Data Analysis and Management at the APS: Successes and Challenges

Nicholas Schwarz
Principal Computer Scientist / Group Leader

Software Services Group
APS Engineering Support Division
Advanced Photon Source
Argonne National Laboratory

15 August 2013

APS Software Services Group

Sprint Team

Arthur Glowacki
Software Engineering Specialist

John Hammonds
Principal Application Developer

Faisal Khan
Software Engineering Associate

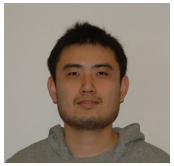

Ke Yue
Software Engineering Associate

APS-U

Sinisa Veseli
Principal Software Engineer

EPICS Team

Janet Anderson
Computer Scientist

Andrew Johnson
Computer Scientist

Tim Mooney
Physicist

Nicholas Schwarz
Group Leader

Marianne Binetti
Administrative Secretary

The Advanced Photon Source

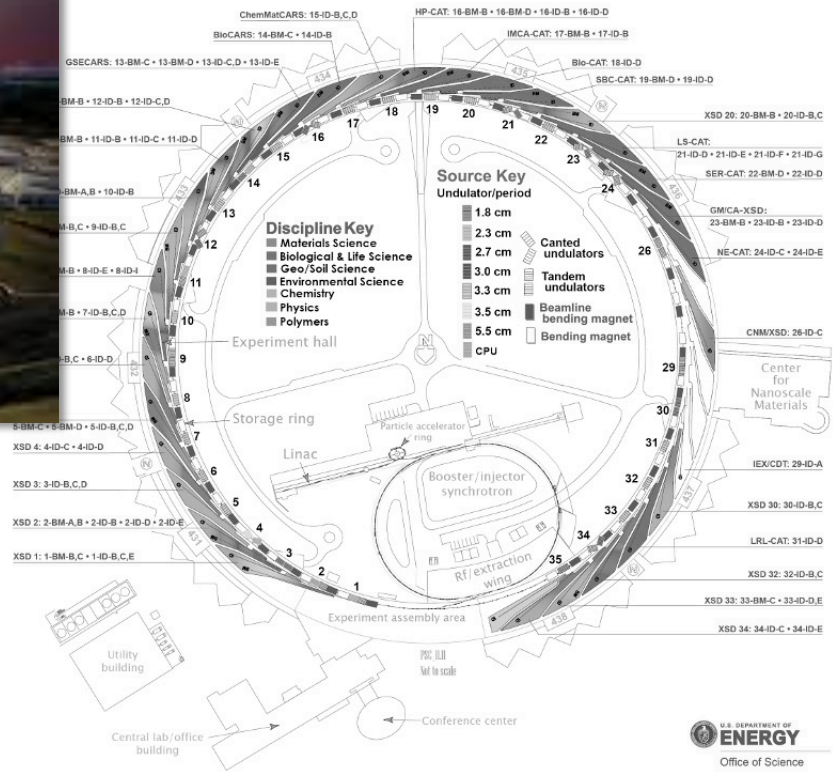

Increasing Amounts and Complexity of Data

Detectors

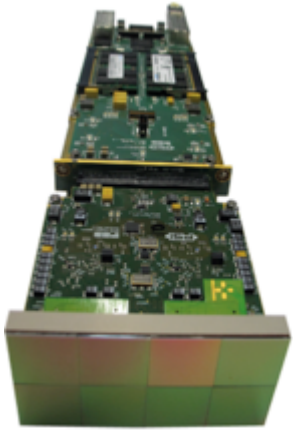

Experiment Automation

Techniques

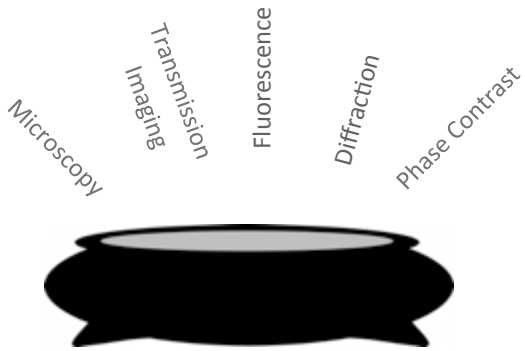

APS Average Data Rate

Cumulative typical data volume at present: 112 TB/month
Cumulative maximum data volume: 368 TB/month

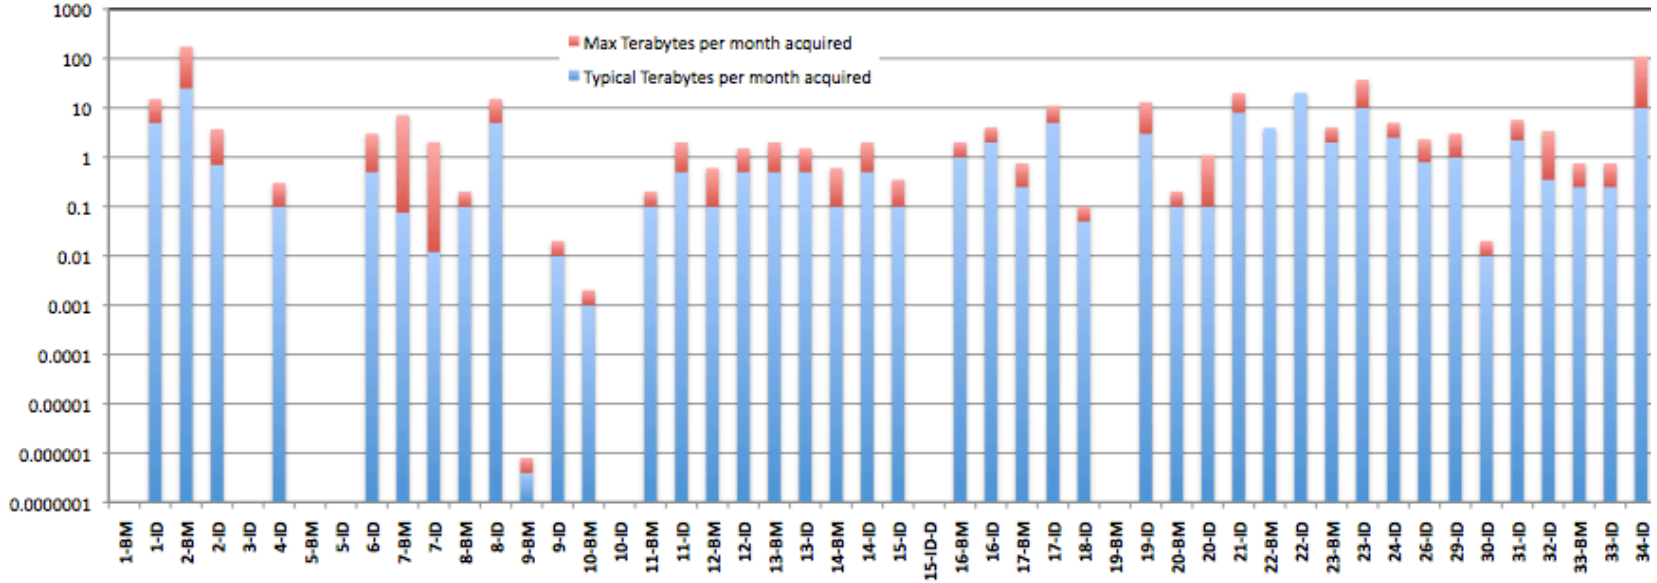

F. De Carlo (ANL)

Increasing User Base

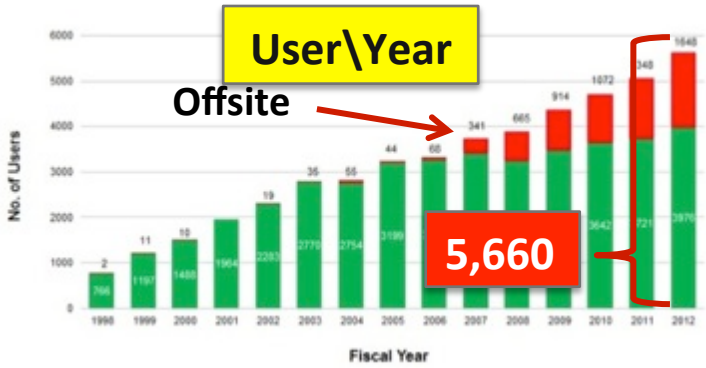

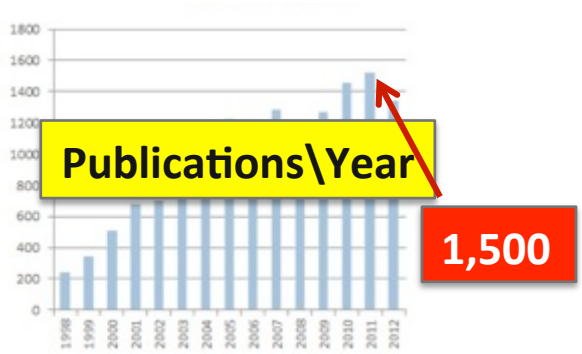

Architecture

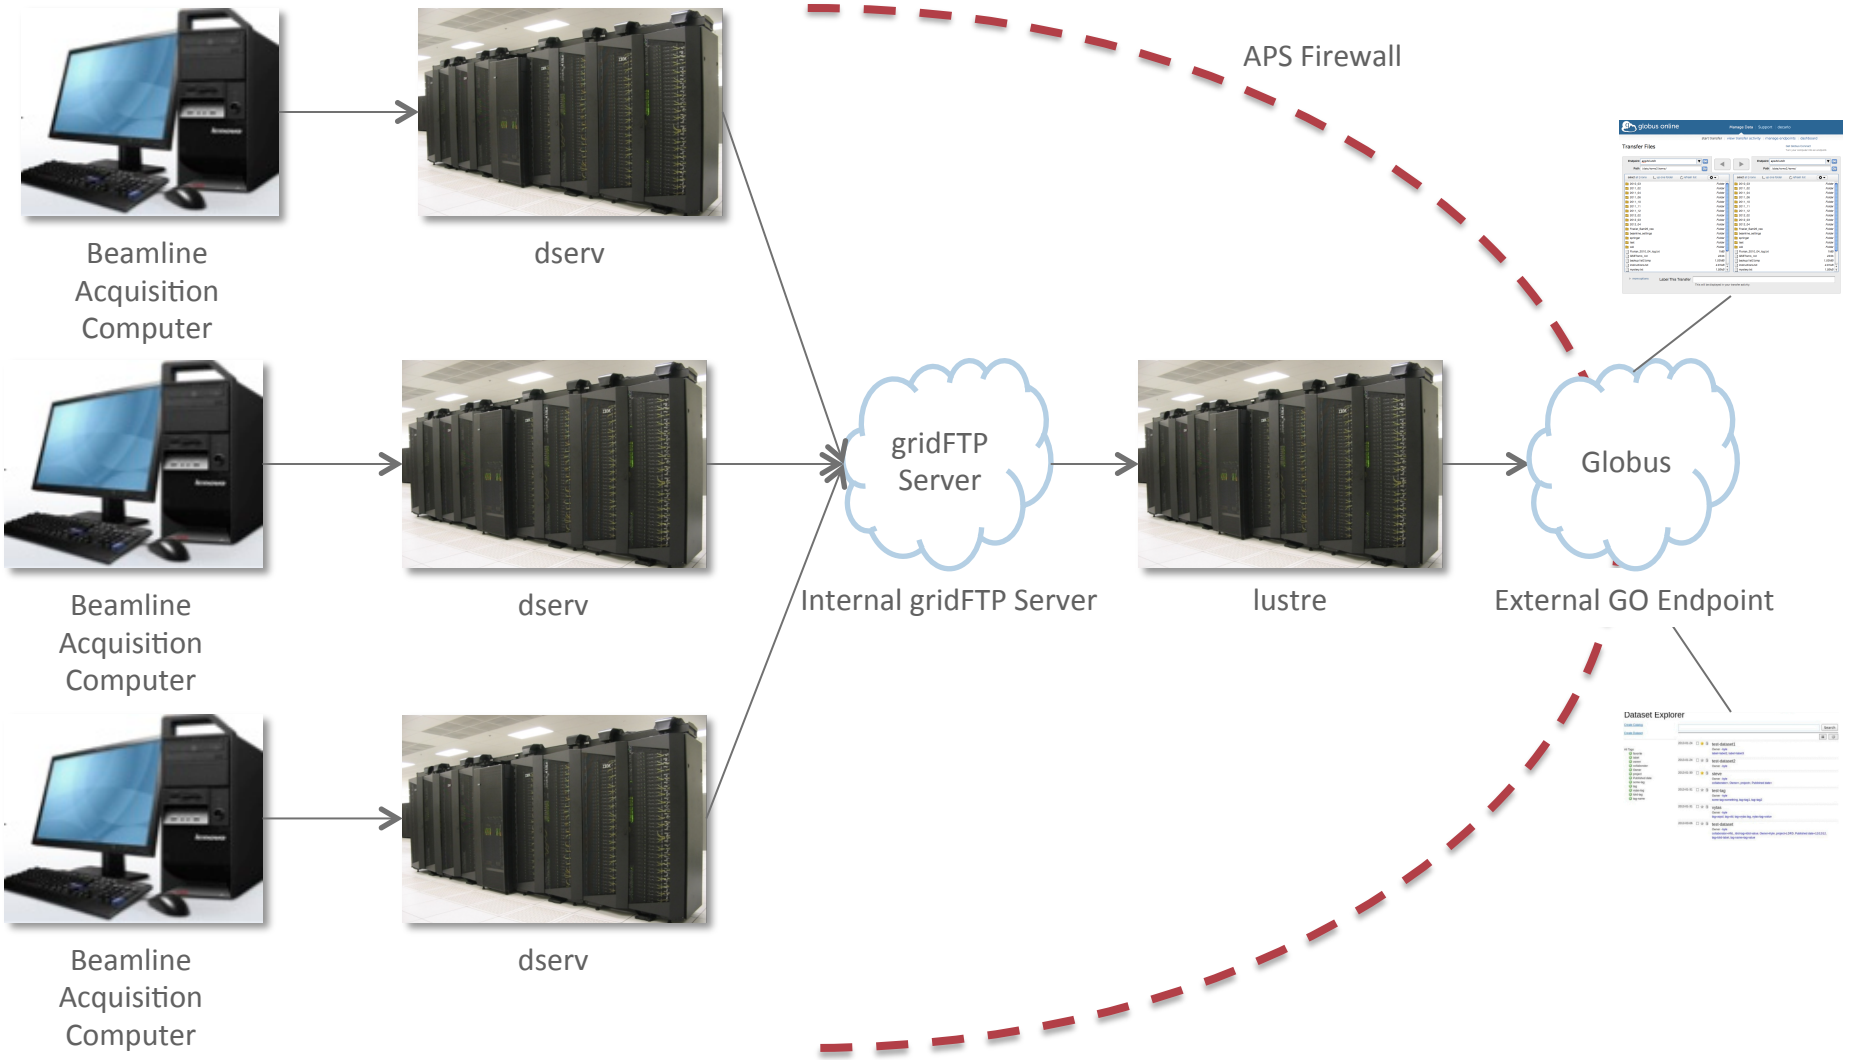

Workflow Pipeline

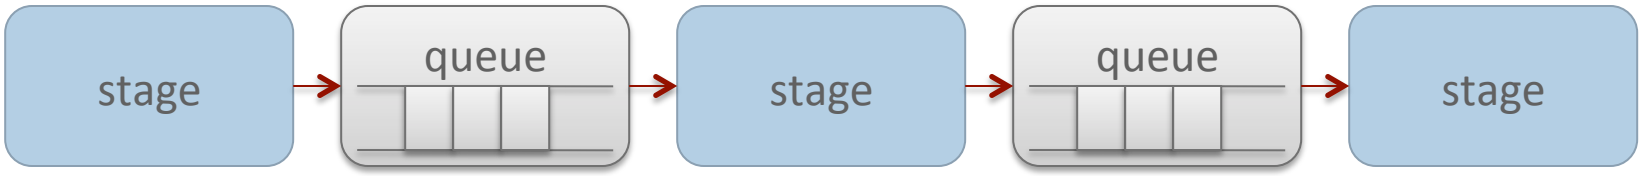

- Series of actors connected by message queues
- Stage
  - Acquisition
  - Run data analysis
  - Transfer files
  - Many languages: C++, Java, Python
- Message queues
  - Pass messages from one actor to the next
  - JMS message queues (ActiveMQ)

Pipeline Stages

Stages are composed of two separate classes:

- Director
  - interface with the message broker
- Actor
  - run scan / analysis / file transfer
  - save results
  - report status

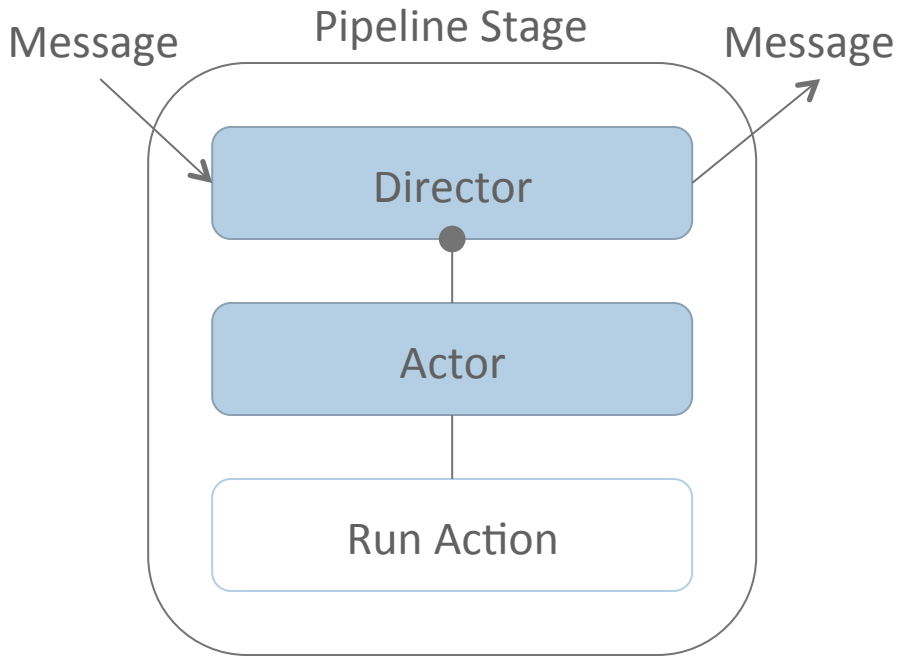

Director

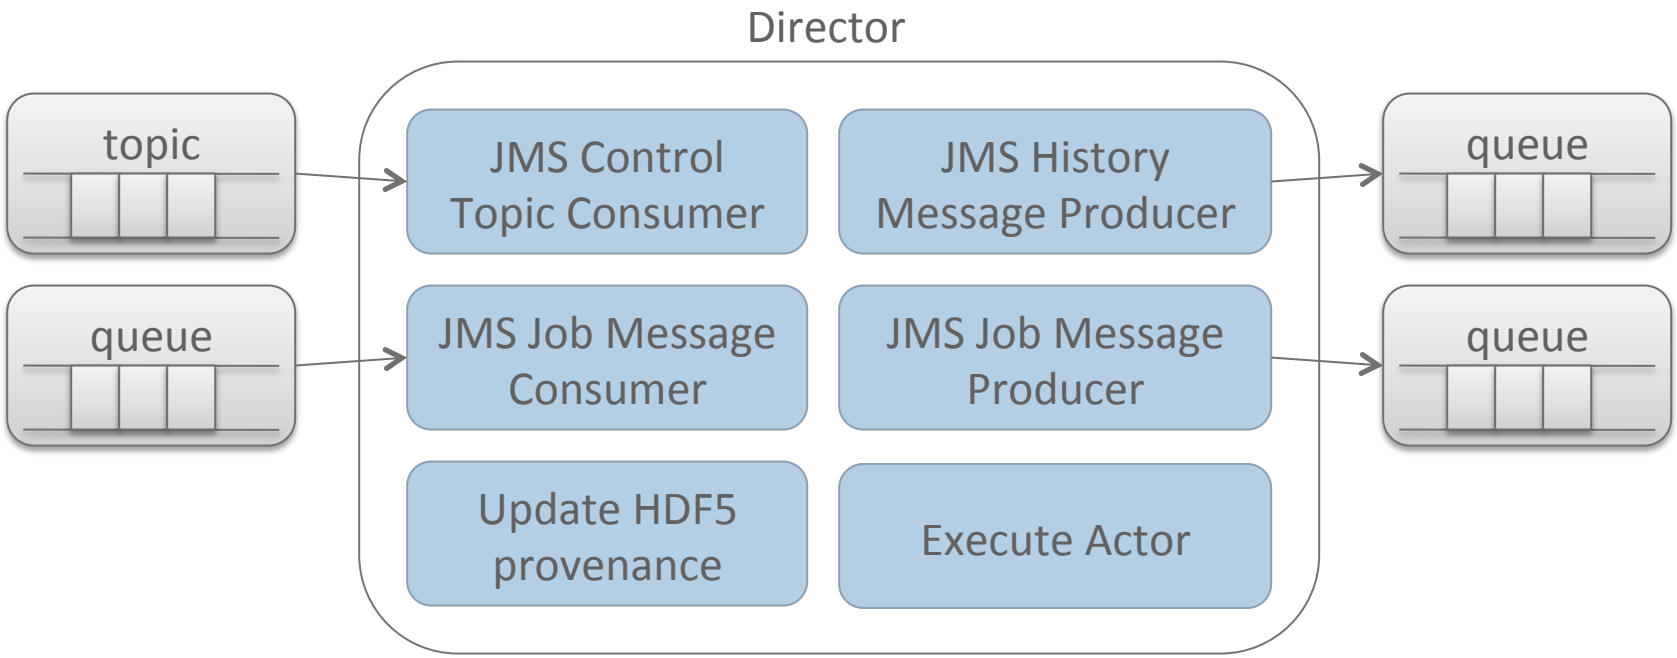

- Broker interface
- Handle job messages
- Control messages
- Update history
- Maintain provenance
- Execute Actor

Actor

- Actor
  - Called by Director
  - Abstract interface for concrete implementation
  - Execute job / analysis
  - Report status
- Analysis
  - External application
  - Internal code
  - Read input data from HDF5
  - Write output data to HDF5

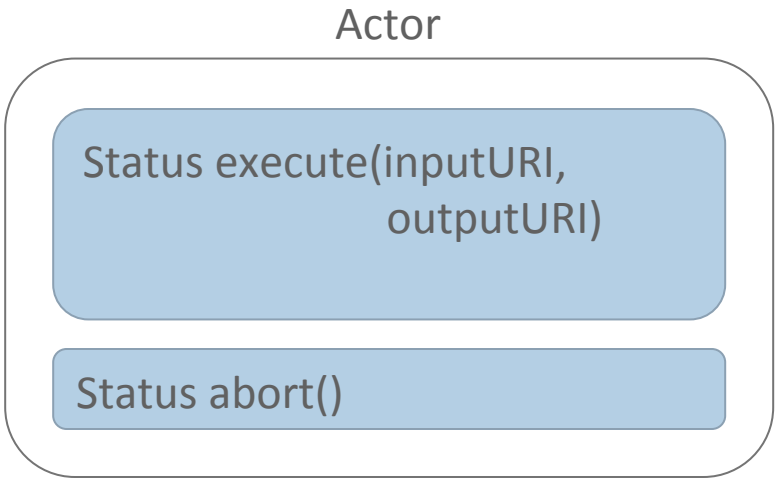

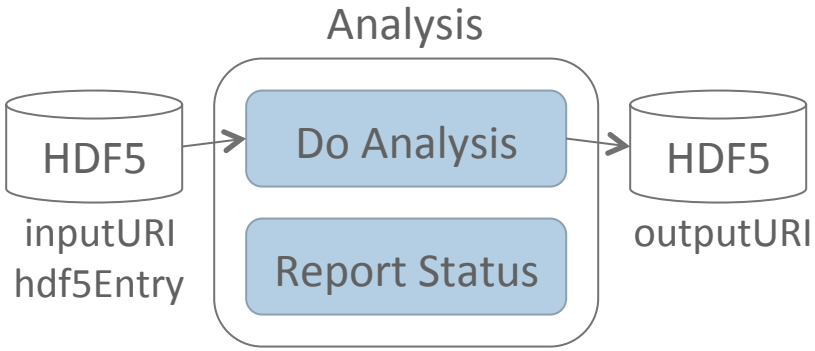

X-ray Photon Correlation Spectroscopy (XPCS)

Suresh Narayanan and Alec Sandy (APS)

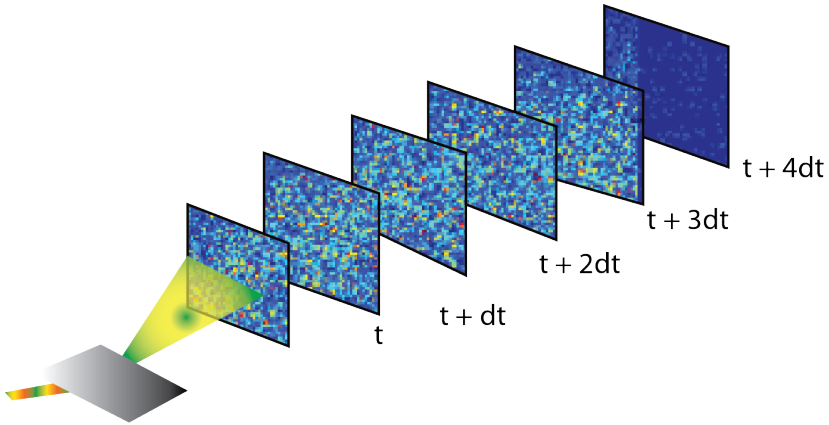

2D small-angle scattering pattern from a suspension of silica spheres (*right*).

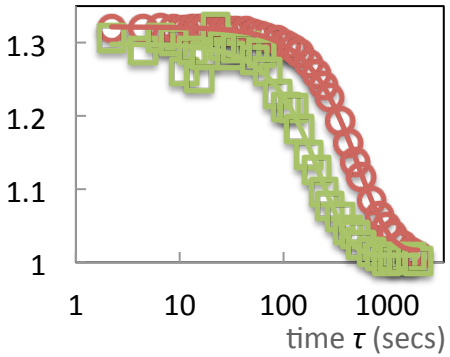

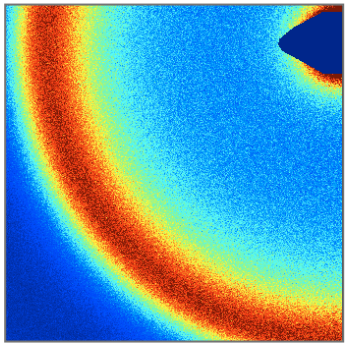

g_2 calculation at two different length scales (*left*).

Beamline and Detector

- 8-ID
- 1K x 1K @ 60 fps
- 1K x 1K @ 200 fps (FastCCD2)

XPCS Data

- Kinetics measurements
- 2D images over time
- A few minutes to collect
- Varies from 0.5 GB to 50 GB raw

Reduction

- Multi-tau autocorrelation

XPCS - Data Workflow

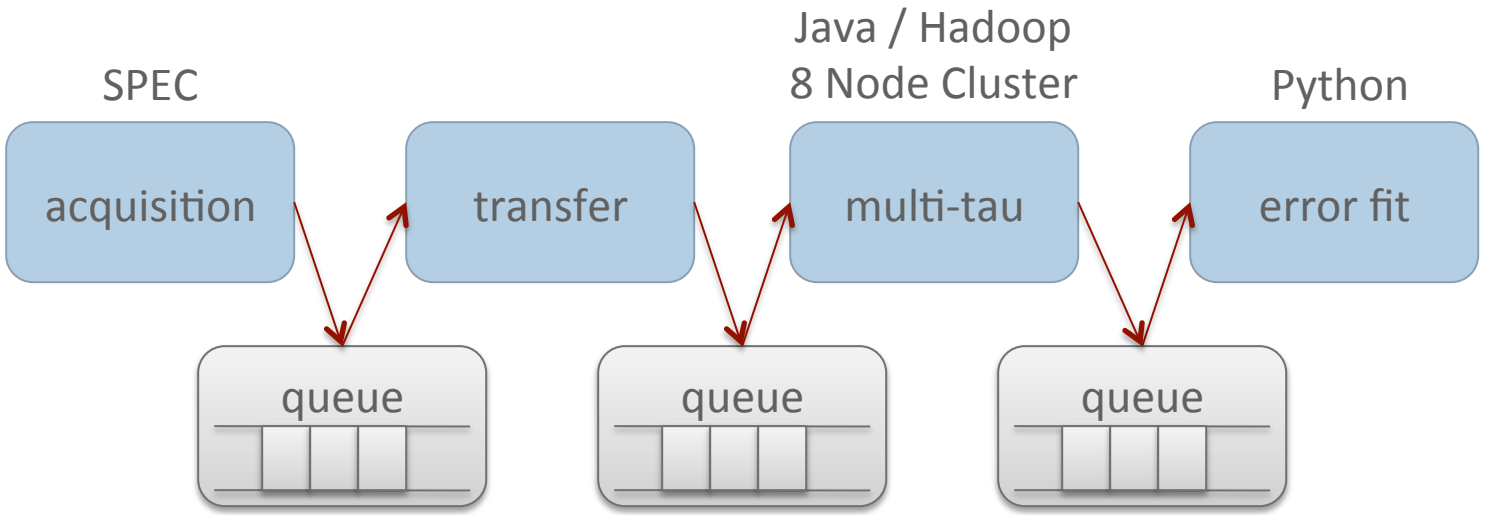

XPCS - Hadoop Multi-Tau Autocorrelation

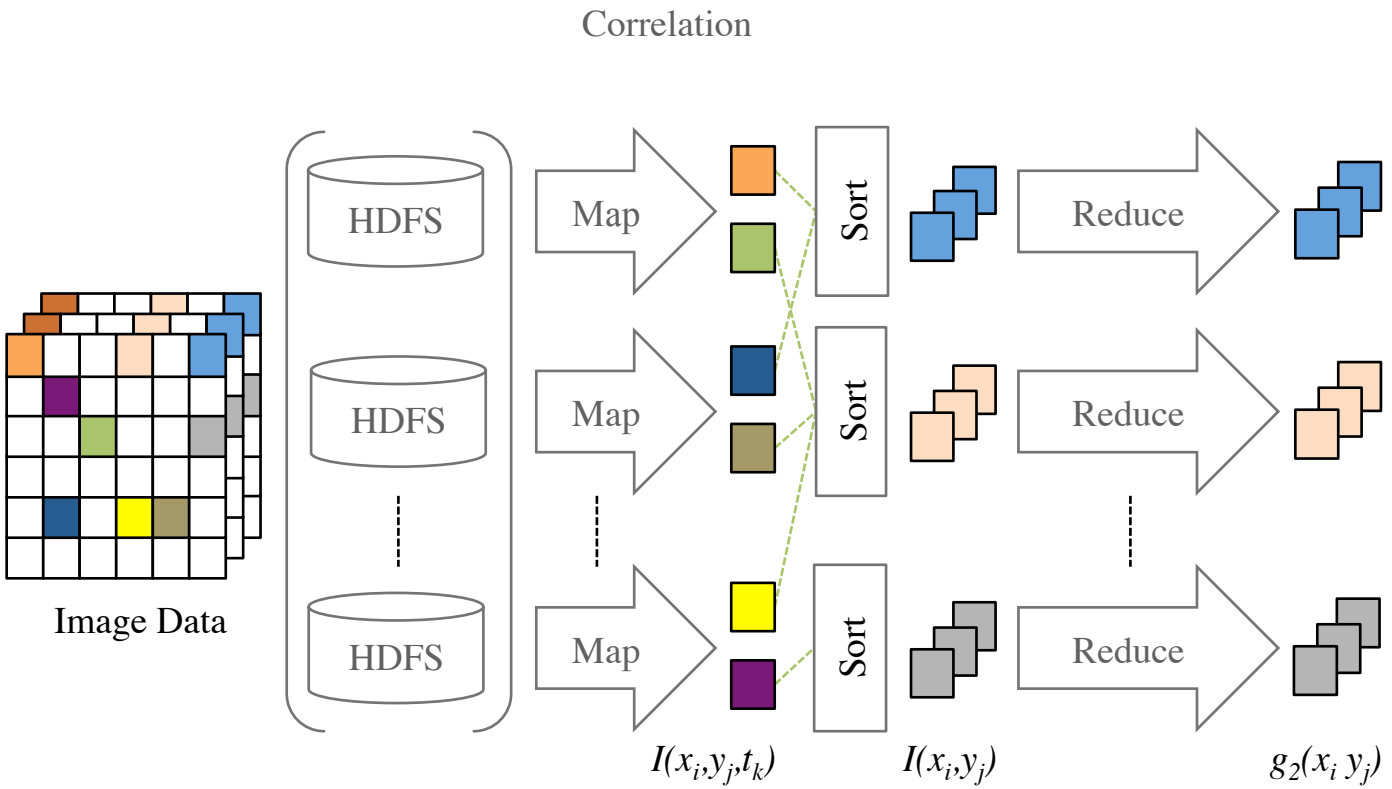

XPCS – Hadoop Normalization

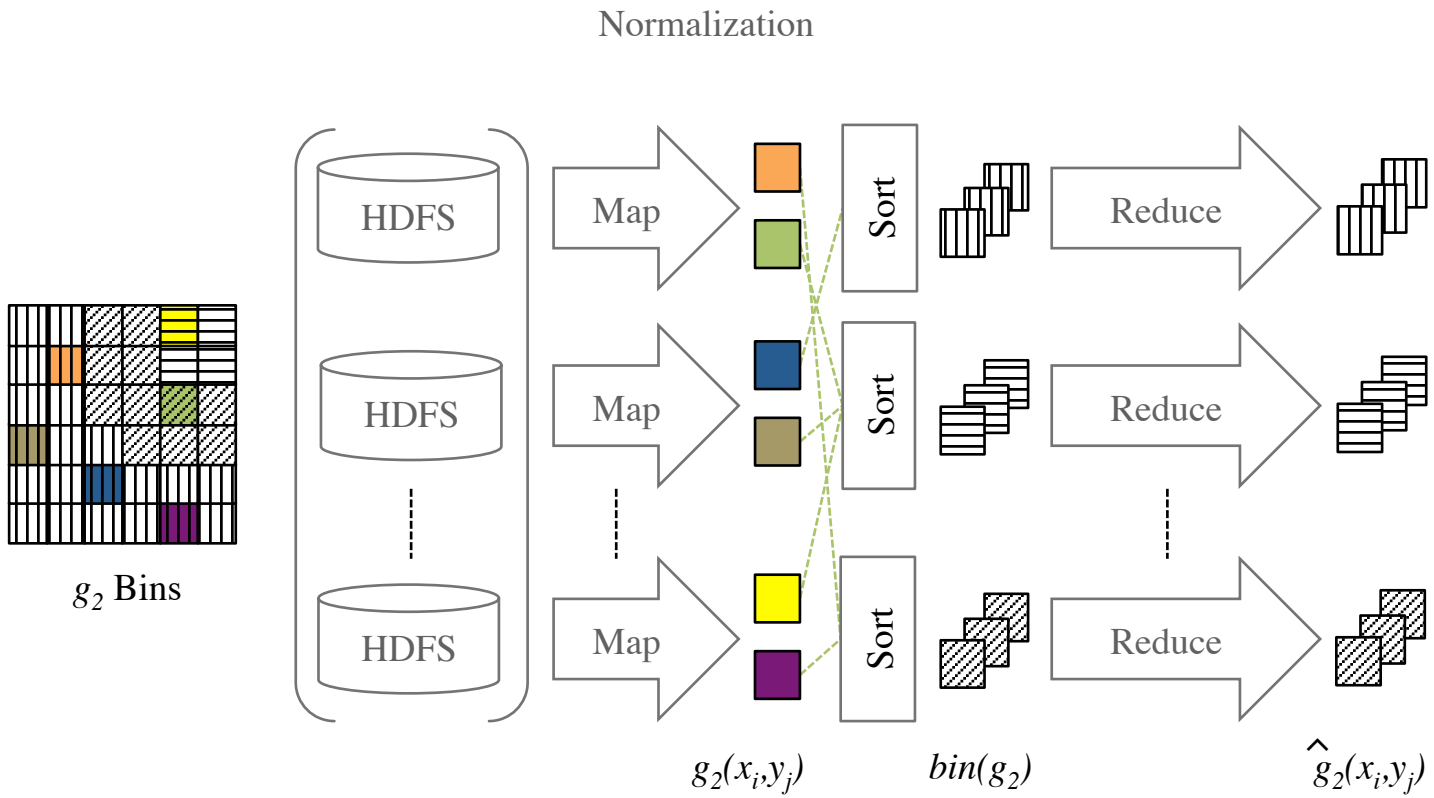

XPCS - Performance Results

Compute Resource

- Distributed-Memory Cluster
- 8 Nodes – 32 Cores
- 2.5 TB of distributed HDFS storage
- On demand

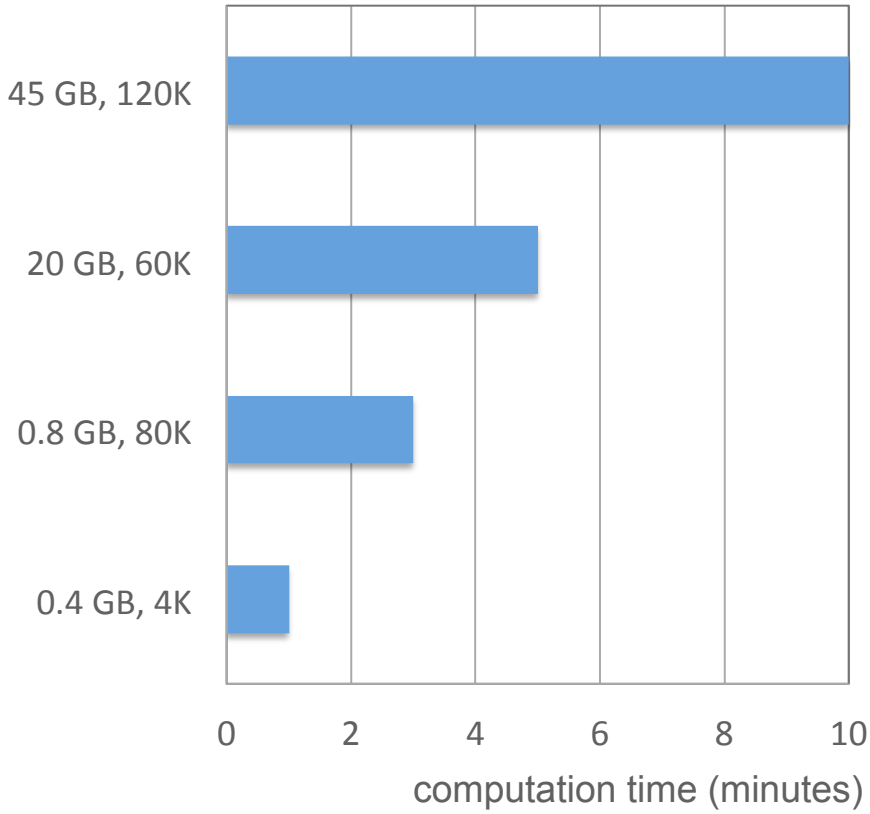

XPCS – FastCCD2

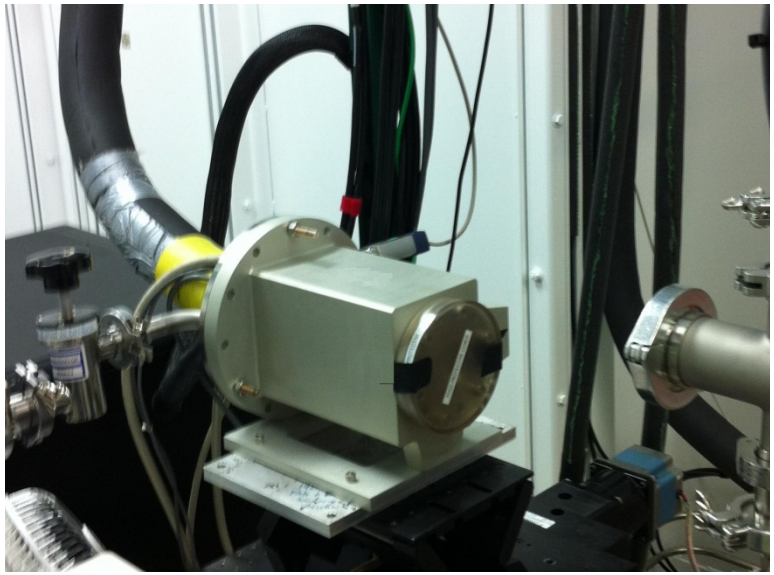

- Custom detector designed by LBNL
- 1K x 1K @ 200 FPS
- Planned to be used by ALS, LCLS, NSLS-II and APS
- ATCA processing crate
- Real-time autocorrelation implementation

Fast Micro-Tomography

Francesco De Carlo (APS)

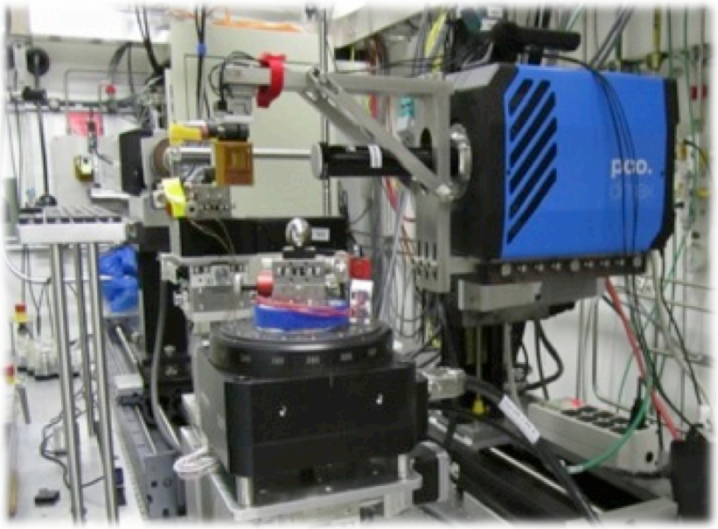

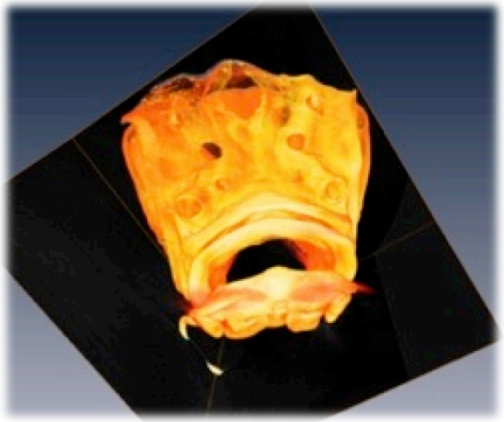

Keith Chang (Penn State College of Medicine)

Beamline and Detector

- 2-BM
- pco.dimax (36GB of onboard RAM)
- 2K X 2K @ 16-bits

Tomography Dataset

- Projections around 180 degrees
- ~1,500 FPS
- A few seconds of data collection
- ~1 minute readout
- Up to 11GB raw data per minute

Reconstruction

- Takes ~5 minutes
- Reconstructed data ~22GB

Fast Micro-Tomography - Data Workflow

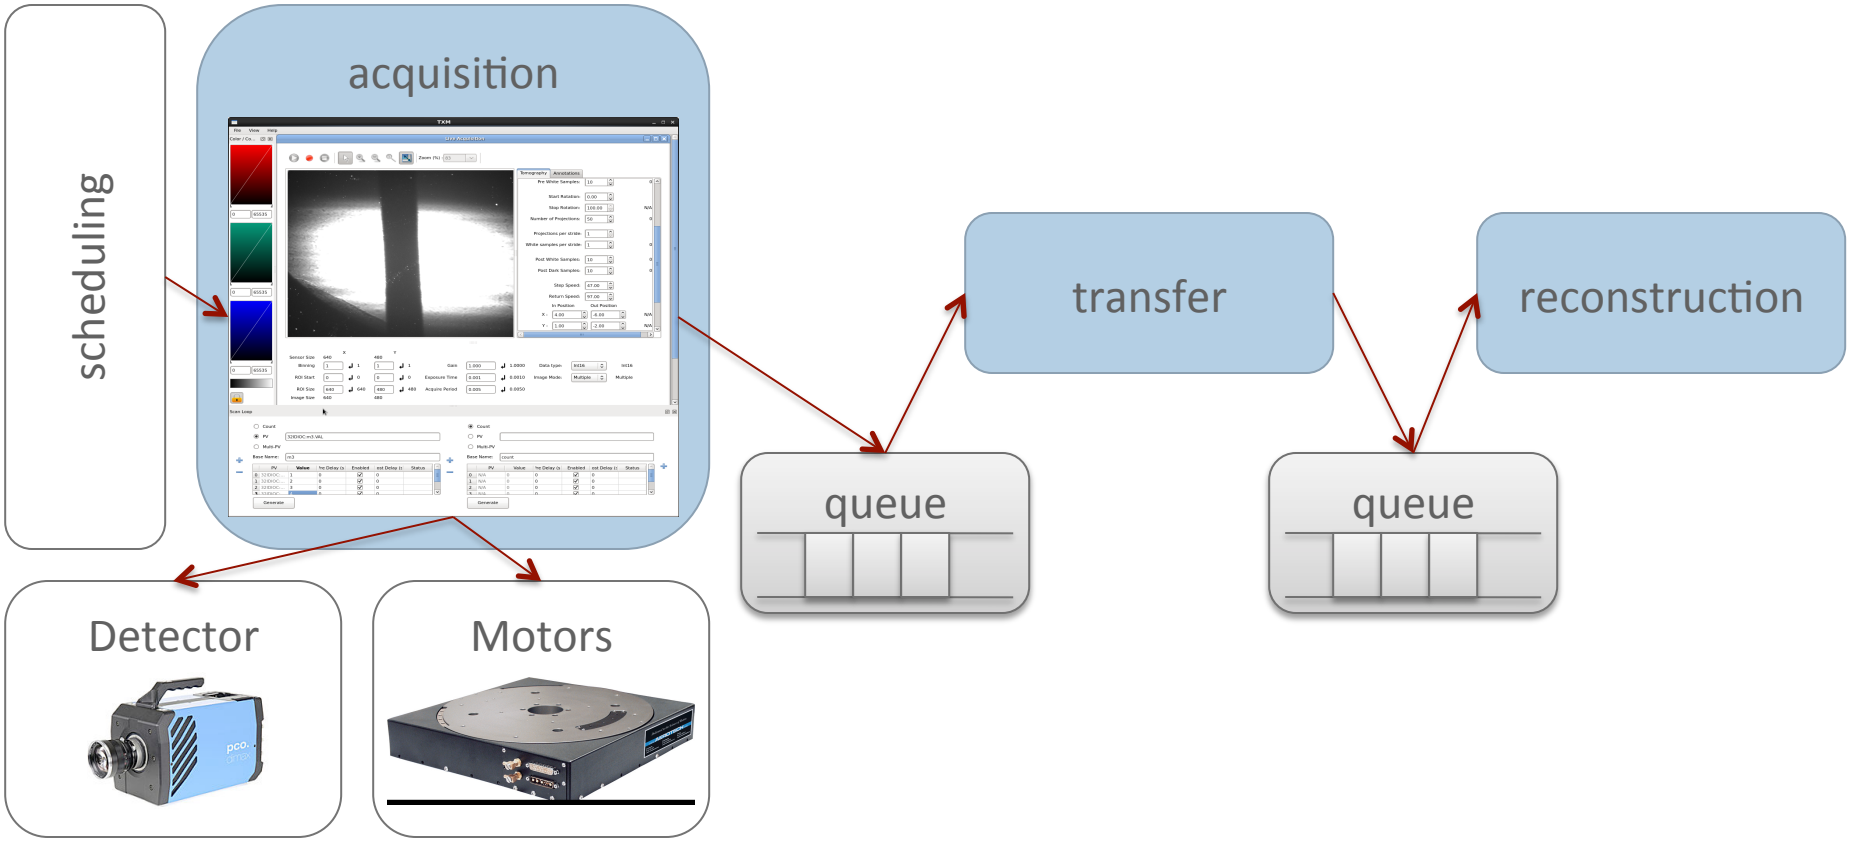

Multimodal ptychographic nanotomography of roll-coated polymer solar cells

Beamline and Detectors

- 2-ID-B, 2.5 keV photons, 150 nm focus
- Small-angle: Cornell MMPAD (266×396×150 µm pixels @ 1 ms/frame)
- Fluorescence: SII Vortex-EX SDD (2048 channels/point)

Tomographic/Ptychographic Dataset

- 70 nm steps, 6 nm pixels (9 µm x 4 µm field)
- 60 projections, recorded every 6°
- 0.1 ms/point, 18 min/projection, 21 h/tomogram
- 129×58 point scan, 7482 diffraction patterns/projection
- **449K diffraction patterns, *176 GB raw data***

Reconstruction

- 50 iterations @ 10 min/iteration
- 12 days/reconstruction (24 cores)
- ⇒ ***308 single-processor compute days!***

D. Vine (APS), J. Andreasen, A. Böttinger (Risø), K. Green, M. Tate, H. Philipp, S. Gruner (Cornell)

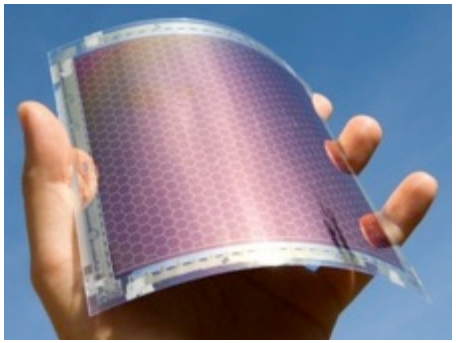

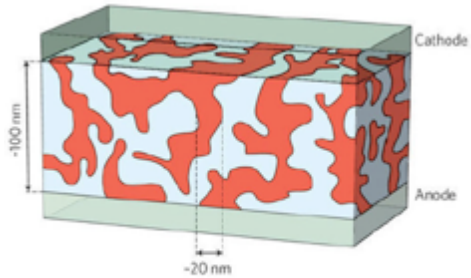

Li *et al.*, Nat. Photon., **6** 153 (2012)

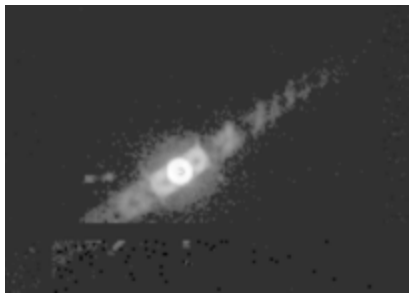

X-ray Laue Diffraction Microscopy

Jon Tischler (APS)

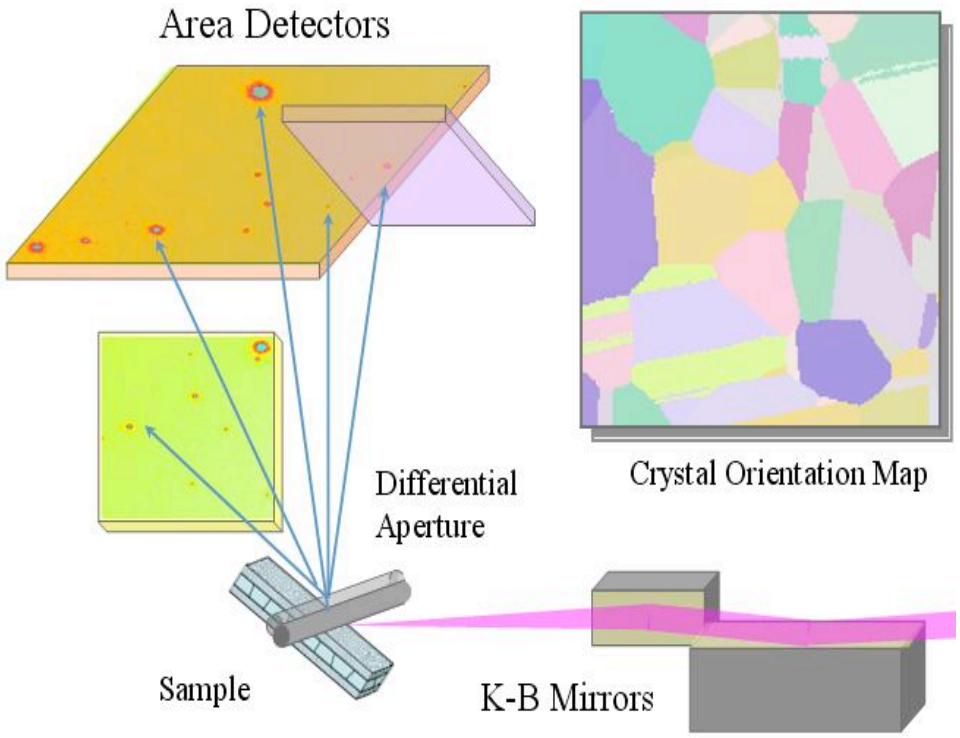

GPU Reconstruction

- Large rewrite / refactor of CPU code / algorithm
- Uses Tesla GPU – 6 GB RAM
- ~10 GB of raw data

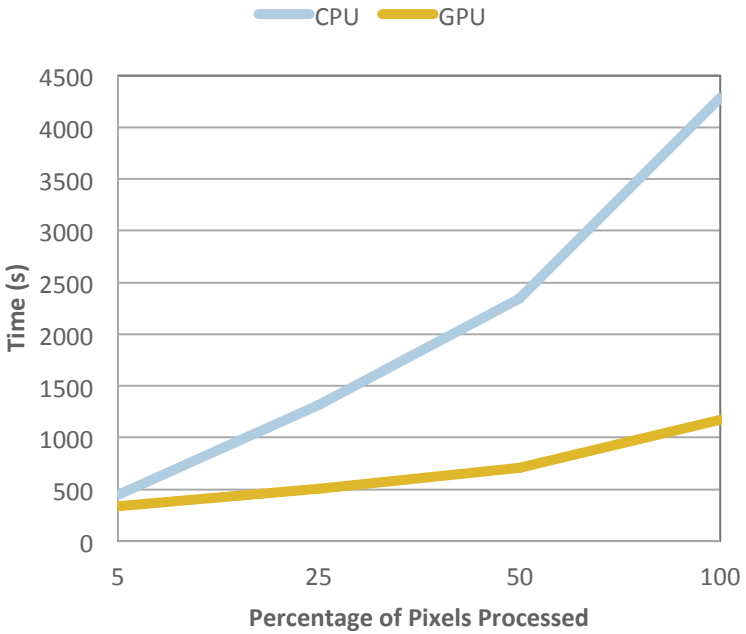

User Access to Data

Globus Online

- Provides access to APS data from the outside
- Used by many user facilities and supercomputing centers
- Optimized bandwidth utilization for faster transfers
- www.globusonline.org

Globus Online

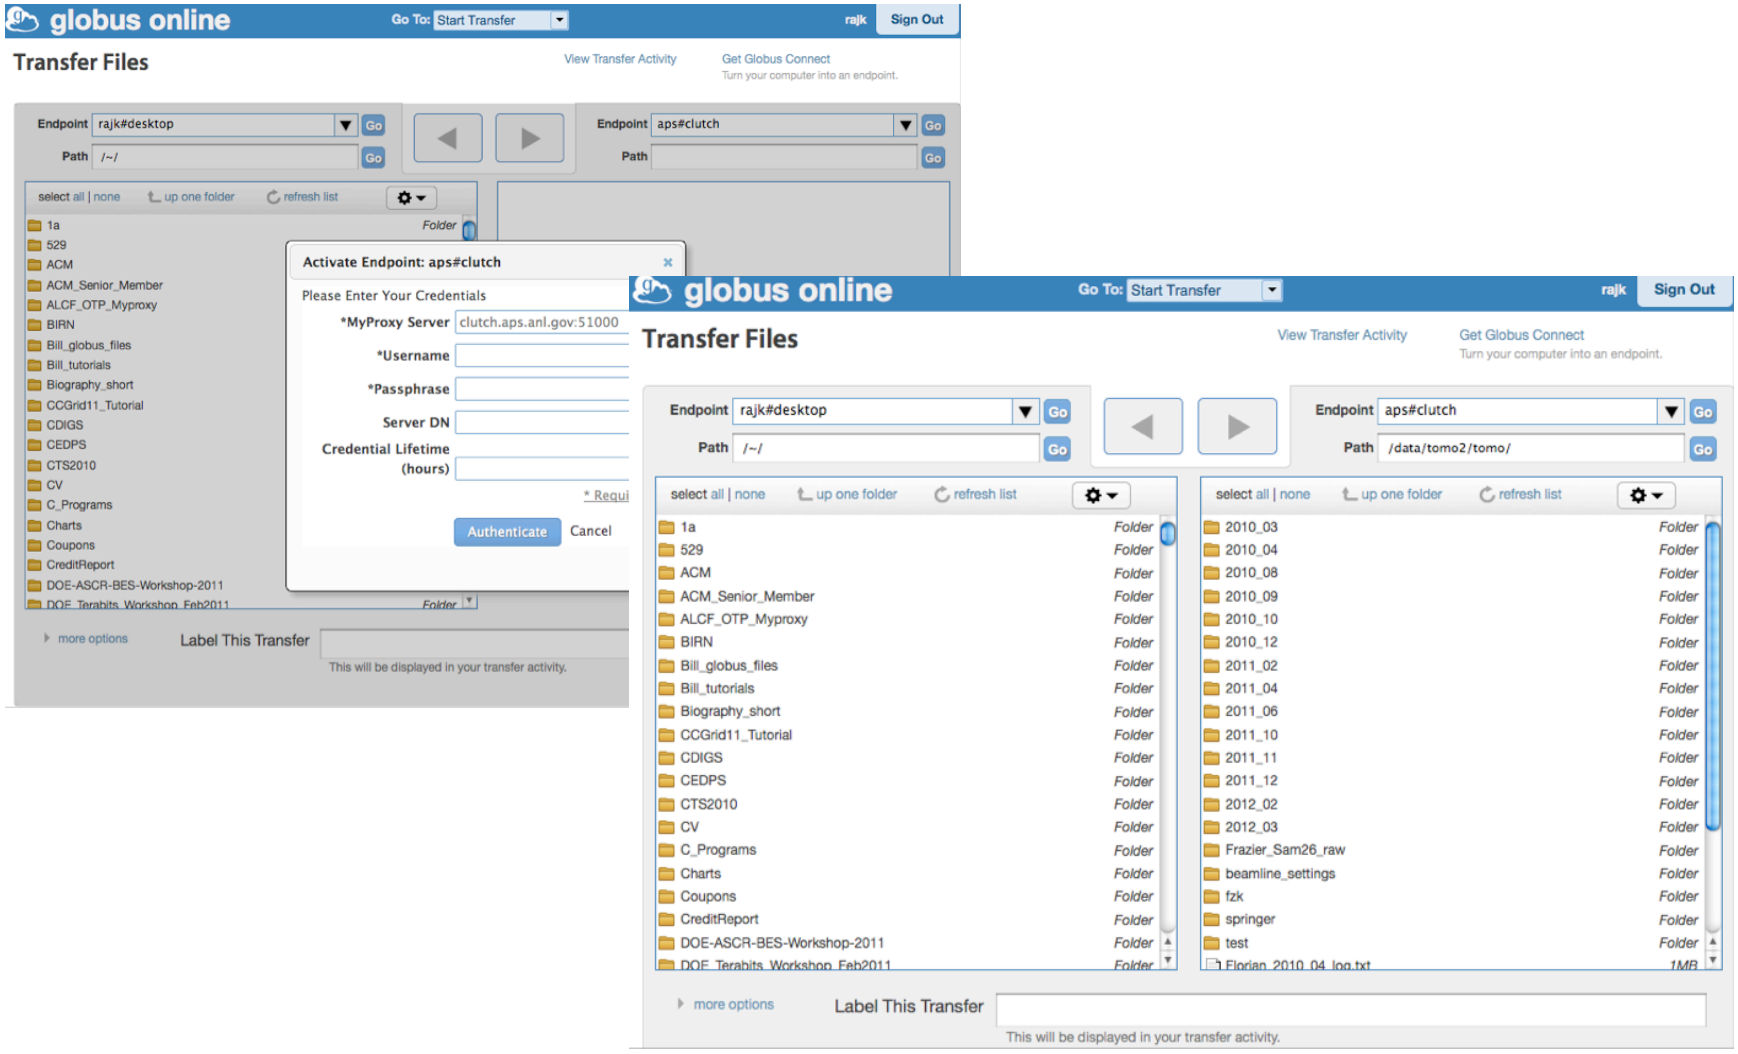

Globus Online Catalog

- Bridge across multiple Globus Online endpoints
- User builds meaningful metadata catalogs
- Not a facility-centric perspective of data

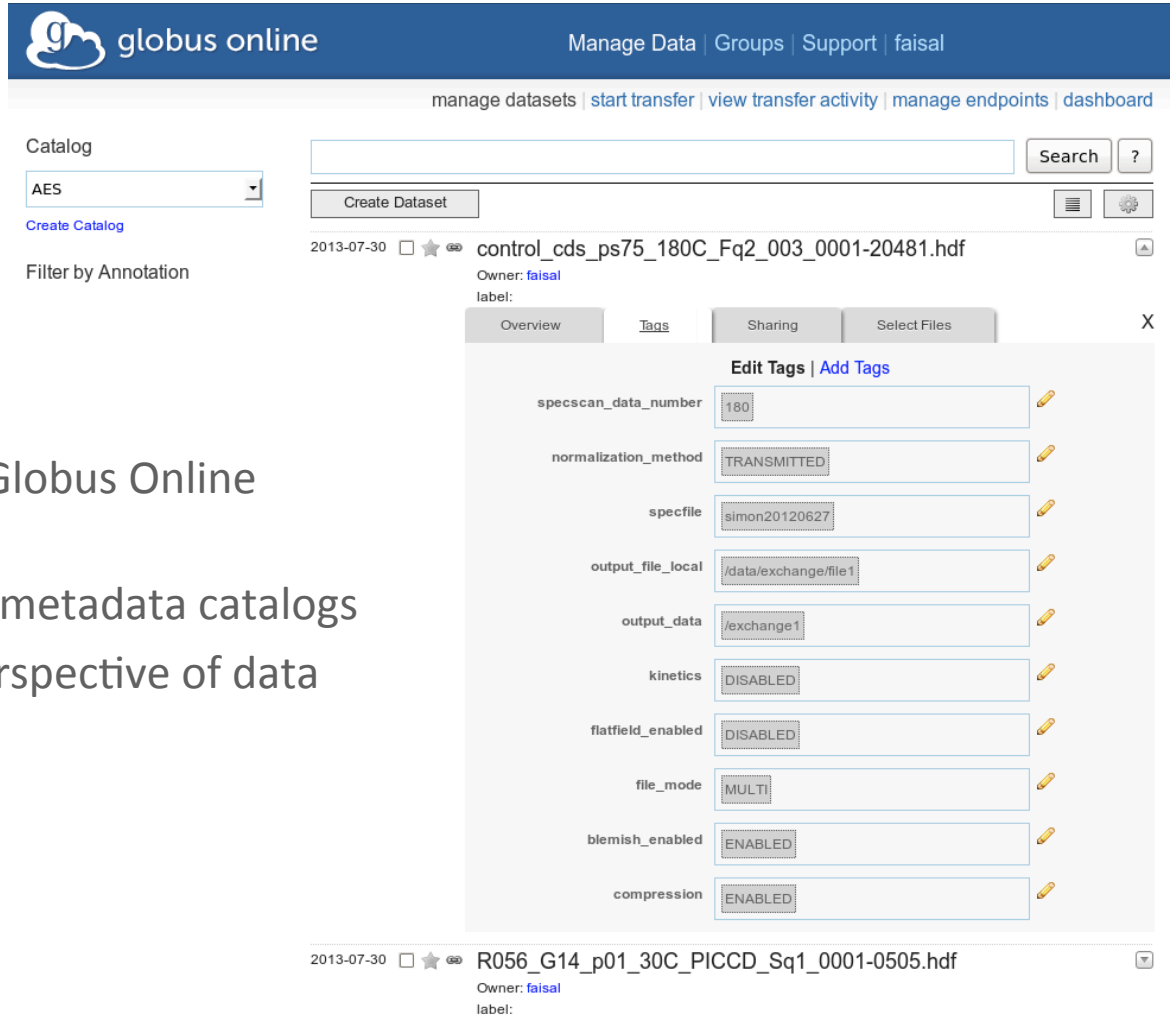

Globus Online Catalog

- Pilot techniques at the APS
- Push data to user specified catalog at end of acquisition
- Read data to a catalog when selected

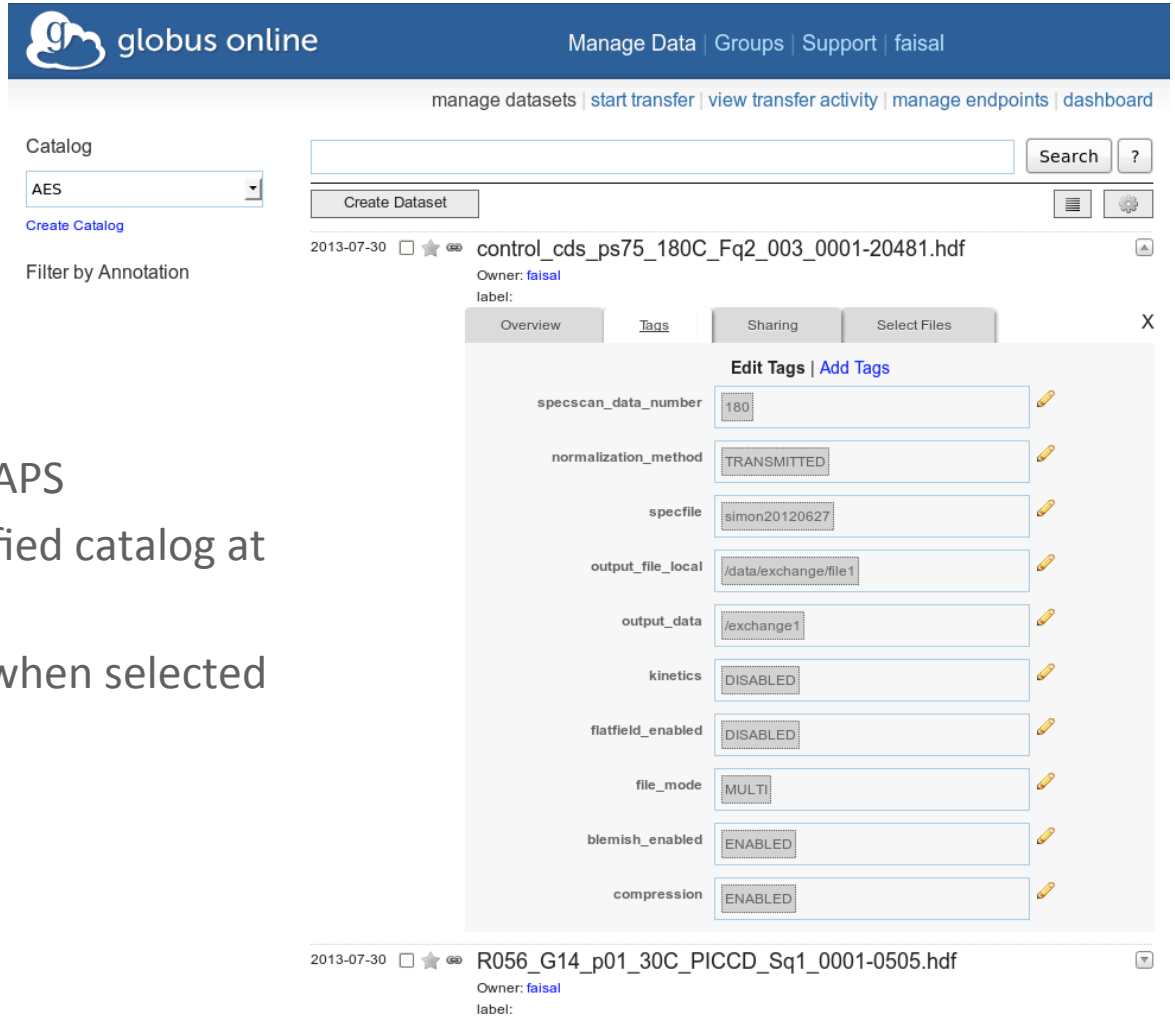

Challenges / Opportunities

- More technique based collaborations
- Tight-knit teams of scientists and software engineers
- Facility integration
- Sustainability – how to keep it working 5 or 10 years in the future?

Thank you!