New Opportunities for Better User Group Software (NOBUGS) 2022

The Scientific Computing Strategy for the Upgraded Advanced Photon Source

Nicholas Schwarz
Principal Computer Scientist, Group Leader
Scientific Software Engineering & Data Management, X-ray Science Division
Advanced Photon Source

September 19 - 22, 2022

U.S. DEPARTMENT OF ENERGY Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

Argonne NATIONAL LABORATORY

APS Upgrade: The Ultimate 3D Microscope

A next-generation synchrotron light source for science and industry

$815 M project to update and renew the facility

Re-uses $1.5 B in existing infrastructure

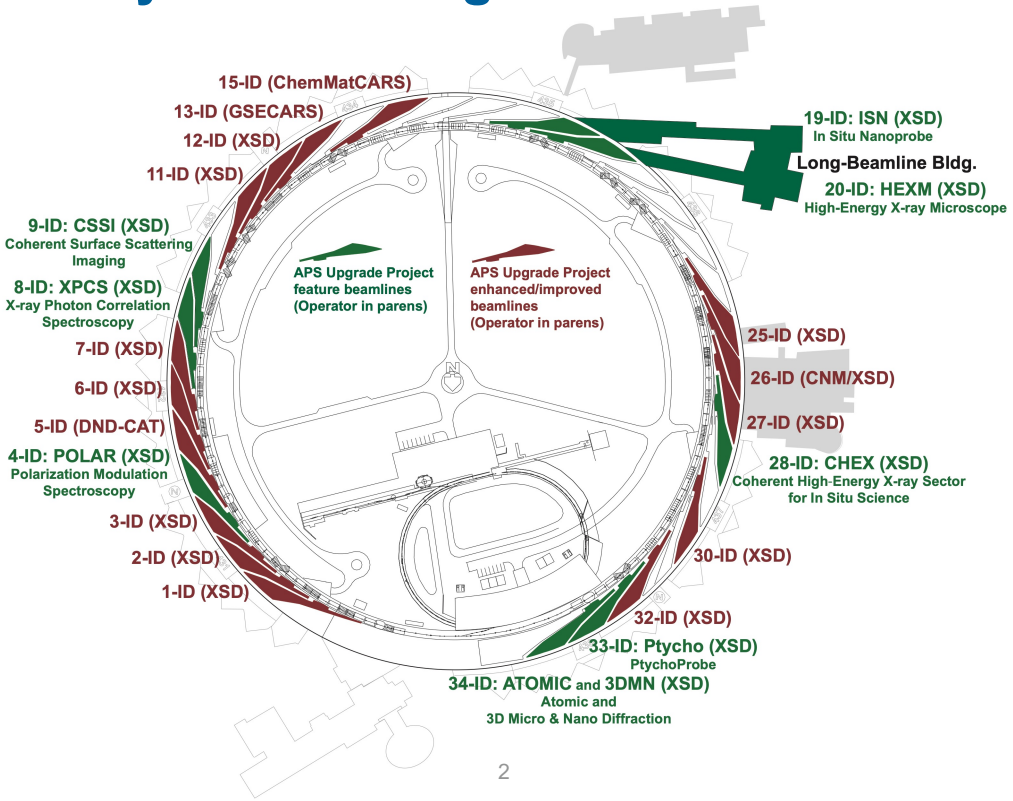

- Completely new storage ring, 42 pm emittance @ 6 GeV, 200 mA
- New and updated insertion devices
- Combined result in brightness increases of up to 500x
- 9 new feature beamlines (green)
- 15 beamline enhancements (red)

Scientific Computing at the APS

Data management & computing are critical to the success of the APS

- Enable cutting-edge research and high-impact scientific output for APS users
- Align with the major science thrusts of the upgraded facility

- Needs are driven by new scientific opportunities, which are enabled by the new source, new measurement techniques, advances in detectors, multi-modal data utilization, and advances in data analysis and Artificial Intelligence / Machine Learning (AI/ML) methods
- Detailed planning for each beamline
  - Anticipated techniques, detectors, and data rates, and networking, beamline controls, and data management and analysis needs

- Over the decade, the APS anticipates collecting 100s of petabytes of raw data per year requiring 10s of PFLOP/s (peta floating point operations / second) of on-demand computing power per year

Scientific Computing at the APS

Strategy and R&D plan for networking infrastructure, advanced data acquisition, data management, computing resources, and algorithms, analysis software, and AI/ML

Key priorities:

- Upgrading **networking infrastructure** within the APS and between the APS and the Argonne Leadership Computing Facility (ALCF)
- Deploying state-of-the-art experiment control software to enable **adaptive feedback**
- Expanding the capabilities and use of common **data management and workflow** services
- Utilizing **high-performance computing resources** for on-demand data processing
- Developing high-speed **data processing and analysis software**, and extensively applying novel mathematical, algorithmic, and **AI/ML** methods prioritizing high-energy, high-brightness, and coherent x-rays

https://www.aps.anl.gov/files/APS-Uploads/XSD/XSD-Strategic-Plans/APSScientificComputingStrategy-2021-09-24-FINAL.pdf

Network Architecture

Updates underway to support APS-U Era data and computing requirements

- 3-tier network infrastructure: facility, sector, hutch
- Supervisory Control and Data Acquisition (SCADA) architecture to better support controls, data, and regular network traffic
- Installed a new fiber plant for all APS beamlines; 768 pairs of new single mode fiber from the APS computer room to beamline networks
- Installed new core network switches capable of 100 Gbps links
- Procuring new sector and hutch switches for APS sectors capable of 100 Gbps links
- Recently upgraded the APS <-> ALCF network connection to 200 Gbps; upgrade to a terabit/s network in the future
- Adding wireless access points inside hutches and installing CAT 6A 10 Gbps copper cable at beamlines

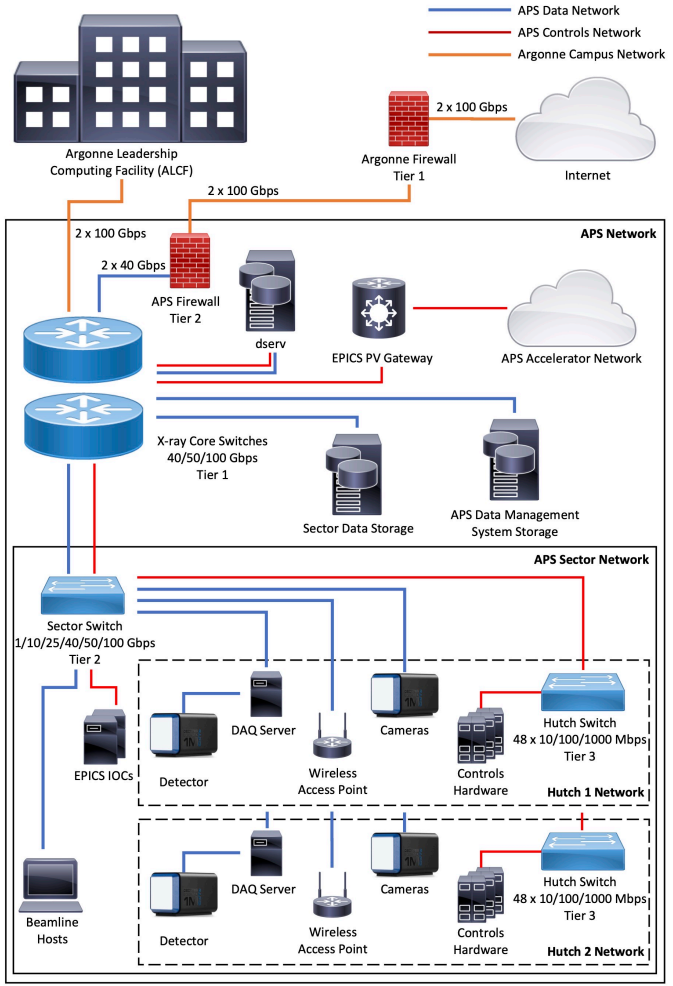

Beamline Controls & Data Acquisition

Fast Scanning and Feedback

- New instruments, especially the nano- and micro-probe scanning instruments, require fast scanning and feedback capabilities
- APS developed *softGlue Zynq* provides configurable FPGA capabilities for data acquisition and feedback
- Standardizing on D-TACQ FPGA devices

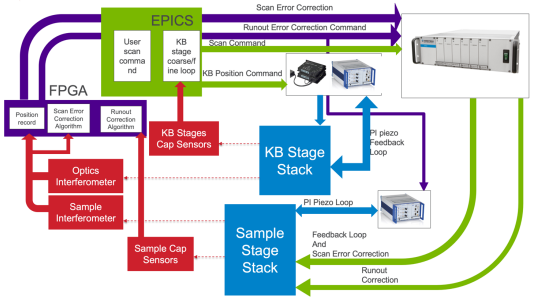

Schematic of fast scanning and interferometry on the APS Upgrade InSitu Nanoprobe Feature beamline

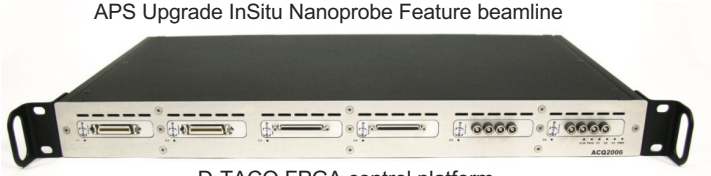

D-TACQ FPGA control platform

Fast Detector Data Handling

- New high data-rate detectors can generate thousands of frames per second using tens of Gbps of bandwidth
- EPICS V7, areaDetector, and newly developed *pvaPy* streaming data framework can process data at thousands of frames per second
- Integration with the APS Data Management System and Globus infrastructure

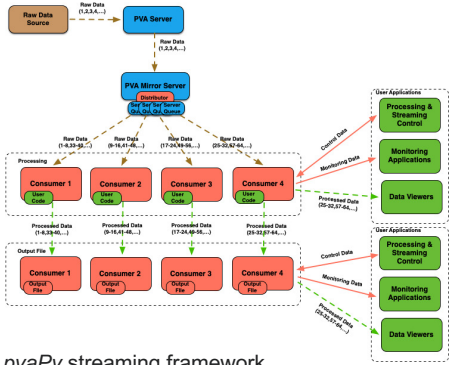

pvaPy streaming framework

Real-time AI-assisted ptychography reconstruction using the pvaPy framework

https://github.com/epics-base/pvaPy/blob/master/documentation/streamingFramework.md

APS Beamline Data Pipeline Project Team: Henke, S., Lang, K., Bicer, T., Guruswamy, T., Jemian, P., Leibfritz, D., Sersted, R., Babu, A., V., Veseli, S.

bluesky Experiment Control System

Leveraging *bluesky* as the future experiment control system

Goal: Deploy bluesky for use on APS-U feature beamlines

Deployed beam

- 2-BM, 4-ID-D,

Beginning depl

- 2-ID-D, 9-BM,

Light source community collaborations

- The APS benefits extensively from NSLS-II's effort
- Collaborations formed across the 5 US light sources

Refer to talk #69

Bluesky at the Advanced Photon Source

Presenter: Pete Jemian

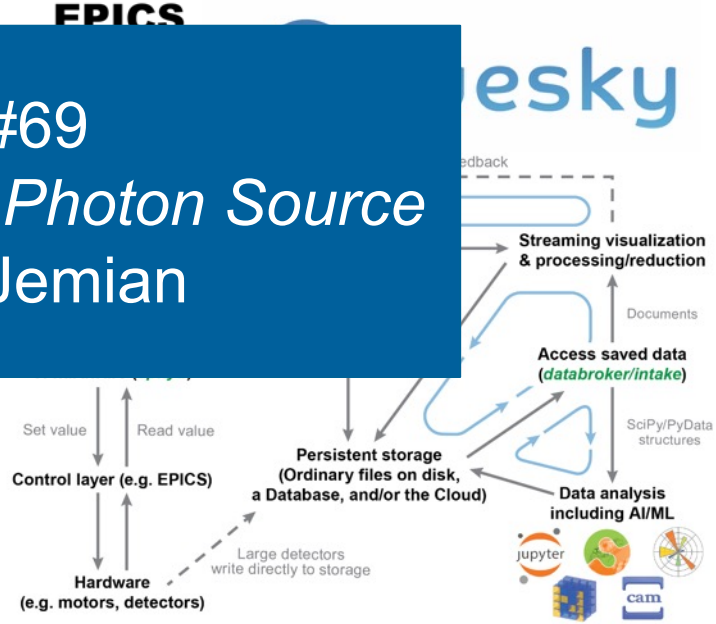

The APS Data Management System

Facility-wide software and hardware system for managing data and computing workflows

- Tools to automate transfer of data between acquisition devices, computing resources, and data storage sys
- Wo pro
- Ow use
- Users can download data at their home institutions using Globus Transfer (globus.org) or SFTP
- Integrated at > 50 APS beamlines (XSD and CAT)

Refer to talk #73

Managing Experiment Data with Ease at the Advanced Photon Source

Presenter: Hannah Parraga

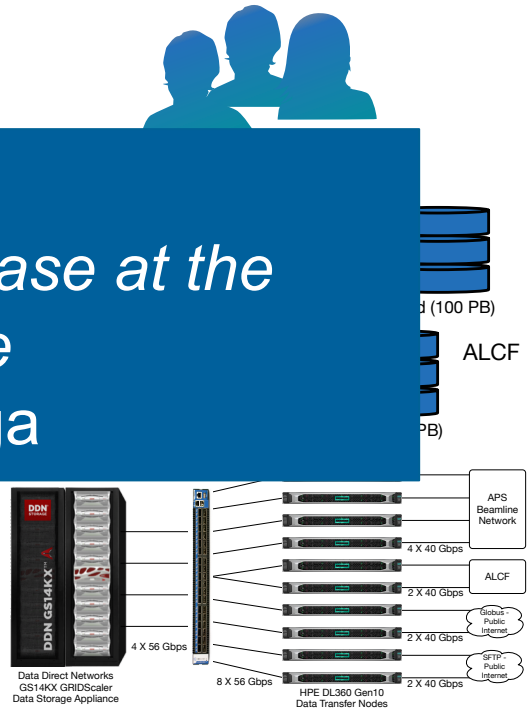

Veseli, S., Schwarz, N., Schmitz, C., "APS Data Management System," J. Synchrotron Rad. 25, 1574-1580 (2018).

Computing Resources

Multi-tiered approach spanning local and remote resources

Local compute resources

- Perform pre-analysis/data reduction (including compression and running ML models) to a form that allows quality control and experiment steering
- This may include, for example, a GPU workstation at a beamline, or the APS computing cluster

High-end compute resources

- Large data processing tasks, ML training, post-processing, and data refinement
- The APS has facility allocations at the Argonne Leadership Computing Facility
- The Argonne Leadership Computing Facility now provides a resource allocation queue/policy that suites APS job size and frequency profiles

Advanced Photon Source (APS)

Orthros – General purpose distributed-memory compute cluster
~27 TFLOP/s CPU cores
Sayre – Single node GPU system for Bragg CDI reconstructions
~111 TFLOP/s
5 x Ti 2080 | 2 x P100 | 1 x Ti 1080 | 1 x Quadro RTX 8000 GPUs
Axinite – Single node GPU system for CSSI and XPCS data processing
~155 TFLOP/s
4 x A6000 GPUs
Monas – 4 node GPU cluster for ptychography reconstructions
~430 TFLOP/s
8 x Ti 2080 GPUs per node

Argonne Laboratory Computing Resource Center (LCRC)

BeBop
~1,750 TFLOP/s
43,344 Intel Broadwell cores | 65,536 Intel Phi cores
Swing
~925 TFLOP/s
48 NVIDIA A100s | 768 AMD EPYC cores
Blues
~198 TFLOP/s
6,000 compute cores

Argonne Leadership Computing Facility (ALCF)

Theta & Theta GPU
Theta: 281,088 Intel Phi cores (~11.3 PFLOP/s)
Theta GPU: 192 NVIDIA A100s

Polaris
~44 PFLOP/s (~4 PFLOPS/s for exploring use by experimental facilities)

Aurora
Anticipated 2023
Intel CPUs / GPUs
> 1 EXAFLOP/s

Polaris – Instrument 2 Edge

Tightly coupling APS instruments with supercomputers to accelerate scientific discovery

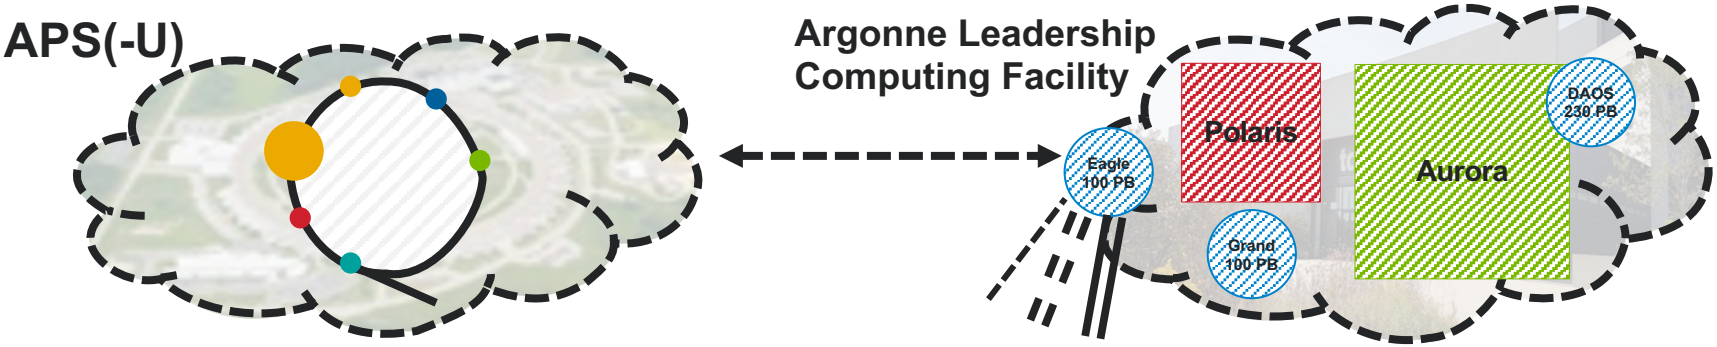

Polaris:

- The latest Argonne supercomputer
- Hybrid commodity CPU / GPU system
- ~44 PFLOP/s peak performance
- ~4 PFLOP/s procured to prototype on-demand use by experimental and observational facilities like the APS

Work Currently Underway:

- Test pre-emptive scheduling queues to provide immediate, on-demand access
- Deploy gateway nodes to allow APS detectors to stream directly to Polaris, avoiding local file I/O
- Developing end-to-end data workflows to connect APS instruments to Polaris, starting with ptychography, high-energy diffraction microscopy, and AI/ML methods

Use of the New Argonne Polaris Supercomputer During APS Beamtime

Integration of APS 8-ID-I XPCS operations with supercomputing resources

A team comprising of staff at the Advanced Photon Source (APS), the Argonne Leadership Computing Facility (ALCF), and the Data Science & Learning (DSL) division have successfully demonstrated the first use of the new ALCF Polaris supercomputer for APS data processing.

This work lays the path to utilization of this new system, and the soon to be delivered Aurora supercomputer, for routine use at APS instruments for facile real-time and post-experiment data processing and analysis by both APS staff and APS Users (on-site and remotely).

- APS developed high-performance computing software, XPCS-Boost, is deployed on the Polaris supercomputer for data processing
- The APS Data Management System is integrated with Globus Gladier/FuncX workflow tools to provide a single end-to-end data pipeline
- A convenient web-based data portal enables staff and users (on-site and remote) to view data as it's acquired and processed in near real-time

Artist rendering of the new Polaris supercomputer.

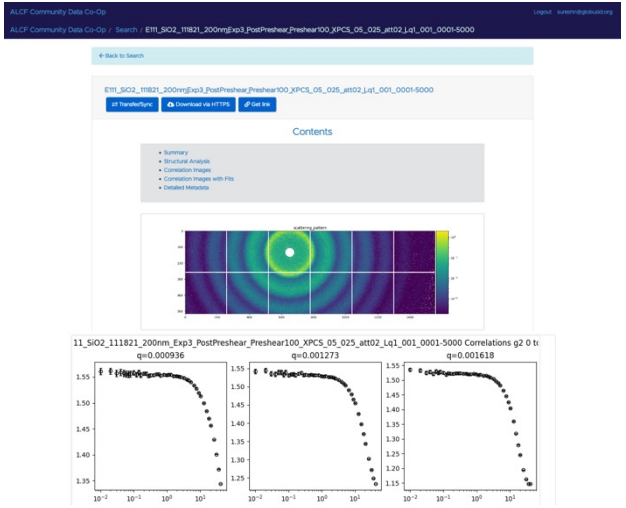

Web-based data portal accessible to on-site and remote users showing processed XPCS data from APS 8-ID beamline.

High-Performance Software for APS-U Era Needs

cohere
Bragg Coherent Diffraction Imaging (BCDI) reconstruction

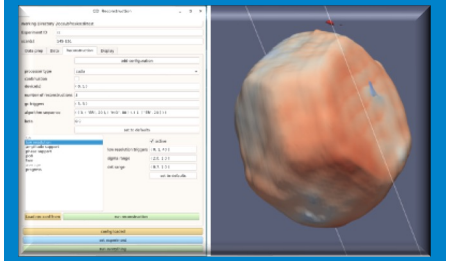

MIDAS
Microstructure for Diffraction Analysis Software

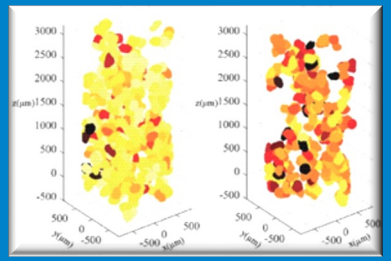

RSMap3D
Rapid Reciprocal-Space Mapping

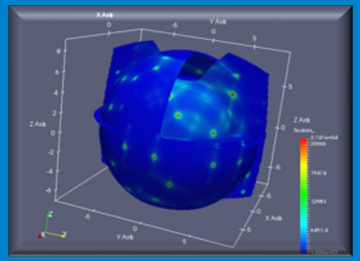

tike
Comprehensive Ptychographic Reconstruction Toolkit

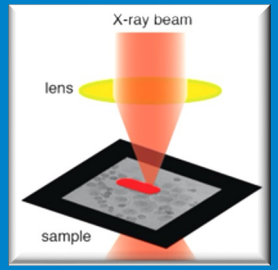

TomoPy
Tomographic Reconstruction Software

TomoPy
Tomographic reconstruction in Python

XPCS
Scalable X-ray Photon Correlation Data Reduction Software

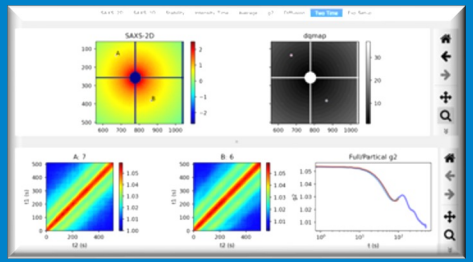

XRF-Maps
Scalable X-ray Fluorescence Elemental Mapping Tools

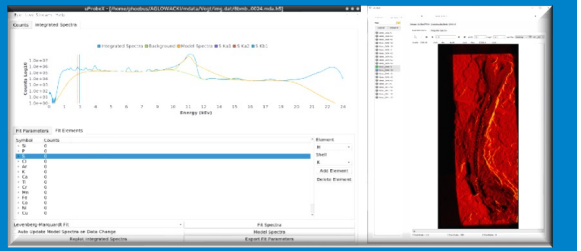

https://www.aps.anl.gov/Science/Scientific-Software
https://github.com/AdvancedPhotonSource

AI/ML Enabled Science at the APS

Data Reduction

PtychoNN: Machine learning ptychography reconstruction

- 100s of times faster and requires up to 5 times less data than conventional iterative approaches

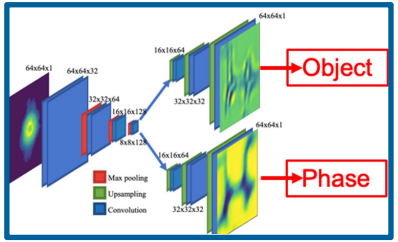

BraggNN: Machine learning method for determining Bragg peak locations from far-field high-energy diffraction microscopy data

- >200 times faster than conventional pseudo-Voigt profiling approach

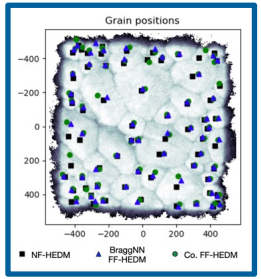

Cherukara, M., Zhou, T., Nashed, Y., Enfedaque, P., Hexemer, A., Harder, R.J., Holt, M. V., "AI-enabled high-resolution scanning coherent diffraction imaging," Applied Physics Letters 117, 044103 (2020).

Liu, Z., Sharma, H., Park, J. S., Kenesei, P., Miceli, A., Almer, J., Kettimuthu, R., Foster, I., "BraggNN: fast X-ray Bragg peak analysis using deep learning," IUCrJ 9, 104-113 (2022).

Experiment Steering

Smart Data Acquisition: Machine learning optimizes acquisition scanning path in real-time

- Using Google OR-Tools, AI calculates next points to scan
- Motor movement is reduced by 80%

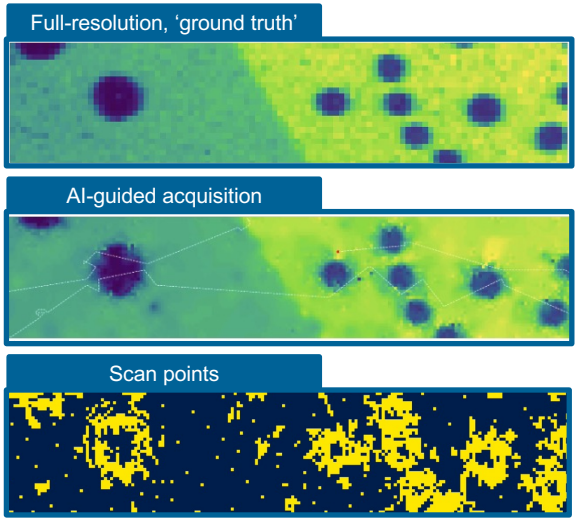

Contacts: Saugat Kandel, Tao Zhou, CD Phatak, et al.

Zhang, Y., Godaliyadda, G. M., Ferrier, N., Gulsoy, E. B., Bouman, C. A., & Phatak, C., "SLADS-Net: supervised learning approach for dynamic sampling using deep neural networks," Electronic Imaging, 2018(15), 131-1.

Knowledge Extraction

TomoGAN: Generative adversarial network improves the quality of tomographic reconstructions

- Uses up to 1/16th less dose or projections

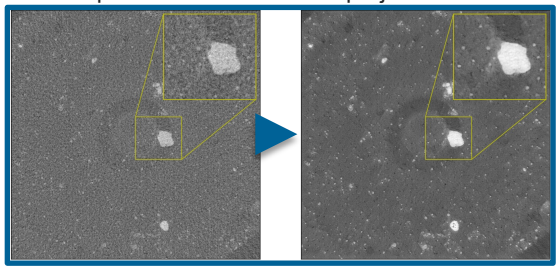

Generating interatomic potentials: Unsupervised machine learning generated interatomic potentials for a refractory oxide

- Diffraction measurements initialize an active-learner that iteratively improves an ML model

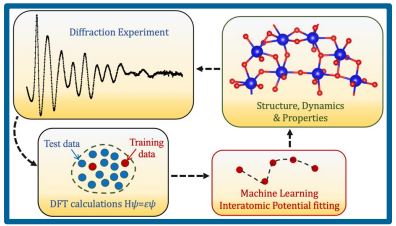

Liu, Z., Bicer, T., Kettimuthu, R., Gursoy, D., De Carlo, F. and Foster, I., "TomoGAN: low-dose synchrotron x-ray tomography with generative adversarial networks: discussion," JOSA A, 37(3), pp.422-434 (2020).

Sivaraman, G., Gallington, L., Krishnamoorthy, A. N., Stan, M., Csányi, G., Vázquez-Mayagoitia, Á., Benmore, C. J., "Experimentally driven automated machine-learned interatomic potential for a refractory oxide," Physical Review Letters, 126(15), 156002 (2021).

Argonne NATIONAL LABORATORY

AI@Edge Enables Real-Time Ptychography

Train AI @ ALCF, deploy AI @ beamline

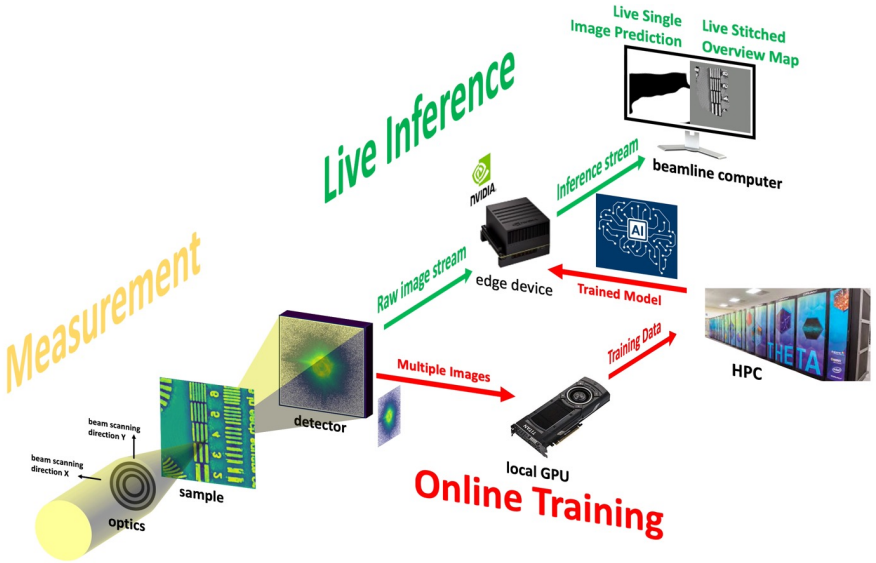

Live detector images

Cumulative image

Live inference

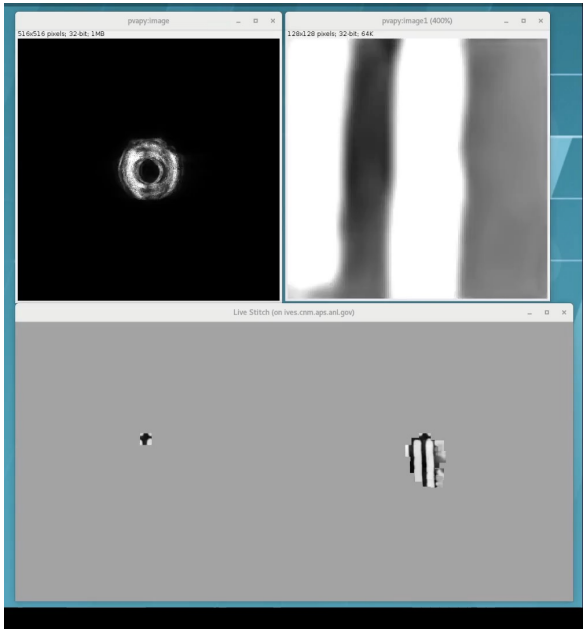

Computationally intensive phase retrieval methods are replaced with neural network models (PtychoNN) that learn to invert raw coherent imaging data to sample amplitude and phase in a single shot

- Neural network models are trained on many GPUs at the Argonne Leadership Computing Facility
- Trained models are deployed on an edge computing device (NVIDIA Jetson)
- The edge device makes live predictions from the detector data stream

Babu, A. V., Zhou, T., Kandel, S., Bicer, T., Liu, Z., Judge, W., Ching, D., Jiang, Y., Veseli, S., Henke, S., Chard, R., Yao, Y., Sirazitdinova, E., Gupta, G., Holt, M. V., Miceli, A., Cherukara, M. J., "Deep learning at the edge enables real-time, streaming ptychography," Paper Forthcoming.

Light Source Data and Computing Steering Committee (LSDCSC)

Steering committee across the 5 US Department of Energy, Office of Science, Basic Energy Sciences light sources

Advanced Light Source (ALS) / Lawrence Berkeley National Laboratory

Advanced Photon Source (APS) / Argonne National Laboratory

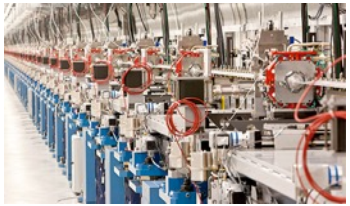

Linac Coherent Light Source (LCLS) / SLAC National Accelerator Laboratory

National Synchrotron Light Source II (NSLS-II) / Brookhaven National Laboratory

Stanford Synchrotron Radiation Lightsource (SSRL) / SLAC National Accelerator Laboratory

- Develop and maintain a strategic plan in the area of computing and data, and assist in the coordination and execution of related work
- Developed reports and guidance on data and computing needs, and a common data policy template
- Coordinated on proposals for the AI/ML awards
- Currently working on joint projects and coordinating future directions

LSDCSC Members

Nicholas Schwarz (Chair) – APS, ANL
Stuart Campbell – NSLS-II, BNL
Alexander Hexemer – ALS, LBNL
Vivek Thampy – SSRL, SLAC
Jana Thayer – LCLS, SLAC

Summary

- Efforts in scientific computing at the APS are driven by the needs of the future facility, and strategic plans and are aligned with the science thrusts of the upgraded facility
- Track record of successful developments aligned with future needs
- Fruitful collaborations contribute to scientific computing needs at many levels

Next steps:

- Continue implementing and refining the APS Scientific Computing Strategy
- Further develop collaborations related to scientific computing
- Continue engaging staff and users regarding APS scientific computing

Acknowledgements

- Anakha Babu
- Antonino Miceli
- Arthur Glowacki
- Arvind Ramanathan
- Ashish Tripathi
- Barbara Frosik
- Brian Toby
- Brian Robinson
- Daniel Ching
- David Leibfritz
- David Wallis
- Deming Shu
- Denjamin Blaiszik
- Dennis Trujillo
- Doga Gursoy
- Fang Zhang
- Felix Lacap
- Francesco De Carlo
- Franck Cappello
- Hannah Parraga
- Hemant Sharma
- Henry Smith
- Howard Yanxon
- Ian Foster
- Ignacia Guerra
- Jeff Kirchman
- Jini Ramprakash
- John Weizeorick
- John Hammonds
- John O'Connell
- Jonathan Tischler
- Jonathan Almer
- Joseph Sullivan
- Justin Wozniak
- Keenan Lang
- Kevin Peterson
- Kurt Goetze
- Kyle Chard
- Lahsen Assoufid
- Luca Rebuffi
- Mark Engbretson
- Mathew Cherukara
- Max Wyman
- Miaoqi Chu
- Michael Papka
- Michel Van Veenendaal
- Mike Hammer
- Nicola Ferrier
- Orlando Quaranta
- Pete Jemian
- Pete Beckman
- Prasanna Balaprakash
- Rajkumar Kettimuthu
- Roger Sersted
- Ryan Milner
- Salman Habib
- Sebastian Strempfer
- Sinisa Veseli
- Skip Reddy
- Stefan Wild
- Steven Henke
- Suresh Narayanan
- Sven Leyffer
- Tejas Guruswamy
- Tekin Bicer
- Thomas Uram
- Ti Leggett
- Tim Mooney
- Todd Munson
- Tomas Walsh
- Valerie Taylor
- William Allcock
- Xianbo Shi
- Xuli Wu
- Yudong Yao
- Zichao Di
- And many others…

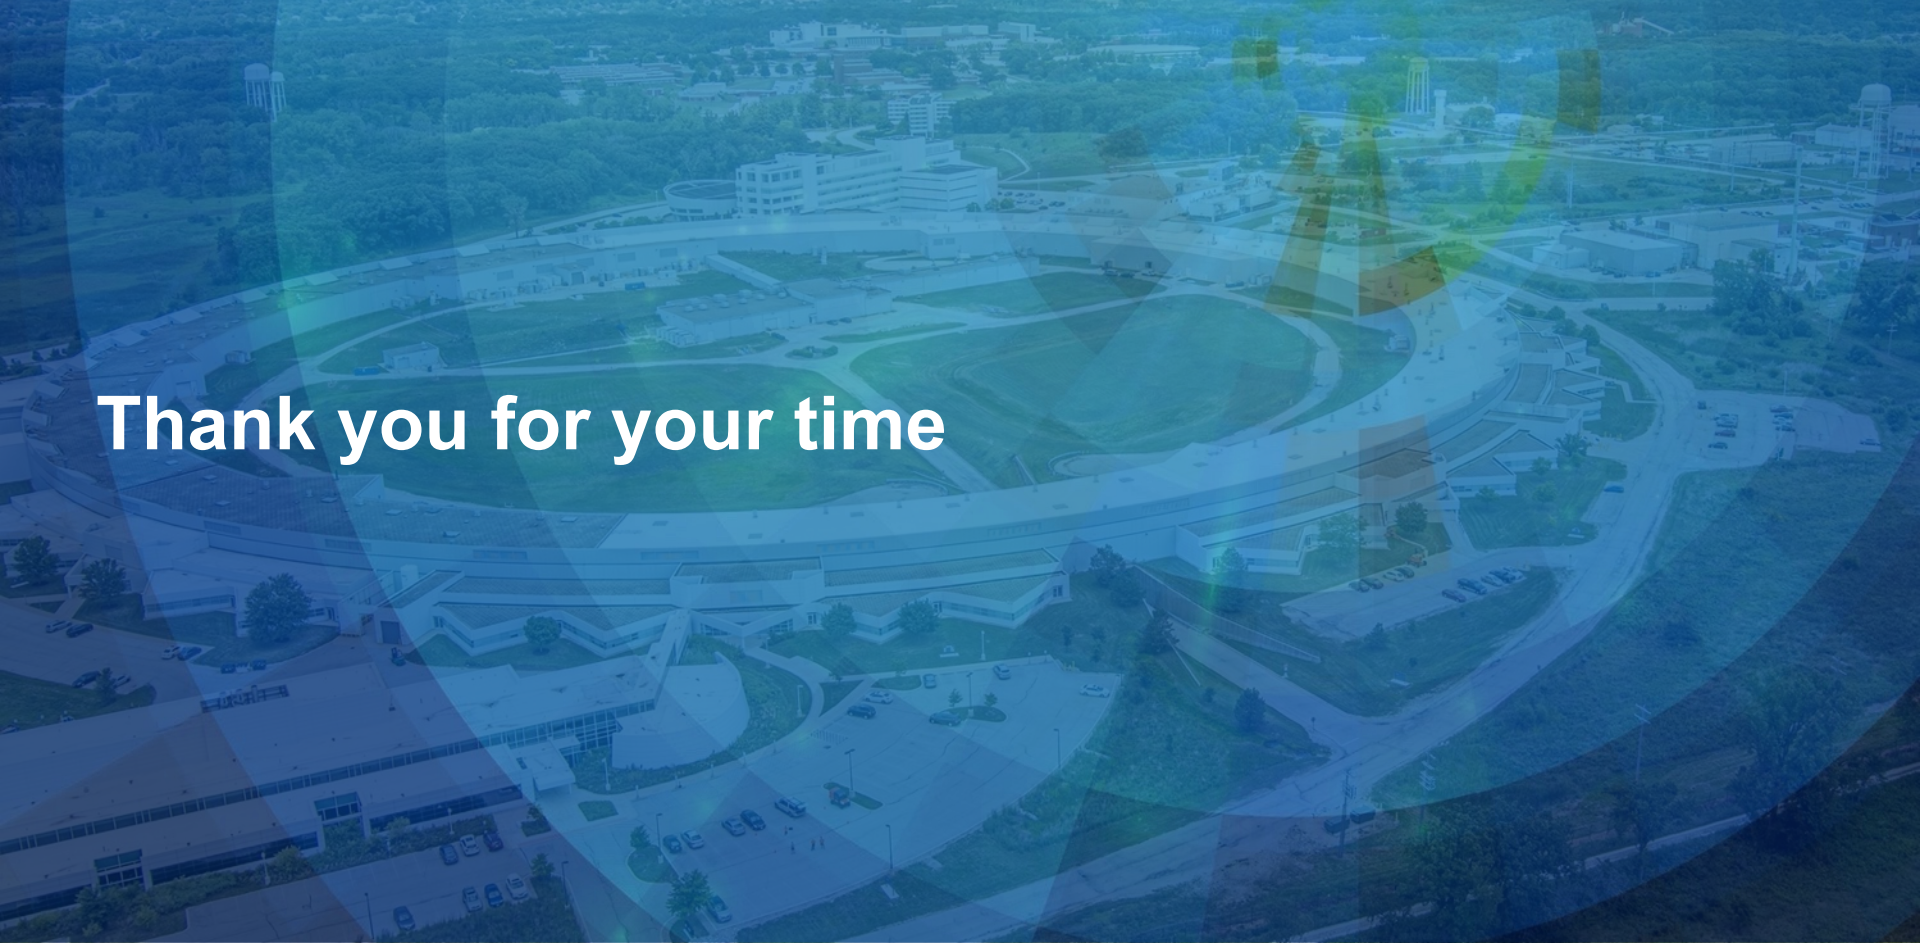

Thank you for your time