Data Reduction Activities at LCLS

European XFEL Users' Meeting 2024

Jana Thayer/LCLS Data Systems Division Director

January 26, 2024

Ric Claus, Dan Damiani, Mikhail Dubrovin, Chris Ford, Wilko Kroeger, TJ Lane, Valerio Mariani, Riccardo Melchiorri, Silke Nelson, Christopher O'Grady, Amedeo Perazzo, Thorsten Schwander, Murali Shankar, Jana Thayer, Mona Uervirojnangkoorn, Matt Weaver, Cong Wang, Clemens Weninger, Seshu Yamajala, Chuck Yoon, Sioan Zohar, TID Electrical Engineering group

Overview

Scientific breakthroughs rely on abundant compute.

Data Reduction helps mitigate challenges in data acquisition, storage, management, and analysis.

Data reduction that also extracts features on the fly can produce actionable information which can be used to build smart(er) experiments.

- Review LCLS-II Data System Challenges and Drivers
- LCLS Data Reduction Pipeline
- Opportunities in Data Science
  - AI/ML at the Edge
  - Edge to HPC workflows
  - Experiment Steering

LCLS-II Challenges

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

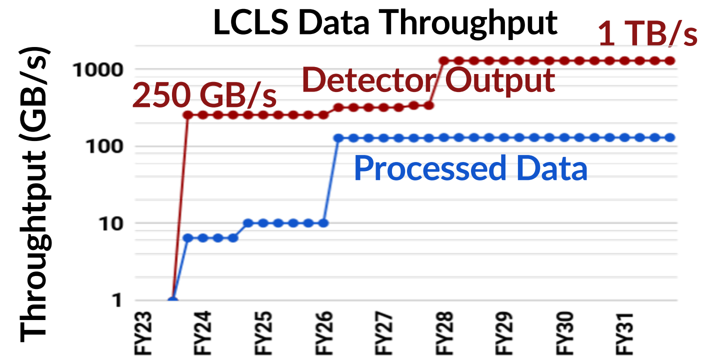

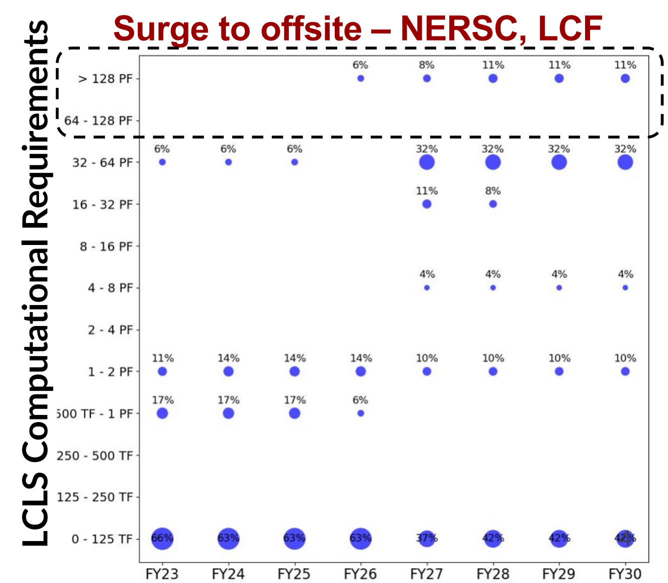

LCLS-II Data Reduction Pipeline challenges

TMO high speed digitizers and ePix imaging detectors represent big data producers

TMO DREAM Reaction Microscope Technique (Coulomb Explosion Imaging, Electron/Ion Correlations)

- Data volume is primarily high speed digitizer (HSD) data, 20 HSD channels with ~10 hits/channel
- FPGA data reduction: save only peaks and area around each peak (deadband); write zero-suppressed data
    - Tunable peak detection threshold, deadband region
    - Data reduced by factor > 50x

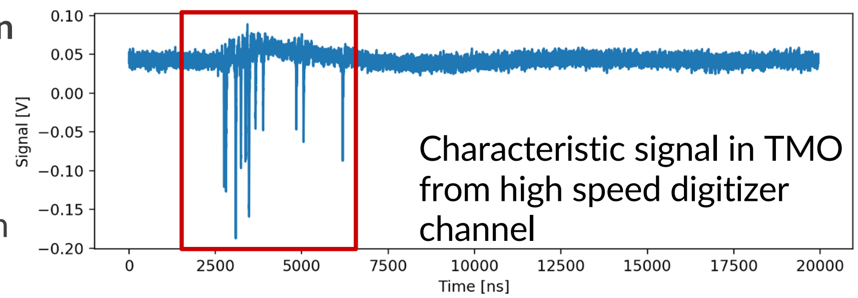

ePixHR imaging detector: 10x reduction required

- ePixHR in MFX (~2028): 16 MP @ 35 kHz = 1120 GB/s
- ePixHR @ 5 kHz expected in TXI in late 2024
- Current LCLS 1 software performance
    - 5 Hz/core for corrections for a 2 MP ePix10K
    - Implies 56K cores for 16 MP 25 kHz detector
- Note: trigger decision has been run at 1 MHz

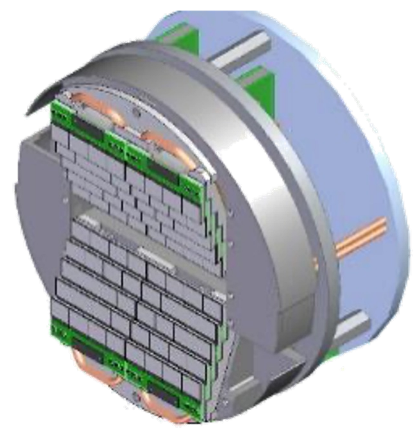

SLAC

LCLS Data System enables & accelerates scientific discovery

LCLS facility provides **access to computing** for massive-scale data analytics in a multi-tiered computing landscape that includes **edge**, **local SLAC**, and **ASCR facilities**.

Users provide the last mile: develop their own analysis on top of this stack.

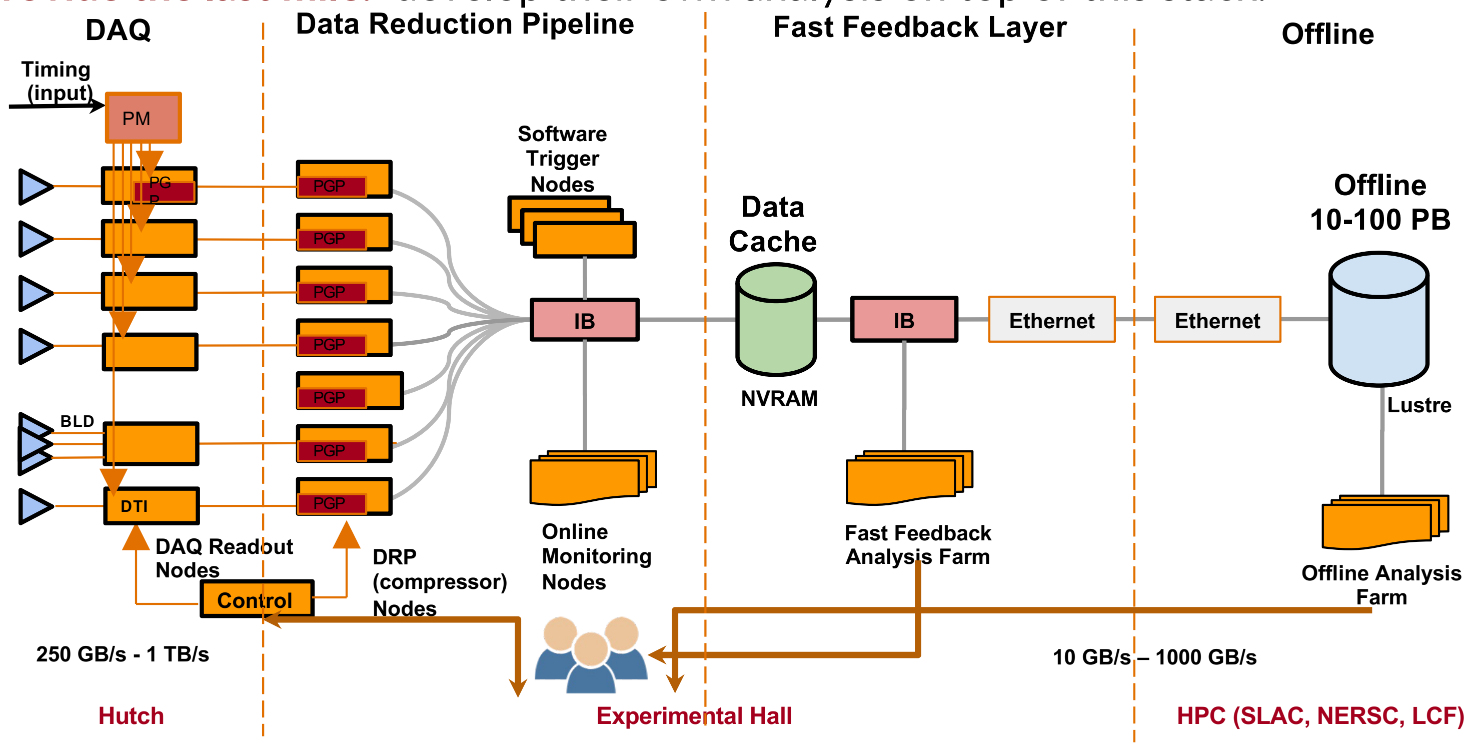

SLAC

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

Data Reduction Pipeline (DRP)

All current experiments reduce data when processing data offline → now do it in real time

On-the-fly data reduction

- mitigates network, storage, computing bottlenecks
- enables streaming to HPC

Users select from toolbox of data reduction algorithms

- Parameterized data reduction algorithms run on the DRP compute layer
- Algorithms: Compression, feature extraction, trigger/veto, multi-event reduction
- Validation: save a programmable fraction of unreduced data

Software Trigger Nodes perform online event build collecting data from multiple detectors from the same event.

Two decisions per event, per shot:

- Store or not?
- Send data to online monitoring?

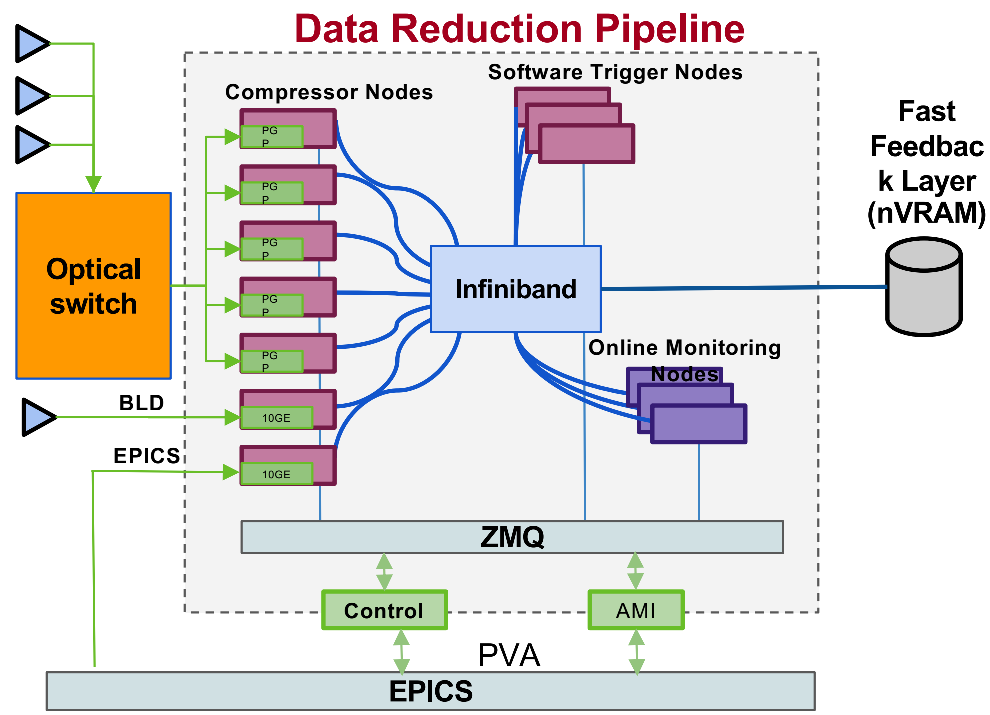

Streaming feature extraction can provide actionable information for experiment control!

Expected data reductions for XPP, DXS, MFX, CXI

Algorithms	**XPP**	**DXS**	**MFX**	**CXI**
Veto		X	X	X
SZ Compression	X	X	X	X
“Cube”: Average image binning	X	X	X	
Pixel binning		X	X	X
ROI/Projection	X			
Angular integration (and Pie-Slicing)	X		X	
“Parallelizable MPEG-style”	X		X	
Peak-finding/Thresholding	X	X		

Data Reduction Pipeline

Future real-time data reduction in FPGA and/or CPU/GPU

1 MHz capable DAQ with real-time data reduction

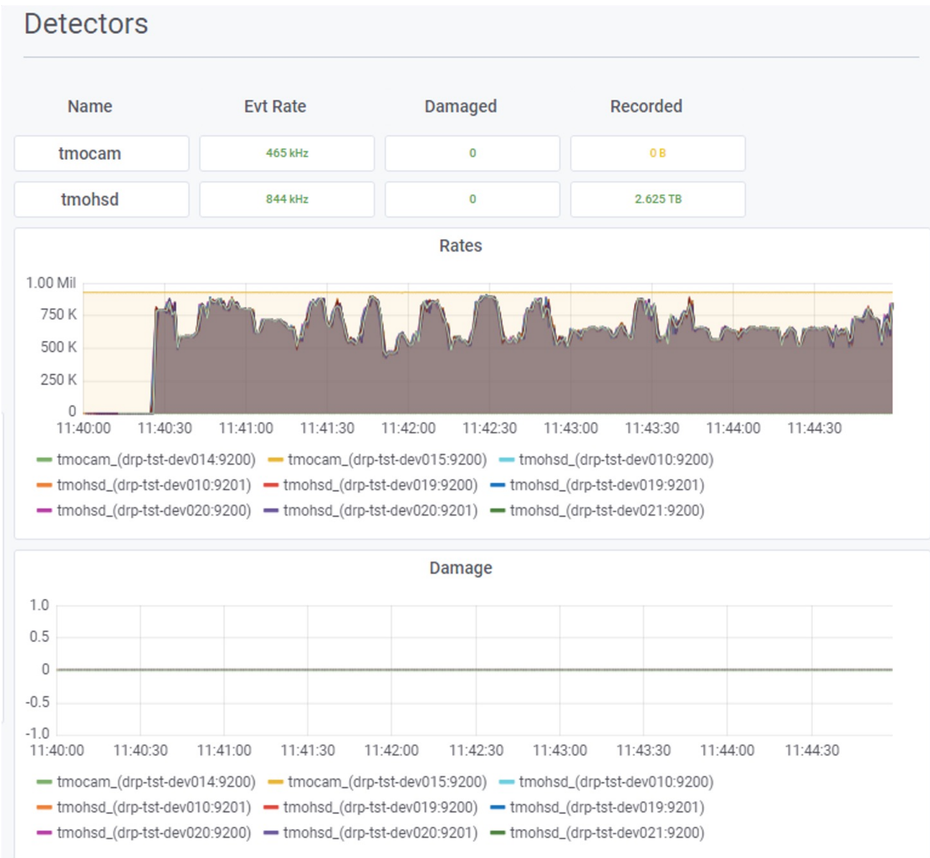

DAQ and Data Reduction Pipeline tested at 120 Hz in TMO.

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

Status of data reduction:

- Deploy data reduction algorithms appropriate to an instrument as the instruments come online.
- TMO high-speed digitizers capable of producing 200 GB/s of data.
- Tested data reduction for waveforms in FPGA
- High-speed cameras not yet deployed.

How does the DRP work in practice?

TMO and RIX use data reduction for High Speed Digitizer Data

- Since 2020, we have been practicing data reduction at 120 Hz rates with the LCLS-II data system in TMO and chemRIXS, recording raw and reduced data to evaluate performance of Data Reduction Pipeline.
- At high rate (> 120 Hz), turn on data reduction, set prescale to 100 Hz (unreduced data), gather unreduced data alongside reduced data
  - Can afford to be conservative with DRP parameters in the beginning.
  - Validate data reduction using prescaled data during first day of experiment
  - Goal is to build trust and confidence between facility and users
  - End goal is good science not good data reduction

Real-time Feedback to validate data reduction performance

AMI2 online monitoring framework attaches analyzes a subset of the most recent DRP data

Graphical user interface for developing new analysis on the fly without writing code

View a selectable fraction of events that meet user-specified criteria.

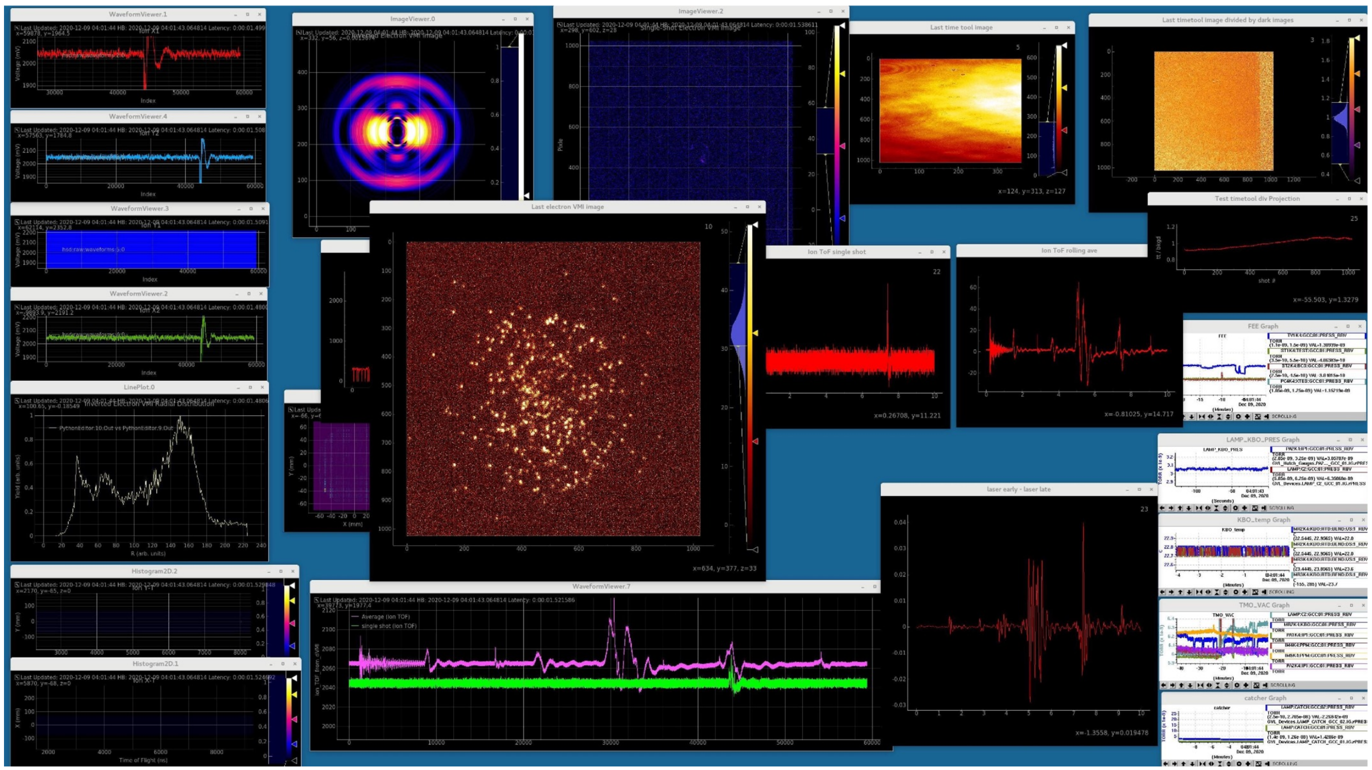

Note:

- AMI2 has a **GUI** and **scripted** interface
- **Uses psana2** under the hood
- psana2 python scripts can be adapted to run in AMI2

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

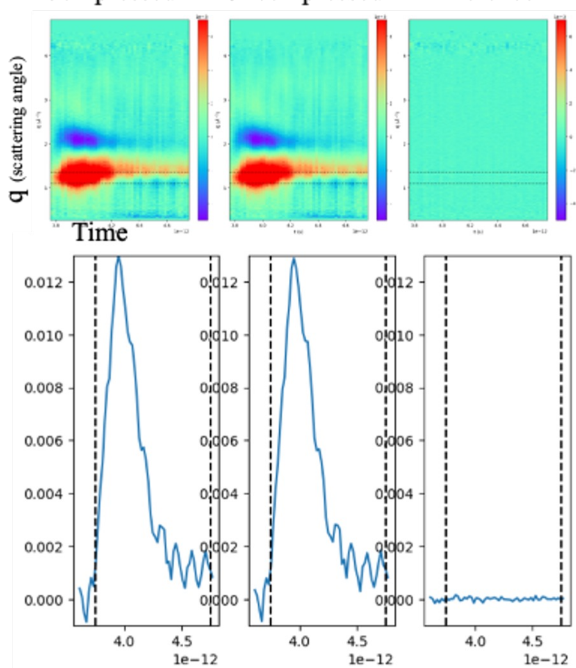

Credit: Stefano Marchesini

SLAC

BOLD PEOPLE VISIONARY SCIENCE REAL IMPACT BOLD PEOPLE VISIONARY SCIENCE REAL IMPACT

Data Reduction at the Edge

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

SLAC

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

MRCO reconstructs attosecond pulses using ML at the Edge

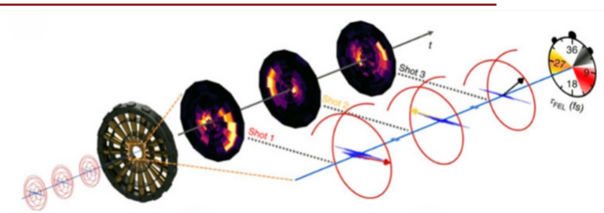

Gain insight into attosecond electron dynamics:

- 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

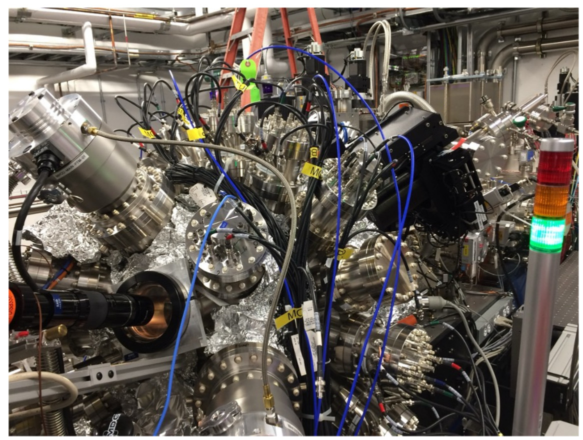

MRCO/Cookiebox

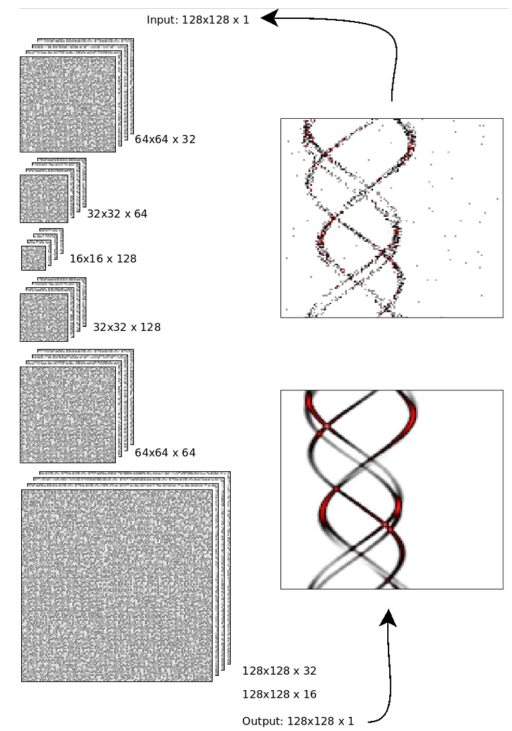

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number FWP-100643 and FWP-35896.

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

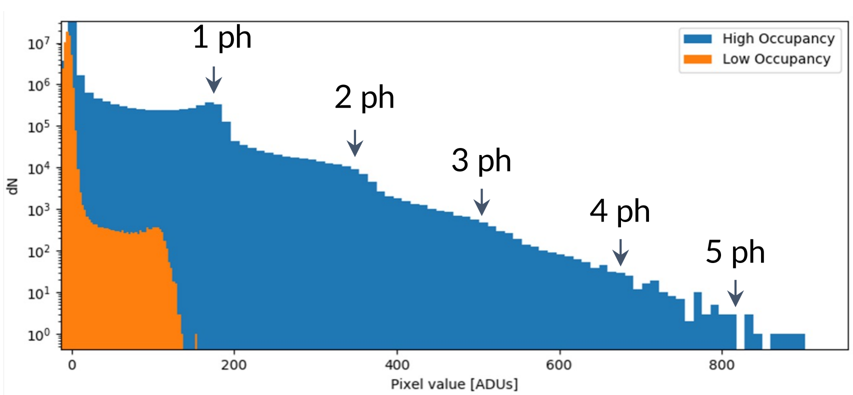

SLAC

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

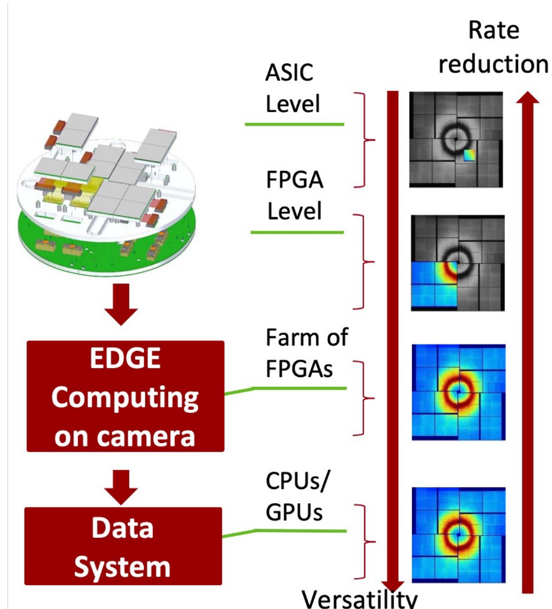

1st generation DRP: Veto for Crystallography and Single Particle Imaging

Experiment Description

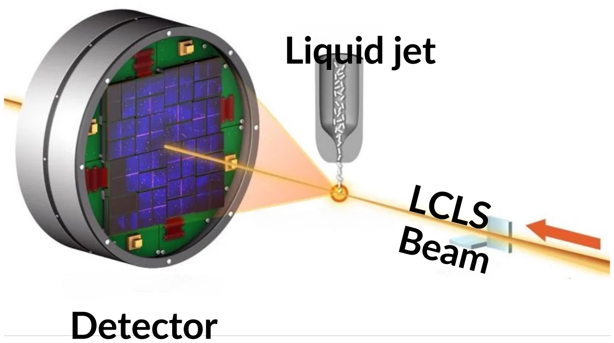

- Individual nanocrystals are injected into the focused LCLS pulses
- Diffraction patterns are collected on a pulse-by-pulse basis
- One exposure per crystal
- Each image processed independently
- Crystal concentration dictates "hit" rate

Megapixel Detector

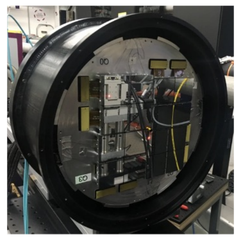

- 4 MP@5 kHz in 2024
- 16 MP@40kHz in 2028

60 GB/s → 1 TB/s

X-ray diffraction image

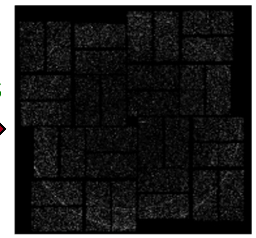

Data Reduction

- Remove"no hits"
- >10x reduction

3 TFlops
16 TFlops

6 GB/s → 100 GB/s

Intensity map from multiple pulses

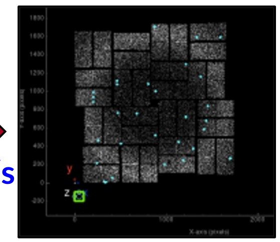

Interpretation of system structure / dynamics

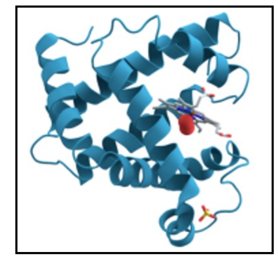

Data Analysis

- Bragg peak finding
- Index / orient patterns
- Average
- 3D intensity map
- Reconstruction

4 PFlops
20 PFlops

autocorrection, calibration

PeakNet

Next generation DRP: write peaks

Indexing, averaging, 3D intensity map, reconstruction

PeakNet: A 1 MHz AI-based Autonomous Bragg Peak Finder

Credit: Cong Wang

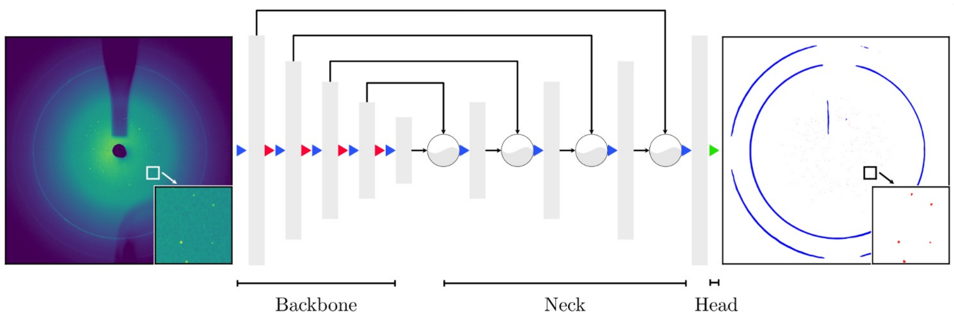

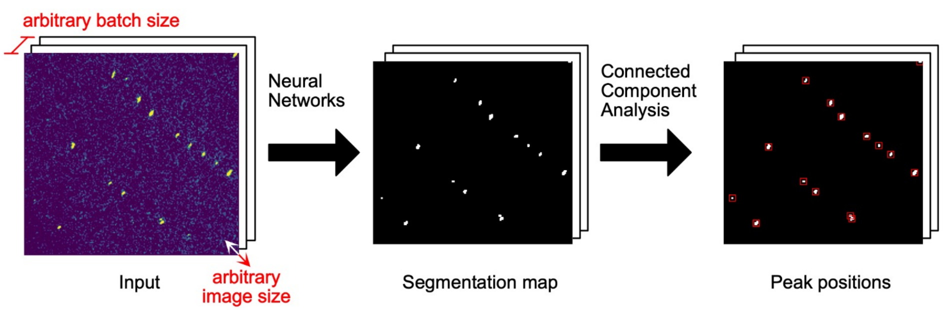

Deployed in stealth mode in CXI+MFX, running alongside a traditional peak finding-algorithm to compare performance

Wang, C. et al., 2023 (https://doi.org/10.48550/arXiv.2303.15301)
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.

Significance and Impact

- *Once proven, use PeakNet in Data Reduction Pipeline to write peaks instead of raw images to disk.*
- SFX produces vast amounts of data, posing computational challenges.
- PeakNet is a deep neural network for
  - Autonomous Bragg peak detection in real-time
  - Adapts in real-time to shot-to-shot background changes without manual tuning

Features

- Autonomous pixel segmentation into 1) Bragg peaks, 2) artifact scattering, and 3) background, requiring no user parameter tuning.
- Currently transitioning to a new modular model architecture for
  - faster training/retraining
  - easier deployment to new use cases
  - enabling easier migration to FPGA

Figures: Greg Stewart at SLAC

Quasi-Real Time Workflows using High Performance Compute

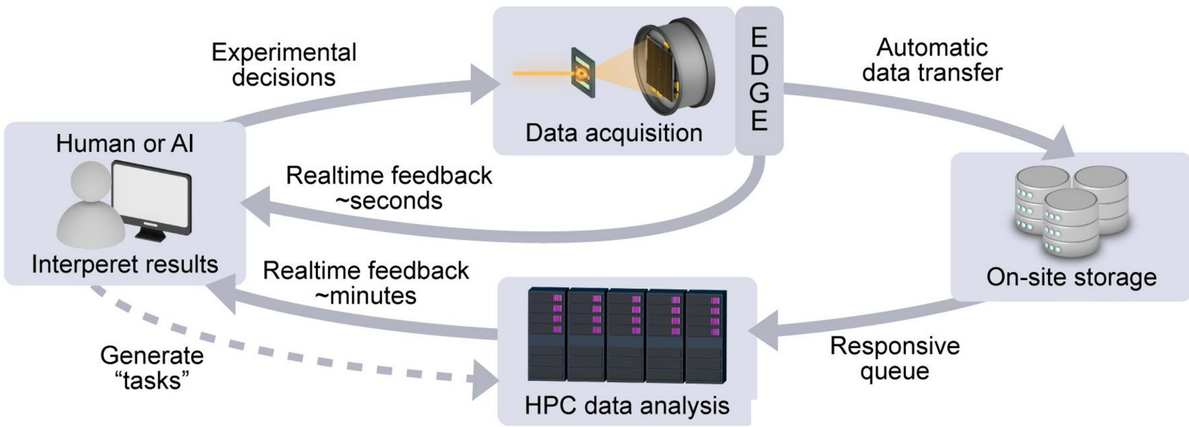

Analyze data at the rate of production by providing seamless access to network, compute, and storage.

Stream data to remote HPC for prompt analysis

AI/ML at the Edge must be capable of fast adaptation to changing conditions

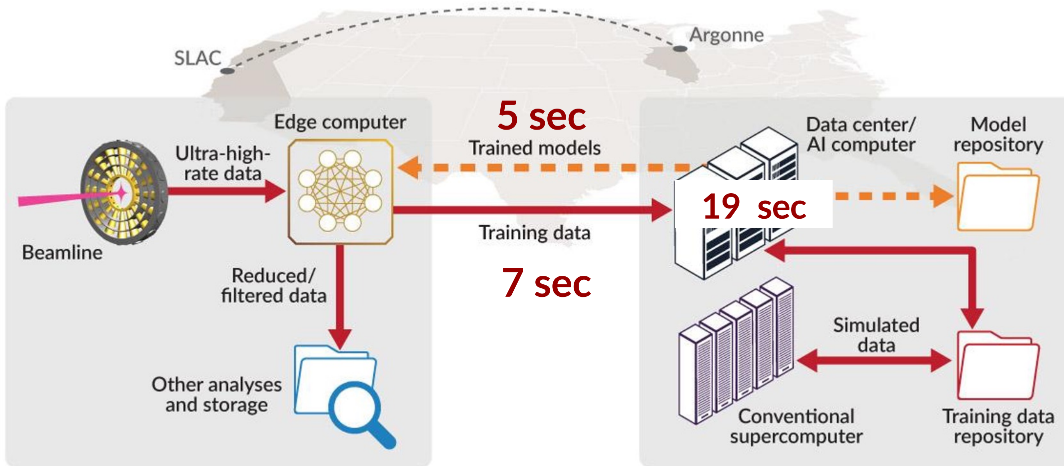

More good information, faster → better decisions → better data → experiment success!

Summary

Advances in computational power and analysis methods that leverage massive data quantities will maximize the science output from LCLS

but need to implement data reduction to mitigate bottlenecks in network, storage, and compute!

LCLS is supporting the development of a data system infrastructure capable of handling the demands of Big Data:

- Real-time data analysis capabilities (data reduction, complex workflow orchestration)
- On-demand utilization of super-computing environments
- Strategic development of AI/ML for targeted applications
- Ability to automate experiments (execution to analysis)

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

Backup

Data compression at different stages of analysis pipeline

Earlier data reduction is usually a win: less data volume to transfer, store, and process.

- **Earlier data reduction (online):**
  - **Pros:** reduces volume, network bandwidth allowing transfer of data (out of detector or across ESnet), reduces storage cost, may reduce downstream computing needs (if reduction method is equivalent to a pre-processing step)
  - **Cons:** introduce latency in the pipeline and some risk that information is lost if lossy compression method is used. If you do something wrong, you may not be able to reprocess to get it right.
- **Later data reduction (offline):**
  - *Non-starter for LCLS-II-HE*: network transmission costs, storage costs, ability to write to disk, cost to analyze data are all astronomical
  - **Pro:** You can repeat your analysis if, for example, your calibration is wrong

SLAC

The Future of Data Reduction at LCLS-II / LCLS-II-HE

Data compression can be implemented at different stages of the pipeline.

- First generation DRP implementation is deliberately minimal, transparent and non-ML to build trust while still achieving the minimum necessary data reduction.
  - Implemented in a single place in Data Reduction Pipeline
  - Key tools: prescale (raw+compressed data at 100 Hz), online monitoring, and fast feedback for validation, buildup of workflows; learn where the pain points are.
- Next generation DRP implementation: more ambitious feature extraction algorithms
  - Provide actionable, real-time information to drive experiment steering (baby steps)
  - Move data compression and algorithms closer to the detector
    - bit-level compression (SparkPix-RT)
    - auto-correction techniques (pedestal correction, gain on SparkPix-RT)
    - begin incorporating ML for well-characterized diagnostics that are used throughout the facility and generate large raw data volumes (ATM, MRCO)
  - Recognize pain points and develop solutions or rapid adaptations:
    - detector characterization, calibration, good/bad/twinkling pixels, beam spot finding
    - work with users to understand the reasons for reprocessing data

SLAC

ILLUMINE - $10M over 5 years for Experiment Steering

ILLUMINE - Intelligent Learning for Light Source and Neutron Source User Measurements Including Navigation and Experiment Steering

LAB 23-3030: Advanced Scientific Computing Research for DOE User Facilities

Authors: Jana Thayer (PI), Ryan Herbst, Vivek Thampy, Chun Hong Yoon (SLAC) Stuart Campbell, Daniel Allan, Andi Barbour, Thomas Caswell, Natalie Isenberg, Phillip Maffettone, Daniel Olds, Max Rakitin, Nathan Urban (BNL) Nicholas Schwarz, Franck Cappello, Ian Foster, Antonino Miceli (ANL) Alexander Hexemer, Dylan McReynolds (LBNL) Jonathon Taylor (ORNL)

Abstract:
This research proposes the development of a **multi-facility framework** to address the challenges posed by the growing volume and complexity of data collected at x-ray and neutron sources.
By integrating advanced computing, algorithms, and analysis, the framework aims to enable **rapid data analysis** and autonomous experiment steering.
It will leverage real-time compression, machine learning inference, and **decision support techniques** to optimize data collection and explore experiment phase space.
The framework, built upon the Bluesky data collection platform, will provide accessible and reusable components to enhance the efficiency and quality of experiments, unlocking new scientific possibilities.

SLAC

ILLUMINE - SLAC-led 5 light source + neutron source $10M, 5Y effort for Experiment Steering Infrastructure

A modular framework to close the loop between fast analysis, machine-assisted decision-making, and data acquisition to drive experiments on the timescales of seconds, minutes, or hours

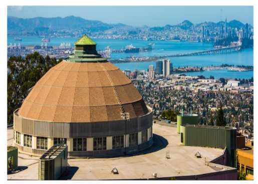



Brookhaven National Laboratory

Argonne NATIONAL LABORATORY

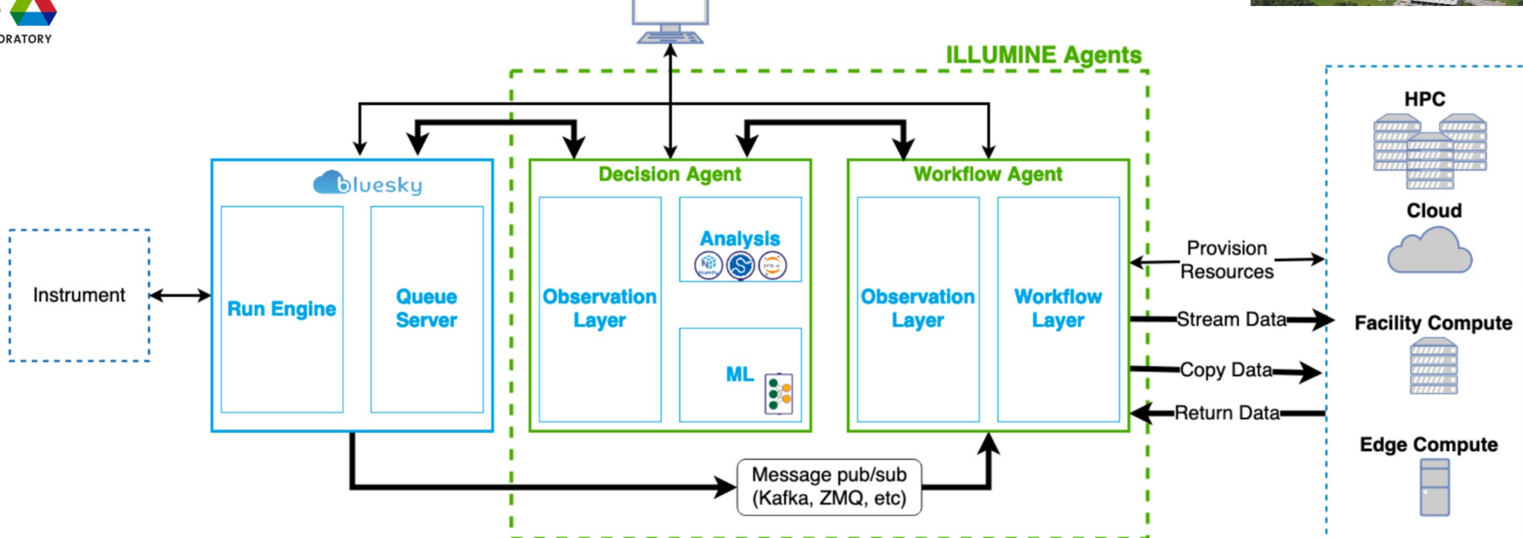

SLAC

LCLS Automated Data Movement and Run Processing (ARP)

Data In → Science Out; use ARP to orchestrate user analysis workflow on local or remote HPC

User defines:
- ❑Script/ workflow
- ❑Run #s
- ❑HPC choice*

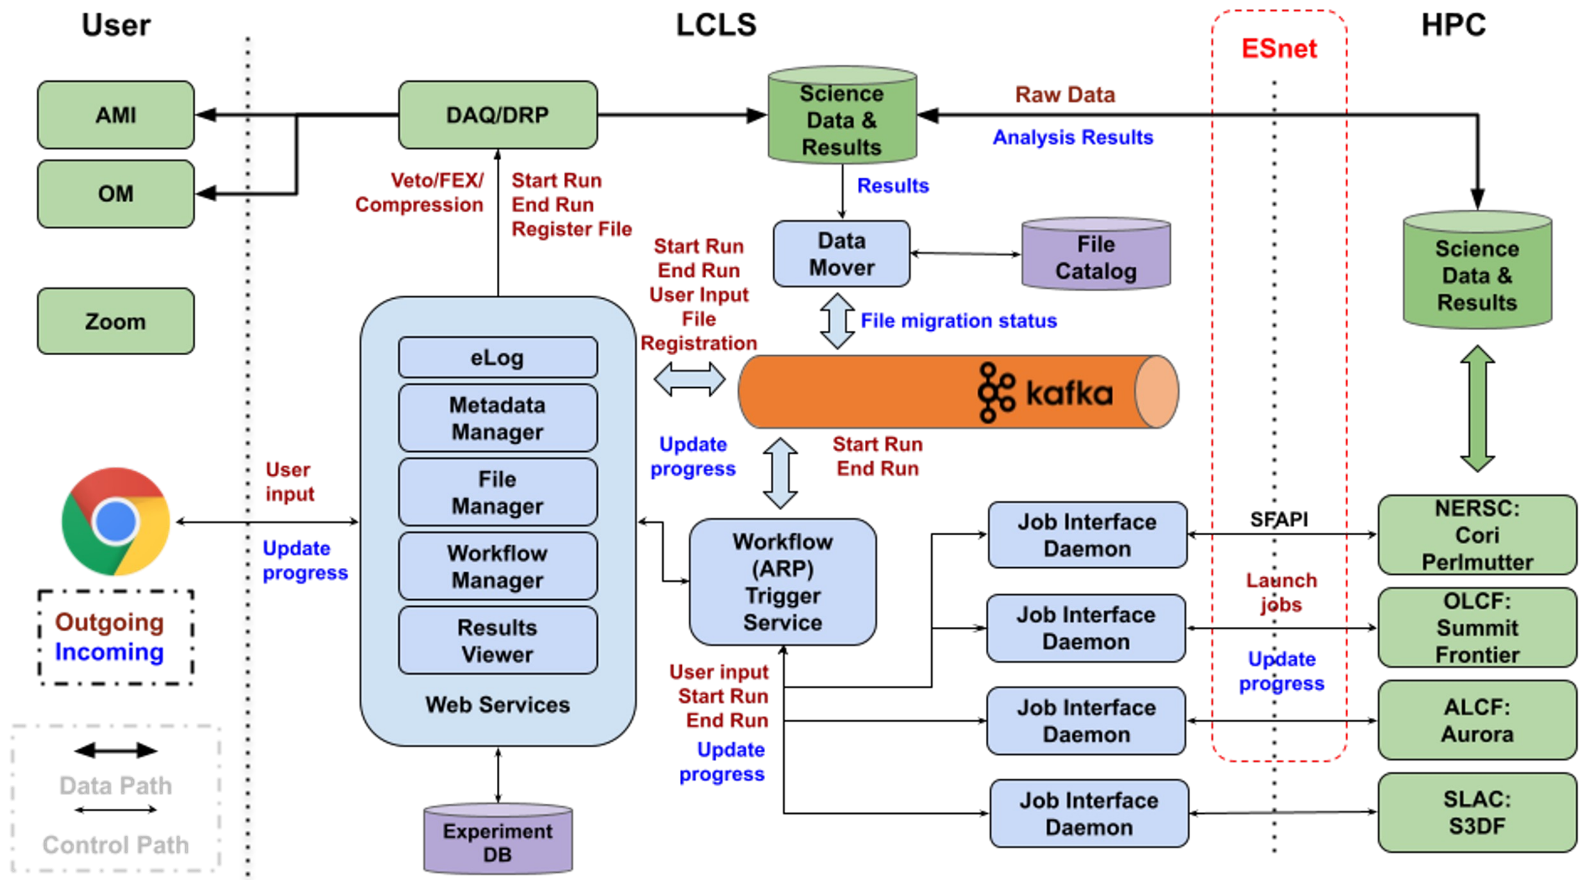

*agreed upon and reserved in advance

Quasi-real-time analysis using NERSC

https://doi.org/10.48550/arXiv.2206.11992

ExaFEL project streams data to NERSC for analysis results within minutes

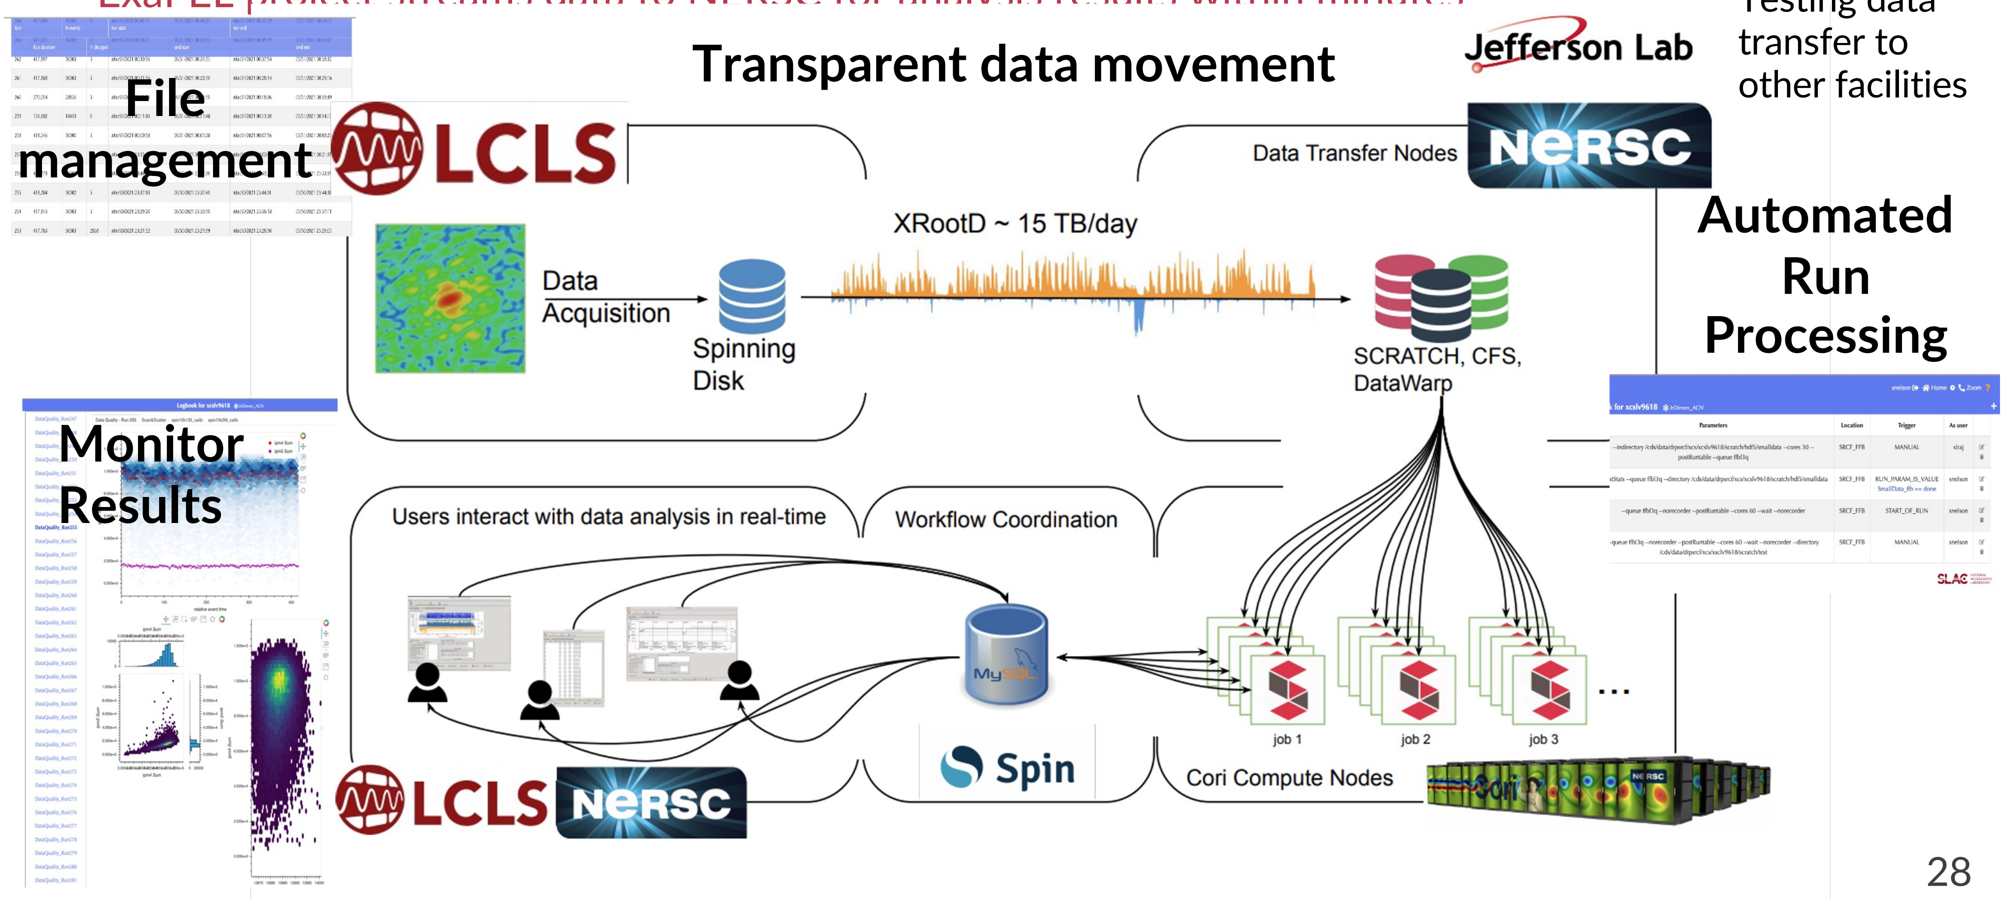

Actionable Information from Sensor to Data Center (AISDC)

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

Develop feature extraction for TMO CookieBox, Serial Femtosecond Crystallography (SFX), Single Particle Imaging (SPI), and High Energy X-ray Diffraction Microscopy (HEDM)

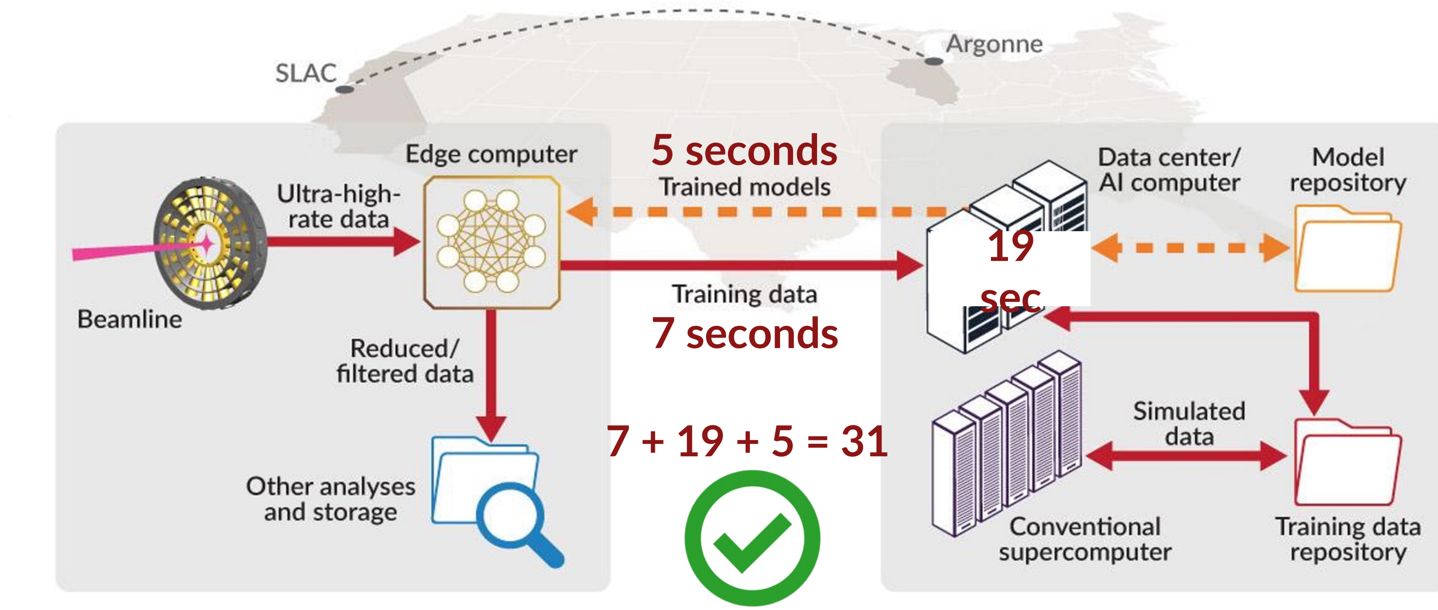

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number FWP-100643 and FWP-35896.