APS Scientific Computing Strategy

19 June 2017

Introduction

The Advanced Photon Source (APS) at Argonne National Laboratory (ANL) is a U.S. Department of Energy (DOE) Office of Science-Basic Energy Sciences (BES) scientific user facility. The core mission of the APS is to serve the scientific community by providing experiment facilities utilizing x-rays (beamlines) to allow users to address the most important basic and applied research challenges facing our nation.

All aspects of APS operation depend on computation, but data analysis software and computing infrastructure are of particular importance for facility productivity. Demands for increased computing at the APS are driven by new scientific opportunities, which are enabled by new measurement techniques, technological advances in detectors, multi-modal data utilization, and advances in data analysis algorithms. The priority for the APS is to further improve our world-class programs that benefit most from high-energy, high-brightness and coherent x-rays. All of these require advanced computing. The revolutionized high-energy synchrotron facility proposed in the APS-U project will increase brightness and coherence, leading to further increases in data rates and experiment complexity, creating further demands for advanced scientific computation.

The APS and ANL are poised well to employ advanced computing to maintain a world-leading position in the synchrotron community. The APS has a world-class photon science program with a large and diverse user base, and ANL is home to world-leading supercomputing infrastructure and computer science expertise in the Computing, Environment, and Life Sciences directorate (CELS). This colocation provides an unprecedented opportunity for collaboration.

The APS has organized the core groups required to achieve these goals under the X-ray Science Technologies (XST) umbrella within the X-ray Science Division (XSD). The XSD Beamline Controls (BC) group is responsible for beamline data acquisition, through control and operations systems and software. The XSD Computational X-ray Science (CXS) group is mainly responsible for the development of theory, mathematical models, algorithms and software for interpreting x-ray measurements. The XSD Scientific Software Engineering & Data Management (SDM) group is responsible for software engineering for data analysis applications and data management tools, enabling high-performance computing (HPC). The management and support of information technology resources within the APS is handled by the APS Engineering Support (AES) division Information Technology (IT) and Information Solutions (IS) groups.

This document identifies the core challenges facing the APS in this area, and presents a plan for addressing these challenges over the coming years within the anticipated resource environment at the APS, and by leveraging synergies with other parts of ANL, across the BES and DOE complex, and with the outside community. Each group maintains strategies and work plans listed at the end of this document that describe activities in greater detail. This document will be reviewed and updated at least once a year.

Beamline Control and Operations Software

Powerful and up-to-date beamline control and operations software is key to realizing the potential of the APS. In order for the next generation of instruments in the APS-U era to help answer the most pressing scientific challenges, the facility must support the most advanced and flexible control and beamline operation systems. Despite the need for flexibility, employment of standardization and pre-existing solutions are required to take advantage of economies of scale.

The Experimental Physics and Industrial Control System (EPICS) is the primary control system for beamlines (and the accelerator) at the APS. Maintained by the APS, as part of a world-wide collaboration with contributions from many facilities and user groups, EPICS versions will continue to be updated regularly as new features become available, and the facility will adopt EPICS V4 and V7 where appropriate. The synApps package serves as the main collection of control system tools for running experiments at the APS, and the areaDetector package serves as the state-of-the-art EPICS interface for detector data acquisition. Both are used almost universally at the APS. The APS will maintain a similar update and release schedule for these components as it does for the core control system.

The current synoptic display software in use at the APS, MEDM, is nearing end-of-life. This class of software is critical for the stable operation of the facility. The APS has evaluated three possible successors for MEDM at beamlines: CSS-BOY, EPICSQt, and caQtDM, and has selected caQtDM as the future successor. Over the coming year, the APS will deploy caQtDM for use at all APS beamlines.

Mail-in automation has proven to catalyze very high levels of productivity by decreasing demands on beamline staff and by lowering the barriers for scientists who are not routine facility users. A standardized facility-wide mail-in support system would allow many more beamlines to offer this service where user participation in data collection is not needed. Where user participation in data collection is required, but only to oversee semi-automatic operations, remote operation has also shown to be productive and cost-effective by MX beamlines.

Experiment control tools at the APS vary depending on particular technique and experiment requirements. SPEC is by far the most widely used experiment control software. Other tools such as APSPy, JBlueIce, TXM, and uProbeX serve a specific user base. For the new APS beamlines planned as a part of the APS-U project and for new beamline improvements, the APS will adopt the NSLS-II bluesky control system, and collaborate closely with its core development team.

With increased scientific exploration into the fundamental processes of atomic interaction, there is present demand on the control systems to handle data collection at kilohertz frequencies and there will be an increased demand to respond to signals at frequencies orders of magnitude greater than megahertz. The APS will continue current and develop new efforts in the development of FPGA and softGlue FPGA frameworks to meet these needs.

Goals for the next five years:

- Keep the beamline control systems up-to-date with new developments from EPICS, synApps, and areaDetector.
- Replace the MEDM synoptic display manager with caQtDM.
- Develop general-purpose mail-in automation tools, and allow more remote access.
- Implement bluesky as the primary experiment control software for the APS-U beamlines, and at other beamlines where significant development will be done.
- Develop high-speed signal handling capabilities that can be, integrated with motion, acquisition, and detector controls and deploy where needed across the APS.
- Integrate high-speed data acquisition with HPC data analysis software and resources to enable real-time feedback and adaptive experiment control.

Data Analysis Algorithms and Software

The APS is focusing data analysis algorithm and software development in the areas needed to answer the novel scientific inquiries enabled by the future APS and the planned APS-U project, which are the techniques driven by coherence, imaging, and high-energy, as well as multi-modal techniques. The table below lists the areas targeted for data analysis algorithm and software advancements critical for operation of the APS-U beamlines.

Scientific Requirement	Beamline	Efforts / State	Funding
Bragg Coherent Diffraction Imaging (CDI) Software	Atomic, InSitu, Ptycho, CHEX	BCDI project	Ops
Correlation Software	CSSI, XPCS, CHEX	XPCS project	Ops, McGill Univ.
General-Purpose Reciprocal-Space Mapping (RSM) Software	3DMicroNano, XPCS, HEXM, CHEX	RSMap3D project	Ops
X-ray Fluorescence Mapping (XFM) Software	3DMicroNano, InSitu, Ptycho, VelociProbe	XRF-Mapper project	Ops
CDI Ptychography Software	InSitu, Ptycho, CHEX, PRISMA	Ptycholib project	ASCR, Northwestern Univ., Ops, CAMERA
Tomography Software	CHEX, HEXM	TomoPy project	LDRD, Ops
Multi-modal High-Energy Diffraction Software	HEXM	MIDAS project	Industrial Partners, Ops
Coherent Surface Scattering Imaging (CSSI) Software	CSSI	Algo. dev. required	LDRD, CAMERA
Multi-modal XRF Tomography	Ptycho	Algo. dev. required	LDRD
Multi-modal Diffraction Tomography	HEXM	Algo. dev. required	LDRD

Multi-modal XRF Ptychography	InSitu, Ptycho, VelociProbe	Algo. dev. required	LDRD
New Laue Diffraction Mode	3DMicroNano	Algo. dev. required	Ops
Data Quality Tools	All	DQV project	Ops
Real-time Streaming and Automated Feedback Infrastructure	All	MONA project	LDRD, Ops

Requirements for data analysis algorithms and software for the APS-U era. Table describes APS-U beamlines that will require particular algorithms and software, current efforts including the names of tools, if they exist, that meet the requirements, and how the efforts are funded.

These techniques already number amongst the most computationally intensive techniques performed at the APS and throughput demands are expected to grow by as much as multiple orders of magnitude, due to improved detectors and the upgraded source. Data reduction and analysis will rely heavily on the use of high-performance computing (HPC), utilizing appropriate technologies such as multi-threading, MMPs and GPUs, and distributed computing environments to obtain results with near real-time completion, so that results allow user-driven or even automated steering of experiments.

Most software will largely be developed as open source and will be made available with user community code contributions encouraged. A graded approach according to impact and priority will be applied to development. Packaging and active support either as distributable applications or as Software-as-a-Service (SaaS) will be provided for software systems that have been deemed to be most important for the success of APS users. Beamlines not directly driving the APS-U will also benefit from the reuse of tools developed for priority applications.

Key accomplishments have been made in this area over the past few years: The TomoPy reconstruction toolkit, the Hadoop MapReduce correlation software for XPCS, the HPC-enabled Microstructural Imaging using Diffraction Analysis Software (MIDAS) for high-energy diffraction microscopy, the multi-threaded XRF-Mapper software for elemental mapping, and the distributed GPU-based Ptycholib software for ptychography are some examples that form the staple applications critical to the APS-U era facility. These programs will continue to be developed for improved performance and algorithms. A more complete list of software produced at the APS can be found at https://www1.aps.anl.gov/Science/Scientific-Software.

New efforts are currently underway to address the development of new algorithms and HPC software for multi-modal analysis, including fluorescence tomography, fluorescence ptychography, tomography diffraction, and Bragg CDI and ptychography, and for a new approach to Laue diffraction reconstructions. The MONA (Monitoring, Optimization, Navigation, Adaptation) project will prototype data streaming coupled with real-time data analysis and automated feedback.

One way the Laboratory supports the computational efforts at the APS is via LDRD funding. Beginning in FY11, the *Tao of Fusion* LDRD helped seed the TomoPy application; likewise, the FY13 *Next Generation Data Exploration: Intelligence in Data Analysis, Visualization and Mining*

LDRD was aimed at multi-modal analysis. Currently funded LDRDs include *Modeling, Analysis, and Ultrafast Imaging (MAUI)*, *Multimodal Imaging of Materials for Energy Storage (MIMES)*, *Enabling Nanometer-scale X-ray Fluorescence Tomography*, and *Coherent Surface Scattering Imaging*.

In addition to LDRDs, the APS receives funding and personnel support from the ALCF Data Sciences Program (ADSP) for *Large-Scale Computing and Visualization on the Connectomes of the Brain*, and from the NERSC Exascale Science Applications Program for *Optimization of data-intensive tomography workflows at light sources*. Effort for CDI ptychography was initially funded by ASCR and now via the Intelligence Advanced Research Projects Activity (IARPA) and Northwestern University; early efforts for the MIDAS software were funded from APS industrial partners. These activities are directly aligned with facility priorities and the core requirements listed in Table 1.

The APS has actively collaborated with other facilities and organizations, and members of the user community to develop data analysis algorithms and software. As examples, most ANL-funded LDRDs in these areas involve collaborators from ANL's Mathematics and Computer Science Division. Select user groups have contributed greatly to analysis algorithms and software. LBNL's CAMERA provides assistance and software for modeling and GISAXS. In addition, the APS has been involved in the NOBUGS conference community, and maintains active participation in the series of hack-a-thons organized by the Experimental Facilities Computing (ExFaC) Working Group. Most recently, the APS and the NSLS-II at BNL have developed a comprehensive computing collaboration plan so as to best utilize our scare resources.

Goals for the next five years are directed towards enabling the full benefit of the portfolio of anticipated future beamlines, including the APS-U beamlines:

- Further refinement of computational needs for APS-U Scientific Requirements (Table 1) through topical workshops with beamline developers and selected users.
- Research and development of new algorithms for multimodal data analysis methodologies needed by future beamlines (see table above).
- Creation and deployment of a robust set of high-performance computing (HPC) enabled software tools utilizing the next-generation Xeon Phi processors that address cross-cutting critical technique domain areas in the areas of coherence, imaging, high-energy, and multi-modal techniques (see Table 1).
- Integration of general purpose data streaming, feedback, adaptive control, and verification tools with beamline control software, HPC data analysis software, and data management resources.

Data Management and Distribution

The need for data management and distribution tools and data storage resources continues to grow. Currently the APS collects on the order of 2 PB of raw data per year. In the coming years, the data storage needs of the APS are anticipated to increase by at least two orders of magnitude, where some beamlines are anticipated creating circa 2 PB/day. Great strides have been made in this area over the past years. The APS continues to rollout a multi-tiered data management and distribution system for APS beamlines.

During FY13-FY15, the APS piloted facility-wide data management and distribution tools and resources with effort and funding from the LDRDs. These activities provided R&D effort and seeded connections between the APS and the Argonne Leadership Computing Facility (ALCF), the Mathematics and Computer Science (MCS) division, and the Globus Services team at the University of Chicago.

Through APS operations effort, Data Management (DM) tools that integrate with beamline data workflows, and large data storage systems are available. These tools help automate the transfer of data between acquisition devices, computing resources, and data storage systems. Ownership and access permissions are granted to the users signed-up to perform the experiment. A metadata catalog allows beamline staff to populate experiment conditions and information for access via a web portal. Users can download data at their home institutions using Globus Online (globus.org). Short-term data storage is available within the APS; larger longer-term storage systems totaling approximately 3 PB are hosted by the ALCF (see figure below). These resources are deployed at a number of XSD beamlines: 1-ID, 2-BM, 2-ID, 6-ID, 7-BM, 7-ID, 8-ID, 11-ID, 23-ID, 26-ID, 32-ID, 33-ID, and 34-ID.

The current APS Data Management Policy can be found at https://www1.aps.anl.gov/users-information/help-reference/data-management-retrieval-practices.

Goals for the next five years:

- Continue to roll-out these capabilities to more beamlines, and continue collaborative plans with the ALCF and the Globus Services team to develop further functionality.
- Provide additional "hot" storage as needed in cooperation with the ALCF, and when the capacity requires, will integrate these storage systems with ALCF tape backup systems.
- Experiment metadata collected using

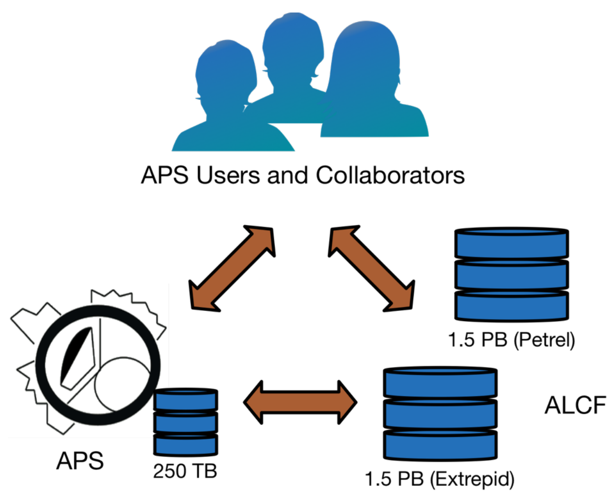

Storage available at the APS. A 250 TB data storage system located at the APS serves short-term needs. The Petrel system housed in the Theory and Computing Sciences building is managed by the Argonne Leadership Computing Facility (ALCF) and provides 1.5 PB of storage. The Extrepid system, run by the APS but housed in an ALCF computing center in Building 369, provides an additional 1.5 PB of storage space for the APS.

bluesky will be directly integrated into the DM system.
- Identify or develop an electronic notebook tool and implement across the facility.
- Integrate data workflows with outside services, such as the Materials Data Facility (materialsdatafacility.org) and the NIST Materials Data Repository (mgi.nist.gov/materials-data-repository), to provide full data lifecycle management.

Computing and Network Infrastructure

As with data storage, the computing resources required by the APS are anticipated to grow by at least two orders of magnitude. To satisfy these needs, the APS has adopted a graded approach to resource utilization. Small-scale resources, such as Intel Phi processors and GPUs, local to beamlines will be used when sufficient. For moderate computational needs, the APS currently maintains an on-site computing cluster, and ANL maintains computing resources as a part of the Laboratory Computing Resource Center (LCRC). For the most demanding computational problems, large-scale computing facilities must be used.

The APS has been involved in many activities aimed at using centralized and large-scale computing resources over the past few years. Notable activities include:

- In cooperation with CELS, the APS successfully deployed Hadoop MapReduce based XPCS analysis on ANL's Magellan virtualized cloud-computing resource located in building 240. This resource has been fully integrated into the automated XPCS data reduction workflow.
- Efforts are currently underway at the APS to routinely utilize NERSC for high-energy diffraction microscopy and tomography reconstructions.
- With computer scientists in MCS, the APS is researching the use of virtualization and containers to pre-emptively utilize computing resources in the LCRC for on-demand beamline use. Prototype on-demand applications have been demonstrated using ALCF and NERSC.
- Applications such as MIDAS (high-energy diffraction microscopy) and TomoPy have been run at scale with up to 30,000 cores, allowing data analysis during beam time.

Networking is another key consideration when planning both computing and storage resources. Sufficient bandwidth and latency constraints are required both within the APS, and within and outside of ANL. These demands will evolve with the architecture and utilization of computing and storage needs. Over the previous few years that APS completed a 100 Gbps network connection to the central laboratory network, a 40 Gbps connection to the ALCF computing center in building 240, and established a direct 20 Gbps network connection to data storage in building 369. The core network within the APS and between servers is 10 Gbps. Working with ESnet and ANL, the APS implemented a Science DMZ for particular outside user groups in order to reduce latency and better enable remote access and experiment control.

In coming years, costs for use of internal vs. externally managed computing resources will be reviewed to ensure that the APS is most wisely using laboratory funding.

Goals for the next five years:

- Continue to adopt a graded approach to computing resources at the APS.
- Work with ALCF, NERSC and ALCF to develop strategies for on-demand utilization of leadership facilities that have minimal impact on the core mission of those facilities allowing the APS to utilize them for moderate-scale and large-scale computing problems.
- Maintain and expand capacity of network connections to computing resources at ANL as needed.
- Increase core network bandwidth to 100 Gbps or greater, depending on need.

SWOT - Data Analysis Algorithms and Software

Strengths	Weaknesses
• World-leading computational efforts in a number of scientific areas. • World-class beamline staff and user groups contribute new algorithms and software. • Close collaborations among APS users and staff, and the XSD-BC, XSD-CXS, and XSD-SDM groups. • ANL has world class expertise in applied math, computer science, HPC as well as data and computation facilities.	• Current funding situation does not allow for the APS to meet all of its data analysis algorithm and software needs. • Most current generation data analysis tools are not suited to stream data in HPC environments needed to keep up with anticipated data rates. • Many scientist-developed packages lack professional software engineering. • Lower facility productivity due to lack of data analysis tools.
Opportunities	**Threats**
• Collaborations with ANL expertise will help bring state-of-the-art HPC applications to the APS. • Having shared staffing with MCS will help increase interactions with non-APS experts. • Collaborations with DOE facilities and resources could amplify development efforts, and provide needed software in a cost-effective manner for the entire DOE complex. • Other light sources wish to collaborate on data analysis activities. • The APS Upgrade-enabled techniques may be fully realized, answering new scientific questions; the APS maintains its position as the most productive light source.	• Without further investment and collaboration in this area, the APS will not fully realize its scientific potential. • User groups may seek to perform cutting-edge experiments at other light sources. • Collaborative projects with external organizations may not produce code meeting APS needs. • Other domestic and international light sources have considerably larger and more active software and algorithm development programs that can leapfrog APS leadership.

SWOT - Data Management & Distribution

Strengths	Weaknesses
• World-leading expertise at ANL in data sciences, data management and transfer (e.g. Globus Services team). • APS is one of the DOE's largest data collecting user facility, producing a wealth of scientifically valuable data. • Collaborative efforts continue to form between the APS and expertise elsewhere at ANL.	• Preponderance of existing unique solutions at beamlines involving manual, inefficient management steps; no common user experience. • Current manual methods cannot keep pace with increasing data rates. • Lowered productivity due to time taken away from staff and users to address tasks that may be automated.
Opportunities	**Threats**
• Leverage expertise from CELS, UoC, and the Globus Services team. • Reduce cost by leveraging outside software resources and expertise. • Consistent data management user experience. • Increase scientific productivity through automation of data management tasks.	• The full potential of the APS-U cannot be realized without managed data workflows. • Lowered scientific productivity due to an inability to keep up with increases in data. • International light sources that have invested heavily in data management software may overtake the APS in terms of scientific productivity.

SWOT – Controls and Beamline Operations Software

Strengths	Weaknesses
• World-leading expertise at ANL in data sciences, data management and transfer (e.g. Globus Services team). • APS is the DOE's largest data collecting synchrotron user facility, producing a wealth of scientifically valuable data. • Collaborative efforts continue to form between the APS and expertise elsewhere at ANL.	• Patchwork of solutions to critical problems exists due to a lack of funding in this area; not a cost-effective use of beamline resources. • Lowered beamline productivity due to a lack of proper software for beamline operation.
Opportunities	**Threats**
• Development effort from new beamlines that are developed in the APS Upgrade can supply reusable software modules to be employed in existing beamlines. • Collaboration with other DOE light	• The benefits of the APS Upgrade, which require advanced in data collection and beamline operation, are not fully realized. • APS beamline capabilities improve but usability stagnates, and beamlines are perceived as no longer world-class.

sources is a cost-effective way to maintain cutting-edge beamline operations software. • Easier to use instruments will make users more productive and lower barriers to research groups that are not currently synchrotron users. • Cutting-edge user groups are drawn to the APS, and further deeper collaborations with beamline staff. • Benefits of the APS Upgrade, which require beamline operation software, will be fully realized, maintaining the APS as the world-leading synchrotron.	• Cutting-edge user groups are drawn to other facilities where experiments are easier. • Experimental techniques and data collection advances at other facilities outpace advances at the APS.

Group Strategies and Work Plans

XSD-BC Strategy:
https://www1.aps.anl.gov/sites/default/files/BCDA%20group%20strategy%2020160120.pdf

XSD-CXS Strategy:
https://anl.box.com/v/CSX2017StratGoals

XSD-SDM Strategy:
https://www1.aps.anl.gov/sites/default/files/2017-SDM-Strategy-Goals-2017-06-19.pdf

XSD-SDM Detailed Work Plan:
https://www1.aps.anl.gov/sites/default/files/SDM-FY17-FY18-Plan-2017-06-19.pdf