LCLS-II Online Processing, Data Reduction, and Future Opportunities in Data Science

IFDEPS 2024

Jana Thayer/LCLS Data Systems Division Director

March 19, 2024

Many thanks to the people doing the honest work: Ric Claus, Ryan Coffee, Dan Damiani, Dionisio Doering, Gabriel Dohrliac, Chris Ford, Mikhail Dubrovin, Conny Hansson, Ryan Herbst, Wilko Kroeger, Xiang Li, Stefano Marchesini, Valerio Mariani, Riccardo Melchiorri, Silke Nelson, Chris O'Grady, Amedeo Perazzo, Frederic Poitevin, Lorenzo Rota, Thorsten Schwander, Murali Shankar, Monarin Uervirojnangkoorn, Matt Weaver, Seshu Yamajala, Cong Wang, Zhantao Chen

Detectors, Data Reduction, and Online Data Processing

Detectors at light sources are not standalone. They are part of an edge to HPC pipeline

Earlier processing of data happens for many reasons:

- Data Reduction
  - To get the data out of the camera (bandwidth constraints)
  - To reduce downstream network, storage or computing constraints
- Pre-processing
  - Auto-calibration of data
  - Split a multi-step algorithm (calibrate + peak finding) into two steps and do the first in the FPGA layer and the second in the online CPU layer
- Feature extraction: provide actionable information for the purposes of
  - Data Quality Monitoring
  - Experiment Steering or Control

Repercussions:

- Data flow in this pipeline may be bidirectional
- Detector firmware may not be static - it may change with experiment or parameters may change as a function of time

SLAC

LCLS-II Data Challenges

- **LCLS-II Upgrade:** greater data *velocity*, *volume*, and *complexity*
  - Data Rates: 120 Hz to 1 MHz (10000x)
  - Raw Data Volumes: 2 GB/s to 200 GB/s (100x)
  - Recorded Data Volumes: 2 GB/s to 20 GB/s (10x)
  - Computational Requirements: 80% ~1 PF, 20% ~1 ExaFLOP
- **Fast Feedback:** real-time analysis (sec/min) is essential to the users' ability to make informed decisions during experiments.
- **Variability:**
  - Wide variety of experiments with turnaround ~days
  - Large dynamic range: device readout 0.01 Hz - 1 MHz
  - Data Complexity: Variable length data (raw, compressed)
  - Access patterns to data vary by experiment and detector
  - Analysis is a mix of tried-and-true & innovative techniques
- **Time to Science:** Development cycle must be fast & flexible
- **No user left behind:** alleviate the pressure on users to gather resources to mount a significant computing effort.

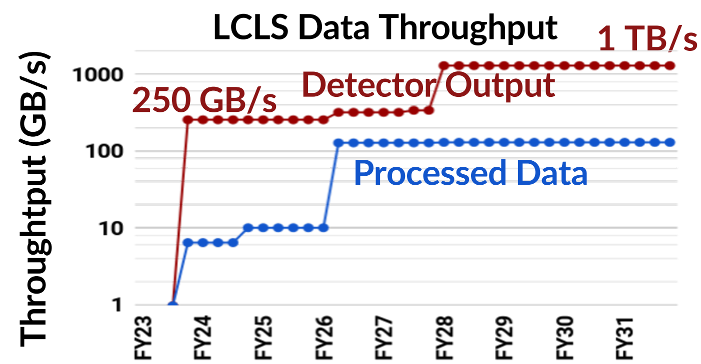

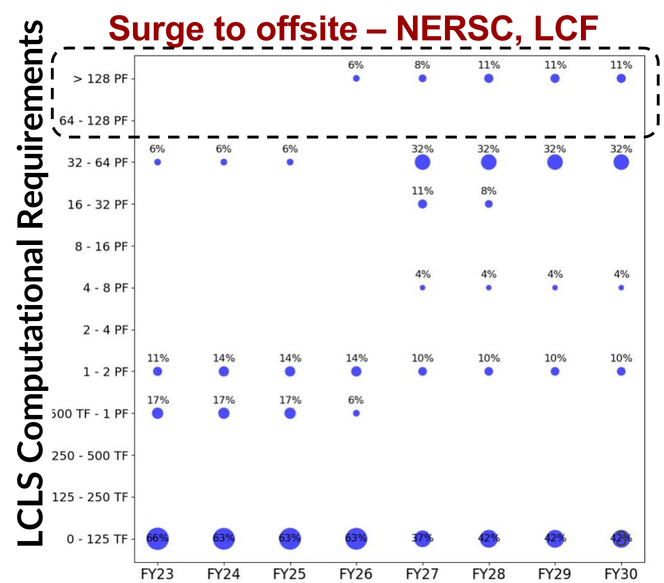

Wide variety of experiments that need to modify analysis during experiments

LCLS Data System enables & accelerates scientific discovery

Detectors are not standalone - they are components in an edge-to-HPC pipeline and used to do data reduction, feature extraction, and as part of experimental control loops

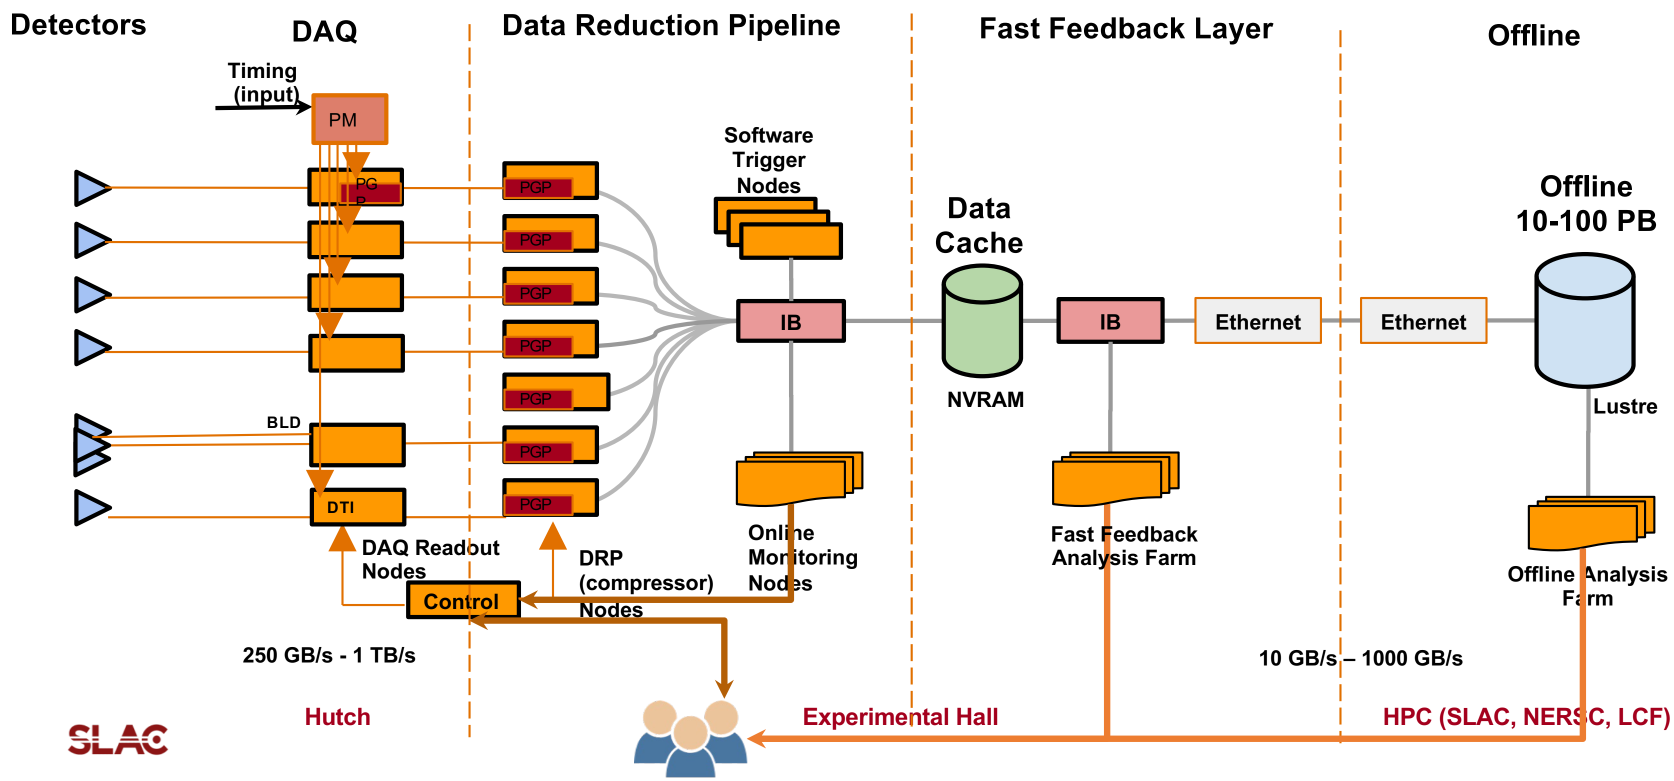

SLAC

LCLS Data System enables & accelerates scientific discovery

Detectors are not standalone - they are components in an edge-to-HPC pipeline and used to do data reduction, feature extraction, and as part of experimental control loops

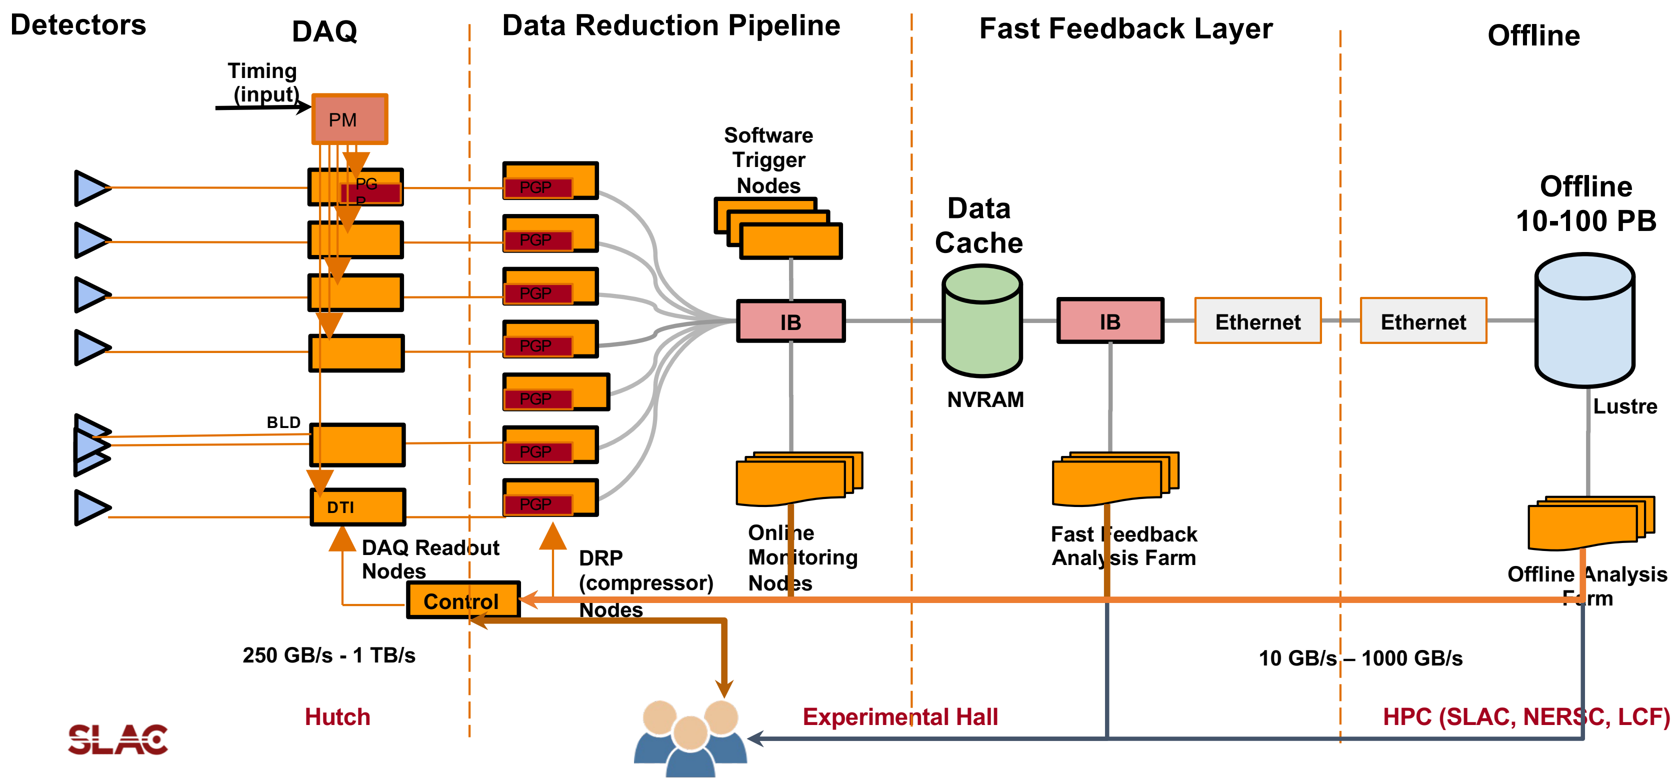

SLAC

Real-time Information Extraction and Detectors

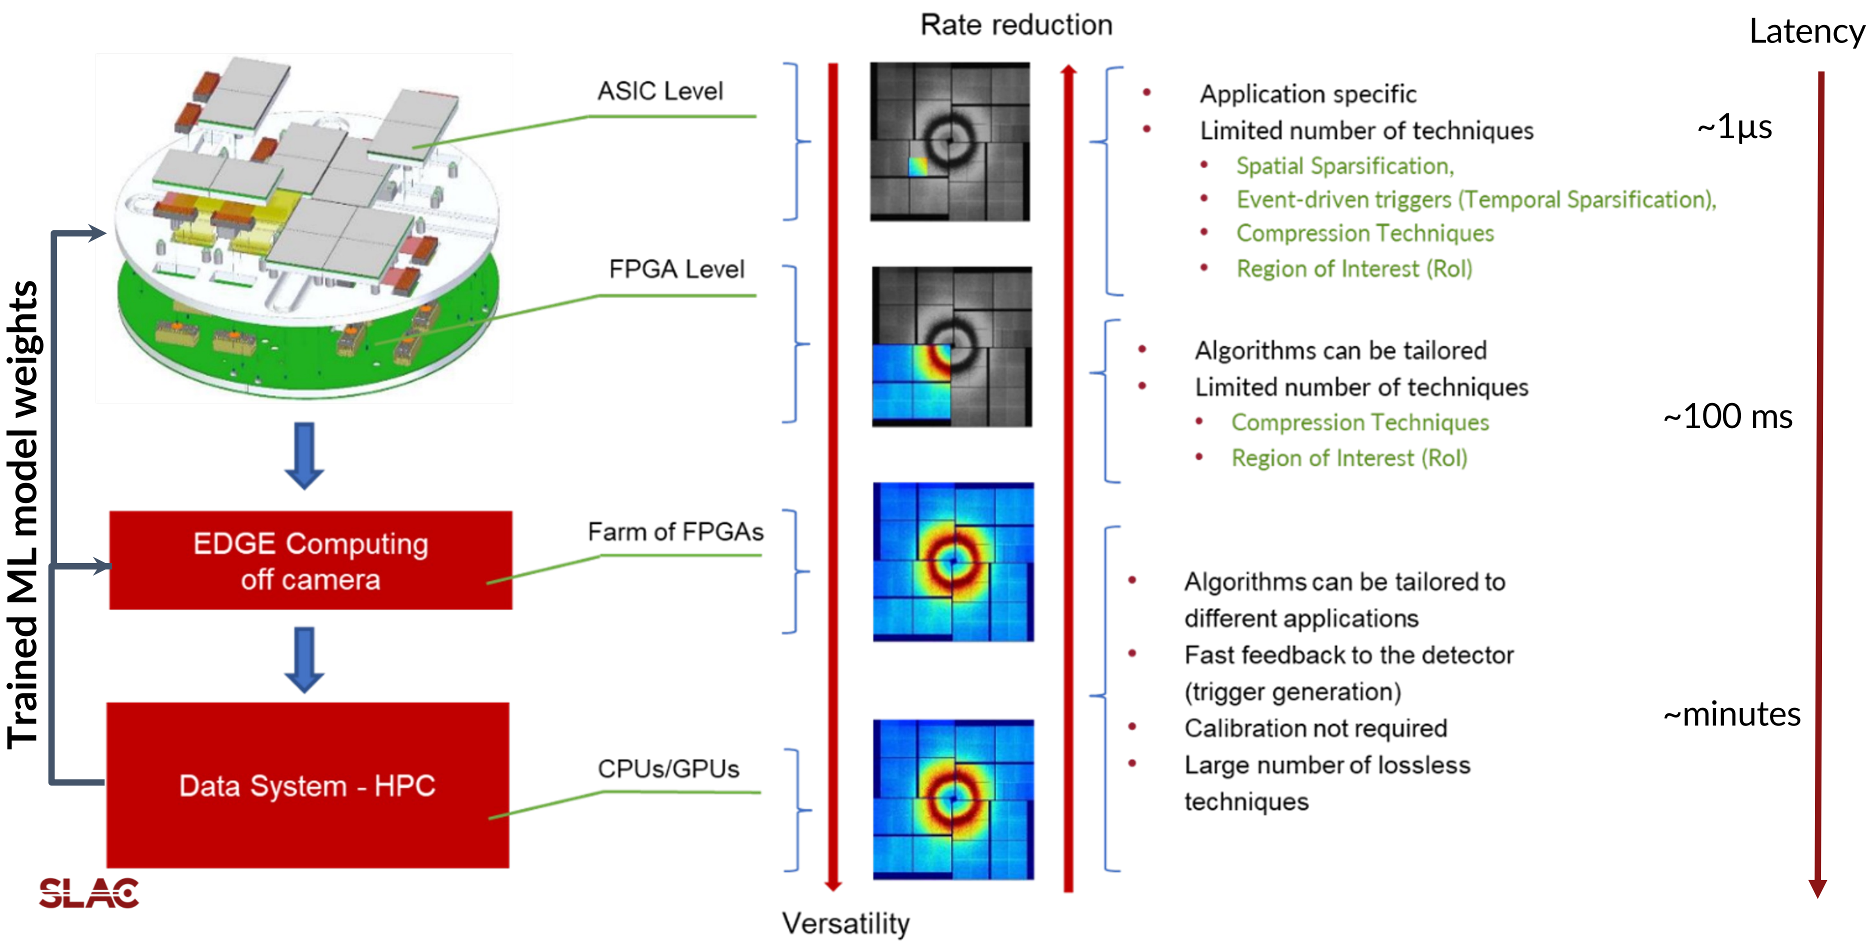

Data Reduction

Produce actionable information with low latency for fast feedback and experiment steering

Undulator	Instrument	Technique	Detector	Data Reduction Type	FFB Algorithm Type	Offline Algorithm Type
HXU	NEH 1.2	X-ray/X-ray	SXR Imaging	ROI	Peak Finding	Indexing
HXU	NEH 1.2	Imaging	epix100-HR + Digi.	Veto	Fourier Transform	MTIP
HXU	NEH 1.2	XAS / XES	RIXS-ccd	N.A.		
HXU	NEH 1.2	Imaging	ePixUHR	Veto	Fourier Transform	MTIP
HXU	XPP	Scattering	CSPAD	N.A.	Cube / Angular integration	Visualization
HXU	XPP	XAS / XES	ePix100	N.A.	Photonize	Stats Analysis
HXU	XPP	IXS / RIXS	ePix100	N.A.	Photonize	Stats Analysis
HXU	XPP	XRD / RXRD	ePix100	N.A.	Photonize	Stats Analysis
HXU	XPP	Scattering	ePix10k-HR	Binning	Cube / Angular integration	Visualization
HXU	XPP	Scattering	ePixUHR	Binning	Cube / Angular integration	Visualization
HXU	XCS/IXS	XPCS	ePix100	N.A.	Photonize	Stats Analysis
HXU	XCS/IXS	IXS / RIXS	ePix100	N.A.	Photonize	Stats Analysis
HXU	XCS/IXS	XRD / RXRD	ePix100	N.A.	Photonize	Stats Analysis
HXU	XCS/IXS	XPCS	epix100-HR	Compression	Photonize	Stats Analysis
HXU	XCS/IXS	XPCS	ePixUHR	Compression	Photonize	Stats Analysis
HXU	MFX	Xtallography	Jungfrau	N.A.	Peak Finding	Indexing
HXU	MFX	Xtallography	Jungfrau	Veto	Peak Finding	Indexing
HXU	CXI	Xtallography	Jungfrau	N.A.	Peak Finding	Indexing
HXU	CXI	Imaging	Jungfrau	N.A.	Fourier Transform	MTIP
HXU	CXI	Xtallography	ePixUHR	Veto	Peak Finding	Indexing
HXU	CXI	Imaging	ePixUHR	Veto	Fourier Transform	MTIP
HXU	MEC		ePix100	N.A.	TIFF	Animated GIF

1 MHz capable DAQ with real-time data reduction

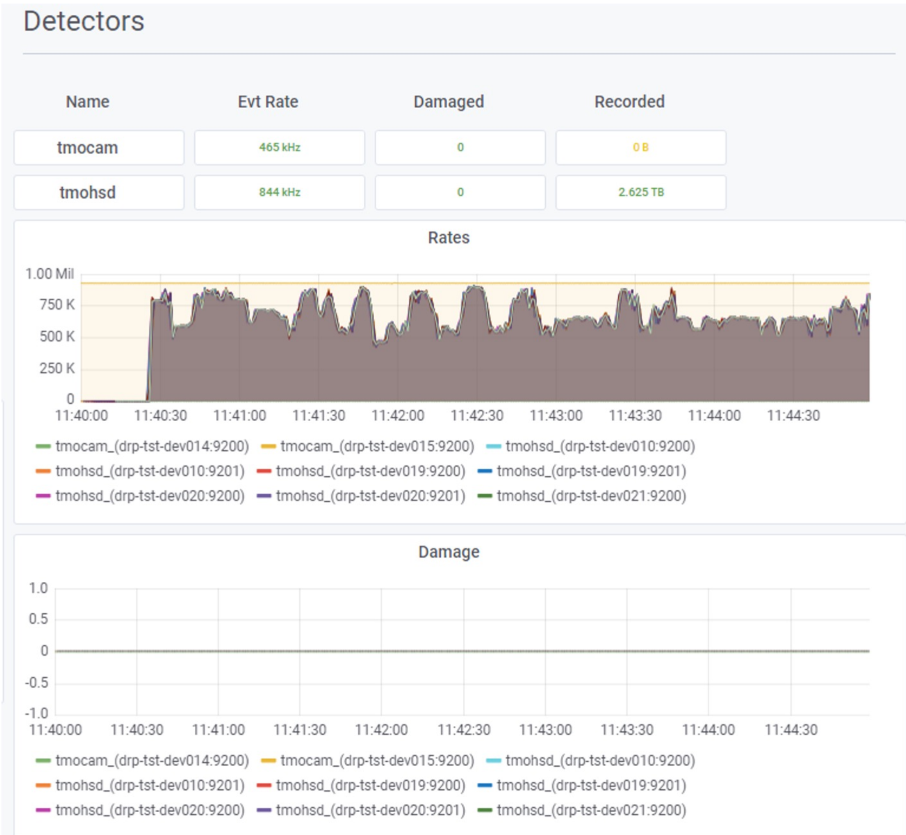

Users select from toolbox of data reduction algorithms

- Parameterized data reduction algorithms run on the DRP compute layer
- Algorithms: Lossless compression, SZ compression, feature extraction, trigger/veto
- Validation: save a programmable fraction of unreduced data (100 Hz)

DAQ and Data Reduction Pipeline tested at 120 Hz → 1 kHz in TMO with data reduction for waveforms in FPGA and ROI for Piranha.

Tested acquisition at 1 MHz without beam using data from 14 high-speed digitizer channels and other instruments such as wave8, Piranha camera.

Lossy compression with fixed error bounds - SZ Compression

SAXS/WAXS is challenging: every shot contains information; hard to distinguish signal

- Demonstrated SZ3 lossy compression with fixed error bounds on single panel emulated ePixHR @ 8 kHz with full calibration in DAQ test stand
    - Data reduced by factor (9x, err= 100), (17x, err=200)
    - No perceptible effect on the science result
- R&D milestones supporting this demonstration:
    - Re-factor calibration software to split segments across many nodes driven by serial number (Mikhail Dubrovin)
    - SZ compression performance improvements (Franck Cappello at Argonne) and segmentation (Stefano Marchesini)
    - Code refactored for highly-parallelized readout
    - Assumptions renormalized: algorithms do not always operate on fully reconstructed, fully calibrated images
- Cons: Does not produce actionable information; need to decompress prior to analysis in offline (there is a computational “penalty”)
- SZ compression has been previously demonstrated on crystallography

Reduced by 17x, err=200

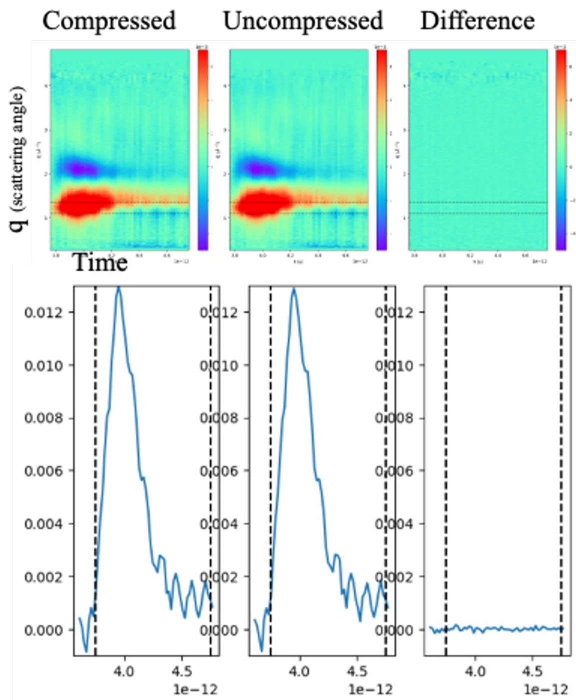

Credit: Stefano Marchesini

Edge to HPC Workflows

Free science from the limits of time and distance by providing adequate access to computing resources

Real-time feedback to validate data reduction validation

View a selectable fraction of events that meet user-specified criteria with ~1s latency

Produce actionable information with low latency for fast feedback and experiment steering

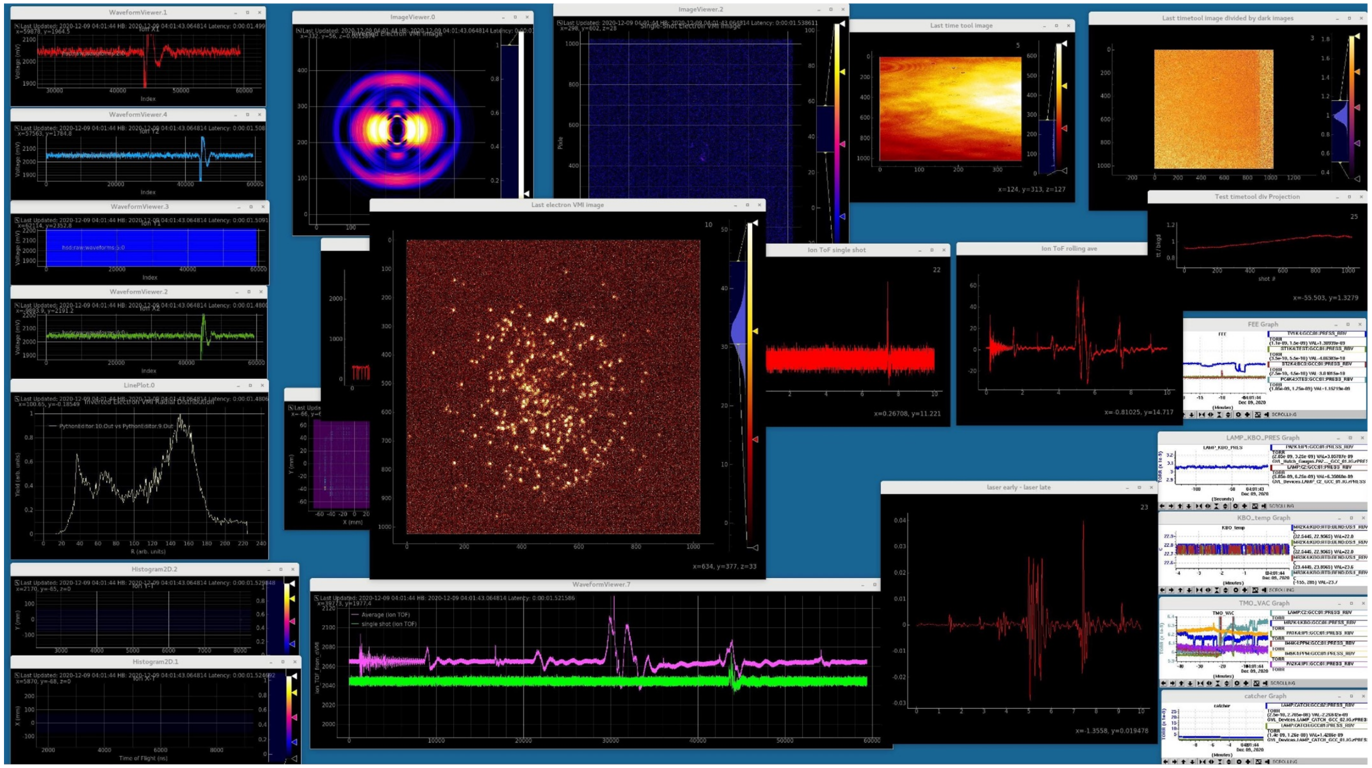

Simple, real-time feedback needed for:

- beamline alignment and tuning
- data reduction tuning/validation
- basic experiment monitoring
- *Can be a source of feature extracted information for experiment steering*

Dedicated workflows for SFX Data Analysis (LCLS ExaFEL)

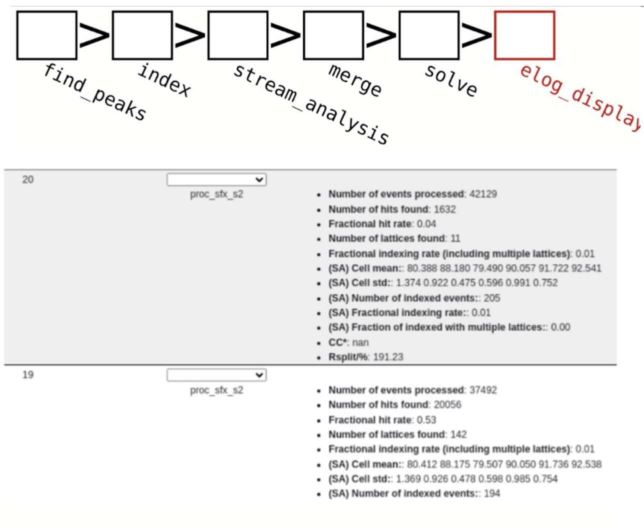

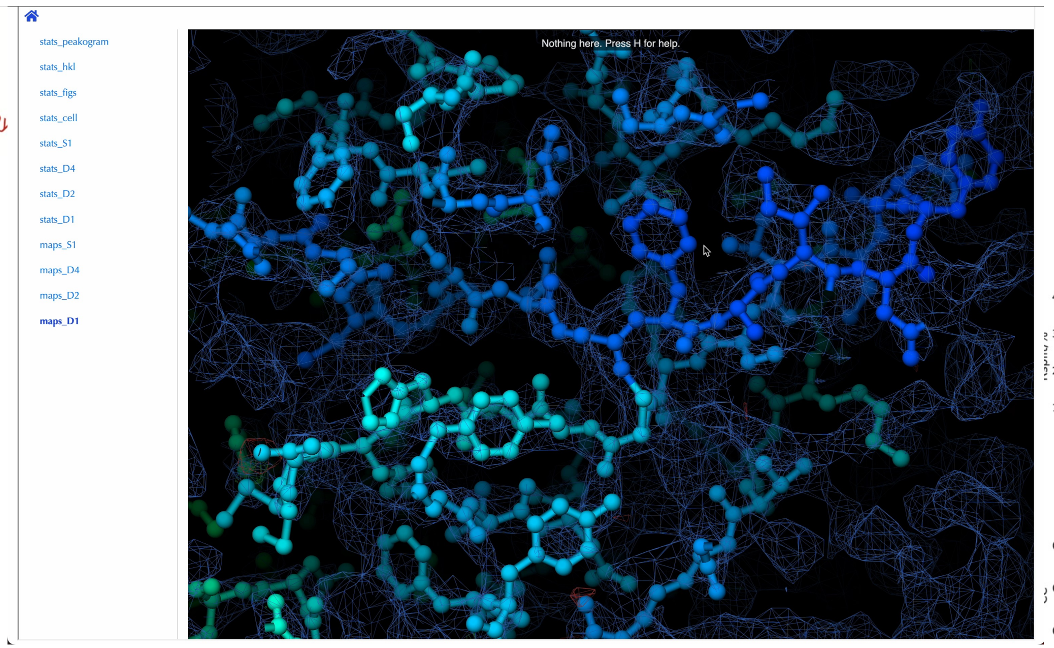

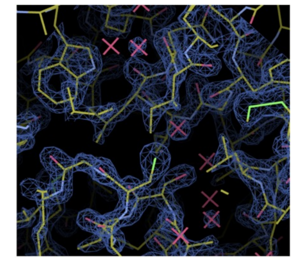

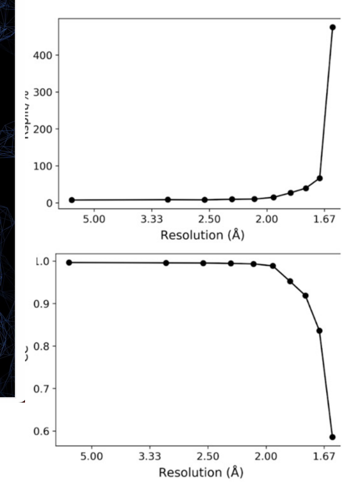

Multistep workflow from raw data processing to final result displayed in the browser. All steps report live (see above).

Users can interact with the ongoing analysis and readily evaluate once the results are good enough to move on.

SLAC

Future Directions

Smart, adaptable detectors for data reduction, fast feedback, and experiment steering lower the barrier to doing science

Smart Sensors: SparkPix-S and SparkPix-RT

Detectors with sparsified readout at ASIC enable leap from 100 kHz detector rates to 1 MHz

SparkPix-S: Pixel-threshold

- Information in both XPCS and XSVS experiments is **"sparse"** and confined in a limited # of pixels/frame, each pixel containing a limited # of photons
- 2D detector with fine spatial resolution, operating at the full rate of the machine, and discriminating between 0, 1, 2, 3.... photons/pixel/frame with high QE

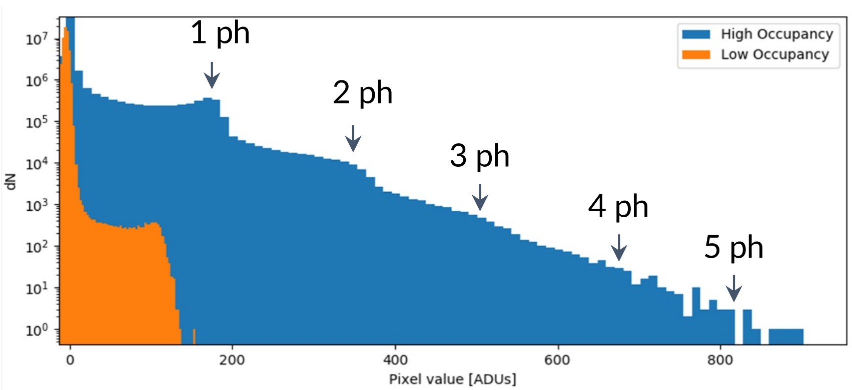

SparkPix-RT

- Solve data transmission bottleneck by implementing compression algorithm solutions in ASIC
  - bit-level compression
  - auto-correction techniques (pedestal)
- R&D needed to deal with calibration and segmentation

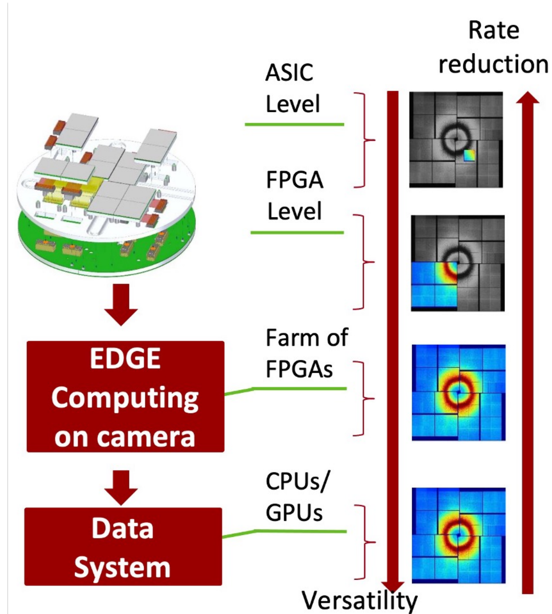

SLAC

AI/ML at the Edge: Data Reduction for TMO MRCO

MRCO reconstructs attosecond pulses using ML at the Edge

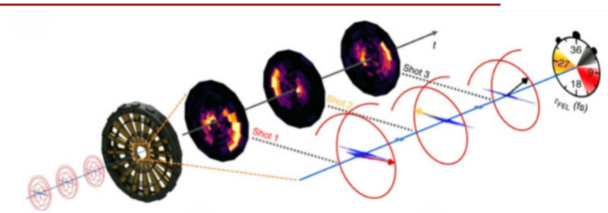

Gain insight into attosecond electron dynamics:

- MRCO/Cookiebox: Angle-resolved Electron Spectroscopy determines photoelectron angular distributions during photochemical processes

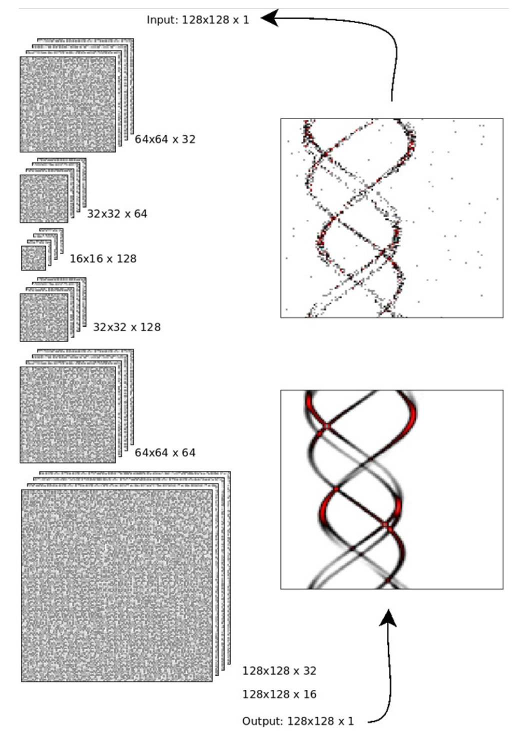

- Deploy AI inference in FPGAs: developed an AI inference library in High-Level Synthesis which enables high rate data processing & low latency feedback
- Implemented CookieNet feature extraction to reconstruct time-energy distribution of an attosecond FEL pulse in real-time to reduce 100 GB/s →~1 GB/s
- Demonstrated in Data Reduction Pipeline FPGA (KCU1500)
- Demonstrated training and inference on Graphcore and SambaNova

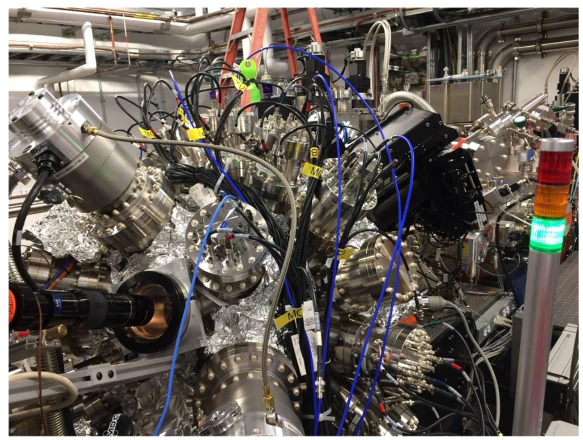

MRCO/Cookiebox

SLAC

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number FWP-100643 and FWP-35896.

ML in FPGA: SLAC Neural Network Library (SNL) Framework

Goal: Provide a set of libraries to synthesize AI inference networks into FPGAs

SNL implementation is targeting scientific instruments (frame rate of 100 kHz to 1 MHz) which must continuously adapt to new data and changing environments.

- Targeted at networks of a medium size, 10 - 20 layers, 100,000s of trainable parameters,
- Dynamic reloading of weights and biases to avoid re-synthesis.
  - Cannot re-synthesize for new training set; cannot risk FPGA implementation failing due to increase in resource usage , timing failure, or change to internal interconnect structure.
- High speed training is needed to support this as are real time bias and weight updates.

Features:

- Supports a Keras-like API for layer definition and configuration, modular and extensible
- Currently supported layer types: Conv2D, MaxPooling, AveragePooling, Dense, Reservoir.
- Current activators: LeakyRelu, Relu

To Do: Quantization, attention layers for transformers (foundation models), global optimization suggestions

SLAC

Figures: Greg Stewart at SLAC

Use ML to analyze data at the rate the production (1 MHz)

Analyze data at the rate of production using ML and providing access to network and compute

- Introduce AI/ML feature extraction at the edge to produce actionable information to feed experiment steering.
- AI-assisted decision making (running offline) uses analyzed information and other inputs to steer experiment.
- Embrace the use of heterogeneous pipelines (FPGA, CPU, GPU) and make them flexible, resilient, and transparent to use and configure

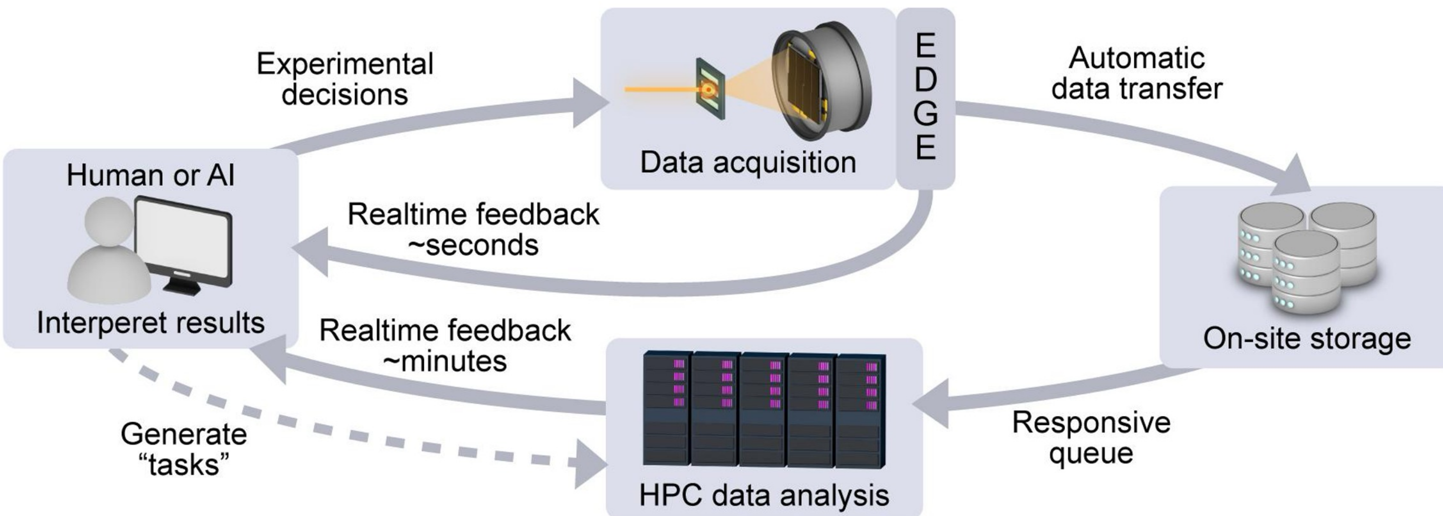

More good information, faster → better decisions → better data → experiment success!

Connect scientific instruments and HPC to create **smart instruments**

Provide actionable information by developing on-the-fly inference at the edge using ML trained remotely on streamed data - rapid (re)training workflows

AI/ML at the Edge can introduce new, compute-intensive workflows, such as those required to re-train a model on streaming experimental data. Experiment conditions can change within 1000 seconds, so rapid re-training necessary.

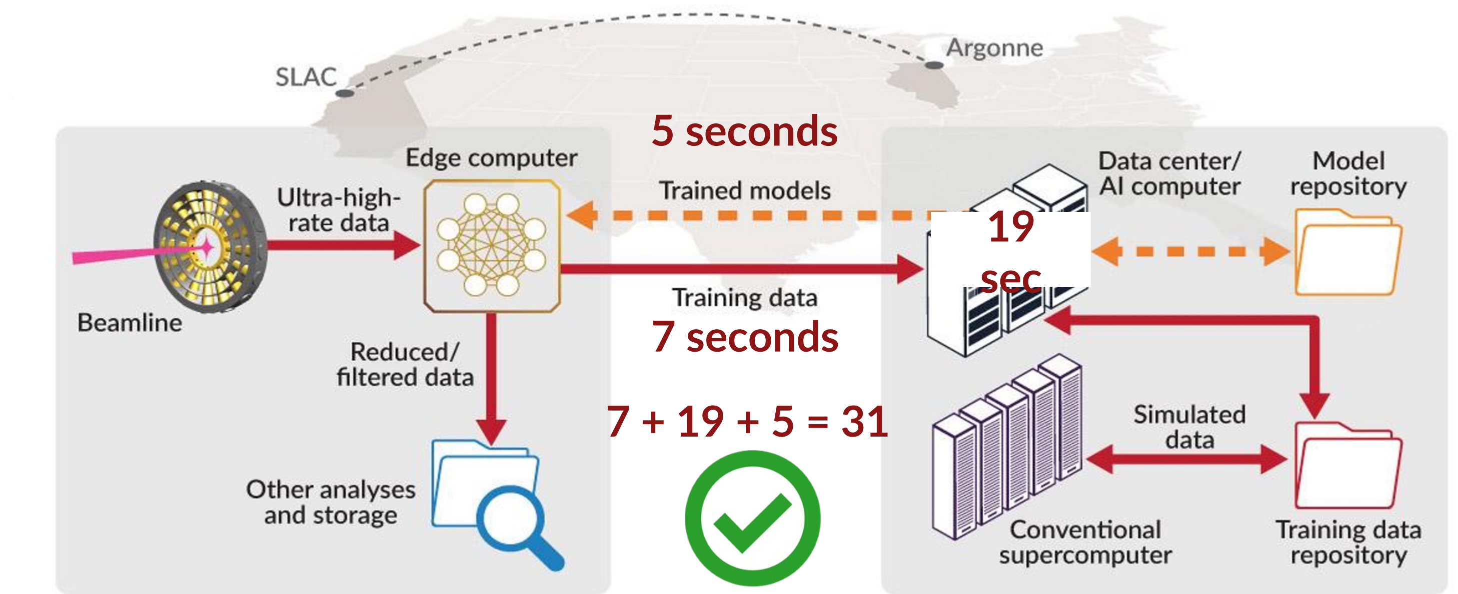

SLAC

This material is based on work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number FWP-

Machine learning enabled real-time experiment steering

Actionable information produced at each layer of computing feeds decision-making algorithms that can drive experiments over seconds, minutes, or hours

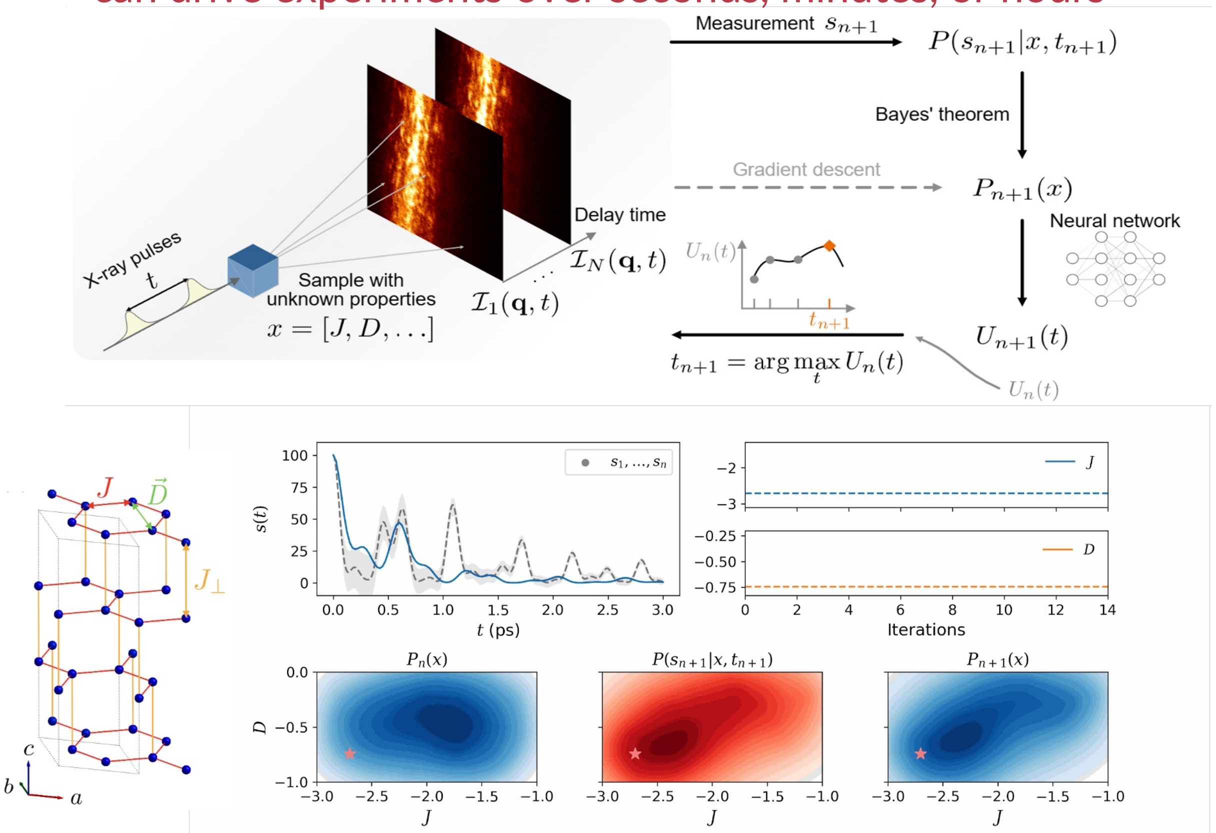

- **Help users make physics-informed decisions during their beam time.**
- Develop a data-driven **experiment steering framework** to suggest next measurement point, time delay *t*, that maximizes information gain
- Uses a surrogate model for spin excitations based on current measurement
- Uses Bayesian design for real-time decision making and parameter estimation
- **Needs access to computing**

Chen, Z. et al., 2023 (https://doi.org/10.48550/arXiv.2306.02015)

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number DE-SC0022216.

More good information, faster → better decisions → better data → experiment success!

Summary Thoughts

AI-powered edge to HPC pipelines are the way of the future, whether we like it or not

- What works for my facility may not work for yours; your mileage may vary.
- At minimum, it would be nice if detectors could auto-calibrate (even partially - pedestal)
- Users should not, in general, be expected to program firmware…but they might change parameters
- As we try to stuff AI/ML, adaptability, and intelligence into ASICs and FPGAs we will run into:
  - Need more resources at the detector (memory, etc)
  - If processing, need to collect and report information/statistics about what the detector is doing
  - Differences of scale between online/offline:
    - When users develop algorithms offline, they usually develop on fully calibrated and stitched together images with infinite computing resources available, batch sizes are large
    - Online data is segmented, every bit is touched once, computing/memory are limited, batch size is one
    - LLMs are the new hotness and they keep getting bigger. FPGAs do not provide enough space.
  - Do we develop something specific and small for the edge that only works for a limited use case? Or do we develop something generic and large in the offline and try to port it online?
- Start with some basics and make them modular: calibration, simple bit-level data reduction.
  - But know that scientists are dreaming about detectors that can adapt to data as it comes in, react to anomalies, and help steer experiments in the most promising directions.