APS UPGRADE FORUM

APS SCIENTIFIC COMPUTING STRATEGY

BRIAN TOBY
Chief Computational Scientist, Senior Physicist, Group Leader
Computational X-ray Science (CXS)

NICHOLAS SCHWARZ
Principal Computer Scientist, Group Leader
Scientific Software Engineering & Data Management (SDM)

28 January 2016

OVERVIEW

Drivers & Challenges

Current State

Strategy

Notable Accomplishments

Next Steps

DRIVERS & CHALLENGES

DRIVERS & CHALLENGES

New scientific opportunities drive demands for increased computing at the APS

These opportunities are/will be enabled by

- new measurement techniques
- technological advances in detectors
- multi-modal data utilization
- advances in data analysis algorithms
- and, of course, all the benefits of the APSU

And they create demand for sophisticated and large-scale

- computational resources (CPUs, storage, and networking)
- algorithms and scientific software

MAPPING OF MESO- AND MICRO-STRUCTURE IN ENGINEERED MATERIALS

Goal: understand function and how to enhance properties of next generation atomically engineered materials

Applications: better high-performance alloys for aerospace, defense, fuel-efficient transport

Measurements: simultaneous high-energy diffraction microscopy (HEDM), ptychography, high-energy micro-tomography and coherent diffraction imaging (CDI)

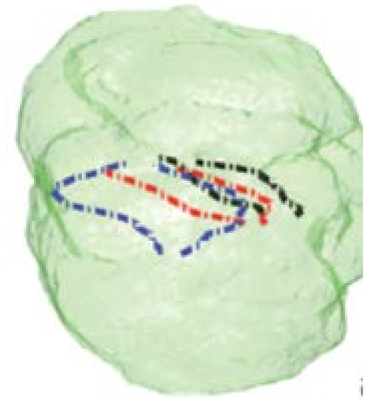

Challenges:

- Big data transport and management (>terabytes/hour)
- Advanced algorithms: image reconstruction with internally consistent experimental artifact correction
- Just-in-time high performance computing to provide real-time feedback to experimenters
- User friendly software. Professionalize existing and new code into a supported package that runs on workstations and HPC.

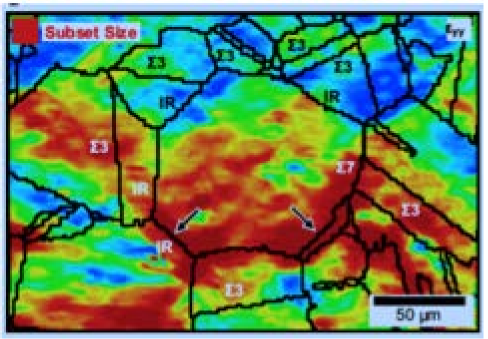

RAVEN: NONDESTRUCTIVE 3D IMAGING OF INTEGRATED CIRCUITS

Chris Jacobsen, et al.

Goal: Image all layers in a 1.5 mm^3 device with < 10 nm resolution (USC/Northwestern/APS proposal for RAVEN project in IARPA)

Applications: National defense & quality assurance

Measurements: ptychography, tomography & possibly elemental mapping

Challenges:

- Huge data volume transmission and management (> 1 petabyte/day!)
- Multimodal image registration algorithm development
- Large-scale HPC adaptation of codes
- Algorithm development for automated reconstruction optimization
- HPC data analysis integration into beamline workflows
- Device-library driven circuit analysis; petabyte-leve ldata interpretation and visualization

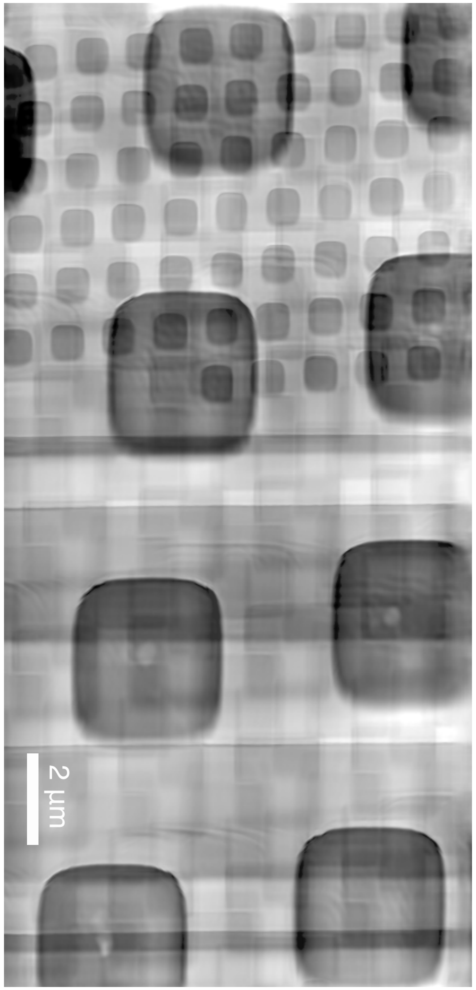

CHARACTERIZE ATOMISTIC DYNAMICS DURING PLASTIC DEFORMATION

Goal: Understand how amorphous polymers and metals fail due to external stress

Applications: New hard and crack-resistant materials for transportation and other high-performance environments

Measurements: XPCS mapping combined with CDI

Challenges:

- Algorithms combining dynamics from XPCS and structure from CDI
- Adaptive control of experiments to capture rare failure events
- Combined use of simulation techniques (metals: DFT or polymers: MD) with experimental data
- Large data volumes (10^6 frames/sec XPCS 10^3 CDI)

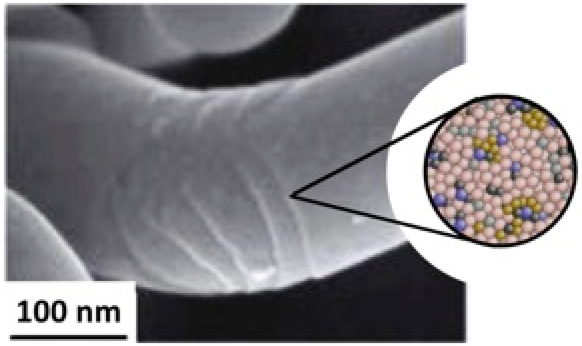

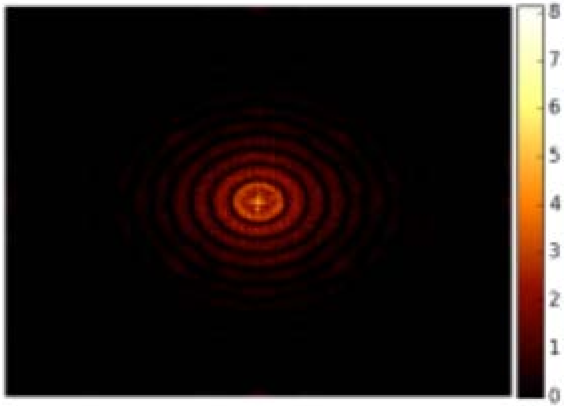

CURRENT STATE

Argonne
NATIONAL LABORATORY

THE APS SUPPORTS OVER 30 SOFTWARE PACKAGES

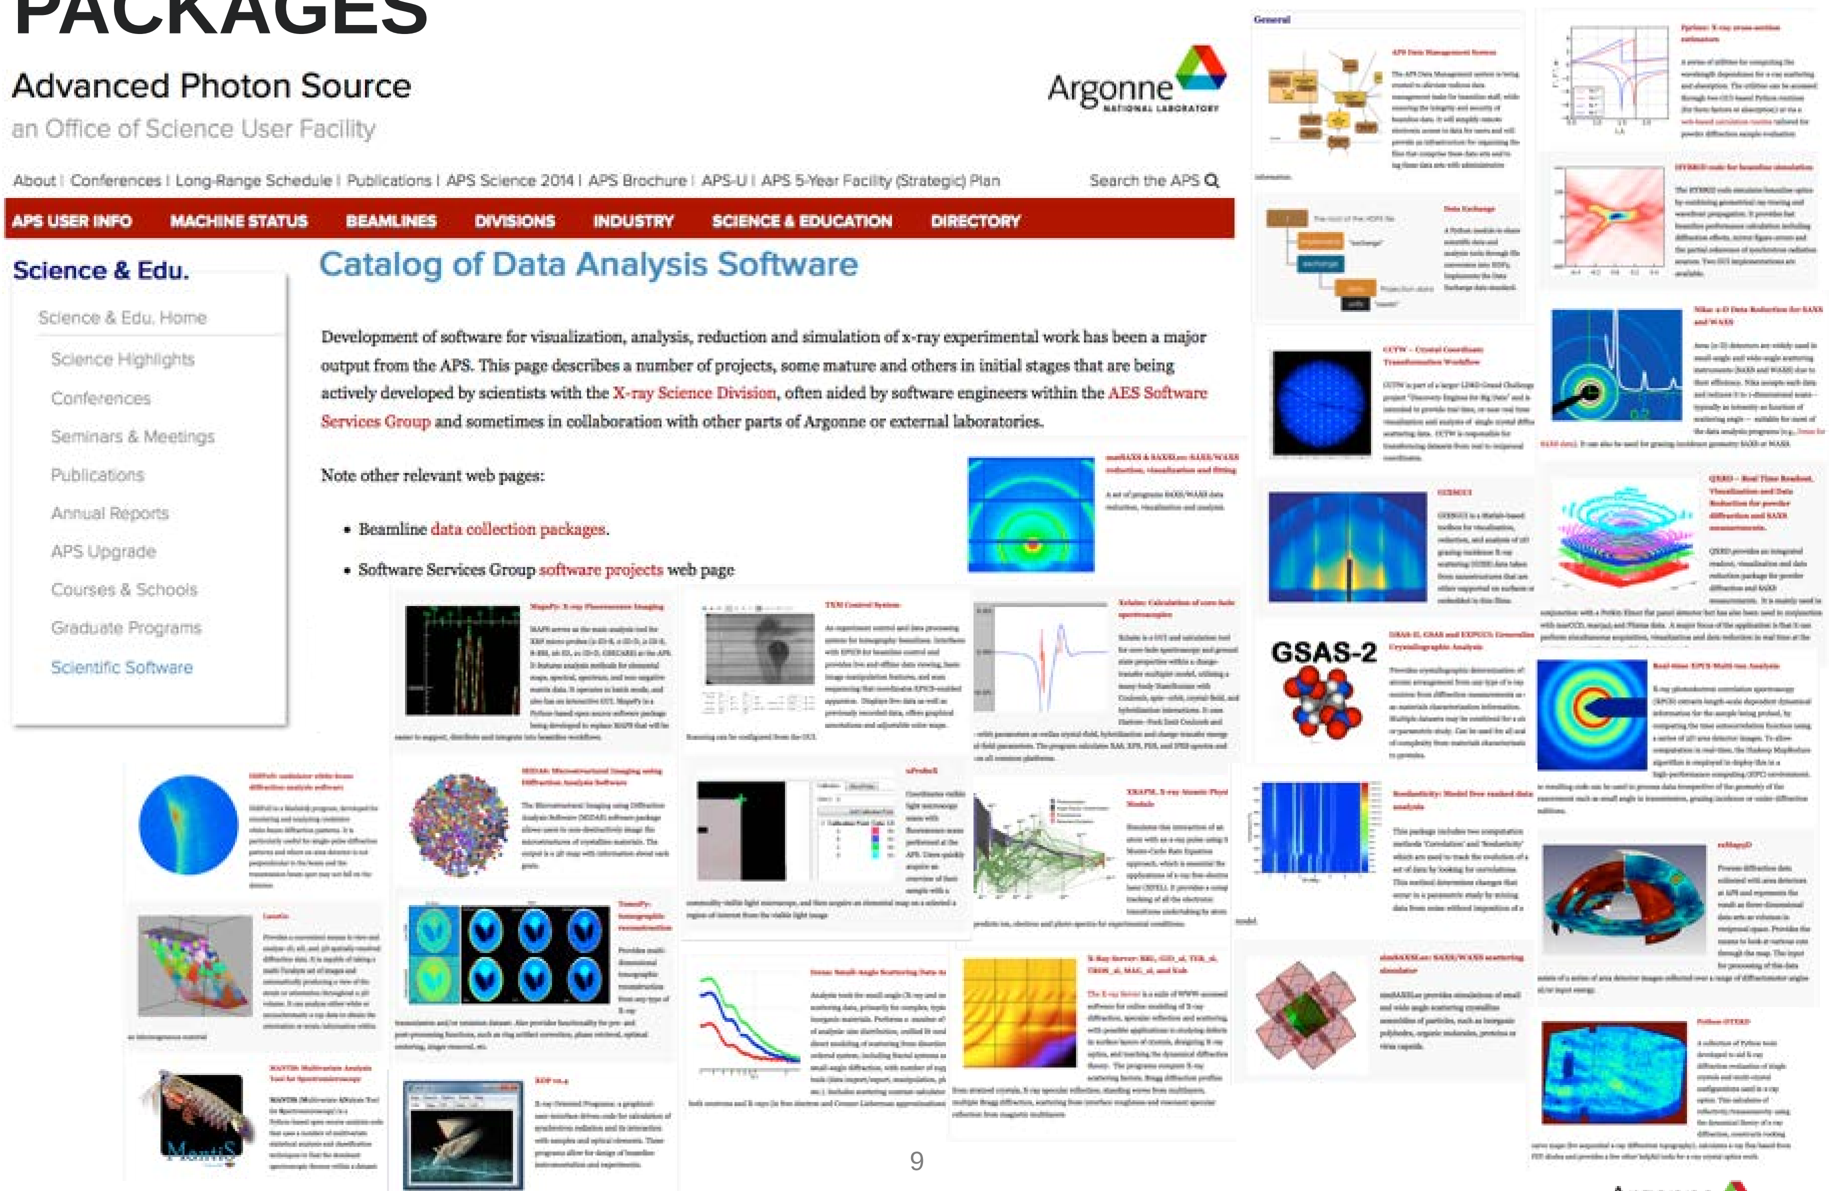

Argonne NATIONAL LABORATORY

SCIENTIFIC SOFTWARE

The APS has a wealth of beamline scientist-generated software

- Highly heterogeneous
- Few packages developed to professional software engineering standards
- Very little is supported to professional standards (regular releases; packaging; user support…)

➔*All XSD groups have unmet software needs – some very extensive*

➔*Migrating existing codes to HPC can be very difficult*

➔*Some of the most challenging/exciting scientific and computational challenges will require significant effort from teams of experts*

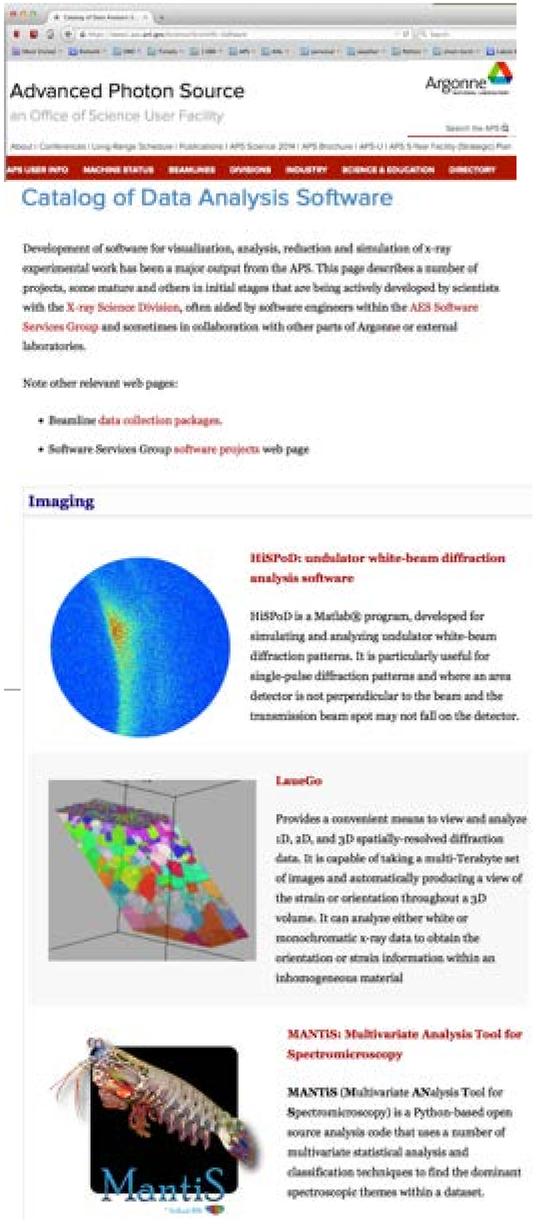

COMPUTING INFRASTRUCTURE

Today: very little beamline computing is done on centralized systems (ALCF & LCRC); most HPC use is in storage ring/accelerator simulation

- The APS has a small cluster (Orthros: 53 nodes, <1000 cores; ~250 Tb) serving 3 beamlines (and APSU).
  - Orthros nodes are configured for and allocated directly to beamlines and are idle when not required
- Significant additional computing done on beamline workstations.

➔*Computing demands are only going to grow, and are beyond the means of the APS to manage*

CHALLENGES OF SHARED RESOURCES

Scheduling: Much APS computing falls in the category of “weather forecasts” (no value if too late); reserving nodes will not scale

- Challenge to CLS: develop preemptive scheduling approaches for LCRC & ALCF

Latency: Must move large amounts of data

- Adaptation of ESNet DMZ concept intracampus to allow firewall-free transfers
- Longer range: algorithms for initial data reduction at beamlines, reducing network footprint

Support: APS runs 24/6; CLS runs best effort from 7 pm to 7 am

- Need to work out an out-of-hours support mechanism

STORAGE MANAGEMENT AND INFRASTRUCTURE

- The APS is currently producing ~2 Petabytes of data/year
- Data from XSD beamlines must be provided to users for archiving
- Beamline scientists need to stage, migrate, clean up and catalog large volumes of data
- Data typically stored in vendor-supplied image formats; lacks metadata & provenance
- Limited on-line storage

BEAMLINE OPERATION SOFTWARE

Most APS beamlines have developed their own versions of user interface/controls

- Idiosyncratic; often not “user-friendly”
- Beamlines need combination of GUIs and scripting; software is continually being adapted

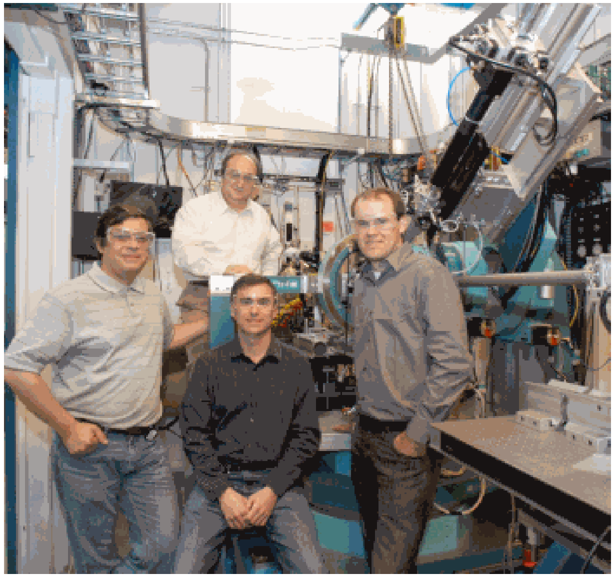

Argonne NATIONAL LABORATORY

STRATEGY

APS SCIENTIFIC COMPUTING STRATEGY

Scientific drivers coupled with technological advances make scientific computing a priority for the facility.

- All techniques at the APS will benefit from advances in this area.
- It’s important to note that without the application of advances in this area, we won’t be able to fully realize the scientific benefits of the APSU.

Participative Process:

- Current landscape
- Technique-by-technique assessment (short- and mid-term needs / overlaps)
- Anticipated technological advances
- APSU first experiments (long-term)

APS SCIENTIFIC COMPUTING STRATEGY

Facility-wide strategy to meet these demands: *APS Scientific Computing Strategy*.

- Prioritization
- Strengthening / aligning internal resources
- Collaboration

Address Four Areas:

1. Scientific & Data Analysis Software
2. Data Management & Distribution
3. Computing Infrastructure
4. Beamline Operation Software

Do all of this with foreseeable budgets.

PRIORITIZATION

Prioritize what's most important to the APS (we can't do it all):

- Primary focus is on APSU enabled and improved techniques
- Improvements that have a facility-wide impact
- Support existing / successful user programs
- Consider initial and ongoing costs, and funding sources

Participative, iterative, *living* process – regular review and discussion

Appendix A – Projects and Priorities

Table A.1 Scientific Software & Data Analysis

Project	Completion	Funding	Priority	Details	Effort	Ongoing Effort / Year
HPC enabled real-time correlation toolkit	2017	APS Operations / CLS	Mission Critical	Development of an HPC enabled application for real-time photon correlation analysis of time-resolved datasets (e.g. XPCS, surface-XPCS, etc.).	2 FTE	0.25 FTE
HPC enabled real-time XAS analysis toolkit	2018	APS Operations	Mission Critical	Development of an HPC enabled application for real-time analysis of x-ray absorption spectroscopy datasets (XANES, EXAFS, XFM).	3 FTE	0.25 FTE
Ultrafast time-resolved imaging with large-scale MD modeling	2018	LDRD	Mission Critical	Integrated ultrafast time-resolved imaging with large-scale molecular dynamics modeling for in situ data analysis and visualization of energy transport.	2 FTE	
Multimodal imaging of materials for energy storage	2018	LDRD	Mission Critical	Integration of multimodal data from x-ray and electron microscopies in order to understand the interaction of materials at multiple length scales from nano to micro.	2 FTE	
HPC enabled real-time CDI analysis toolkit	2019	APS Operations	Mission Critical	Development of an HPC enabled application for real-time x-ray coherent diffraction imaging datasets (CDI, Ptychography).	2 FTE	0.25 FTE
HPC enabled real-time x-ray scattering analysis toolkit	2019	APS Operations	Mission Critical	Development of an HPC enabled application for real-time reciprocal space analysis of x-ray scattering datasets (SAXS, WAXS, HEDM, XPD).	2 FTE	0.25 FTE
Microstructural Imaging using Diffraction Analysis Software (MIDAS)	Ongoing	APS Operations / Industrial Partner / AFRL	Mission Critical	Development and maintenance of a suite of analysis tools for different modalities (near-field, far-field, and very-far-field) of high energy diffraction microscopy (HEDM)		1 FTE
TomoPy	Ongoing	APS Operations	Mission Critical	Package, maintain, and support the TomoPy reconstruction toolkit.		0.5 FTE
General purpose analysis workflow toolkit	Ongoing	APS Operations / CLS	Mission Critical	Leverage best-in-class workflow tools for use at APS beamlines.	2 FTE	0.5 FTE
General purpose	Ongoing	APS	Mission	Development of algorithms for	2 FTE	0.5 FTE

Argonne NATIONAL LABORATORY

STRENGTHEN / ALIGN INTERNAL RESOURCES

Two groups dedicated to scientific software:

- Computational X-ray Science (CXS)
  - 3.5 FTE
  - Algorithm development and external collaborator coordination
  - X-ray scientists and applied computational scientists
- Scientific Software Engineering & Data Management (SDM)
  - 8 FTE
  - Software creation and maintenance, HPC development, and data management
  - Software engineers

- Aligned within a cohort of support groups, including Beamline Controls, Beamline Instrumentation, Detectors, and Optics: Build partnerships with scientists
- *Very* modest staffing increases

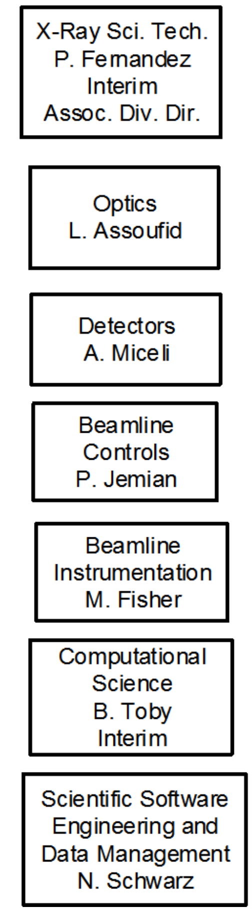

COLLABORATION

Leverage the broader community for expertise and effort. We won't achieve all that we need with only what we have internally.

ANL – ALCF/CLS/MCS

- The APS has the unique advantage of world-class computer science and leadership hardware on-site

BES Facilities Computing Working Group

- Facility Directors' call to develop collaborations
- Avenue to collaborate with other BES SUFs

CAMERA

- Three areas of interest/overlap with APS: Tomography, Ptychography, GISAXS

Scientific Staff & User Base

- Often the main drivers of innovation

SCIENTIFIC SOFTWARE

Algorithm development and software creation

- Aim to provide near real-time feedback that can inform experiment steering
- Interpretation and visualization of large and complex data streams that are beyond unaided human comprehension
- Combined simulations and multimodal data to extend understanding for currently intractable scientific problems

Focus on techniques in support of science enabled by the APSU

- High-energy, multi-modal, microscopy, fluorescence, scattering, coherence imaging, time-dependent

Approach

- Open source
- Graded
- Select high-impact APS packages will be maintained for broader community use (as distributable packages or as software-as-a-service (SaaS)), such as TomoPy

DATA MANAGEMENT & DISTRIBUTION

Expect an order of magnitude increase in data volumes and rates over the coming years, and at least two orders of magnitude increase post-APSU.

Facility-wide tools that provide:

- Automated data transfer (e.g. Globus Services)
- Data ownership
- Metadata cataloging
- Provenance tracking
- Electronic notebook (e.g. OLog)

- Apply best-in-class tools to APS beamline and workflows.
- Prioritizing beamlines based on highest current and anticipated data rates/volumes.
- Work closely with Argonne computing expertise.
- APS provides systems integration.

COMPUTING INFRASTRUCTURE

A multi-order of magnitude increase in required computing power is anticipated over the coming years.

Tiered Approach

- Beamline local resources used when sufficient
- APS operated compute cluster for current mid-range problems
- Shared (LCRC/ALCF/NERSC) resources for most demanding requirements

Continually review costs of internal vs. externally managed resources to ensure the most efficient use of resources.

BEAMLINE OPERATION

The APS wants to provide intuitive yet flexible beamline operation software in order to maximize beam time efficiency.

- Common core of beamline alignment, calibration and experiment operation tools
- Automation elements, e.g. alignment, mail-in, or robotic operation

Due to foreseeable budgets, the APS will *not* devote significant resources to improving beamline operations software for existing beamlines.

- Development projects to create new or improved beamlines (funded by the APSU or Operations) should properly budget beamline operation software as a part of the overall project.

NOTABLE ACCOMPLISHMENTS

PTYCHOGRAPHY RECONSTRUCTION

MCS: Youssef Nashed, Tom Peterka, Rob Ross
APS: Si Chen, Stefan Vogt, Chris Jacobsen, David J. Vine
Other: Junjing Deng, Nicholas W. Phillips

ans: as low
beam
on today,
n future

s: 100%
ilization

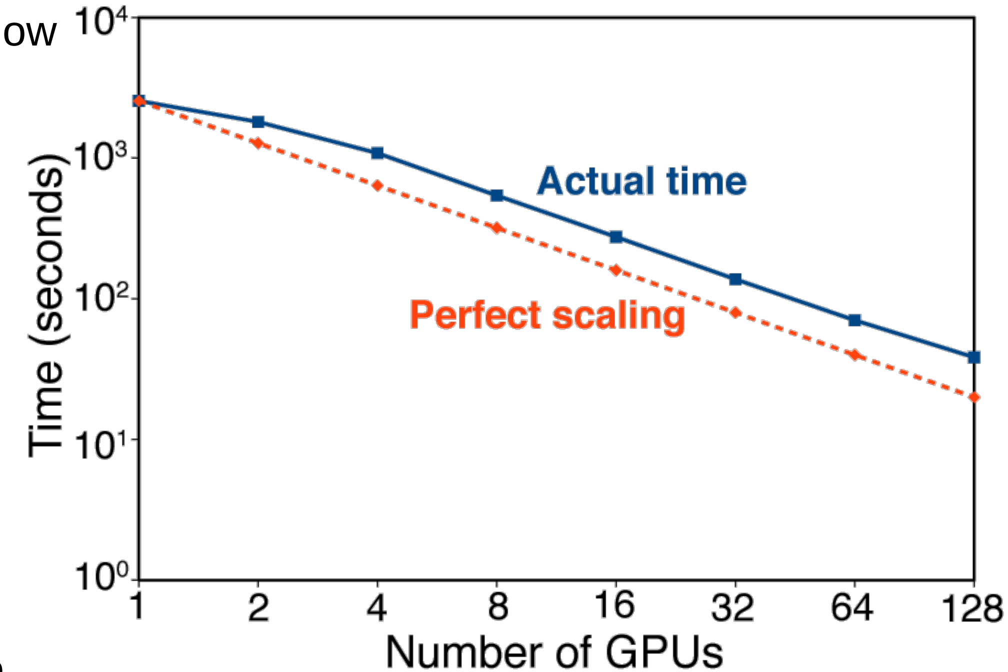

Reconstruction:

- speedup of 10^2 from CPU to GPU
- additional 10^2 speedup from many GPUs in parallel!

Nashed *et al.*, *Optics Express* **22**, 32082 (2014)

MAGELLAN CLOUD RESOURCE FOR XPCS

Magellan

- ANL cloud computing environment
- ~ 1,000 node resource
- Virtualized environment
- Maintained by CLS

XPCS

- XPCS multi-tau Hadoop system set up and running on 10 nodes
- Fully integrated into the XPCS workflow pipeline at 8-ID-I
- In production use this run

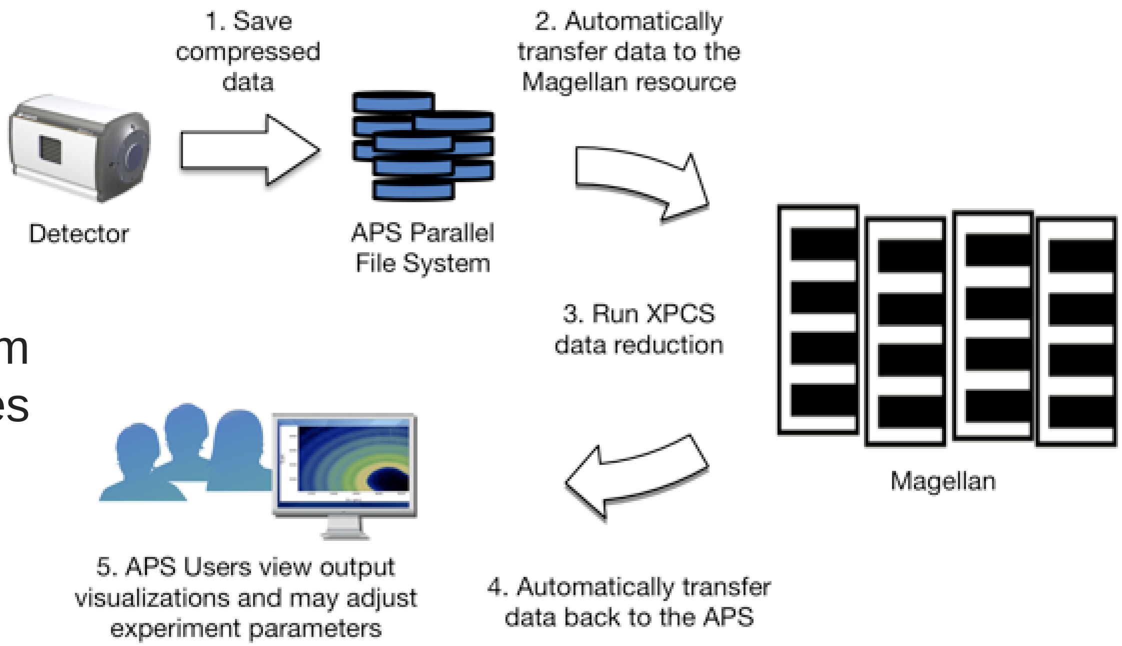

APS: Collin Schmitz, Ben Pausma, Faisal Khan, Suresh Narayanan
CLS: Ryan Aydelottt, Dan Murphy-Olson

PETREL DATA SERVICE

The APS is collaborating with Argonne’s CLS directorate and the Globus Services team to prototype a new data management and distribution service.

- High-speed, high-capacity data store
  - 1.7 PB total storage
  - 100 TB project allocations
- Project-focused, self-managed
- CLS provides the hardware infrastructure
- Globus Services provides the software
- APS provides integration

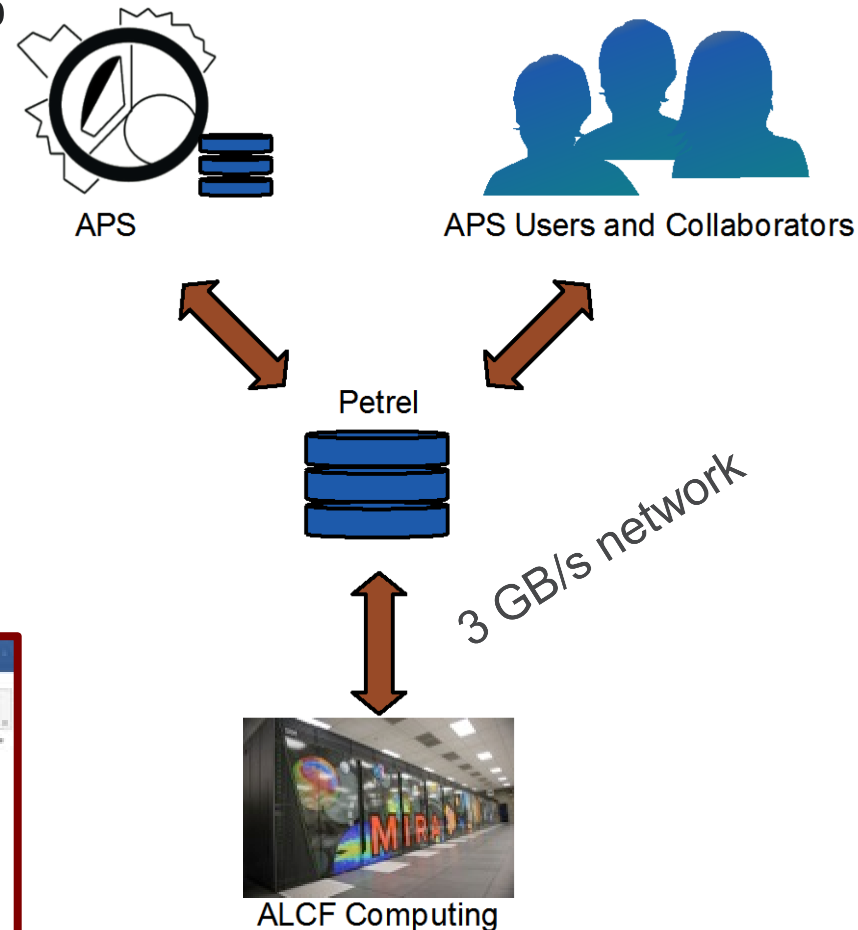

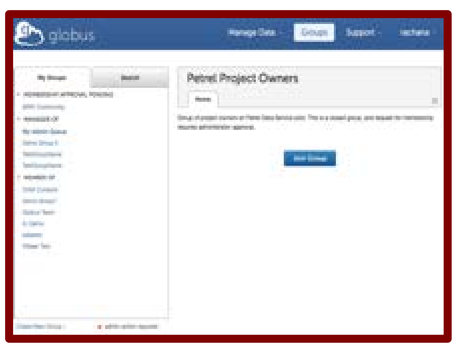

EXTREPID

- Repurposed storage from Blue Gene/P – Intrepid
- 1.7 PB
- Administered by APS
- Hosted by ALCF/CLS
- Online in the coming months

USE OF LARGE SCALE COMPUTING RESOURCES

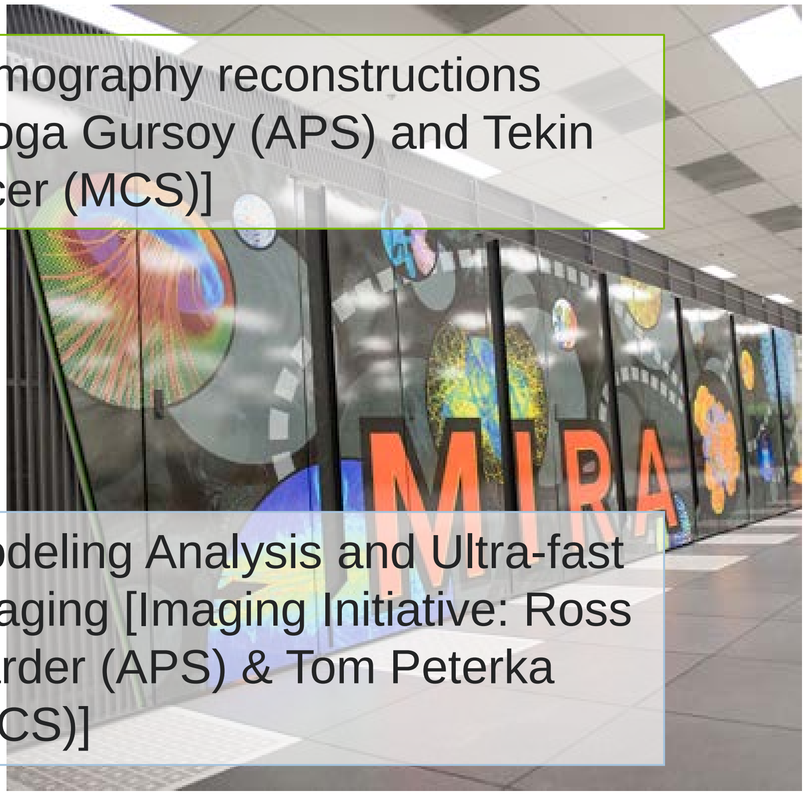

Tomography reconstructions [Doga Gursoy (APS) and Tekin Bicer (MCS)]

Modeling Analysis and Ultra-fast Imaging [Imaging Initiative: Ross Harder (APS) & Tom Peterka (MCS)]

ALCF

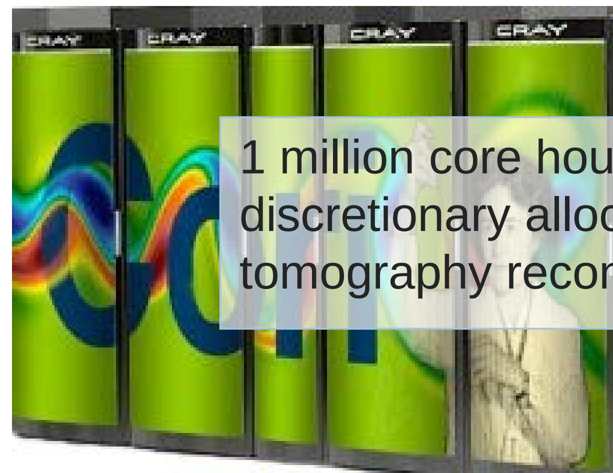

1 million core hour discretionary allocation for tomography reconstructions

NERSC

NEXT STEPS

NEXT STEPS

- Continue to catalog and enumerate all computing needs

- Use APSU's first experiments as a starting point:
- Start developing aligned 5-year and 10-year roadmaps
  - Machine timeline
  - Beamline/technique timeline
  - Detector/beamline hardware timeline
  - Computing hardware timeline
  - Software timeline
  - Algorithm/math timeline

NEXT STEPS

Computing is critical to the scientific success of the APS after the APSU project completes.

Our computing needs must be addressed before the APSU completes.

Strategy Document:

https://www1.aps.anl.gov/X-ray-Science-Division/XSD-Strategic-Thrusts

Thoughts and comments are very welcome!

WE START WITH YES.
AND END WITH THANK YOU.

DO YOU HAVE ANY BIG QUESTIONS?