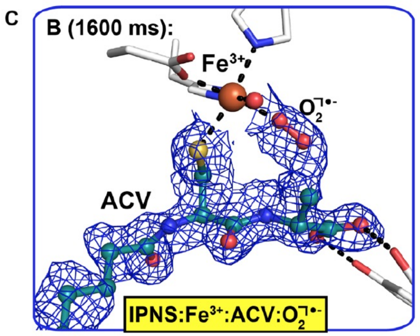

Isopenicillin N synthase:
Production of β-lactam antibiotics

$ACV + O_2 \rightarrow IPN + 2H_2O$

fs / ps — **ns** — **μs** — **ms** — **s**

Drop-on-drop:
Protein crystals: 30 Hz, 160µm drop, 2nL
Substrate: 1.6 kHz, 50µm drop, 60pL
For 5µm crystals the diffusion time is ~5ms

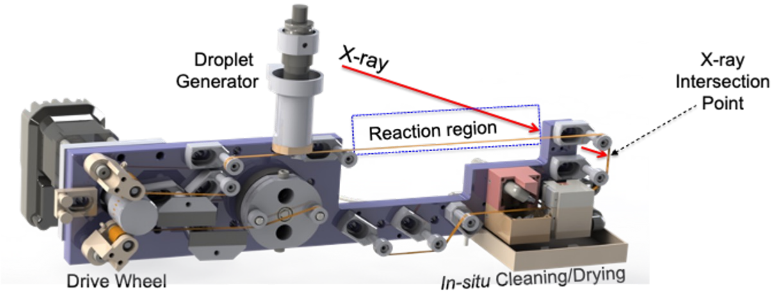

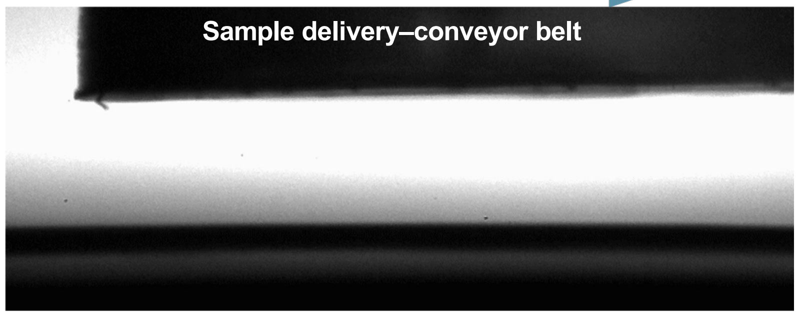

Why is speed a consideration for serial crystallography?

Partial diffraction patterns:
merge data from thousands of crystals

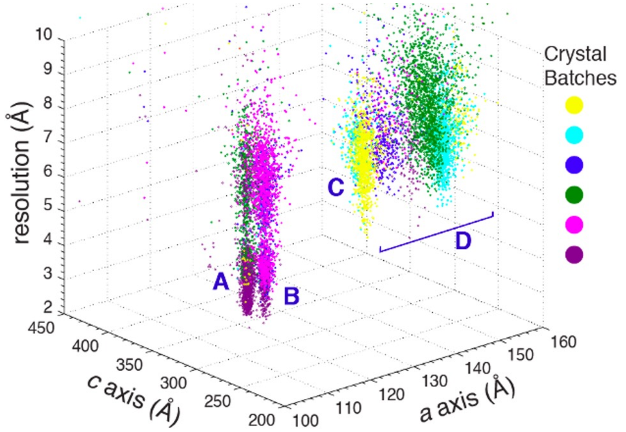

Partial reflections & disperse beam:
apply complex computational models to deduce structure factors

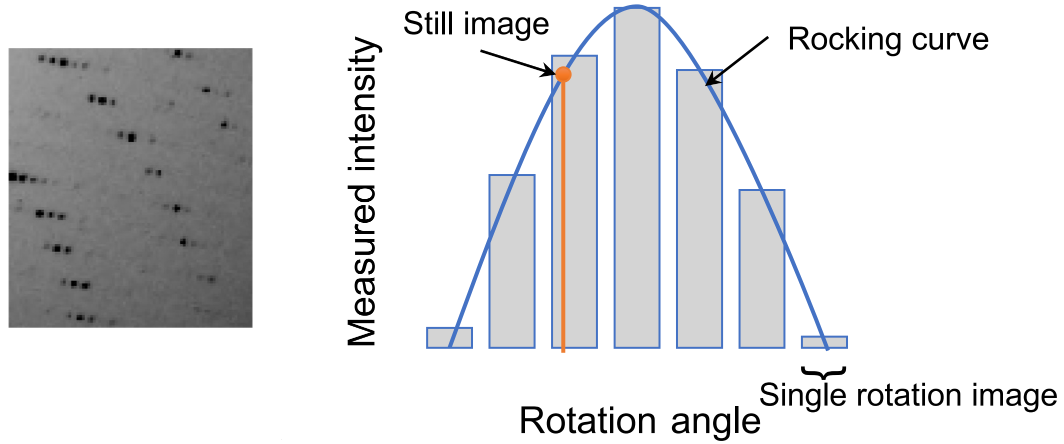

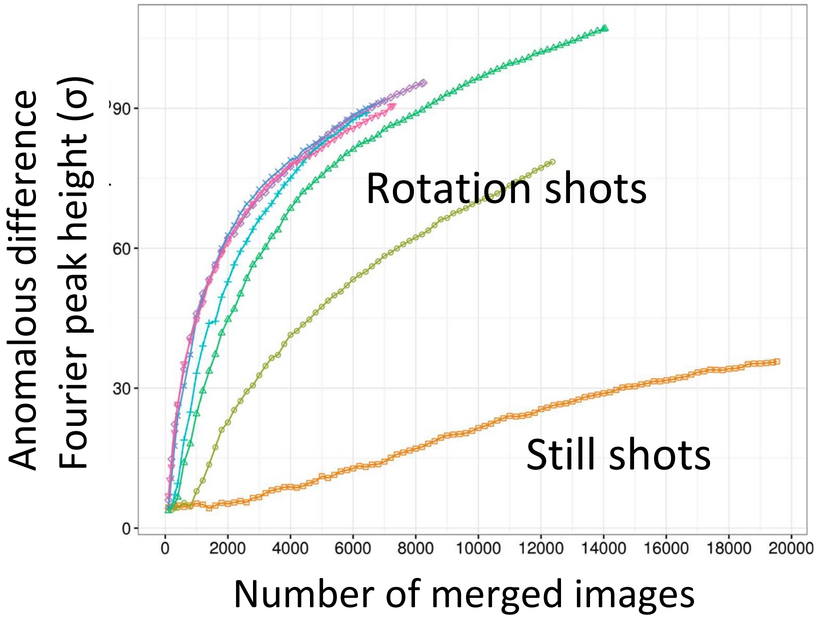

Continual assessment of whole data quality
to decide stopping point,
and compare alternate processing pipelines

Hasegawa, 2017

Why is an MX XFEL Experiment so Stressful?

Jan Kern: "For a 5-shift Photosystem experiment (60 hrs @ 20 Hz) we brought 15 mL sample. Making 1 mL takes 1 week for 1.5 people, so 22 weeks FTE for 5 shifts."

Sample delivery:

- Is the beam hitting the sample?
- Can we further optimize our parameterization

Xray emission spectroscopy:

- Is our sample in the correct oxidation state?

Diffraction data processing:

- Diffraction to high resolution?
- Unit cells isomorphous?
- When can we stop collecting data?

Beamline operations:

- Are there upstream issues?

PI Discussion:

- What sample or time point should we collect next?

Record high-level metadata in Google sheets

Structure solution:

- Do our results support our hypothesis?

Considering the enormous cost of the experiment in terms of sample and beamline operation, the most important mitigation we can offer is immediate data reduction with a 10-minute turnaround time, evaluating the final science metrics in order to know if the experiment is on track and when to move to the next sample.

Collaboration with LCLS, ESNet & NERSC on the assessment pipeline

Long history
- 2014 Edison
- 2017 Cori
- 2021 Perlmutter
- 2027 Doudna

Continual evolution
- Beam
- Detector
- Compute cluster
- Algorithms

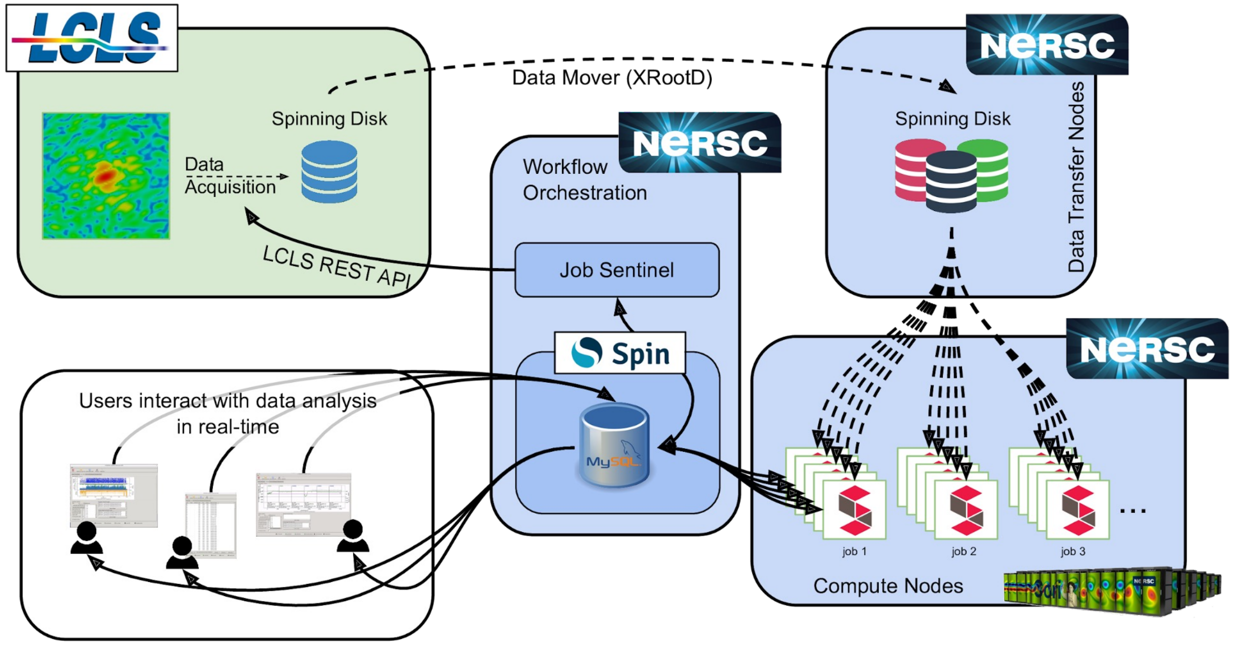

Current scale
- Jungfrau 16 Mpx x 2-byte depth
- Clocked at 120 Hz for 12 hrs
- 150 TB data

We have maintained our 10-minute turnaround for the past 8 years

What we have (and want to keep!):
- Dual implementation for FFB @ S3DF + complete workup at NERSC
- ESNet at 100 Gb/s
- Petabyte disk allocation
- 32 CPU nodes in the realtime queue on demand for 12 hours
- Immediate spinup of MySQL db

Things we would like, please

What we don't have now for serial crystallography at 2 kHz

- Clearer understanding of the goals, and which experiments benefit
- More accurate algorithms
- Much faster feedback to the experiment
- Failover to backup facility
- Data processing portability
- High level standardization
- More things automated

What could we accomplish at 2 kHz?

Current experience at 30 Hz

MFX-LU5017

- 48 hours
- 48% duty cycle
- 594,230 indexed images
- 54 data sets
- 1 mg of protein per data set
- 12 types of enzymes

Chemical crystallography
Mail-in pilot program

- 24 hours
- 44 samples screened
- 14 different user groups
- 11 structures solved

LCLS Mail-in small-molecule serial femtosecond crystallography

Submission Deadline:
1st February 2026

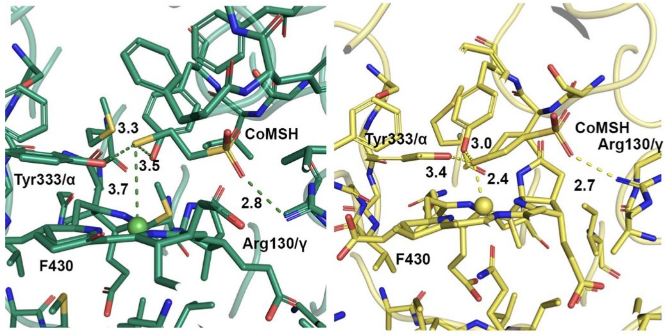

Now scale this to 2 kHz!

Structure Guided Drug Design

Serial crystallography is performed at room temperature without perturbation from cryoprotectants.

- Ligand binding pose is more physiologically relevant.

RT Serial

Cryogenic

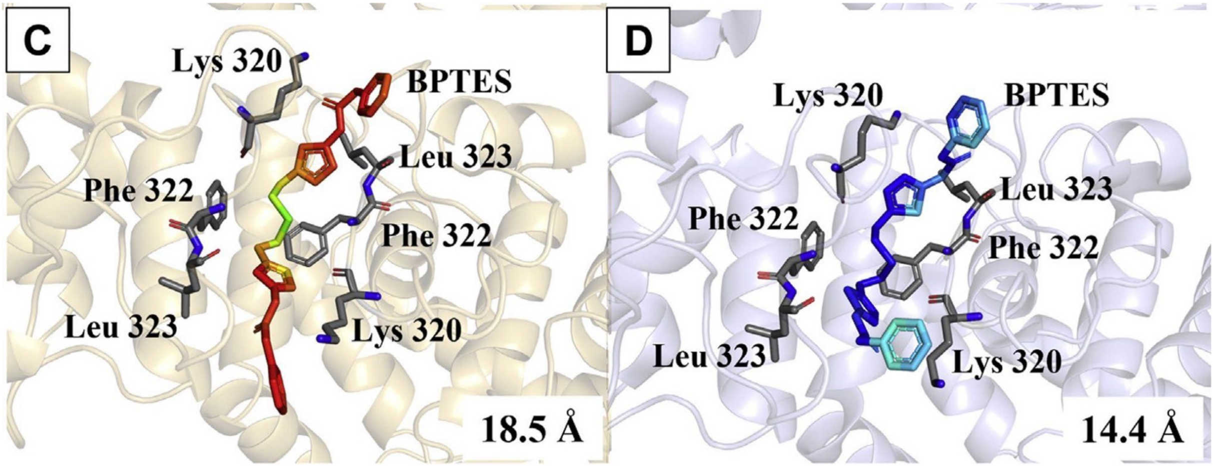

Milano, S, *et al.* (2022). New insights into the molecular mechanisms of glutaminase C inhibitors in cancer cells using serial room temperature crystallography*. J. Biol. Chem.* [DOI: 10.1016/j.jbc.2021.101535]()

Structure Guided Drug Design

Fragment screening starts computationally, then high throughput binding assays.

Serial crystallography at 2 kHz would be a more informative high throughput assay.

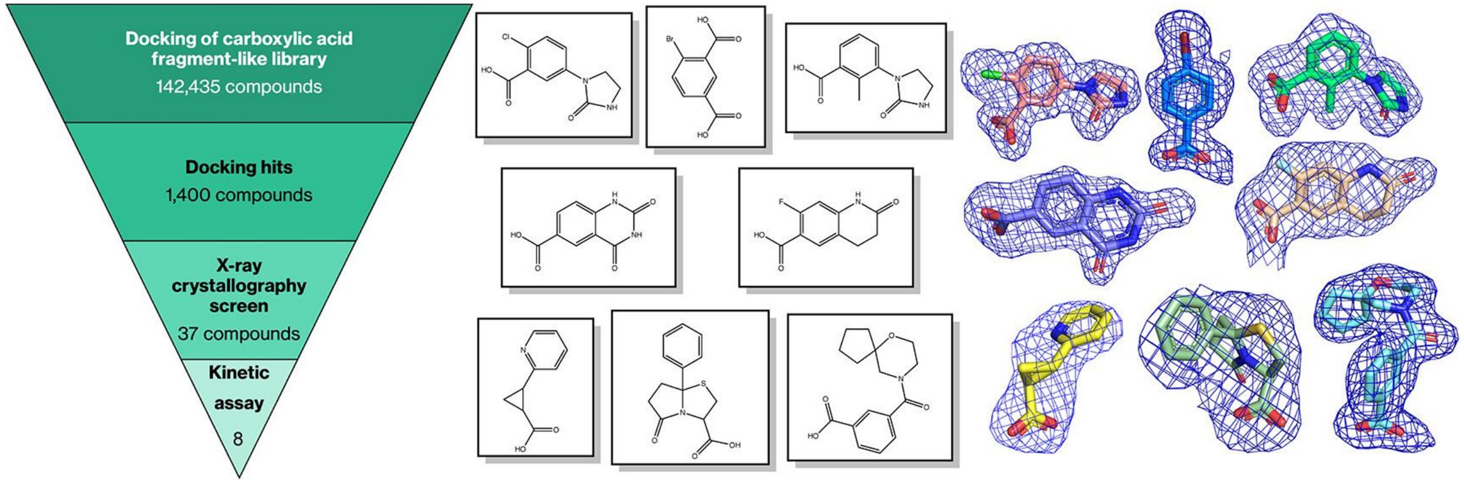

Meeks, K, *et al.* (2024). Novel Fragment Inhibitors of PYCR1 from Docking-Guided X-ray Crystallography. *J. Chem. Inf. Model.* DOI: 10.1021/acs.jcim.3c01879

New Modalities: SAXS/WAXS

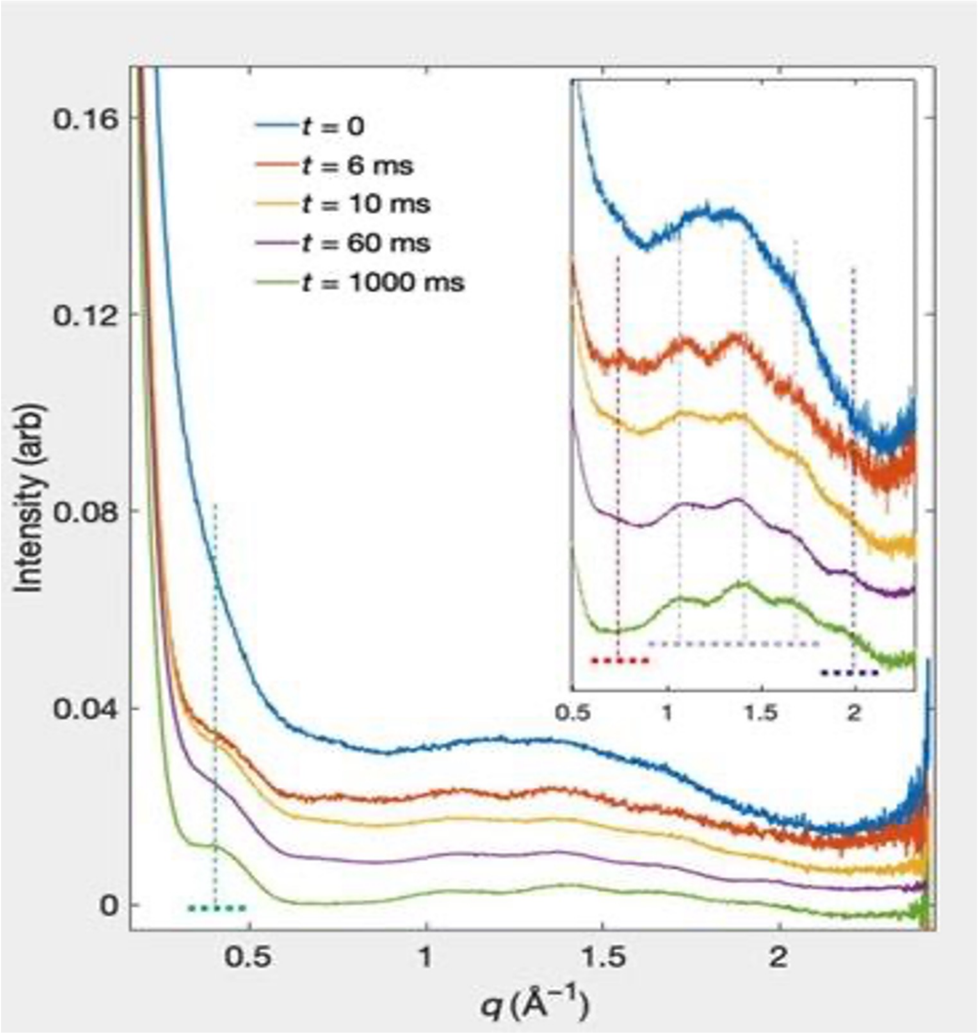

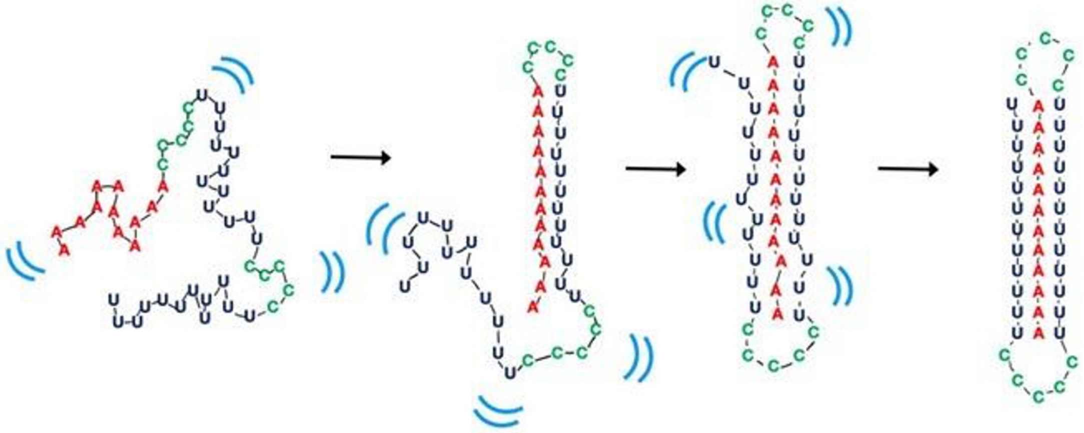

Zielinski, K. *et al.* (2023) RNA structures and dynamics with Å resolution revealed by x-ray free-electron lasers
***Sci. Adv.* 9, DOI: 10.1126/sciadv.adj3509**

New Modalities: X-ray absorption spectroscopy

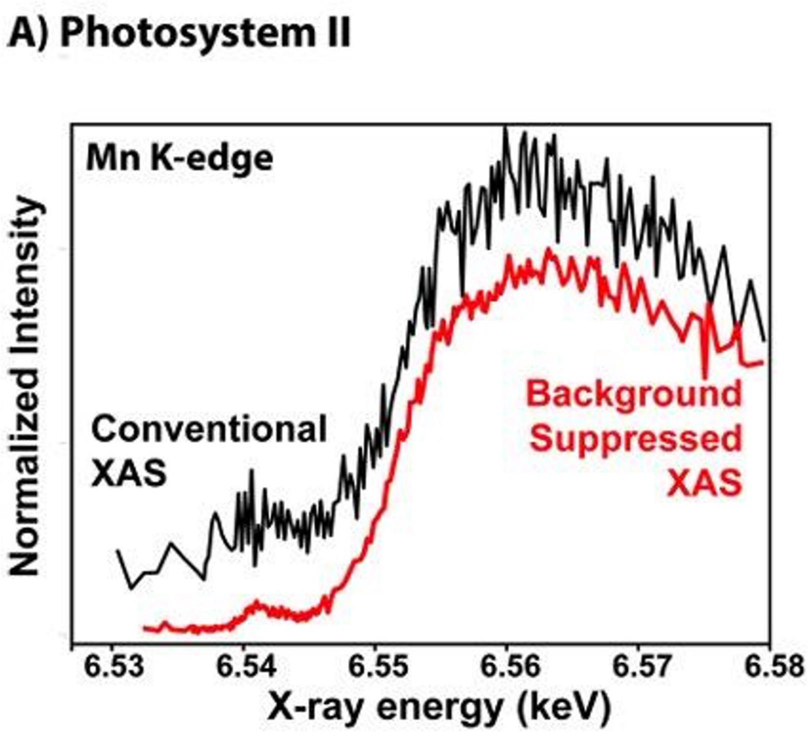

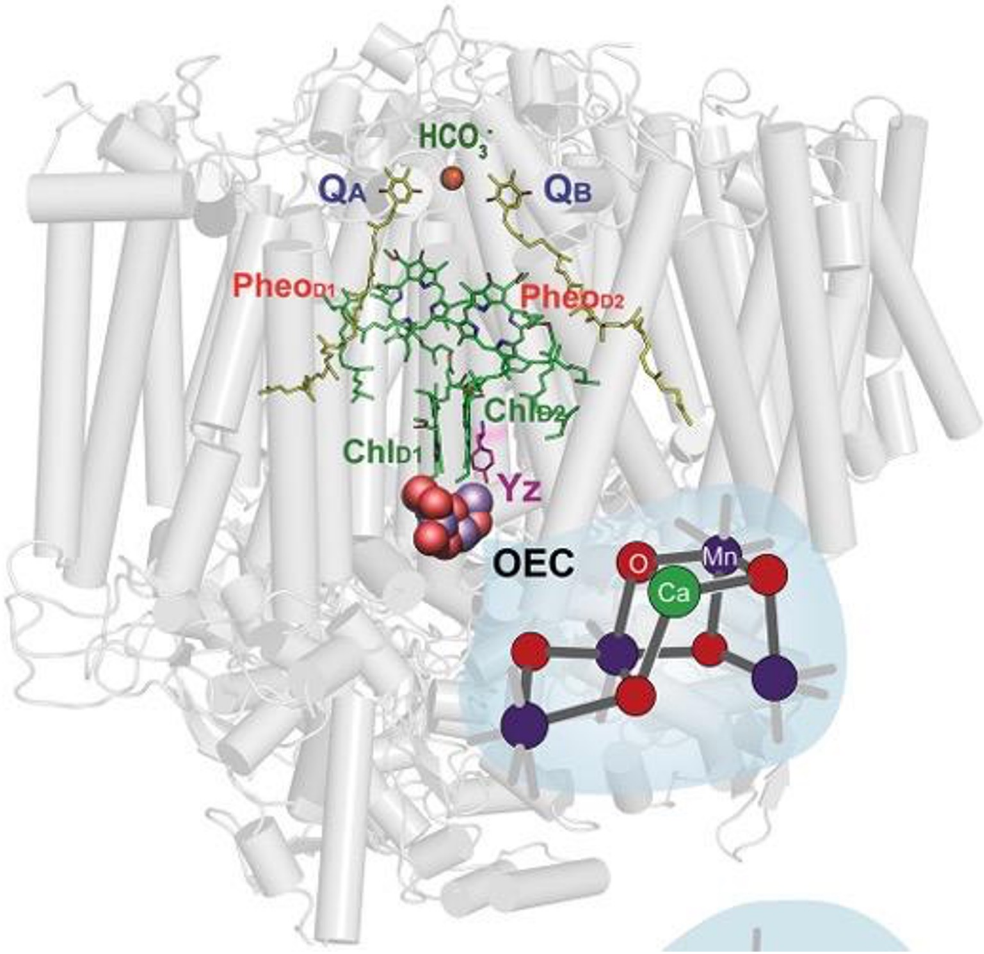

Bogacz, I. *et al.* (2025) X-ray Absorption Spectroscopy of Dilute Metalloenzymes at X-ray Free-Electron Lasers in a Shot-by-Shot Mode. J. Phys. Chem. Lett. DOI: 10.1021/acs.jpclett.5c00399

What experiments will benefit from kHz?

120 Hz, 2 mJ X-ray pulse
LCLS Copper LINAC
Bigger signal, bigger sample

2 kHz, 0.2 mJ X-ray pulse
LCLS-II-HE
Small signal, more repetitions

Better S/N with bigger samples

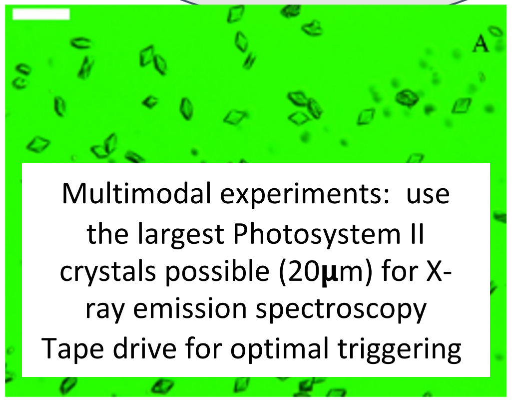

Multimodal experiments: use the largest Photosystem II crystals possible (20µm) for X-ray emission spectroscopy
Tape drive for optimal triggering

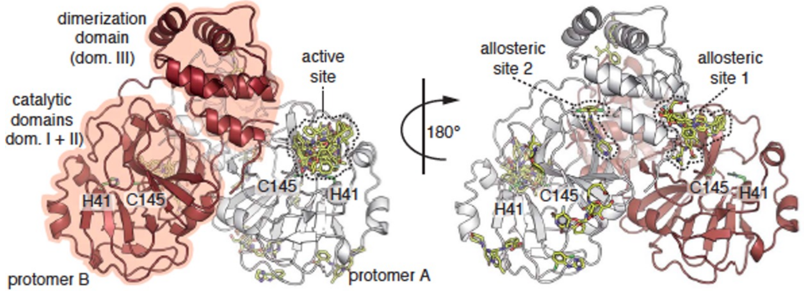

Drug repurposing screen (Günther, Science, 2021)
5953 compounds, 4 weeks at 3 synchrotron beamlines

Overcome heat dissipation challenges

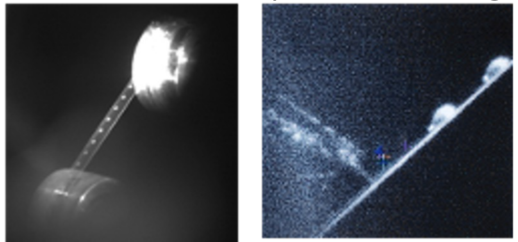

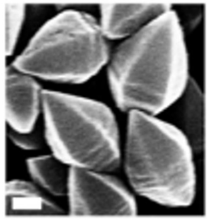

Obligatory submicron crystals

Systems with fundamentally small sample sizes

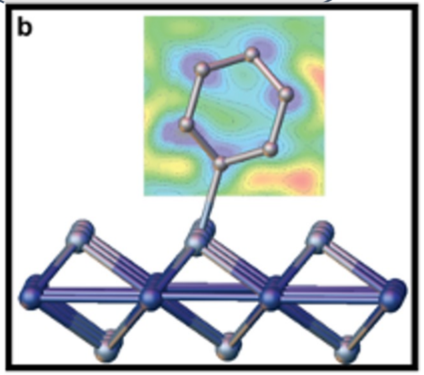

Serial chemical crystallography

Improving accuracy

Partial reflections & disperse beam:
apply complex computational models to deduce structure factors

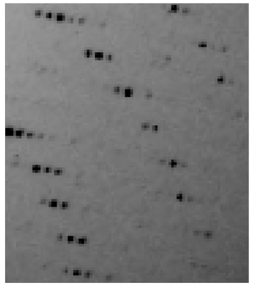

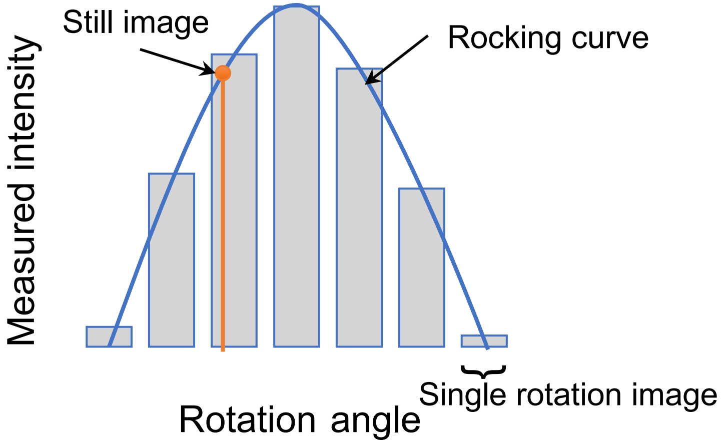

nanoBragg / diffBragg

Use diffraction physics to model one pixel at a time, one photon at a time

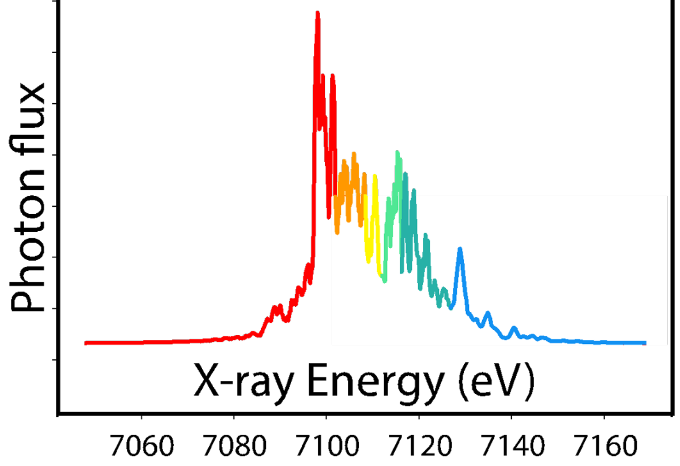

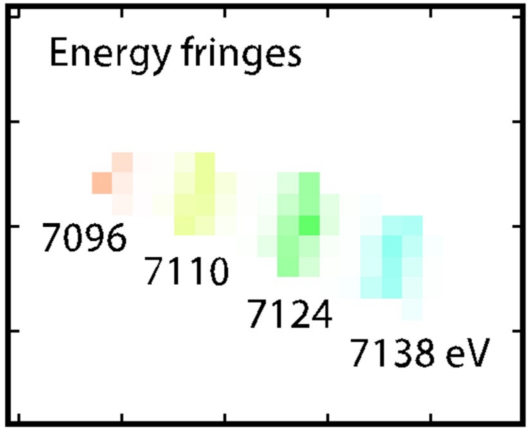

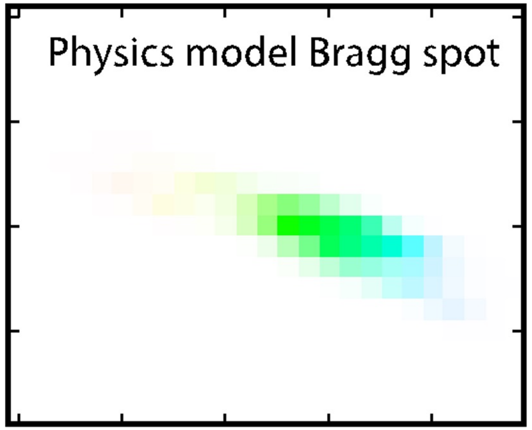

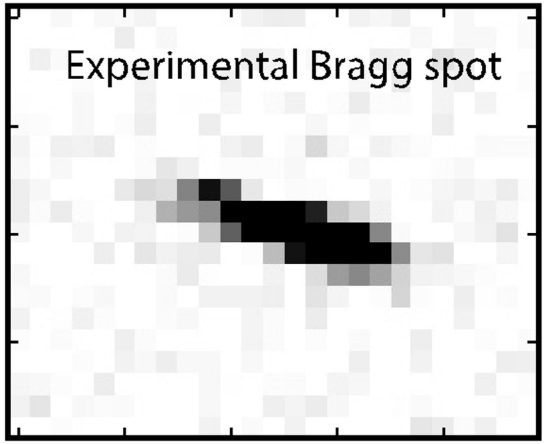

$$I = \sum_{\substack{\text{wavelength}\\ \text{channels}\\ \lambda}} \left(J_0(\lambda) \times |\mathrm{F}_{HKL}(\lambda)|^2 \times \sum_{\substack{\text{mosaic}\\ \text{rotations}}} \mathrm{F}^2_{\text{Lattice}}(\lambda) \right)$$

Pixel Intensity

Incident Spectrum

F_{HKL} = Structure factor

$F_{Lattice} = e^{-(\text{Bragg condition offset } (\theta,\ \lambda)\ /\ \text{reciprocal mosaic size})^2}$

Promises

- More accurate results for tiny biological differences
- Works great for simulated experiment
- Extends to new science, like metalloenzyme redox state

Challenges

- Older code written in C++/CUDA
- Might require larger resources: 64 x GPU nodes
- Real experiments degraded by numerous systematic effects

Systematic effects that potentially defeat our high-accuracy goals

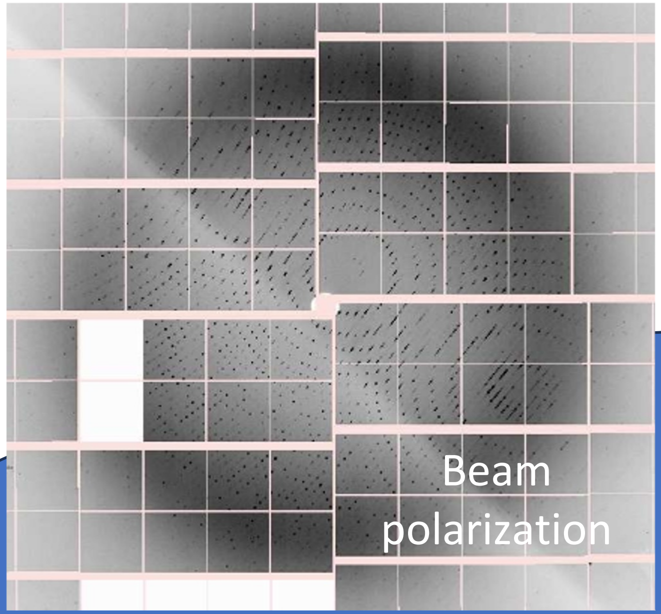

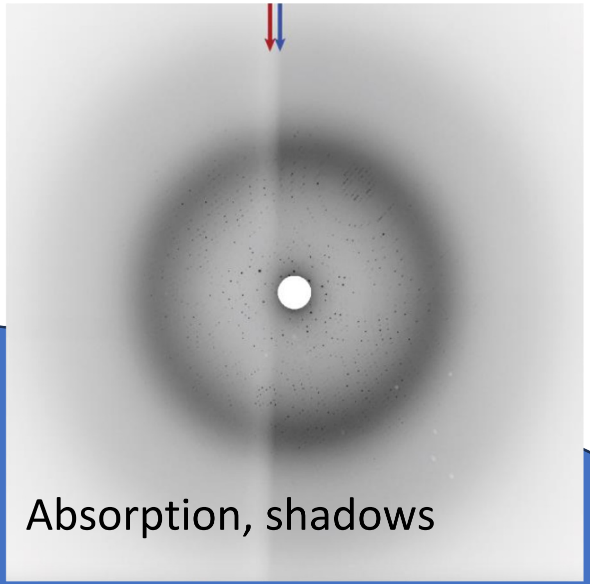

Parallax correction for thick detector

Mosaic crystal, model unknown

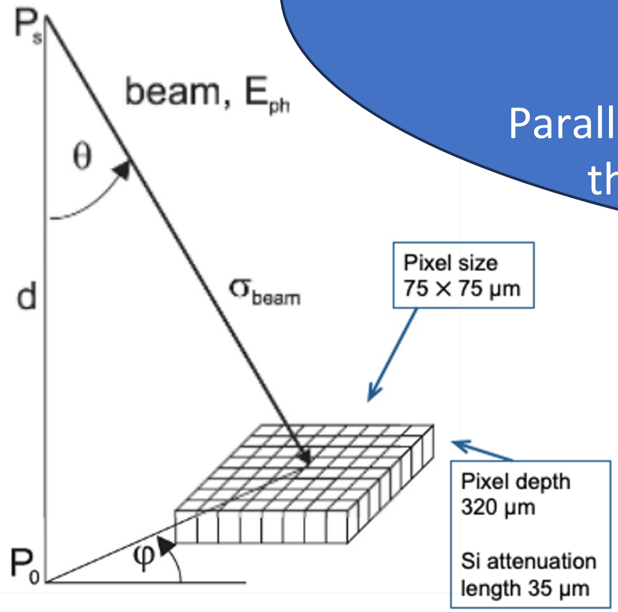

$$I_{pixel} = \sum_{E} \left(J_0(\lambda) \times |F_{HKL}(\lambda)|^2 \times \sum_{domains} F^2_{Lattice}(\lambda) \right)$$

Incident spectrum — Structure factor — Ewald offset, Mosaic domain size

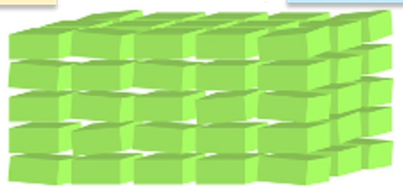

Careless uses variational inference to scale Bragg spots

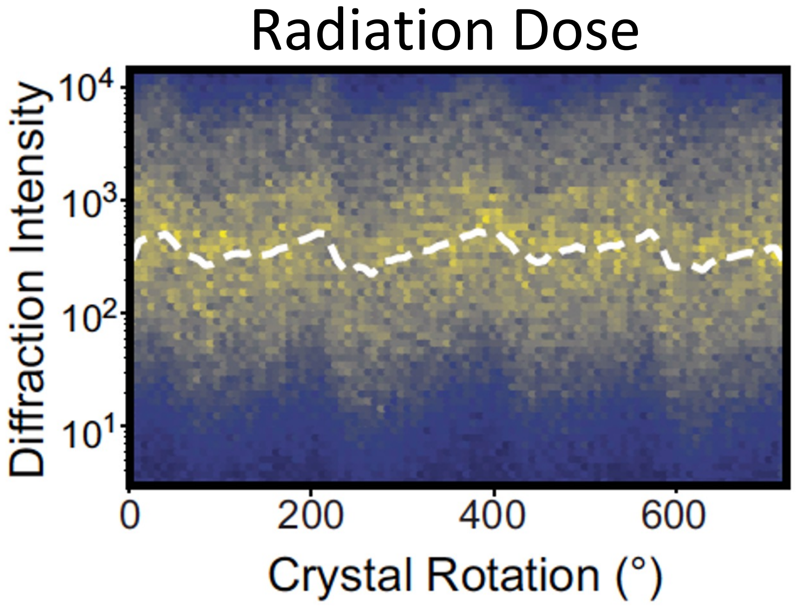

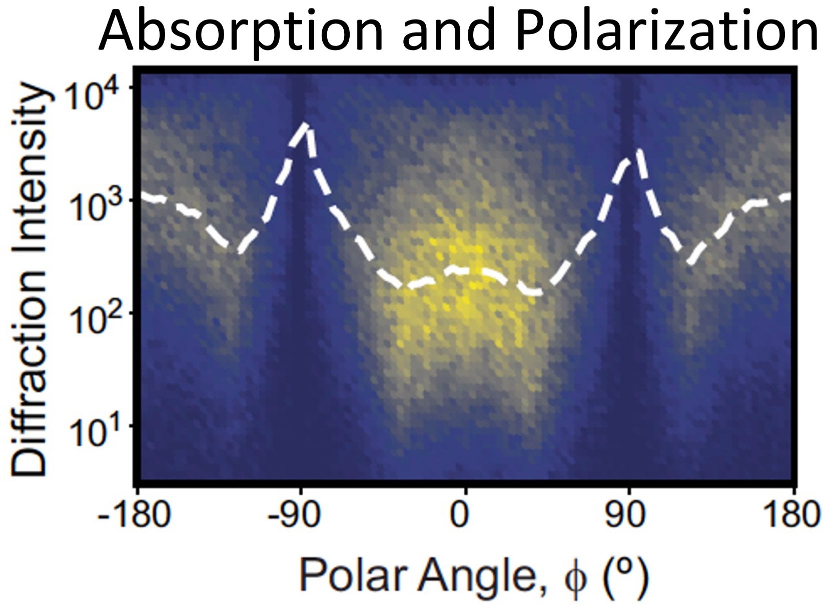

Promises

- NN learns how to scale data together
- Correctly model systematic effects with no physics

Challenges

- Not scalable in memory, only uses one GPU node
- Relies on stochastic training to loop through $>10^4$ images
- No integration trials especially at 2 kHz scale

Dalton et al. (2022) Nature Communications. 13, 7764

Is the experimental feedback fast enough?

Jan Kern:
For a 5-shift Photosystem experiment (60 hrs @ 20 Hz) we brought 15 mL sample.
Making 1 mL takes 1 week for 1.5 people, so 22 weeks FTE for 5 shifts.

Aaron Brewster:
Oh, my word, that's 42 years of sample prep for 5 shifts at 2 kHz.

Nick Sauter:
A 10 minute data collection will burn through 40 days of sample prep.
In the 2kHz regime we are justified in asking for feedback with 1 second turnaround.

Experimental cycle

Data transfer does not occur till the DAQ writes to hard disk

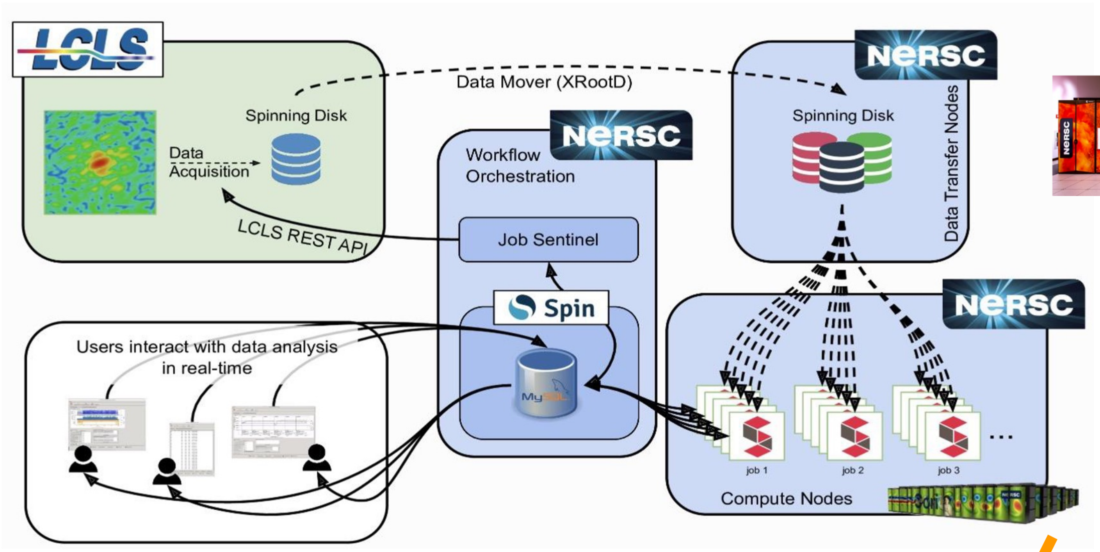

ESnet

Data acquisition
5 minute run

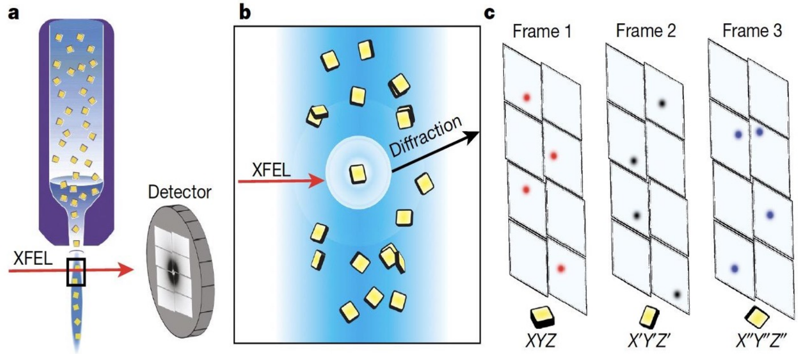

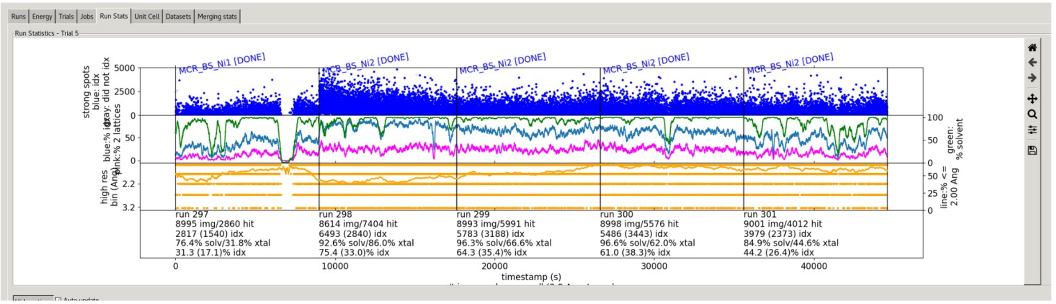

Immediate:
Data quality metrics

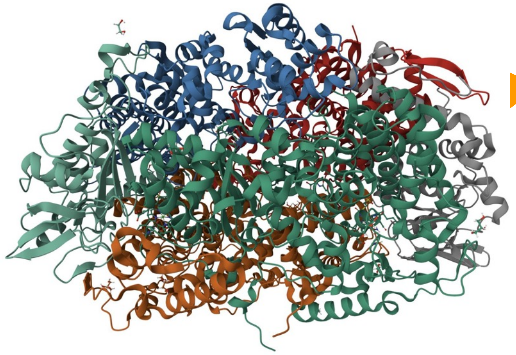

After 5 minutes:
Final science
(Fourier coefficients)

Experimental cycle

Goal: 1 second

Bypass writing the data to file

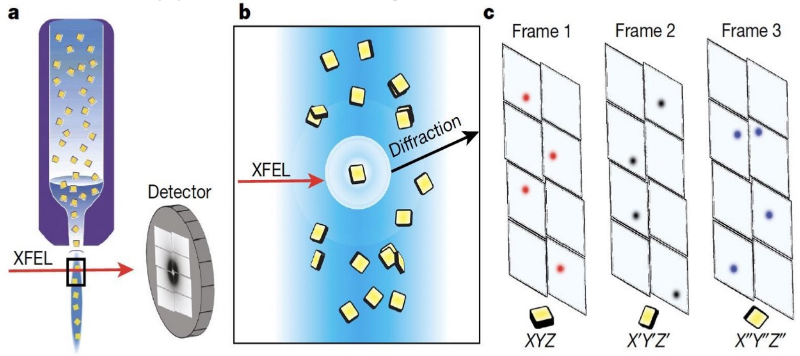

Immediate data quality metrics

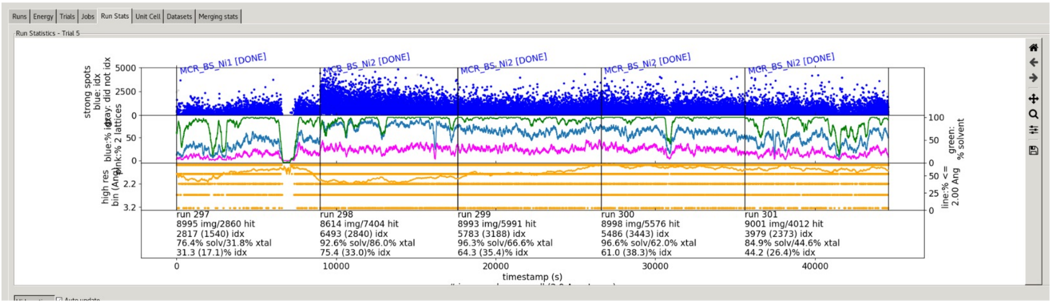

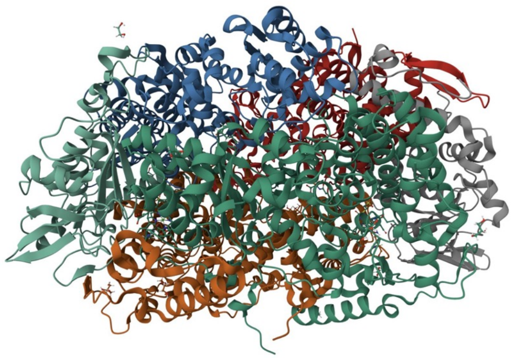

Still possible delay to final science (Fourier coefficients)

Streaming with *dials.stills_process*

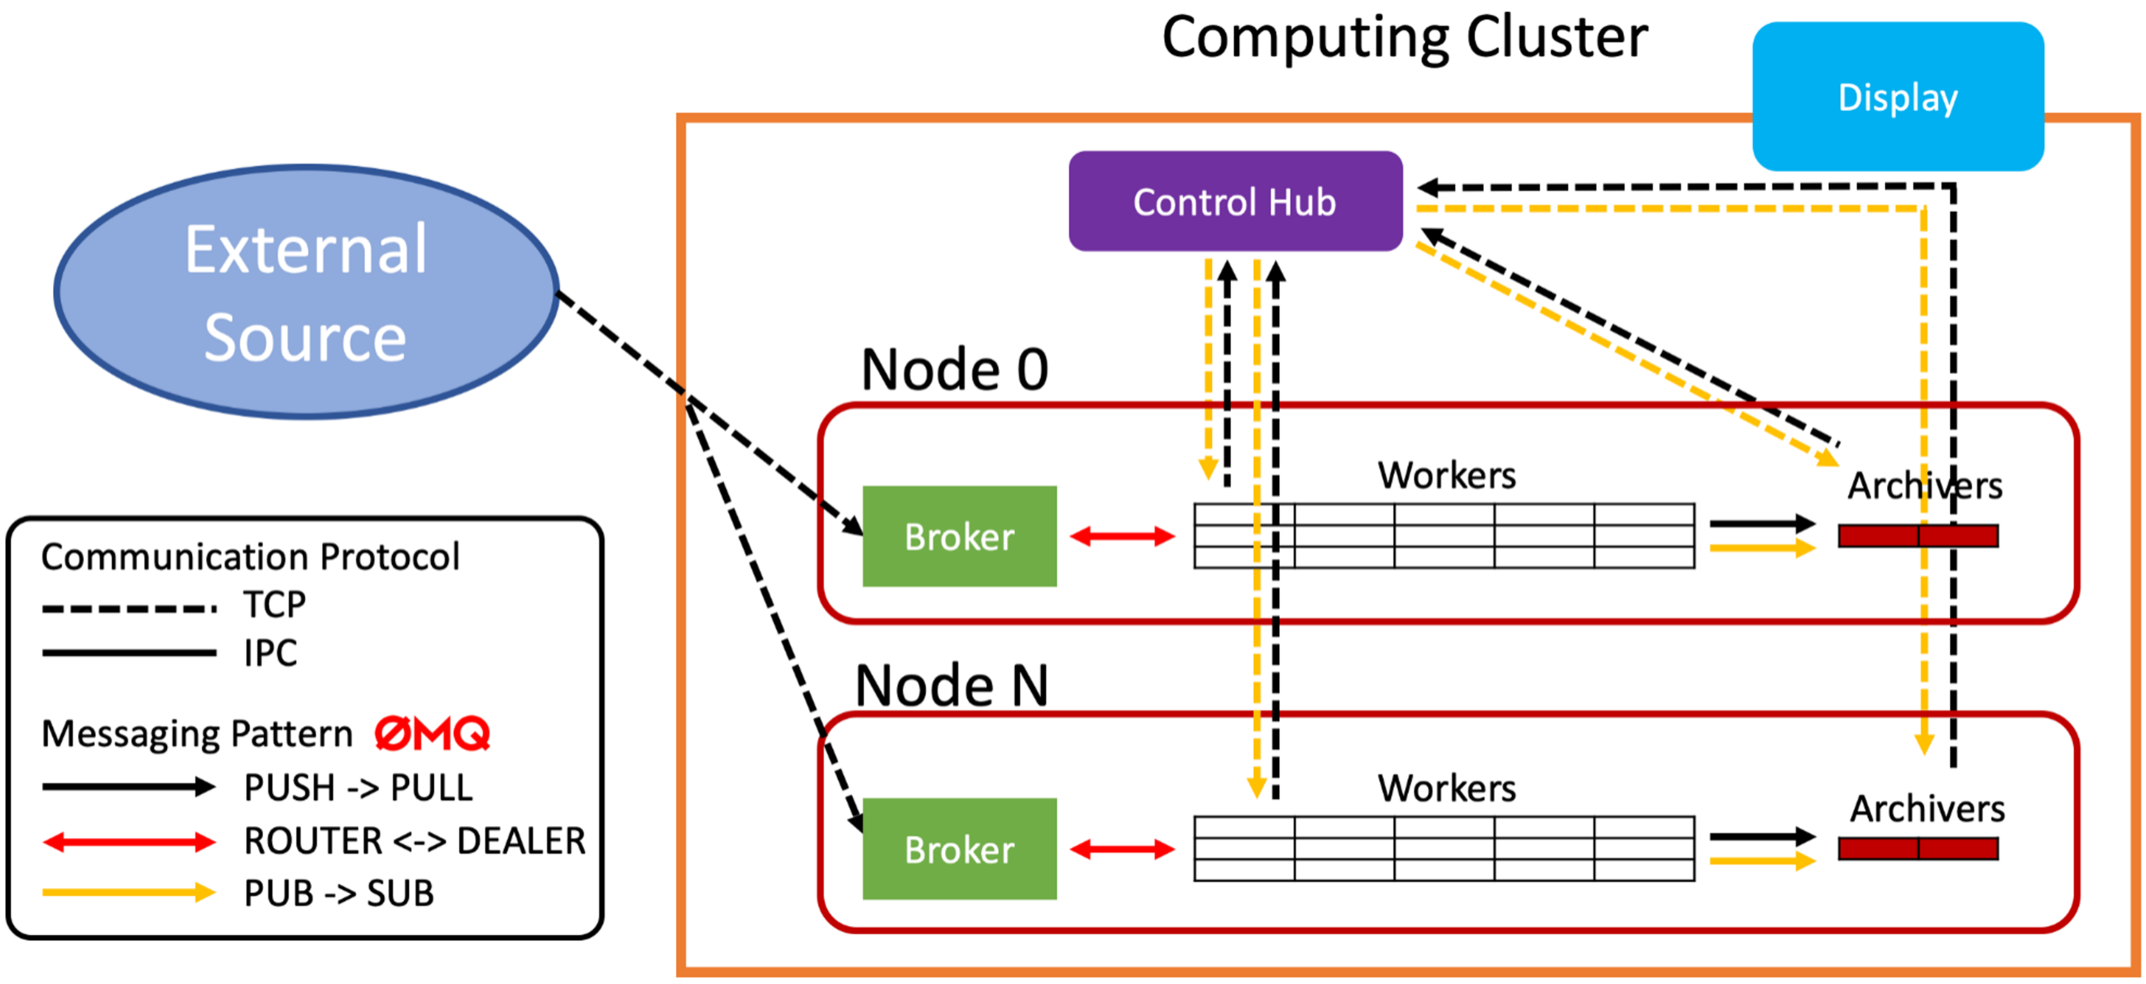

Processing at Scale

Five runs at 100 Hz

4 CPU nodes x 128 ranks

for a total of 512 ranks

Histogram of per-image latency

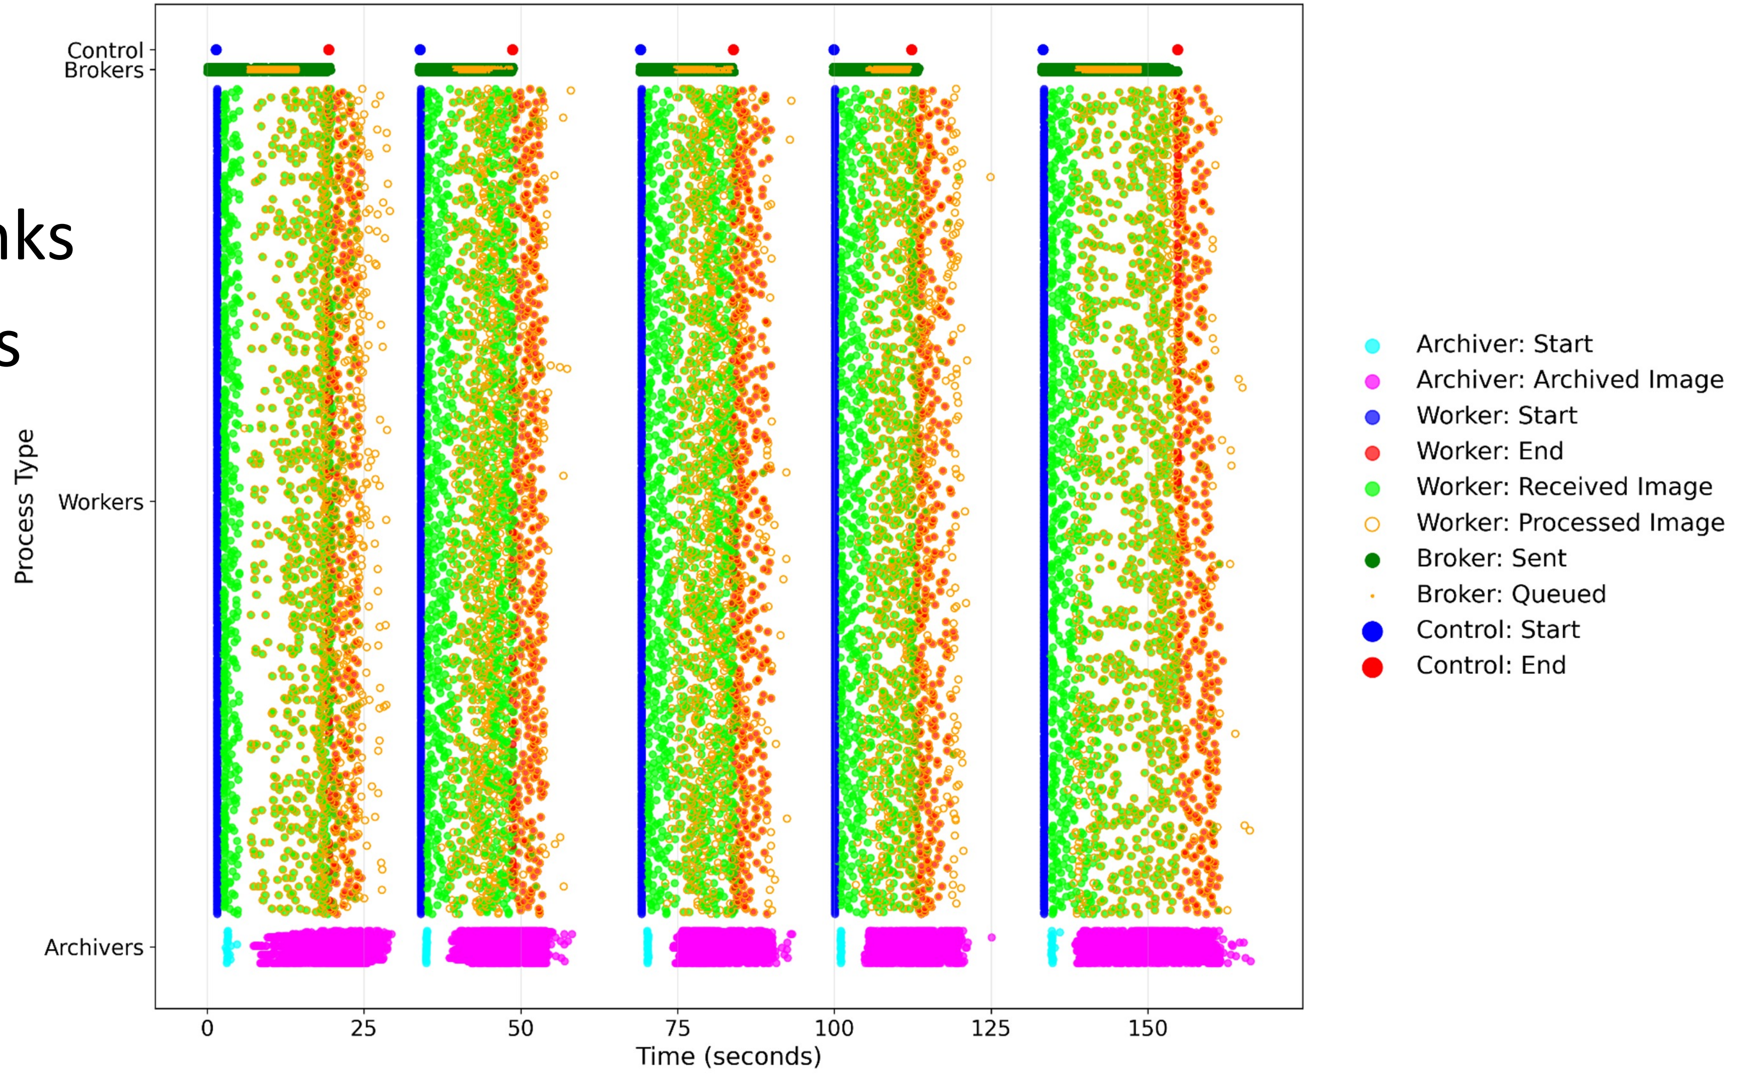

Resonet GPU spot analysis

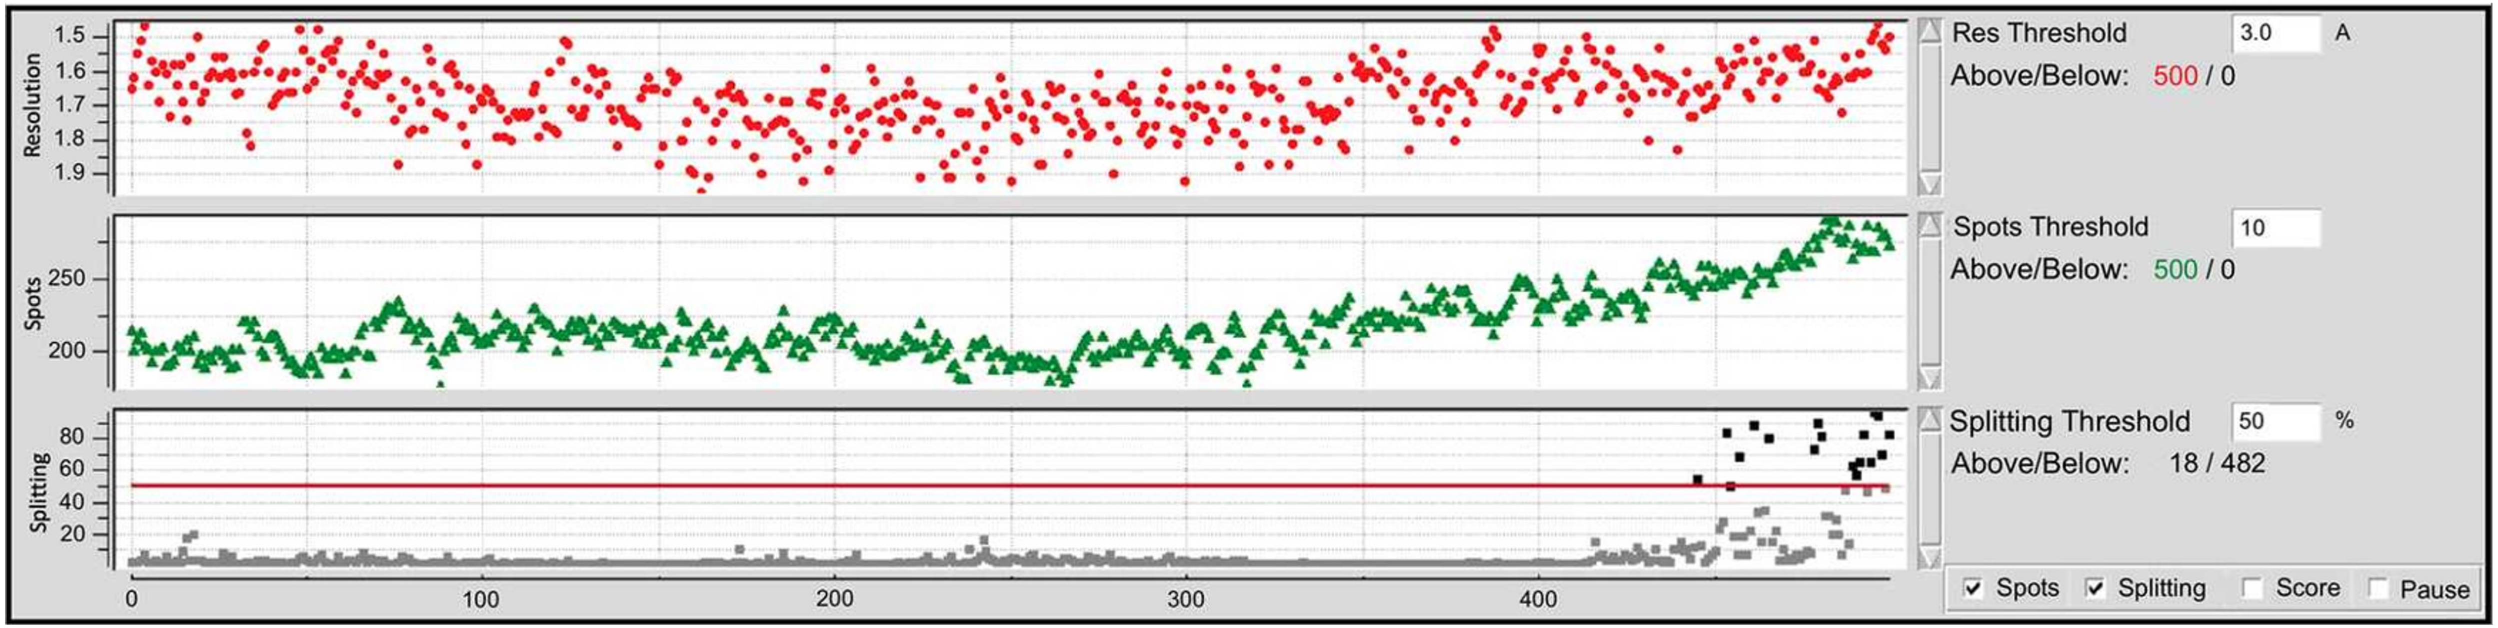

Promises

- GPU image quality metrics with NN
- Determine crystal resolution per pattern
- Identify overlapping lattices
- Suitable for 1-second decision making
- For rotation: helps check rad damage, crystal miscentering; asymmetric diffraction
- Network training is easy with simulated data

Challenges

- No integration trials especially at kHz scale
- Can we convince user it is sufficient for event rejection?
- Might have to retrain network for specific beamlines

Implemented last year by Derek Mendez, SSRL, Acta D. (2024). 80, 26–43

Toro Indexer

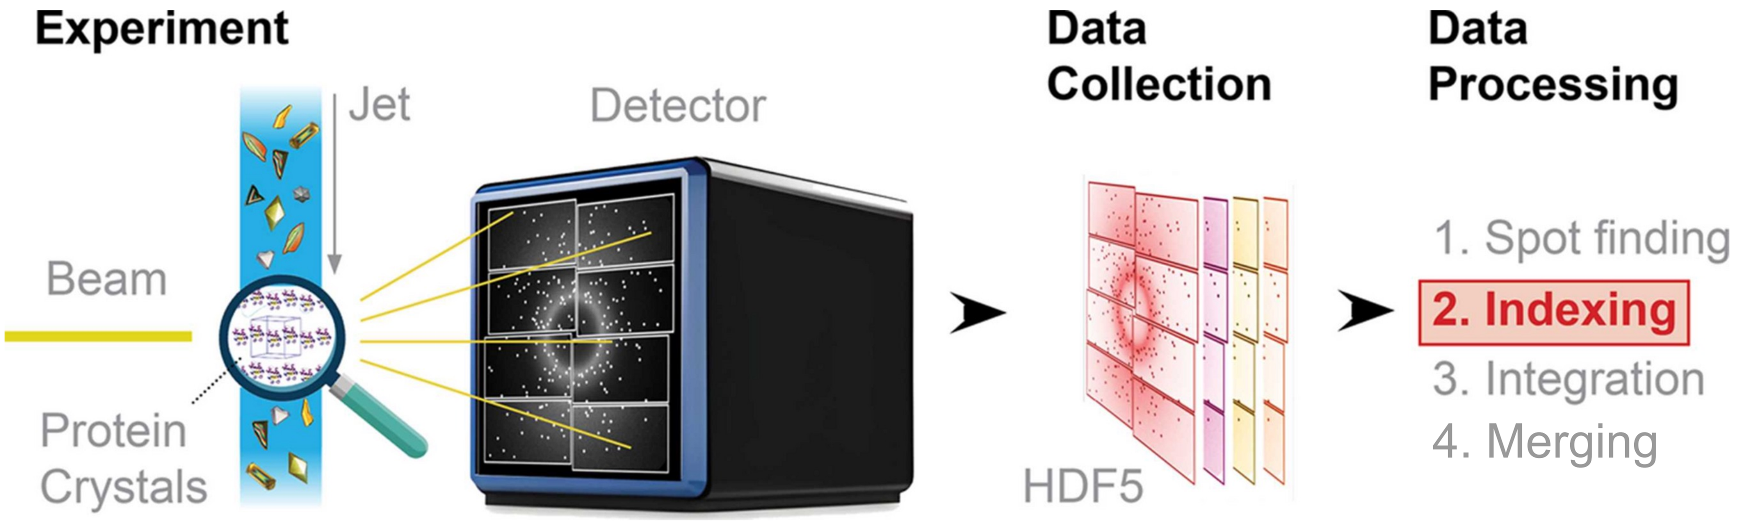

Promises

- Can index images at 1 kHz per A100-GPU node
- Result in fixed time, no sagas
- Suitable for 1-second decision making
- Sample-steering
- Feeds back a final-science metric, not just # spots
- Could be edge computing, could be streaming
- PyTorch so no direct involvement with C++/CUDA
- Robust methods instead of heuristic outlier rejection

Challenges

- Only works if you have GPU spotfinder
- Works in parallel only if you pack 1,000 images
- No integration trials especially at kHz scale
- Can we convince user it is sufficient for event rejection?
- Tunable-might be slower for large unit cells
- Not a general indexer. Only works if you know the cell

Implemented last year by SwissFEL team, J. Appl. Cryst. (2024). 57, 931–944

ASAP::O & CrystFEL

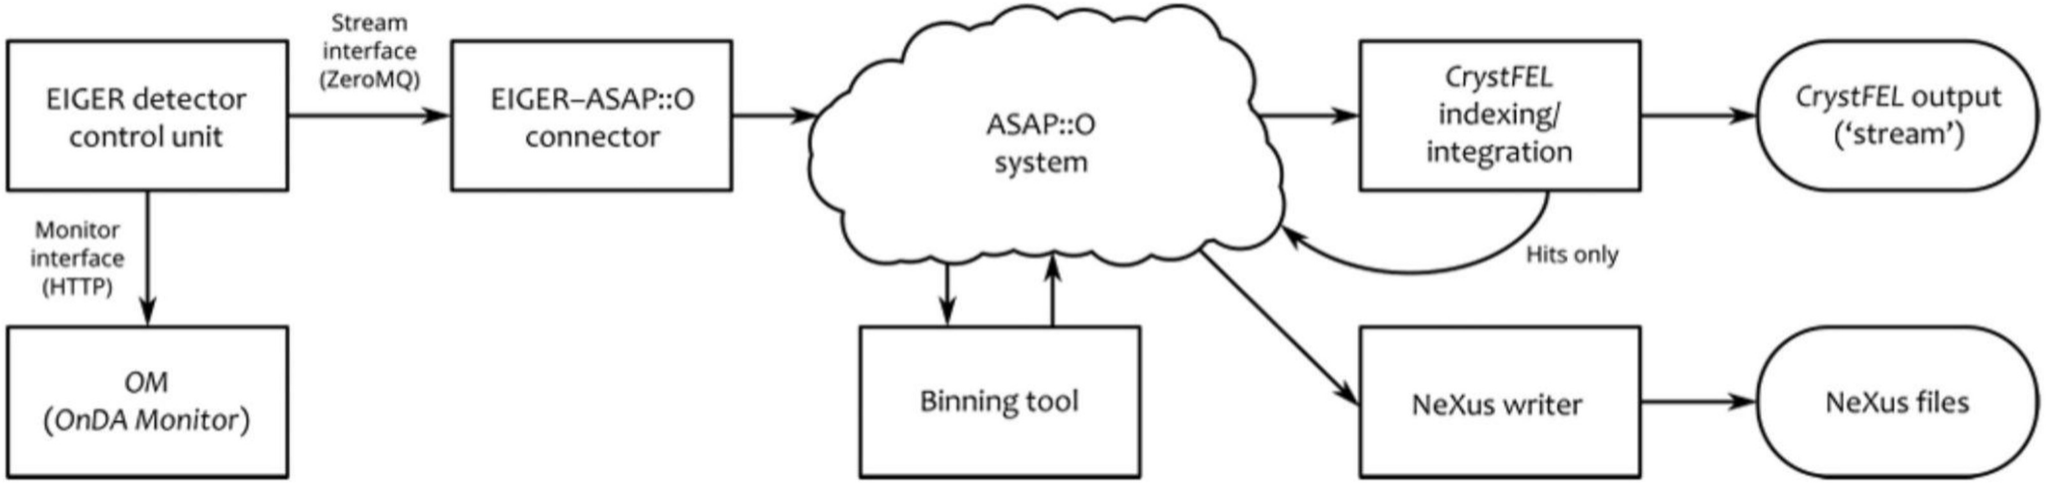

Promises

- Entire CrystFEL pipeline optimized, not just one step
- Process 16 Mpx EIGER2 X at 133 Hz on one CPU node
- Latent time for one image is about 500 ms
- Elimination of the I/O bottleneck
- Dynamic scaling of computing resources.

Challenges

- Sufficient computing resources must be available during experimentation. [Would require 15 CPU nodes at 2kHz]
- Need a robust data streaming platform (ASAP::O).
- Performant crystallographic software.
- Adequate and properly configured networking infrastructure.

White, T. *et al*. IUCrJ. (2025). **12**, 97-108

Failover to backup facility

We want: 55MB compressed image × 2kHz = 820 Gb/s ($66 \mu s$/evt)

ESnet offers: 800 Gb/s (according to Google)

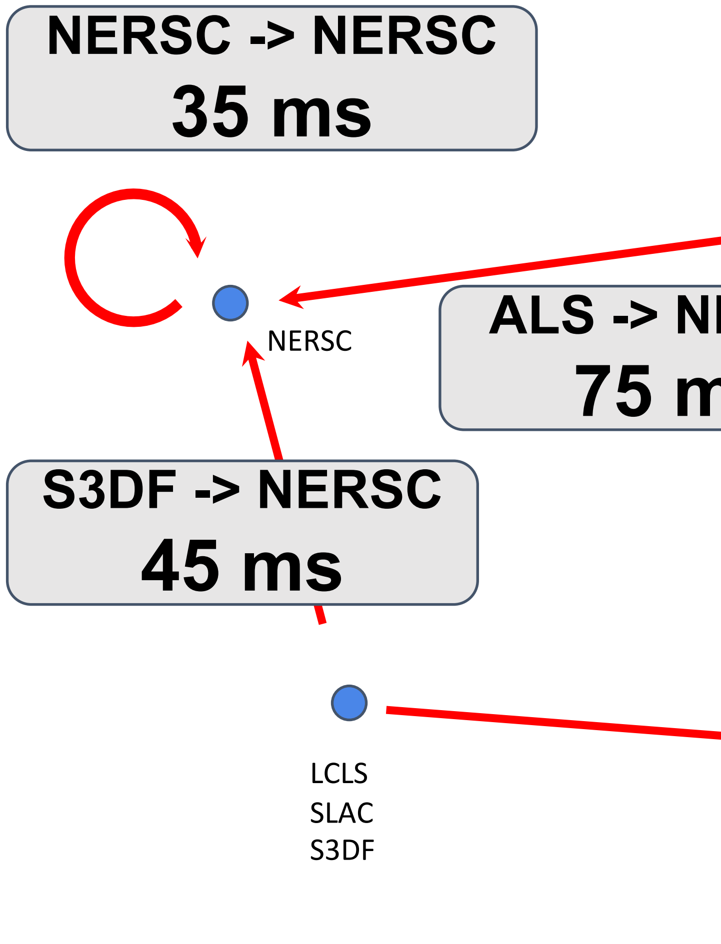

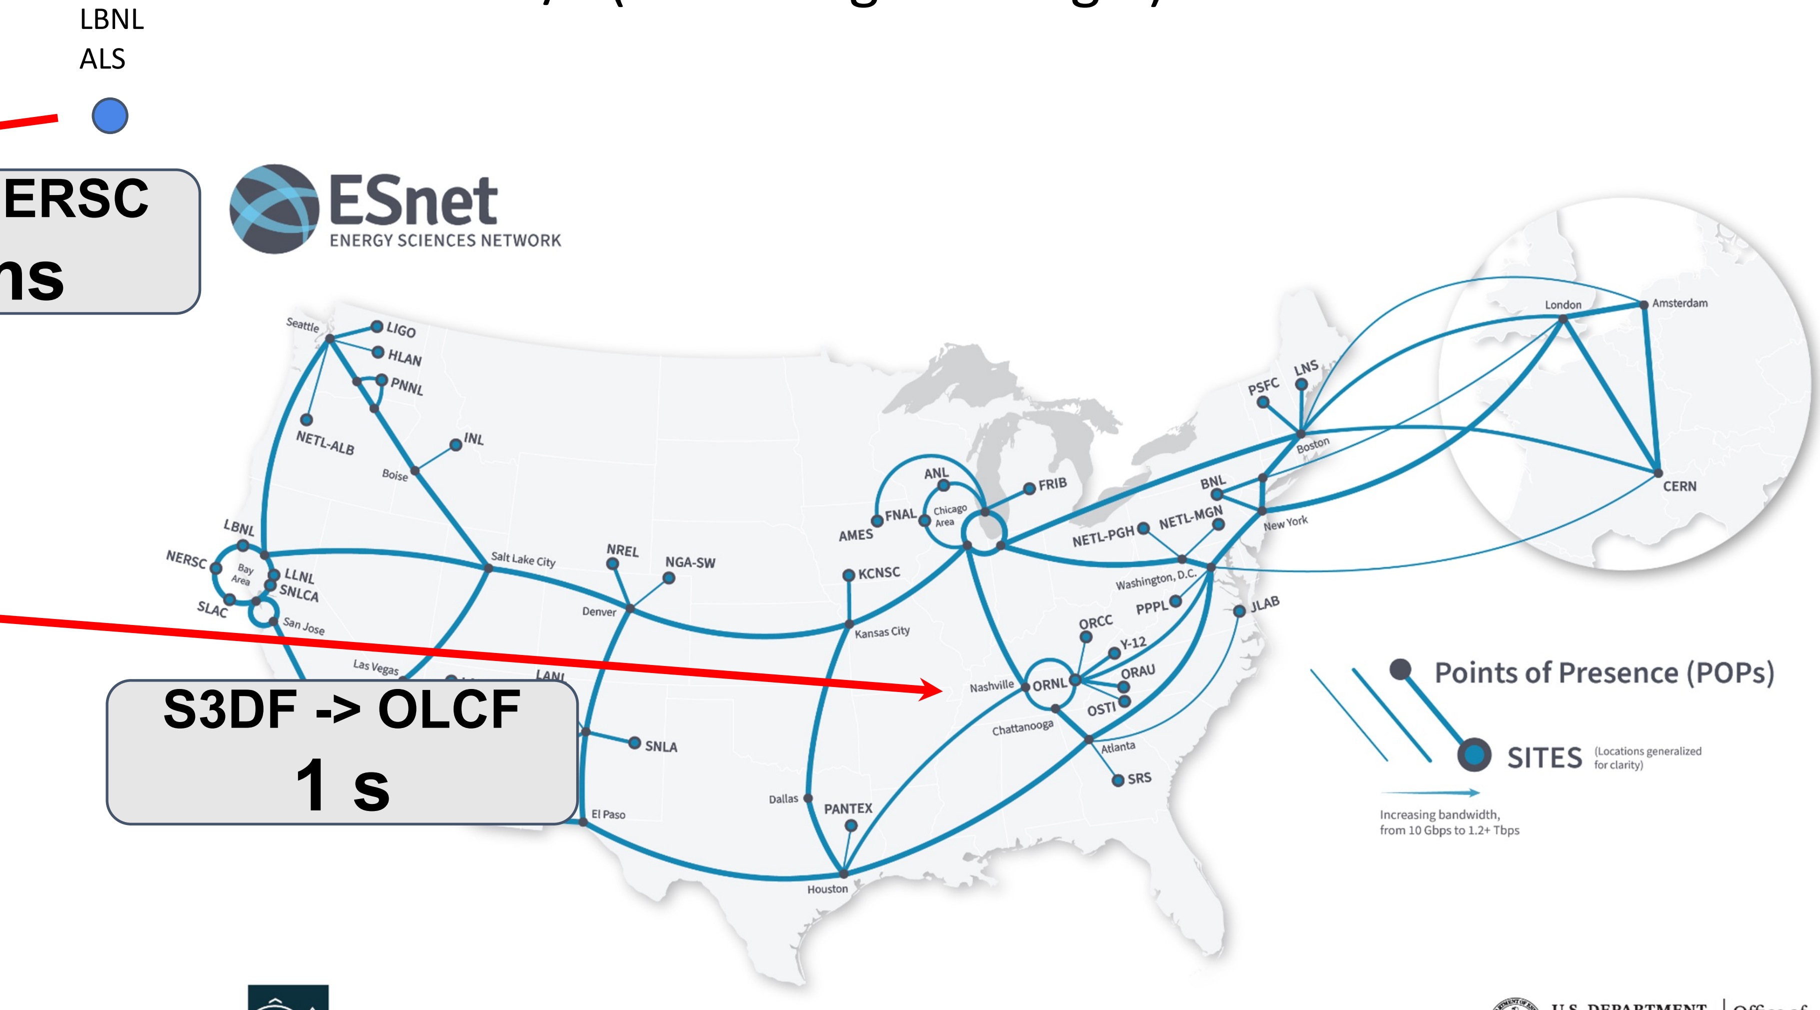

Is data processing portable?

Run	time	sample (batch #, volume, buffer, time loaded, concentra	run TAG	SPREAD conditions
168			PSII_b8_0F	
		reached next batch		
169			PSII_b9_0F	[6521,6571,10 eV / 10s], k
170			PSII_b9_0F	[6521,6571,10 eV / 10s], k
171			PSII_b9_0F	[6521,6571,10 eV / 10s], k
172			PSII_b9_0F	[6521,6571,10 eV / 20s], k
173			PSII_b9_0F	[6521,6571,10 eV / 20s], k
174			PSII_b9_0F	[6521,6571,10 eV / 20s], k
175				water run
176			PSII_b9_0F	[6520,6571,1 eV / 1s], v
177			PSII_b9_0F	[6520,6571,1 eV / 1s], v
178			PSII_b9_0F	[6520,6571,1 eV / 1s], v
179			PSII_b9_0F	[6520,6571,1 eV / 1s], v
180			PSII_b9_0F	[6520,6571,1 eV / 1s], v
	2:00	end of syringe - loaded batch 10		
181-184		not sample		
185			PSII_b9_0F	[6520,6571,1 eV / 1s], v
186			PSII_b9_0F	[6520,6571,1 eV / 1s], v
187			PSII_b9_0F	[6520,6571,1 eV / 1s], v
188			PSII_b9_0F	[6520,6571,1 eV / 1s], v
189			PSII_b9_0F	[6520,6571,1 eV / 1s], v
190			PSII_b9_0F	[6520,6571,1 eV / 1s], v
191			PSII_b9_0F	[6520,6571,1 eV / 1s], v
		end of batch 9 start of batch 10		
192			PSII_b10_1F	[6520,6571,1 eV / 1s], v
193			PSII_b10_1F	[6520,6571,1 eV / 1s], v
194			PSII_b10_1F	[6520,6571,1 eV / 1s], v
195			PSII_b10_1F	[6520,6571,1 eV / 1s], v
196			PSII_b10_1F	[6520,6571,1 eV / 1s], v
197			PSII_b10_0F	[6520,6571,1 eV / 1s], v

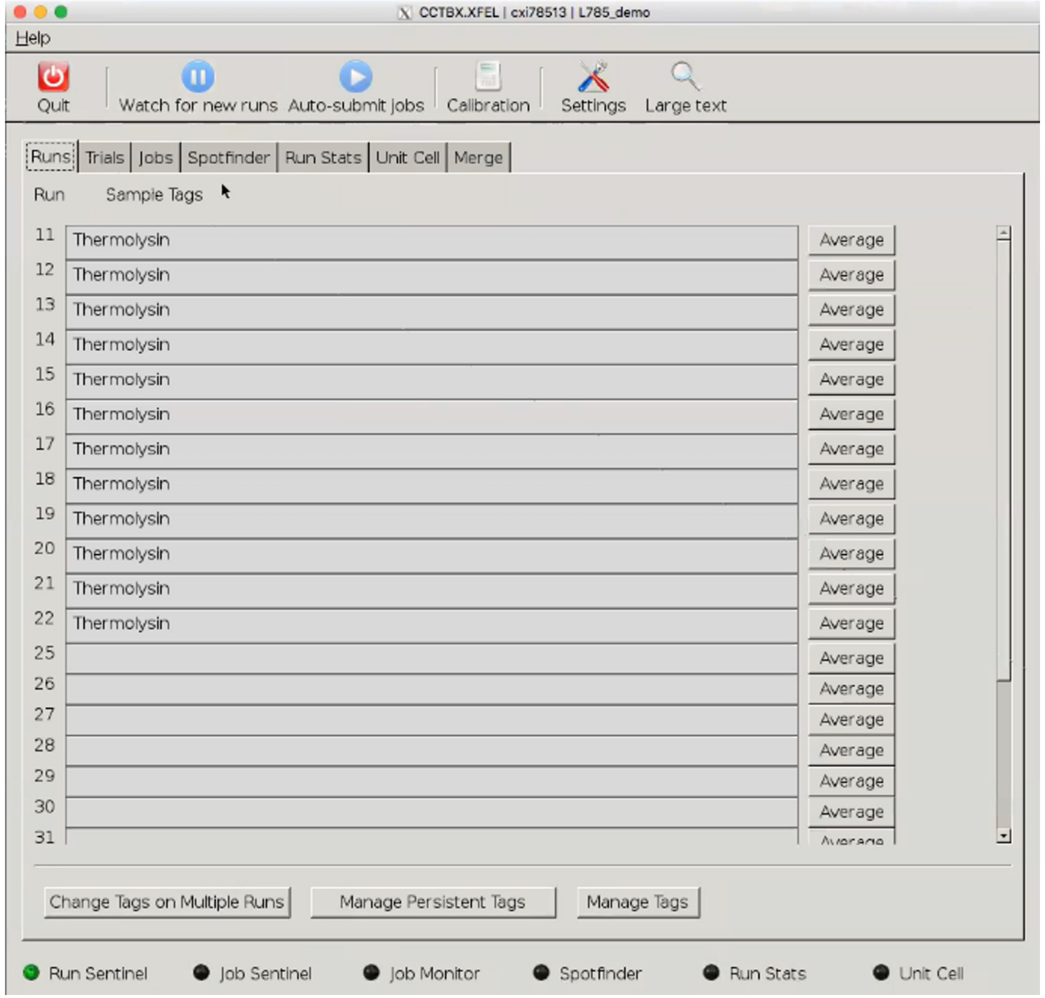

Pros
Google sheet identifies all our runs: protein type, batch number, triggering conditions
MySQL to organizes data reduction & science metrics; with auto-detection of new runs

Cons
Information transferred to the GUI by hand!
Switching to NERSC required doing this all over again
Currently <1000 runs each beamtime. May not scale to 2kHz
High level metadata not portable, doesn't conform to any standard

Resource pointers should be portable too!

```
> more redoscamerge/42937958/PSII_SPREAD_1F_spread_000428.expt
```



```
{
  "__id__": "ExperimentList",
  "experiment": [
    {
      "__id__": "Experiment",
      "identifier": "9653060b-65fd-694a-991d-8378f12bf4dc",
      "beam": 0,
      "detector": 0,
      "crystal": 0,
      "imageset": 0
    },
  "imageset": [
    {
      "__id__": "ImageSet",
      "images": [
        "/pscratch/sd/p/psdatmgr/psdm/mfx/mfx101080424/hdf5/cctbx/mfx101080424_Run0478_nexus_281_285.h5"
      ],
      "single_file_indices": [
        2126
      ],
      "mask": "/global/cfs/cdirs/lcls/mfx101080424/masks/reshuffled_reinspected.mask",
      "gain": null,
      "pedestal": null,
```

XFEL data are multimodal

- Sample metadata
- Sample delivery videos
- FEE spectrometer
- Elastic spectrometer
- XES spectra
- Diffraction images
- Results & conclusions

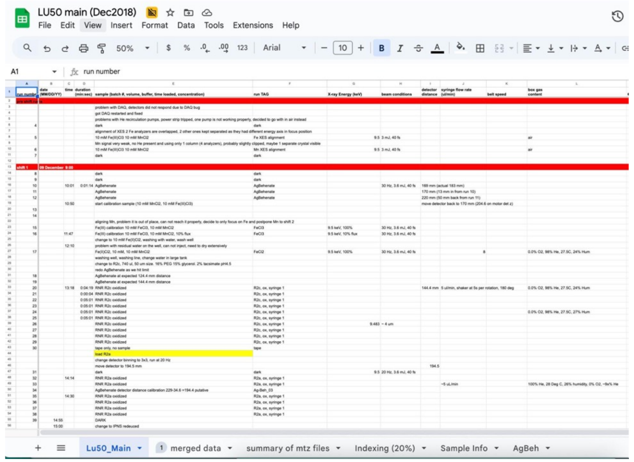

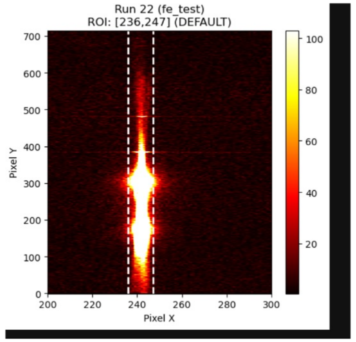

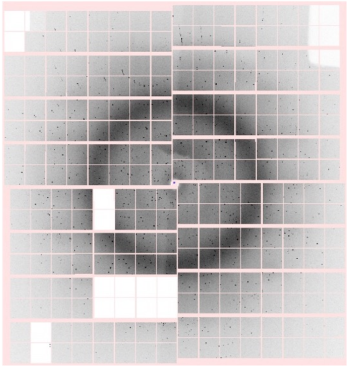

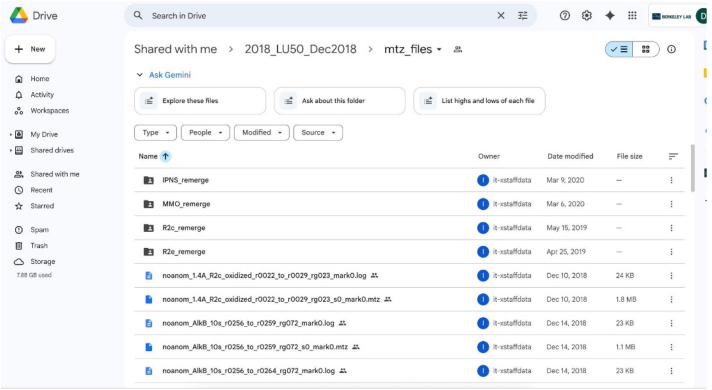

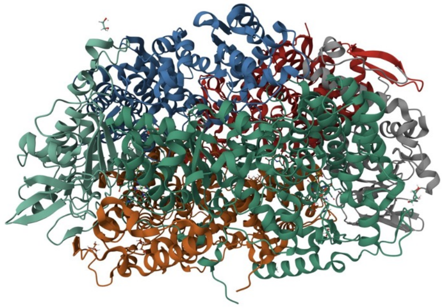

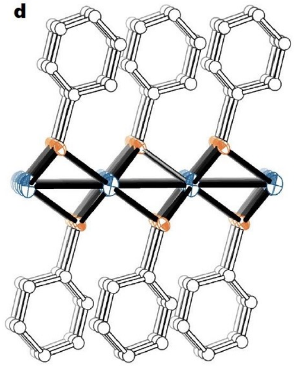

Autonomous experimentation

Streaming data processing

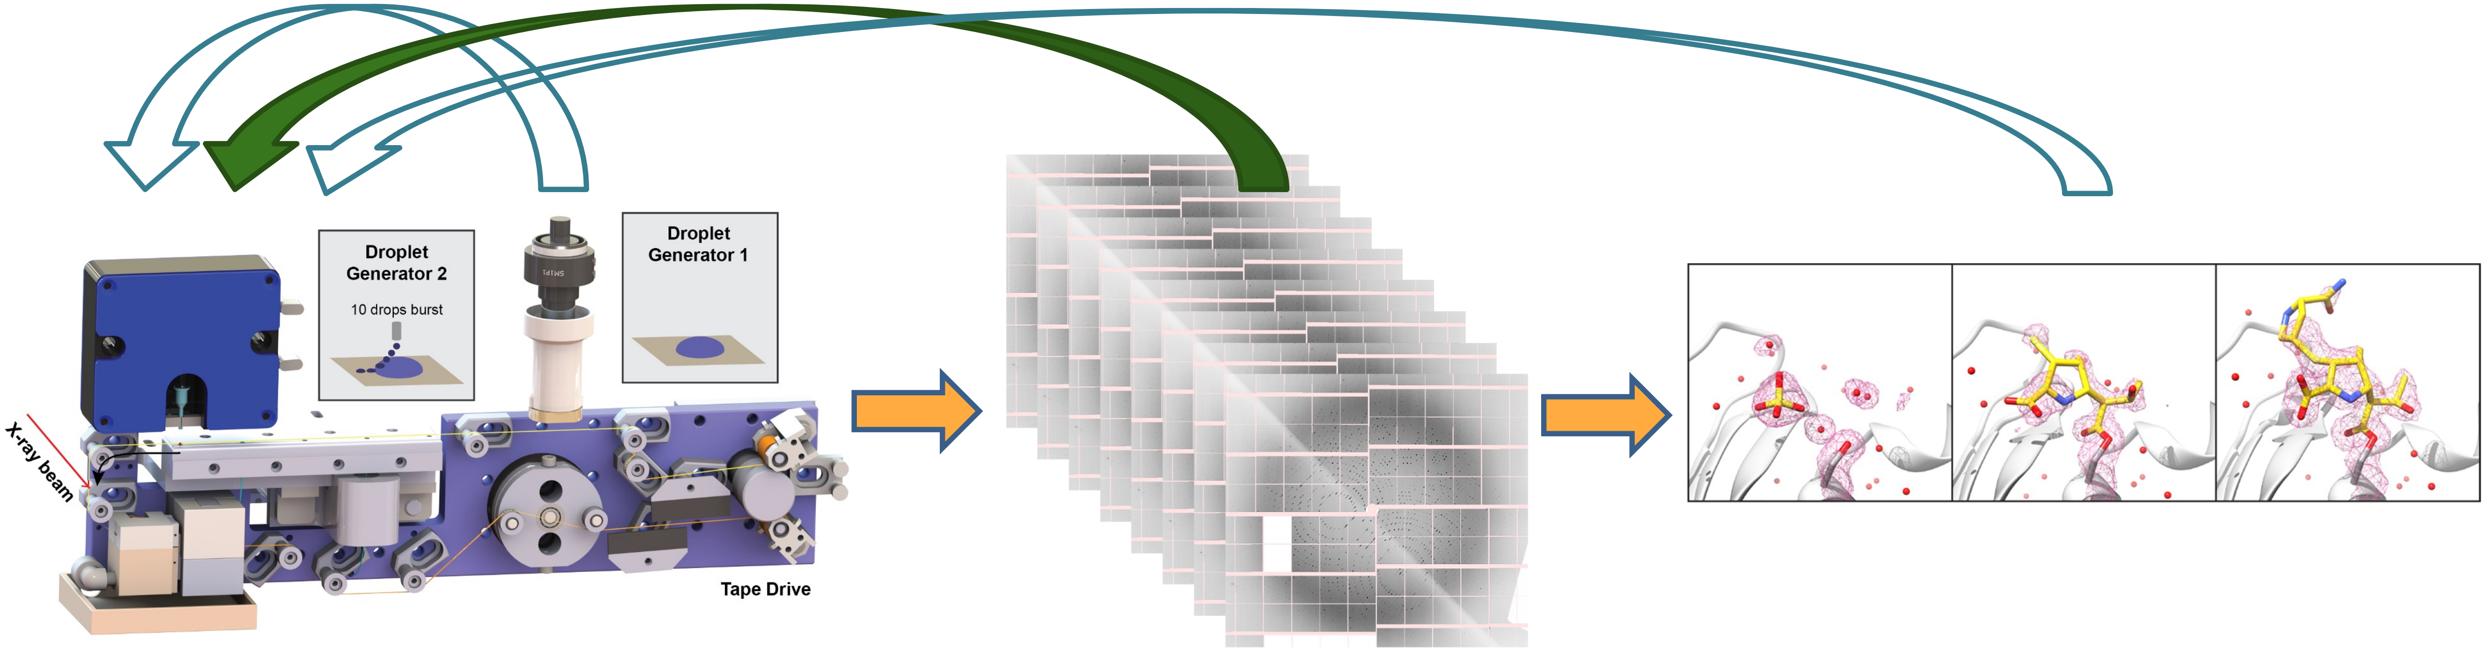

Many calibrations with various levels of reliability

Incident X-ray spectrum, energy scale

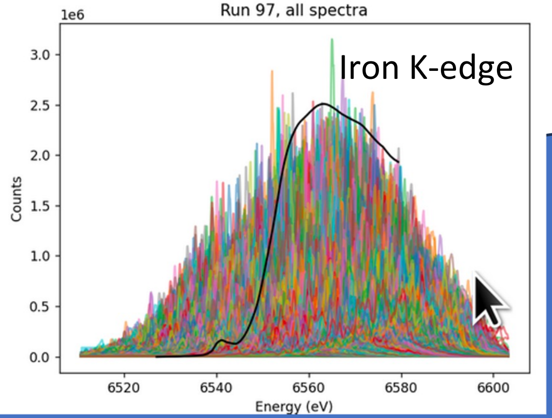

Detector distance

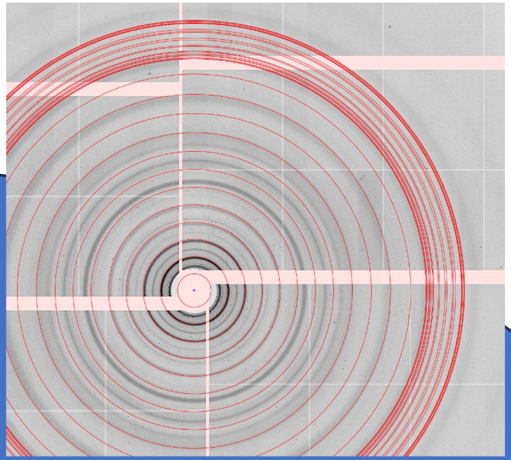

Untrusted pixel mask

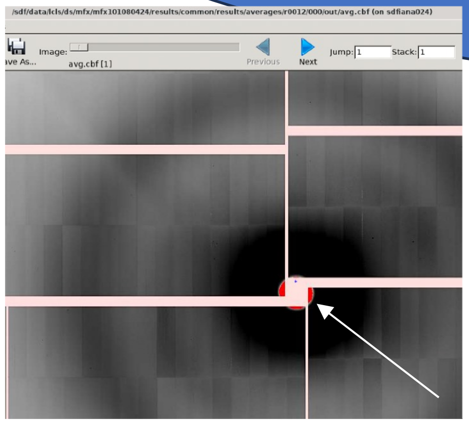

Image data

Pixel gain mode

Patchiness is subtle but visible due to offsets

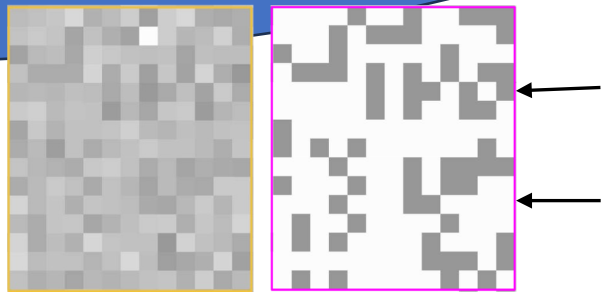

Detector gain switching, inconsistent mode offset

Things we would like, please

What we don't have now for serial crystallography at 2 kHz

- Clearer understanding of the goals, and what experiments benefit
  - Discussion of sample prep and sample delivery
- Better algorithms for better accuracy
- Goal of 1 second feedback to the experiment
  - Data streaming
  - Reliance on GPU and/or ML
- Failover to backup facility
- Database portability, resource portability (like file pointers)
- Better standardization
  - Have a rich metadata layer describing the sample (run,time point, chemical species, conditions, trigger detail)
  - GUI doesn’t scale to 2kHz
- Automate more steps

Things we would like, please

What we don't have now for serial crystallography at 2 kHz

- Clearer understanding of the goals, and what experiments benefit
    - Discussion of sample prep and sample delivery
- Better algorithms for better accuracy
- Goal of 1 second feedback to the experiment
    - Data streaming
    - Reliance on GPU and/or ML
- Failover to backup facility
- Database portability, resource portability (like file pointers)
- Better standardization
    - Have a rich metadata layer describing the sample (run,time point, chemical species, conditions, trigger detail)
    - GUI doesn’t scale to 2kHz
- Automate more steps

Everything here would help us at 120 Hz too!

Acknowledgements

Berkeley Lab Biosciences

Aaron Brewster
Dan Paley
Jan Kern
Hiroki Makita
Philipp Simon
Vittal Yachandra
Junko Yano
Asmit Bhowmick
Natalie Minnetian

LCLS

Jana Thayer
Valerio Mariani
Fred Poitevin
Kevin Dalton
Doris Mai
Katalin Mecseki

NERSC

Johannes Blaschke

OLCF

David Rogers

ALS Gemini Beam Line

Daniil Prigozhin
John Taylor

Current Funding

LBNL/LDRD Self Driving Serial Crystallography
NESAP for Doudna "CCTBX+LCLS+ALS"
NIGMS US DIALS National Resource
NIGMS R35 Serial Crystallography Computation

Discussion time-Nick Sauter, Structural Biology at 2kHz

- What have you **learned**? (Particularly in areas outside your core expertise)
  - Great alignment of goals both within immediate collaboration with LCLS and within larger light source community as expressed by Hannah Parrage (single interface for cross facility awareness, use streaming in live experiment)
- What key **challenges and opportunities** must be addressed to achieve our goals?
  - High-level metadata describing the experiment is recorded (duplicated) in three different places, and none is portable in case a different compute resource is needed.
  - The generalization to GPU and multimode should be easier. Mike Bauer (NVIDIA), “all parallelism should be implicit”. Brad Chamberlain “with Arkouda you don’t have to worry about nodes and ranks.”
  - Wah Chiu, “no funding body to help people process data…this is extremely lacking”. This should include both the computing services and the expert advice.
  - The small calibrations must also be done automatically or else they will consume a large percentage of the time if they must still be done by hand. “We can’t have human beings touching things anymore”.
- What **time-dependent factors** may influence decisions in the near term? (E.g., American Science Cloud)
  - Both the ALS and LCLS will have down time in the next three years, interfering with the ability to test data processing and self-driving experiments on live data.
  - We (David Mittan-Moreau) have one year of LDRD funding to work at the ALS / GEMINI beamline so it would be good to work on a well defined stepping stone, like “Process the data at 100 Hz data acquisition with a processing latency of 1 second, and use the results to automatically drive at least 1 motor control.”