Evolving role of HDF5 at the upgraded Advanced Photon Source

Tejas Guruswamy
Detectors Group, Advanced Photon Source

APS-U Beamline Data Pipelines project

HDF5 User Group Meeting 2023

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

17 August 2023

The Advanced Photon Source

- The **Advanced Photon Source** is a hard X-ray (6 to 150 keV photons) synchrotron light source
- Electrons accelerated to 7 GeV (99.999999% of *c*) generate intense beams of photons for imaging, diffraction, spectroscopy
- Around 70 parallel experimental stations each with unique specialties operating 24/7, 9 months a year
- 10^9 (one billion) times brighter than a medical X-ray
- $>$5,000 worldwide users each year from physics, materials science, chemistry, biology, defence, nuclear safety, microfabrication, electronics, pharma . . .

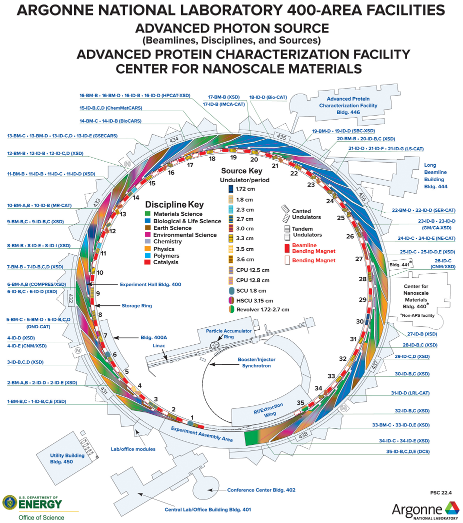
ARGONNE NATIONAL LABORATORY 400-AREA FACILITIES
ADVANCED PHOTON SOURCE
(Beamlines, Disciplines, and Sources)
ADVANCED PROTEIN CHARACTERIZATION FACILITY
CENTER FOR NANOSCALE MATERIALS
16-BM-B • 16-BM-D • 16-ID-B • 16-ID-D (HPCAT-XSD)
15-ID-B,C,D (ChemMatCARS)
14-BM-C • 14-ID-B (BioCARS)
13-BM-C • 13-BM-D • 13-ID-C,D • 13-ID-E (GSECARS)
12-BM-B • 12-ID-B • 12-ID-C,D (XSD)
11-BM-B • 11-ID-B • 11-ID-C • 11-ID-D (XSD)
10-BM-A,B • 10-ID-B (MR-CAT)
9-BM-B,C • 9-ID-B,C (XSD)
8-BM-B • 8-ID-E • 8-ID-I (XSD)
7-BM-B • 7-ID-B,C,D (XSD)
6-BM-A,B (COMPRES/XSD)
6-ID-B,C • 6-ID-D (XSD)
5-BM-C • 5-BM-D • 5-ID-B,C,D (DND-CAT)
4-ID-D (XSD)
4-ID-E (CNM/XSD)
3-ID-B,C,D (XSD)
2-BM-A,B • 2-ID-D • 2-ID-E (XSD)
1-BM-B,C • 1-ID-B,C,E (XSD)
17-BM-B (XSD)
17-ID-B (IMCA-CAT)
18-ID-D (Bio-CAT)
Advanced Protein Characterization Facility Bldg. 446
19-BM-D • 19-ID-D (SBC-XSD)
20-BM-B • 20-ID-B,C (XSD)
21-ID-D • 21-ID-F • 21-ID-G (LS-CAT)
Long Beamline Building Bldg. 444
22-BM-D • 22-ID-D (SER-CAT)
23-ID-B • 23-ID-D (GM/CA-XSD)
24-ID-C • 24-ID-E (NE-CAT)
25-ID-C • 25-ID-D,E (XSD)
Bldg. 441*
26-ID-C (CNM/XSD)
Center for Nanoscale Materials Bldg. 440*
*Non-APS facility
27-ID-B (XSD)
28-ID-B,C (XSD)
29-ID-C,D (XSD)
30-ID-B,C (XSD)
31-ID-D (LRL-CAT)
32-ID-B,C (XSD)
33-BM-C • 33-ID-D,E (XSD)
34-ID-C • 34-ID-E (XSD)
35-ID-B,C,D,E (DCS)
431
432
433
434
435
436
437
438
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
Discipline Key
🟩 Materials Science
🟦 Biological & Life Science
🟫 Earth Science
🟪 Environmental Science
▨ Chemistry
🟧 Physics
🟦 Polymers
🟥 Catalysis
Source Key
Undulator/period
1.72 cm
1.8 cm
2.3 cm
2.7 cm
3.0 cm
3.3 cm
3.5 cm
3.6 cm
CPU 12.5 cm
CPU 12.8 cm
SCU 1.8 cm
HSCU 3.15 cm
Revolver 1.72-2.7 cm
Canted Undulators
Tandem Undulators
Beamline Bending Magnet
Bending Magnet
Experiment Hall Bldg. 400
Storage Ring
Bldg. 400A
Linac
Particle Accumulator Ring
Booster/Injector Synchrotron
Rf/Extraction Wing
Experiment Assembly Area
Lab/office modules
Utility Building Bldg. 450
Conference Center Bldg. 402
Central Lab/Office Building Bldg. 401
U.S. DEPARTMENT OF ENERGY
Office of Science
PSC 22.4
Argonne NATIONAL LABORATORY

The APS Upgrade (Apr 2023 - Apr 2024)

- Facility undergoing major upgrade, replacing all storage ring magnets
- **500 times increase** in coherence and X-ray brightness at most energies
- Enables better measurements of dilute samples, faster scanning, nanoscale focusing
- Entire beamlines, especially detectors, being upgraded to take advantage of increased photon flux
- New network, data storage and compute infrastructure throughout APS

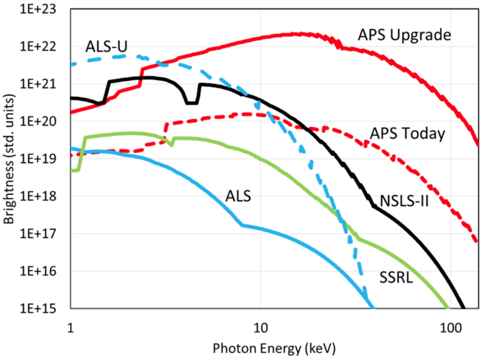

Increased data rate & size

- Order-of-magnitude increase in data generated (**<10 PB/year to >100 PB/year**)
- Updated science goals including real-time feedback, use of ALCF compute (exascale Aurora/Polaris), FAIR data policies
- Overhaul of **batch and stream** processing, analysis, visualization for most scientific techniques (*Beamline Data Pipelines* project)

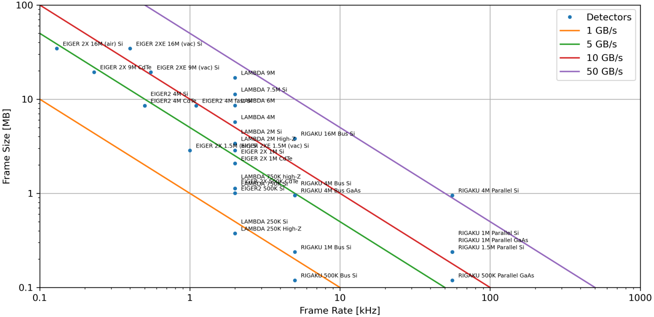

Data format and management tools

- HDF5 preferred as on-disk data format for all new/upgraded data workflows
- Combine all data and metadata in stable format for archiving and offline (re-)analysis
- Leveraging existing solutions where possible
  - **EPICS areaDetector** plugin(s) for direct to HDF5 data acquisition and network streaming of structured data with same compression algorithms as HDF5
  - **Data Exchange** [`https://dxfile.readthedocs.io/`] and **NeXus** [`http://www.nexusformat.org/`] for public, documented schema matching needs of technique
  - **bluesky** experiment orchestration and metadata collection, including **tiled** UI for slicing and exporting HDF5-backed data
  - **globus** cross-institution data management and HPC workflow execution

Improving parallel and GPU use

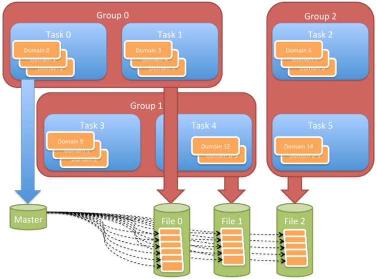

From Mark Miller, 2017
`https://www.hdfgroup.org/2017/03/mif-parallel-io-with-hdf5/`

- Parallel read and (sometimes) write important to keep up with data rates
- *External multiple file links, MPI, or other parallelism?*
- Increasing use of GPU-based analysis; experimenting with NVIDIA Rapids libraries, limited success so far
- *Accelerate HDF5 loading to GPU, network streams?*
- Some techniques, e.g. XPCS, require data layouts efficient on FPGA, CPU and GPU
- *Sparse and/or optimized data layouts for diverse architectures?*
- *Enable AI and data mining to efficiently navigate archived data?*