PaNdata ODI 1st Open Workshop

Dublin 24-25/28th of March 2014

Co-located with the RDA 3rd plenary at Croke Park

https://indico.desy.de/event/1stow

Booklet of collected presentations

PART 3: PaNSIG

Photon and Neutron Science Interest Group meeting

26th and 28th of March at Croke Park / RDA

Thomas Proffen/ Amber Boehnlein	Photon and Neutron Science and facilities - the computing challenges
Tom Griffin:	Data Management solutions – ICAT
Brian Matthews:	Data Management solutions – Tardis
Brian Matthews/ Steve Androulakis	Potential areas for future collaboration and development of common standards and services within RDA
Adam Farquhar:	DataCite
Alun Ashton:	Data analysis issues and frameworks
Eugen Wintersberger:	HDF5
Brian Matthews:	PaNdata
PaNSIG chairs:	Developing a plan of activities for PaNSig

Computing Challenges for Photon and Neutron Facilities

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

Driving Factors: Computing

•Meeting the current science goals for major initiatives such as energy research and materials discovery require improvements in computational tools and techniques.

•Science is driving source upgrades
- Brighter and more precise sources drive detector development
- Detector development drives computing needs—volume, and complexity

•Science also directly drives the computing needs
- Simulations
- Data Analysis and analytics

Facility Challenges

- O(10000) users
  - Diversity of science, needs, skills, longevity
- O(100) beam lines
- O(10) different imaging techniques
- Operational constraints
- Can Do/Make Do Culture
- Source and detector upgrades
  - Challenging in themselves;
  - Environment is changing

Context

It is relatively easy to make a list of common areas of interest for data and computing

Can we collaborate across facilities on those areas productively?

With limited resources, what are good investments?

How do we plan and make the case?

Who, what, where, when and how?

A New Kind of Laser - The LCLS creates X-ray pulses that can capture images of atoms and molecules in motion

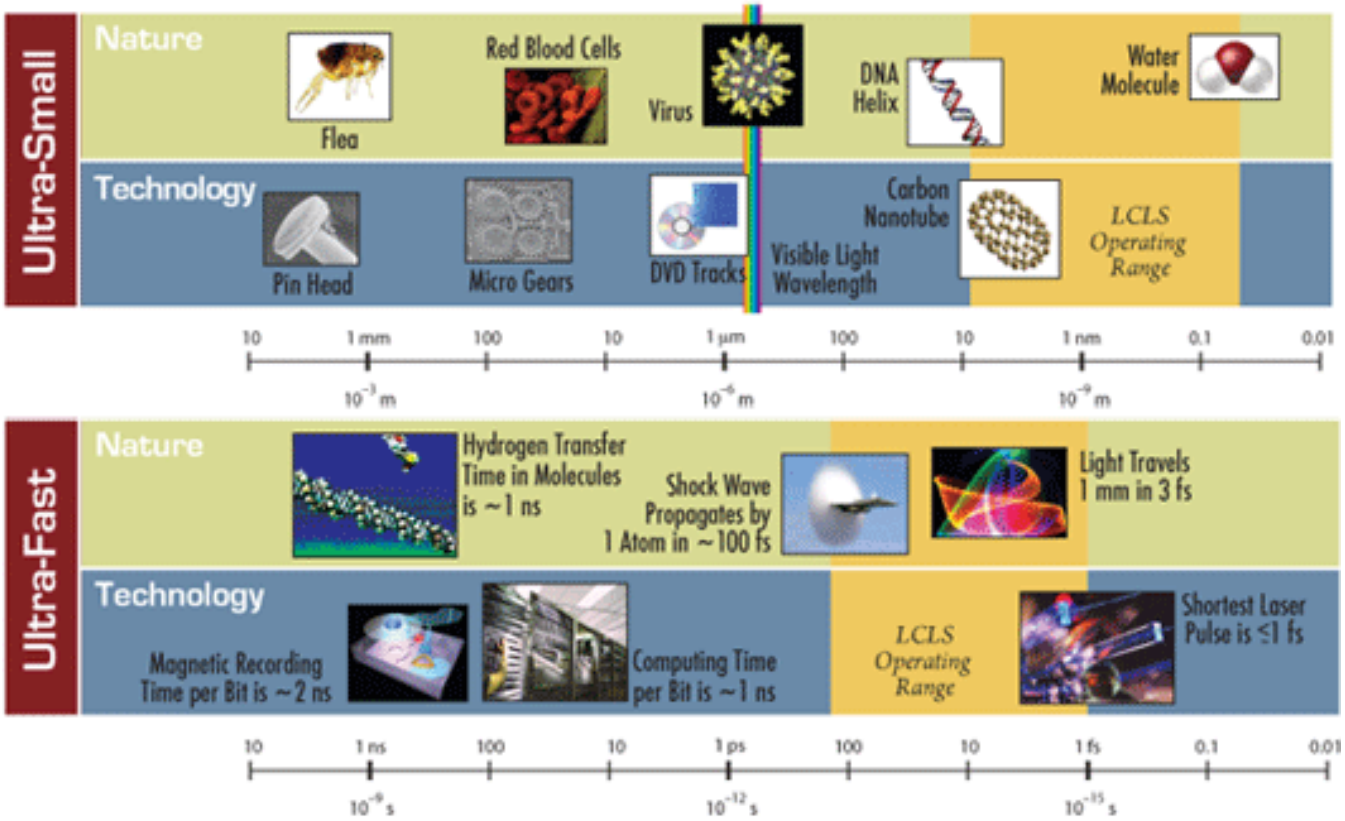

LCLS Data Rates & Volumes

Vary for different instruments/experiments

Three operation modes (X-Ray pulse rate):

- **30 Hz** (first runs on LCLS, AMO/CAMP, Oct-Dec 2009)
- **60 HZ** (second series of runs, May 2010 -> on)
- **120 Hz** (Fall 2010)

Instruments:

- **Fall 2009**: AMO/CAMP
- **Summer 2010**: AMO/CAMP, XPP, SXR
- **Fall 2012:** all 6+1 instruments (station #2 for CXI)

First runs at AMO/CAMP (30 Hz):

- up to **180 MB/s**, **3 TB/day**, **~100 TB** of raw data recorded

At full capacity (120 Hz):

- **Up to 1.5 GB/s** for CXI (event size exceeds 10 MB)
- **30 TB/shift**
- **Up to 3 instruments can take data in parallel (**approaching this condition**)**
- **1-2 PB** of raw data per year

The Data Challenge view from 2012

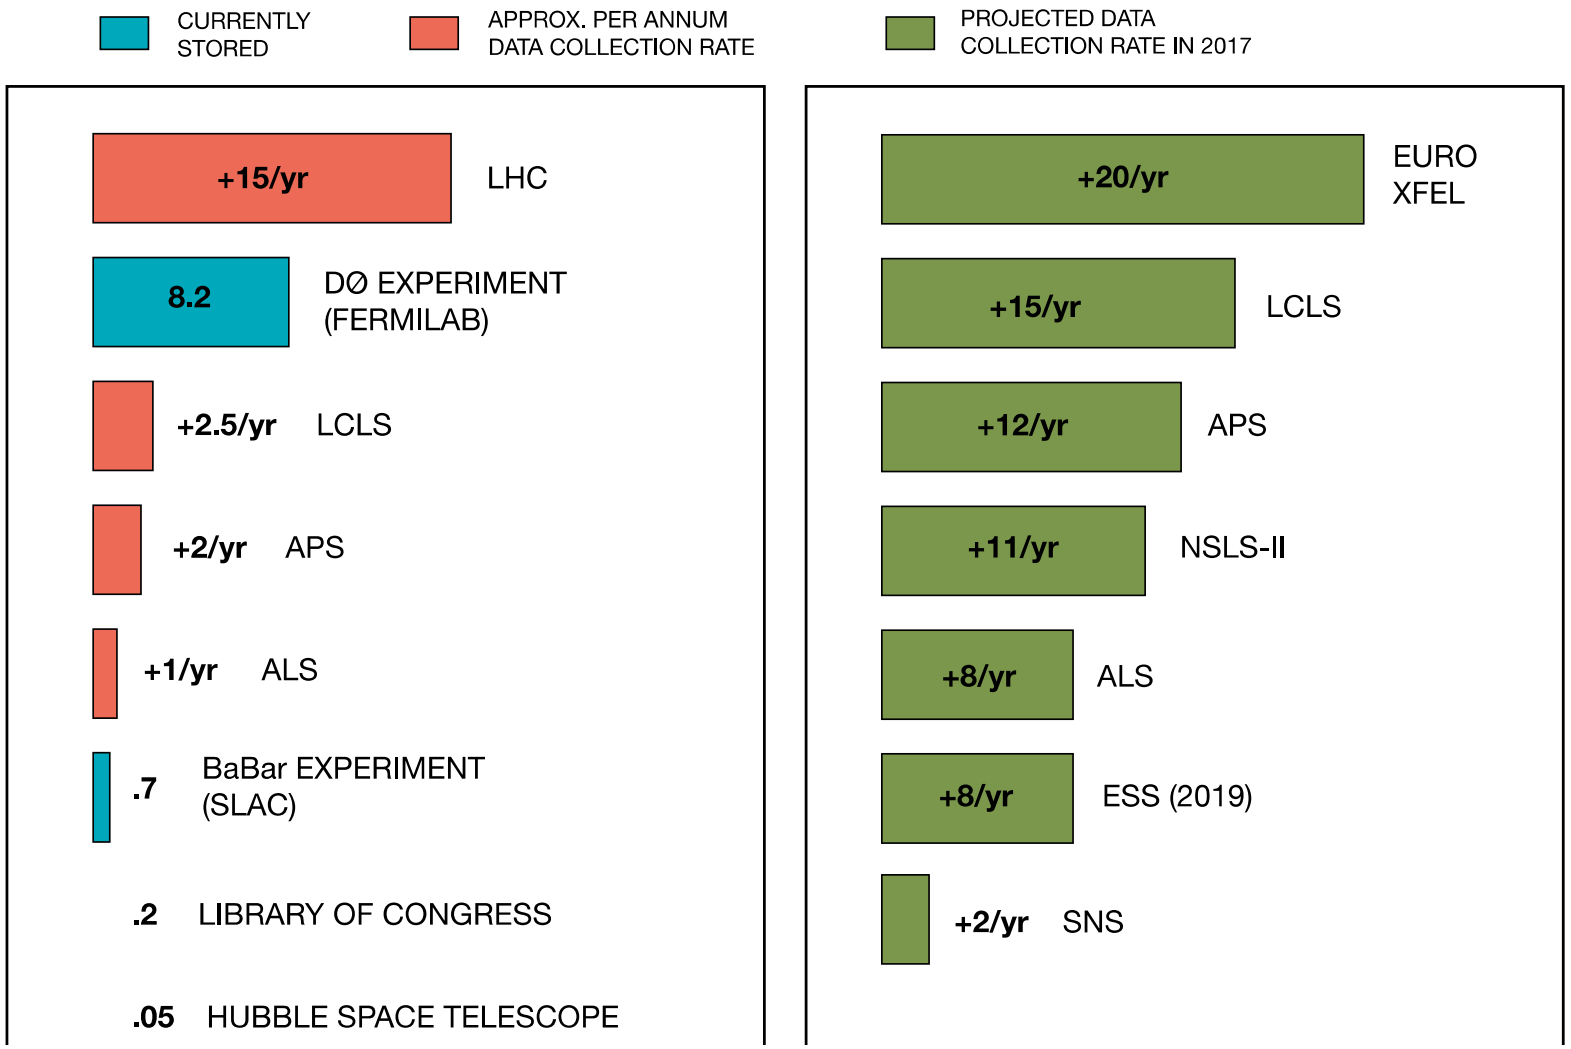

Data Volumes Growing everywhere

SLAC

Data Rates and Data Volumes

Application	Current data rates [GB/h]		Future data rates [GB/h]	
	Peak	Average	Peak	Average
Protein Crystallography	500	50	500	200
Coherent Diffraction Imaging	500	50	4000	400
Tomography	700	50	800	200
Spectroscopy	450	45	18000	1800
Small Angle Scattering	1400	140	14400	4200
Grain Mapping	140	80	800	300

DESY: CFEL/PETRA III+/FLASH → 1.6 PB/year
Cern: ATLAS 100 MB/s → 3 PB/year

- Increasing source brilliance
- Faster detectors with more pixels (e.g. Pilatus)
- New experiment scenarios: Fast time resolved, scanning
- Use of simultaneous Detectors

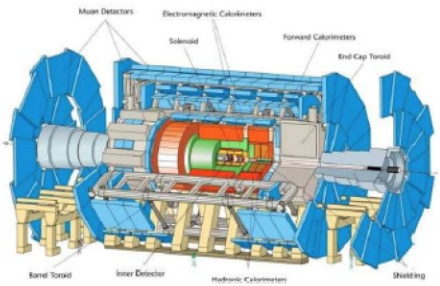

DESY

HELMHOLTZ GEMEINSCHAFT

LCLS Diffract & Destroy

Photo-system II, Nick Sauter (LBL): A single experiment in 2013 generated **120TB** over four days. Data was relayed to NERSC via Esnet, analysis required **135K CPU hours** of computing.

LCLS-II will require > **100M CPU hours** per experiment. Higher resolution and advanced image analysis could grow computational complexity. Some algorithms scale an M*NlogN for M images of N pixels.

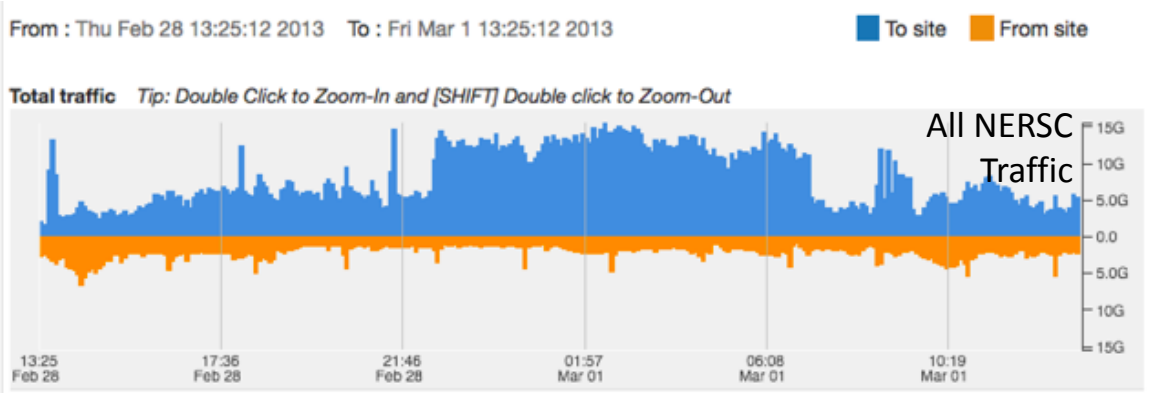

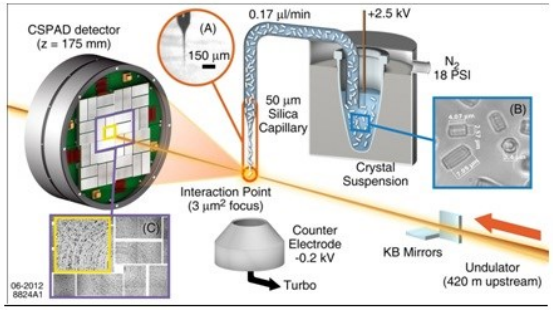

A starter wish list…

Easy to use data management and processing frameworks that scale with data rate
Conditions data storage and management; meta data
Data validation and fast feedback at collection
Data volume reduction/compression techniques
Development and availability of algorithm and algorithm tools
Community developed simulations and simulation tool kits for beamlines and detectors
Computational science

Compute and storage hardware platforms appropriate for the task**S**

Data management

- Life is easier if the data is managed from point of origin
- Detector readouts can have proprietary readouts

Is standardization possible?
Data containers/formats
Metadata
Workflows/tool kits
Visualization tools
Curation/cataloging

Naively, some standardization would make good use of resources and expertise, simplify life, open up many possibilities

Rapid feedback

Time is money

Beam time, experimenter time; instrument scientist time

Experiment simulations

Advance preparation of analysis tools

Could provide better use of beam time, shorten time to publication and refine where to spend resources on computing improvements

http://home.slac.stanford.edu/pressreleases/2011/20110202.htm

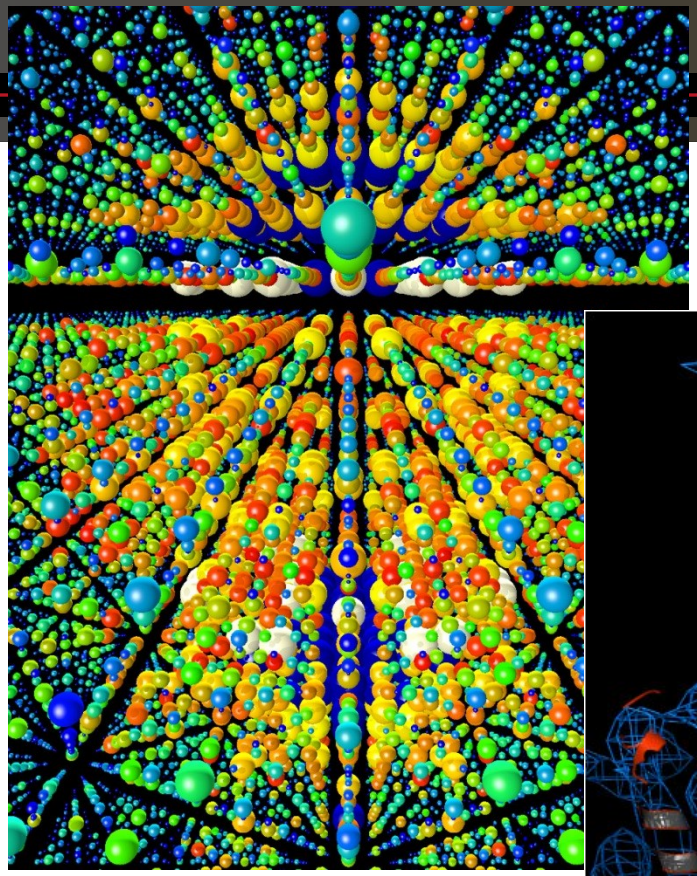

A reconstructed image of t[illegible]osystem I complex. (Image courtesy Raimund Fromme, Arizona State University.)

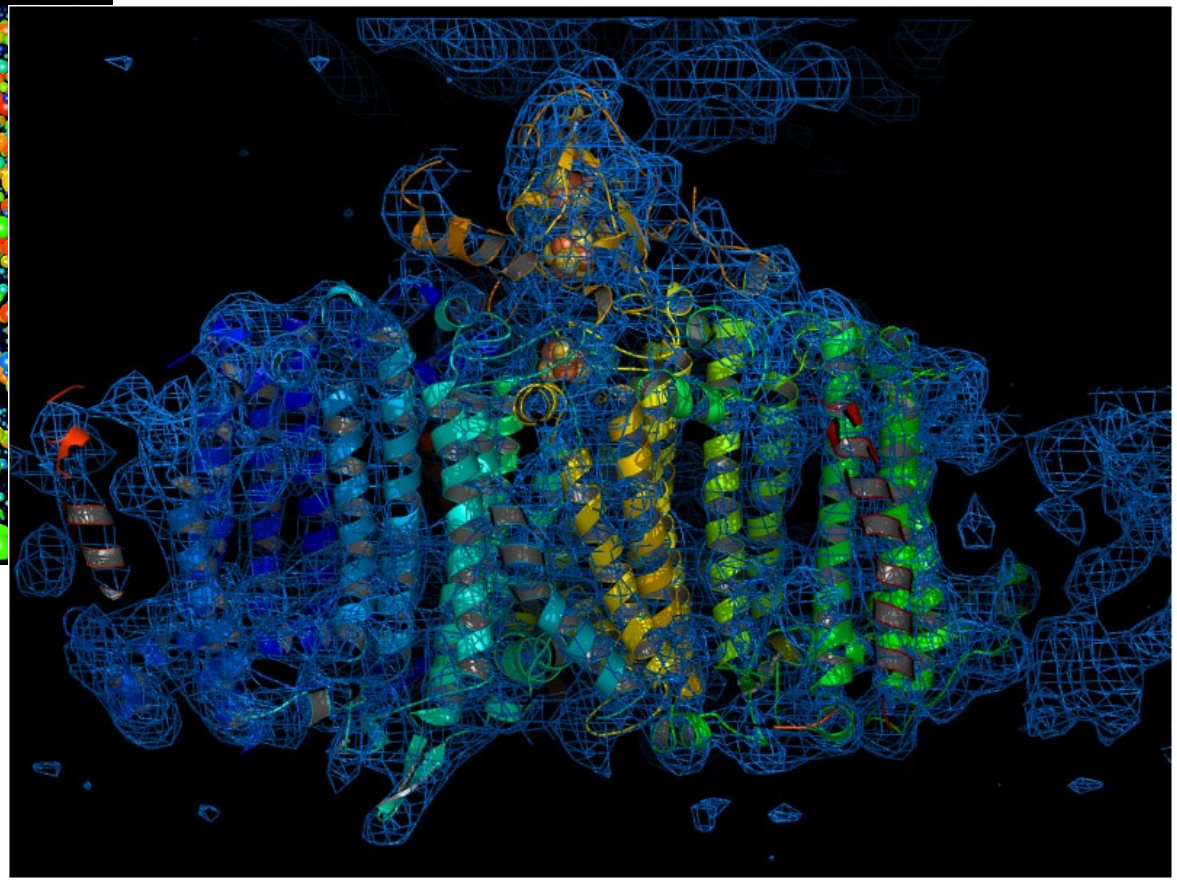

Three-dimensional rendering of the X-ray diffraction pattern for the Photosystem I protein, reconstructed from more than 15,000 single nanocrystal snapshots taken at the LCLS. (Image courtesy Thomas White, DESY.)

High Throughput XRD for Material Discovery

Strategy:

- Ternary and Quaternary Combinatorial Libraries
- Screen for Properties
  - HT Spectroscopy
  - New Materials
- Determine Structural Phase Diagram — Sam Webb
  - HT XRD, HT EXAFS
  - New Phase Diagrams
- Composition-Structure-Property Relations
- Phase Transitions
  - Metastable & Near Equilibrium
- Input for Atomistic Material Modeling
  - New Theory for (Metastable) Materials

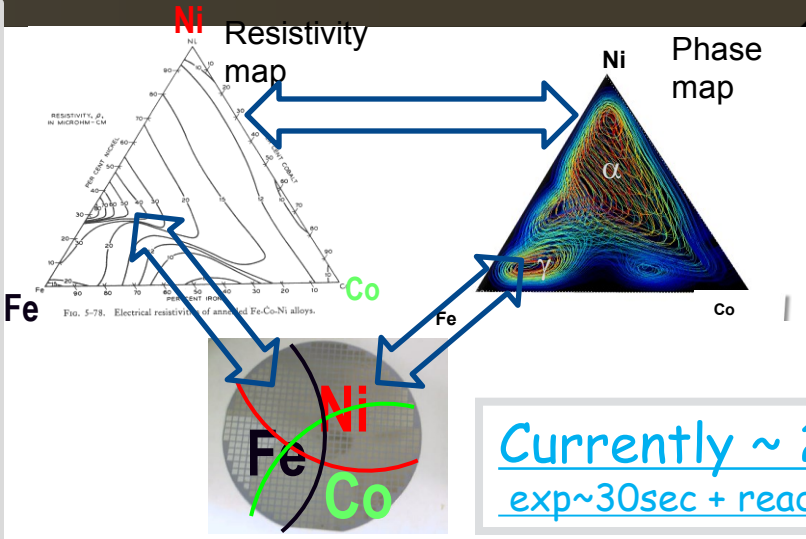

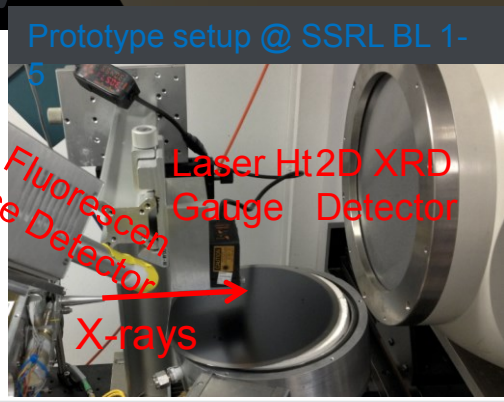

Currently ~ 2000xrd+mca/day
exp~30sec + readout ~10sec

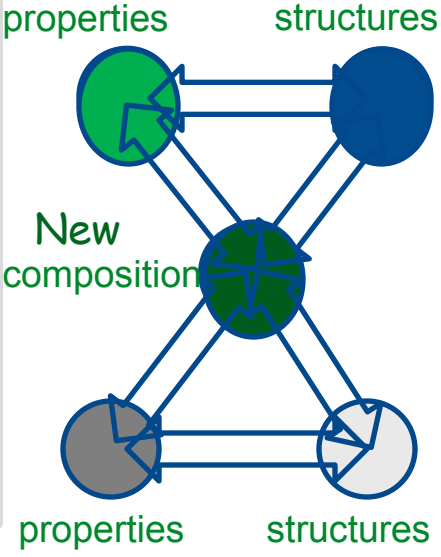

Moving Forward with HT-XRD :

Enhance Capabilities:

- Auto Alignment and Calibrations
- Low Z Detection Sensitivity
- Rapid Annealing

Increase Throughput:

- New X-ray Focusing Optics: exp ~1-5sec
- New 2D detector: readout <1 sec
  - → 20-50,000 xrd+mca/day
  - Need Automation and Robotics

Data Management:

High-Throughput Pipeline:

Robotics, Automation & Machine Learning

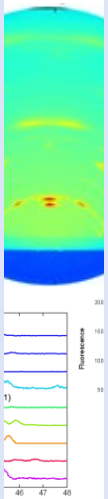

Hubs

SSRL

Hubs

John Gregoire
Ichiro Takeuchi
Matt Kramer
Apurva Mehta

New Composition-Structure-Property Relations

$_{40}Ga_{60}$

New Materials Scientific Communities

New Theory of (metastable) Materials

tabase

CD

COMPUTING CHALLENGES FOR PHOTON AND NEUTRON FACILITIES

Thomas Proffen, ORNL
Amber Boehnlein, SLAC

Creating a Pathway for Scientific Discovery

- Accelerating discovery in materials science
- Enhancing predictive capabilities

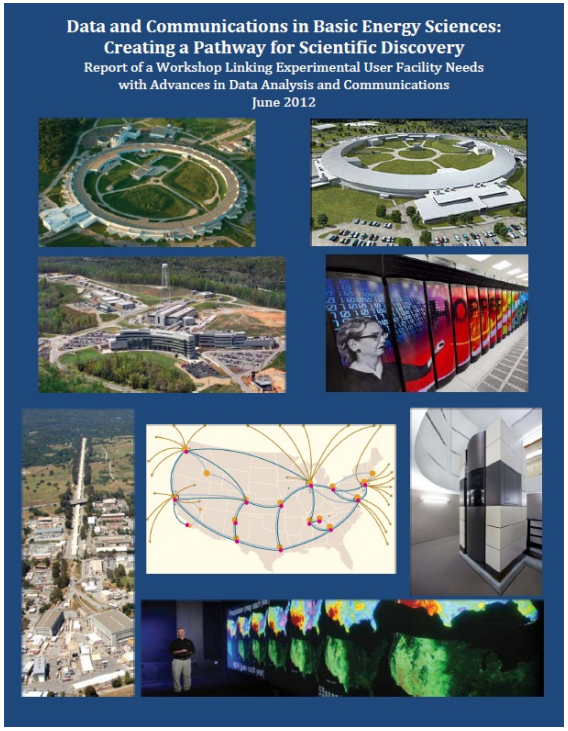

- Theory and analysis components should be integrated seamlessly within experimental workflow.
- Move analysis closer to experiment – future possibility of experiment steering.
- Match data management access and capabilities with advancements in detectors and sources.

Driving Factors: Computing

- Meeting the current science goals for major initiatives such as energy research and materials discovery require improvements in computational tools and techniques.

- Science is driving source upgrades
  - Brighter and more precise sources drive detector development
  - Detector development drives computing needs—volume, and complexity

- Science also directly drives the computing needs
  - Simulations
  - Data Analysis and analytics

Facility Challenges

- O(10000) users
  - Diversity of science, needs, skills, longevity
- O(100) beam lines
- O(10) different imaging techniques
- Operational constraints
- Can Do/Make Do Culture
- Source and detector upgrades
  - Challenging in themselves;
  - Environment is changing

Context

- It is relatively easy to make a list of common areas of interest for data and computing
- Can we collaborate across facilities on those areas productively?
- With limited resources, what are good investments?
- How do we plan and make the case?
- Who, what, where, when and how?
-

Neutron Data Life Cycle

Feedback

DAS
- Neutron events
- Events from sample environment
- Other triggers

Reduction
- Corrected reduced data (histograms, S(Q,E), ..)
- Merging, reconstruction of data
- Instrument/technique dependent
- Need for 'real' time reduction

Analysis
- Multi dimensional fitting
- Advanced visualization
- Comparison to simulation / feedback
- Field dependent, large variety of approaches

Simulation Modeling
- Multitude of techniques (DFT, MD, ..)
- Advanced simulation of experiments
- 'Refinement' using experimental data
- Multiple experiments / probes

User Facility

- Variety of experiments, topics, methods and 'computer literacy' of users are significant challenge.

Example: NOMAD Diffractometer

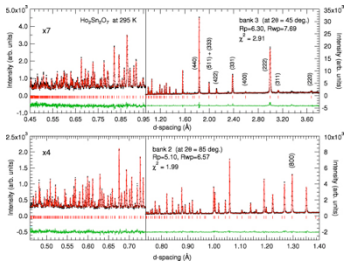

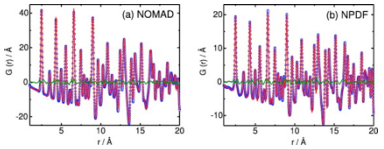

Raw Data: up to 10^12 events per second

Acquisition

Translated Data: Gigabytes to Terabytes

Reduction

Reduced Data: e.g. Powder Diffraction Pattern

Analysis & Simulation

Analysis: PDF, MD simulation, etc.

Feedback guiding changes to the experiment setup

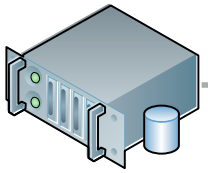

Data captured and stored on multiple systems at the beamline

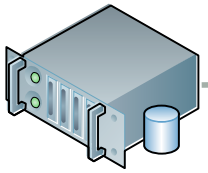

After completion of a “run” data is aggregated on a single system, translation begins

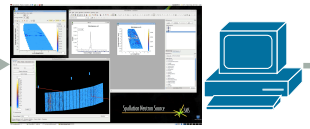

Once data is aggregated reduction begins using a workstation

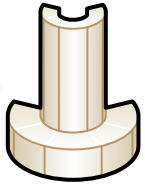

Analysis and Simulation using mid-scale compute

Improving Productivity = Changing the Workflow

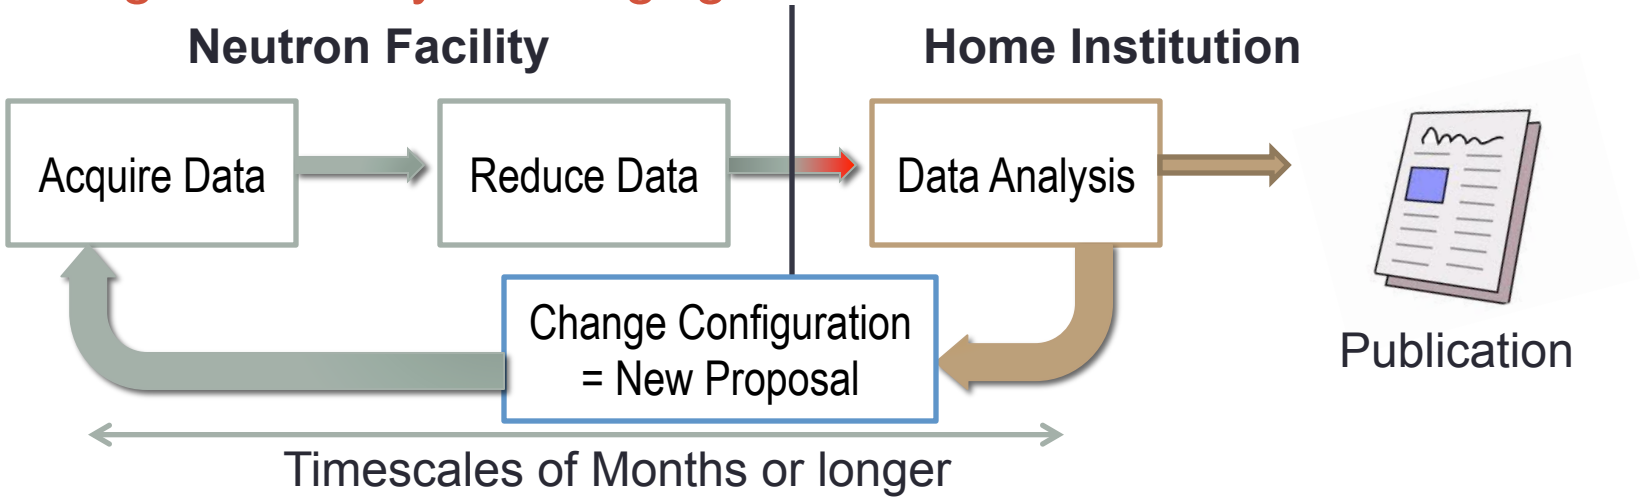

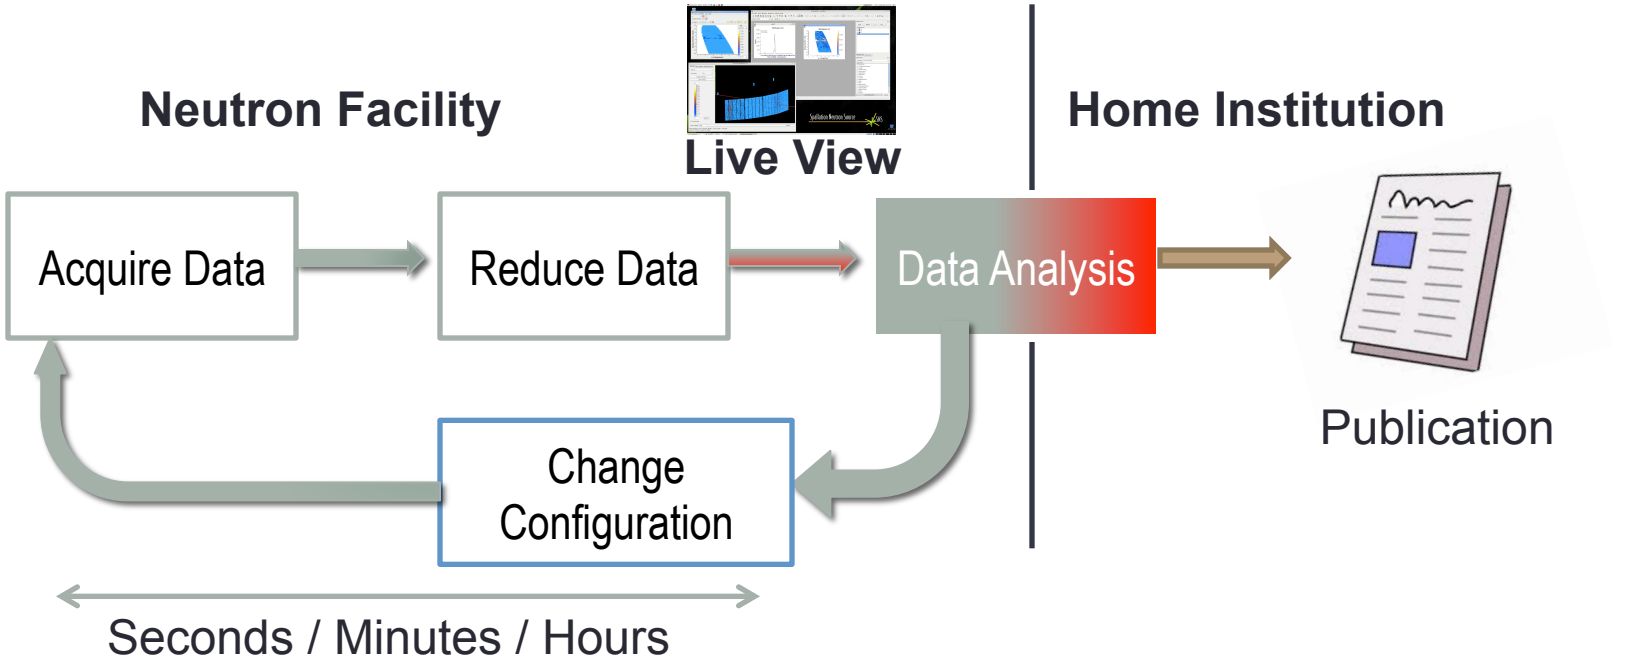

ADARA is enabling real-time feedback from experiment, analysis and computational steering

ADARA

- Leverages our multi-disciplinary capabilities at ORNL coupling Neutron Sciences Directorate with Computing and Computational Sciences Directorate.
- The ADARA Project lets us stream data to computational resources and provide live feedback from experiment in real-time S(Q,E).
- Provides a high performance data backplane for reduction, analysis, and coupling with simulation forming the basis for future work to integrate experiment and simulation.
- Prototype running on HYSPEC instrument. Deployment to other beamlines in 2013/2014.

ORNL has launched the Center for Accelerating Materials Modeling (CAMM)

- The CAMM will integrate materials modeling/simulation (MD/DFT) directly into the chain for neutron scattering data analysis, **offline** and **online** (in near real time)
- Developing workflows for refinement, integration of MD codes, **neutron scattering corrections** ..
- The CAMM is working with ORNL's Materials Science and Technology Division to study coarse grained MD simulations of polymers PEO-AA (CNMS), *ab-initio* MD simulations for ferroelectrics/thermoelectrics

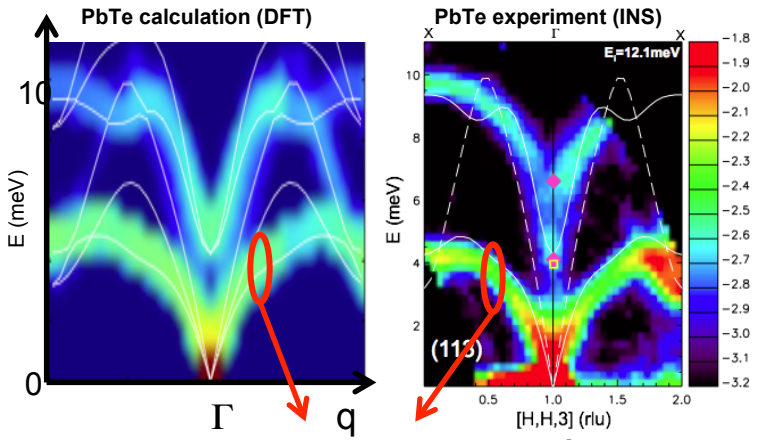

Example: *ab-initio* MD simulations for ferroelectrics/thermoelectrics. Focus on *width* of dispersions

The Center for Accelerating Materials Modeling (CAMM)

- *Partnership between ORNL's Neutron Sciences, Physical Sciences and Computing and Computational Sciences Directorates*
- *ORNL SEED money and DOE funds provided to study force field refinement from quasi-elastic and inelastic neutron scattering data*
- *CAMM formed in response to BES proposal call for* Predictive Theory and Modeling

A New Kind of Laser - The LCLS creates X-ray pulses that can capture images of atoms and molecules in motion

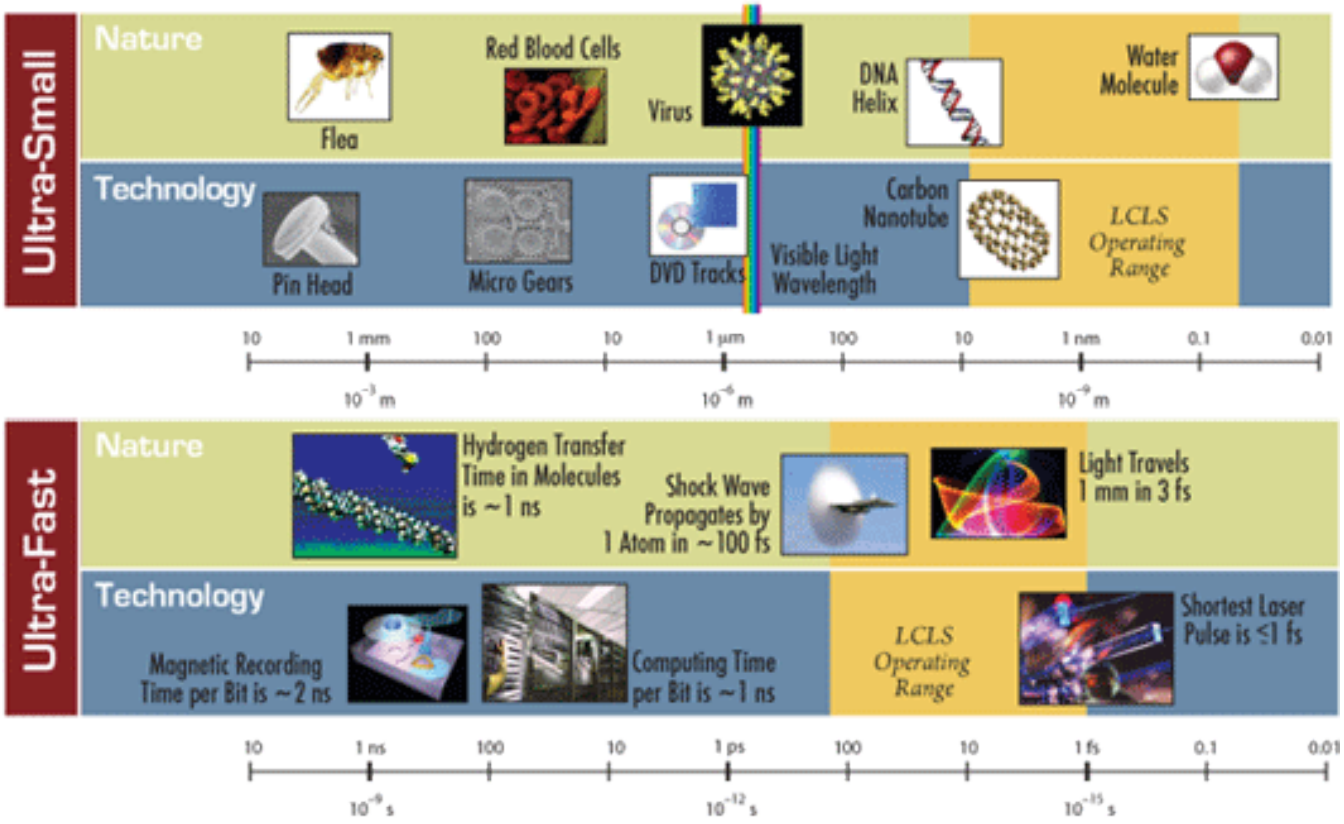

LCLS Data Rates & Volumes

- Three operation modes (X-Ray pulse rate):
  - **30 Hz** (first runs on LCLS, AMO/CAMP, Oct-Dec 2009)
  - **60 HZ** (second series of runs, May 2010 -> on)
  - **120 Hz** (Fall 2010)
- Instruments:
  - **Fall 2009**: AMO/CAMP
  - **Summer 2010**: AMO/CAMP, XPP, SXR
  - **Fall 2012:** all 6+1 instruments (station #2 for CXI)
- First runs at AMO/CAMP (30 Hz):
  - up to **180 MB/s**, **3 TB/day**, **~100 TB** of raw data recorded
- At full capacity (120 Hz):
  - **Up to 1.5 GB/s** for CXI (event size exceeds 10 MB)
  - **30 TB/shift**
  - **Up to 3 instruments can take data in parallel (**approaching this condition**)**
  - **1-2 PB** of raw data per year

Vary for different instruments/experiments

The Data Challenge view from 2012

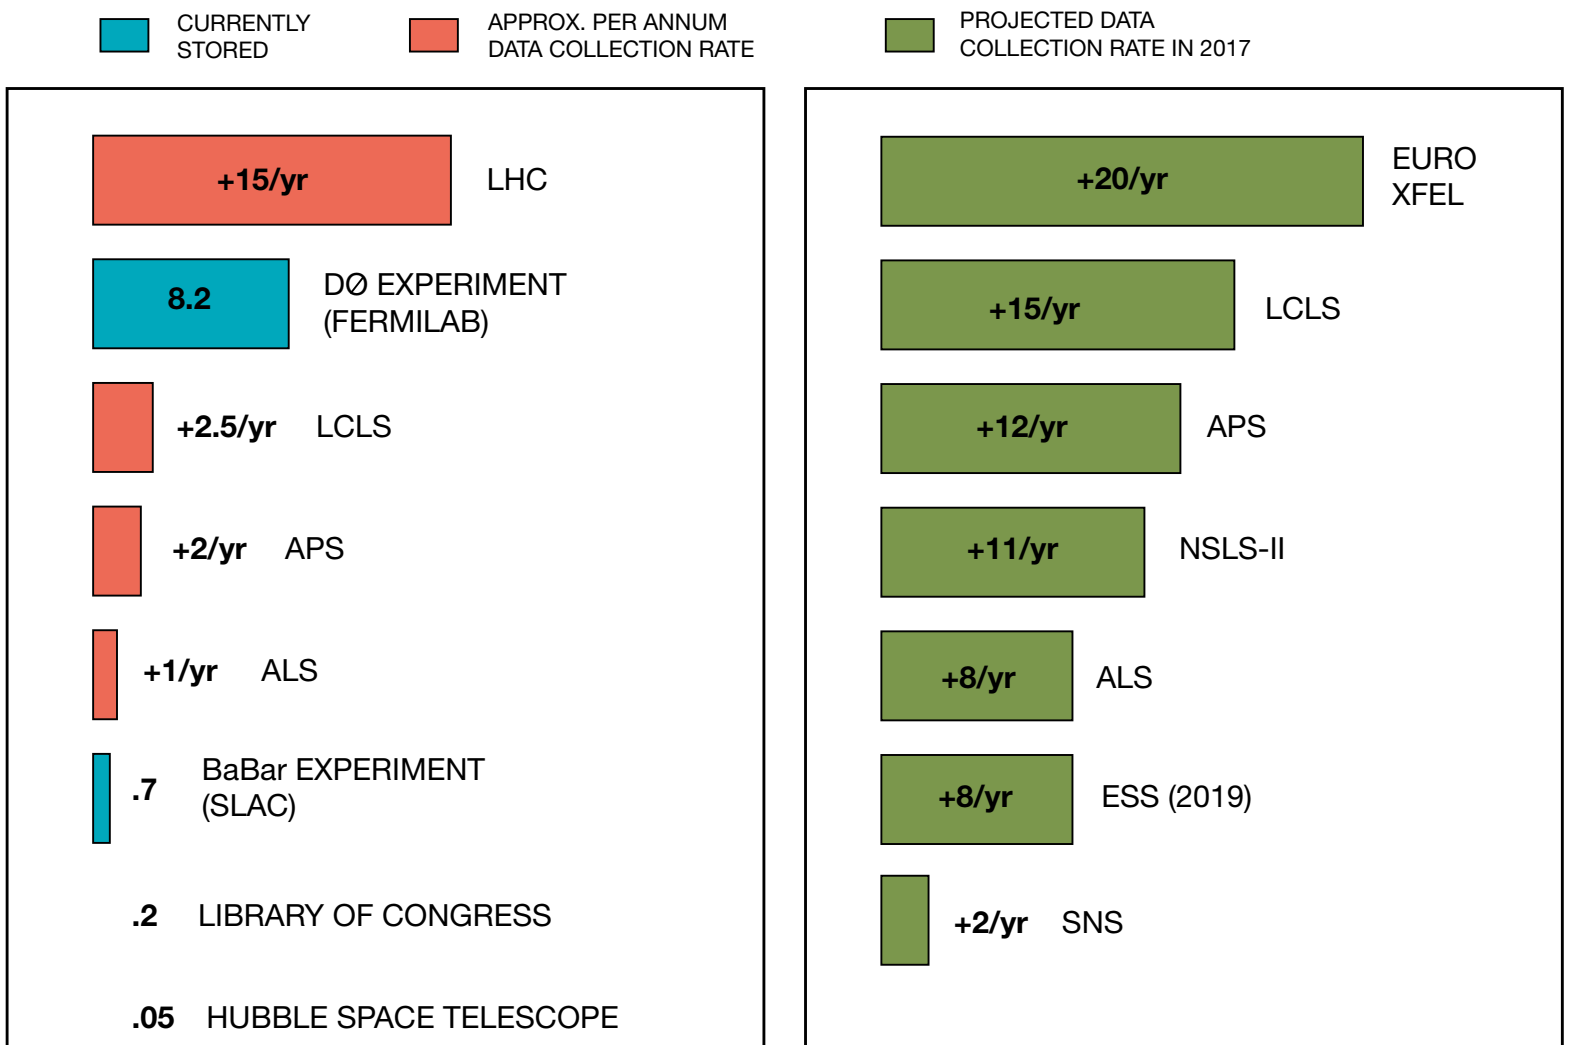

Data Volumes Growing everywhere

Data Rates and Data Volumes

Application	Current data rates [GB/h]		Future data rates [GB/h]	
	Peak	Average	Peak	Average
Protein Crystallography	500	50	500	200
Coherent Diffraction Imaging	500	50	4000	400
Tomography	700	50	800	200
Spectroscopy	450	45	18000	1800
Small Angle Scattering	1400	140	14400	4200
Grain Mapping	140	80	800	300

DESY: CFEL/PETRA III+/FLASH → 1.6 PB/year
Cern: ATLAS 100 MB/s → 3 PB/year

- Increasing source brilliance
- Faster detectors with more pixels (e.g. Pilatus)
- New experiment scenarios: Fast time resolved, scanning
- Use of simultaneous Detectors

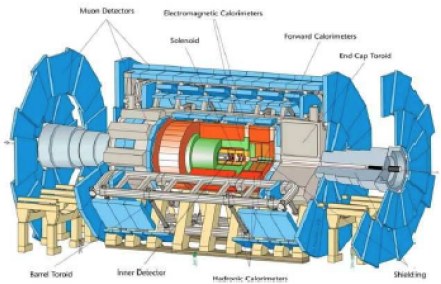

DESY

HELMHOLTZ GEMEINSCHAFT

LCLS Diffract & Destroy

Photo-system II, Nick Sauter (LBL): A single experiment in 2013 generated **120TB** over four days. Data was relayed to NERSC via Esnet, analysis required **135K CPU hours** of computing.

LCLS-II will require > **100M CPU hours** per experiment. Higher resolution and advanced image analysis could grow computational complexity. Some algorithms scale an M*NlogN for M images of N pixels.

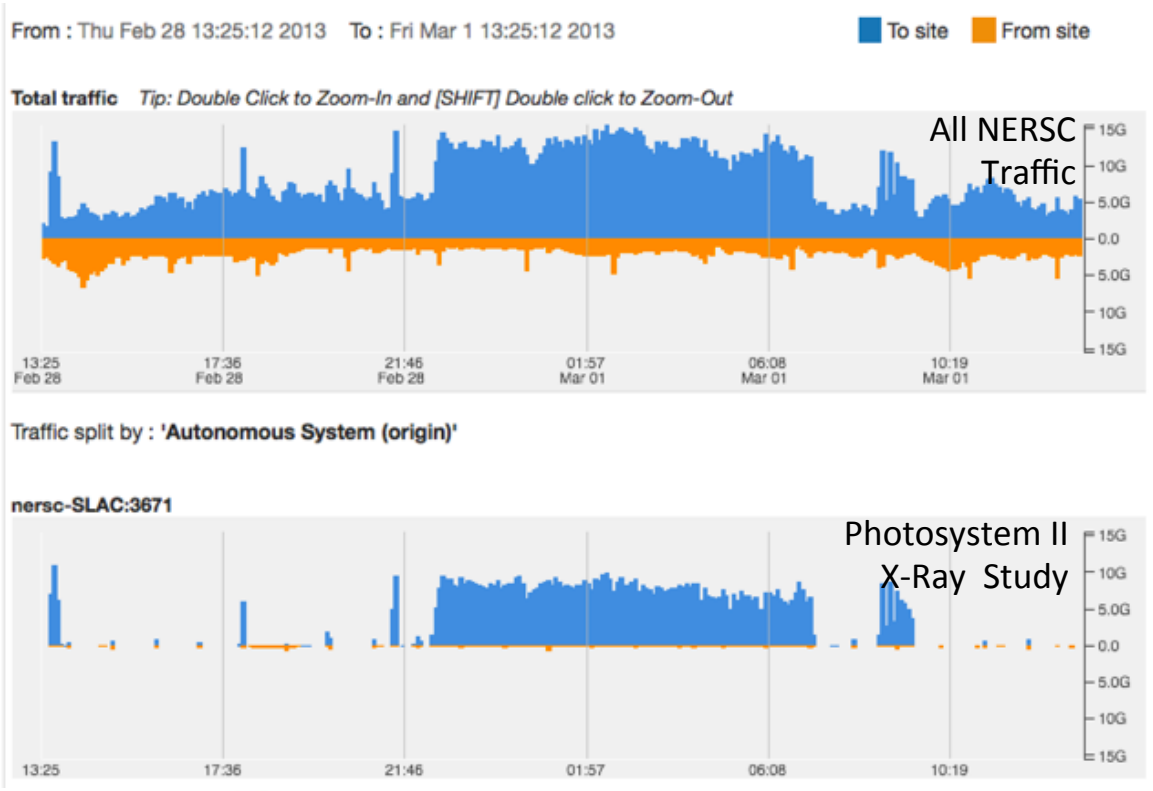

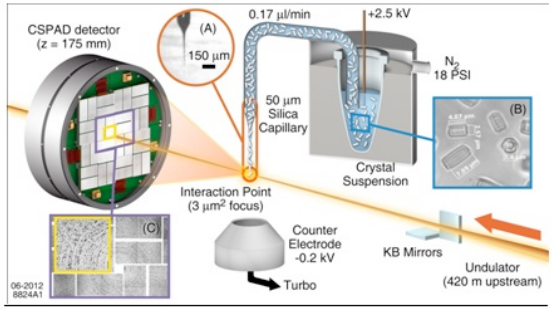

ESnet Energy Sciences Network

A starter wish list…

- Easy to use data management and processing frameworks that scale with data rate
- Conditions data storage and management; meta data
- Data validation and fast feedback at collection
- Data volume reduction/compression techniques
- Development and availability of algorithm and algorithm tools
- Community developed simulations and simulation tool kits for beamlines and detectors
- Computational science
- Compute and storage hardware platforms appropriate for the task**S**

Data management

- Life is easier if the data is managed from point of origin
- Detector readouts can have proprietary readouts

- Is standardization possible?
  - Data containers/formats
  - Metadata
  - Workflows/tool kits
  - Visualization tools
  - Curation/cataloging

- Naively, some standardization would make good use of resources and expertise, simplify life, open up many possibilities

Rapid feedback

- Time is money
  - Beam time, experimenter time; instrument scientist time
- Experiment simulations
- Advance preparation of analysis tools

- Could provide better use of beam time, shorten time to publication and refine where to spend resources on computing improvements

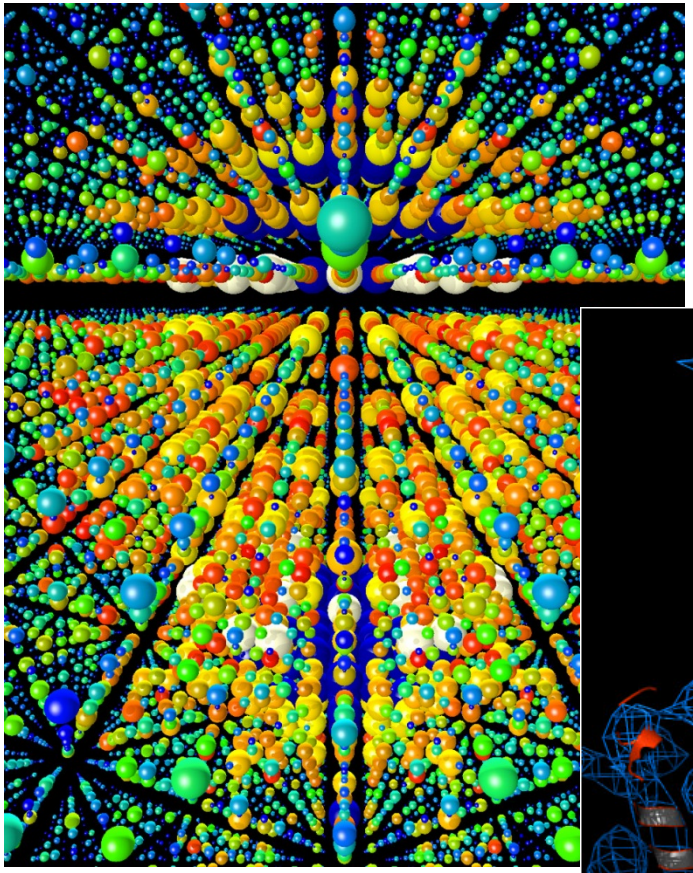

Three-dimensional rendering of the X-ray diffraction pattern for the Photosystem I protein, reconstructed from more than 15,000 single nanocrystal snapshots taken at the LCLS. (Image courtesy Thomas White, DESY.)

A reconstructed image of the Photosystem I complex. (Image courtesy Raimund Fromme, Arizona State University.)

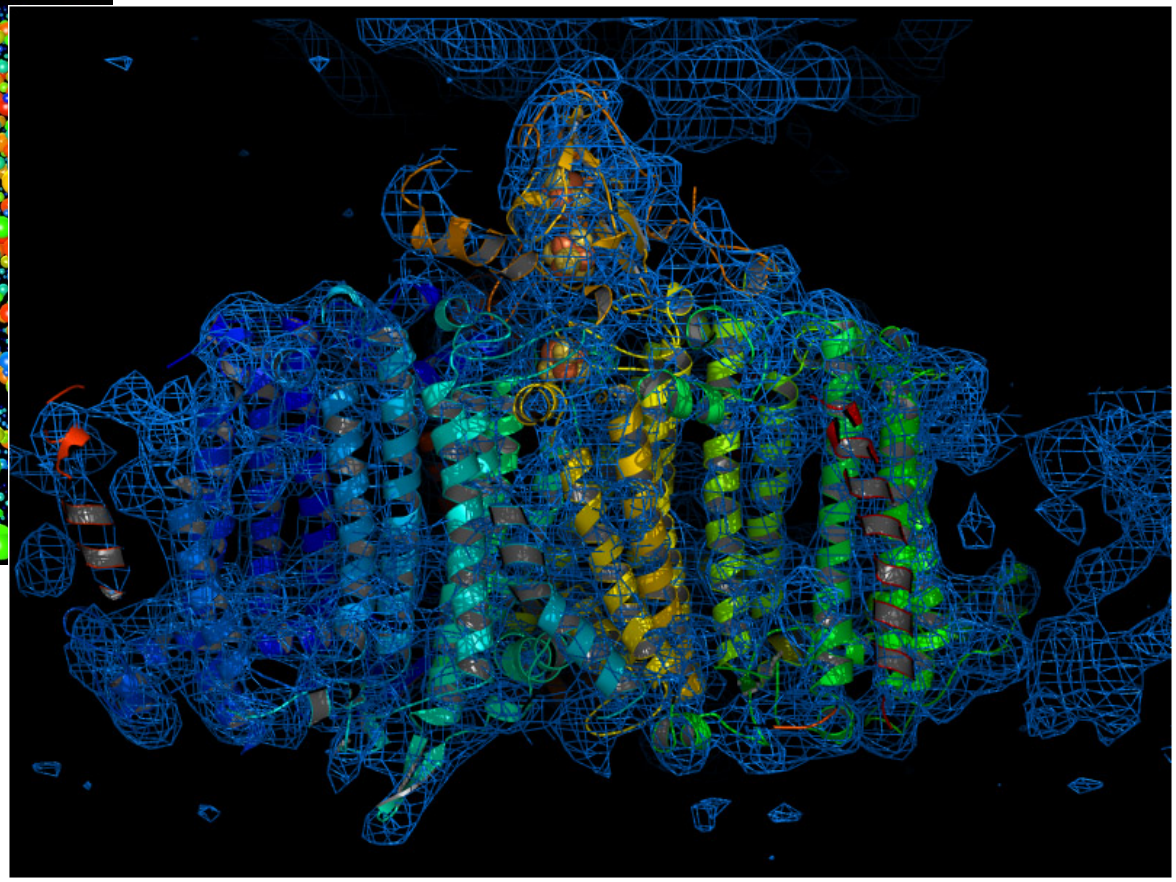

High-Throughput Pipeline:
Robotics, Automation & Machine Learning

John Gregoire
Ichiro Takeuchi
Matt Kramer
Apurva Mehta

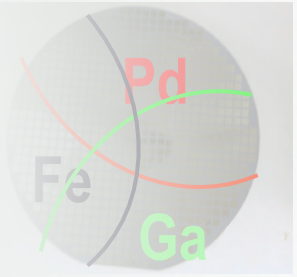

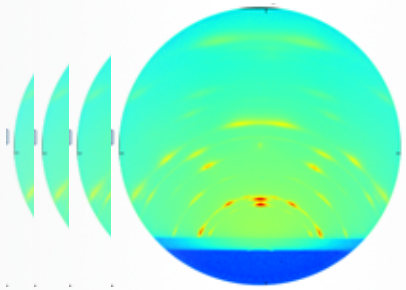

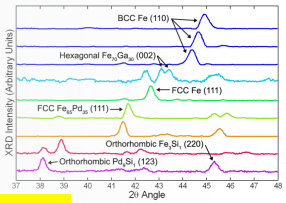

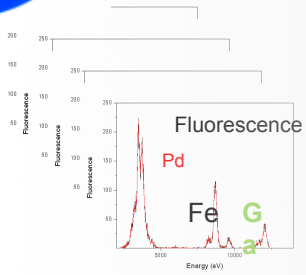

PCA, Density based clustering, Genetic optimization

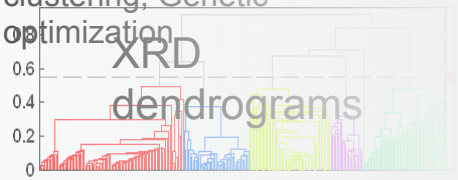

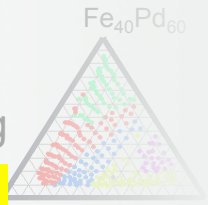

Hubs

Hubs

Library Production

Data Collection

Data Archiving

Property Screening

SSRL

- Auto-alignment and Calibration
- Robotic Library Changer
- Library Tracking

- Multiple Data Stream Archiving
  - Hdf5?

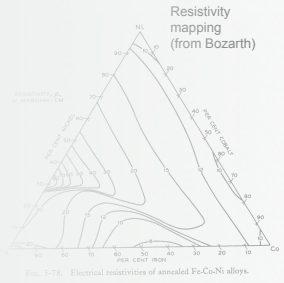

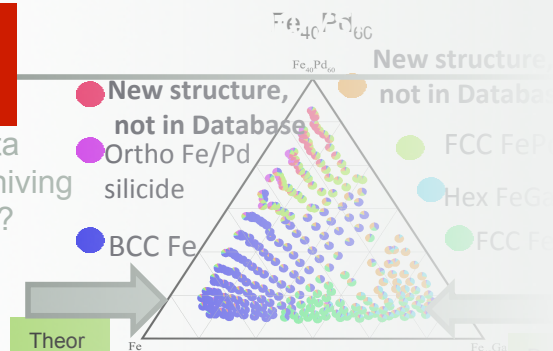

Theor y

Machine Learning

New Composition-Structure-Property Relations

New Materials

Scientific Communities

New Theory of (metastable) Materials

tabase

CD

Science & Technology
Facilities Council

ICAT

Facility Metadata Catalogue

Tom Griffin, STFC ISIS Facility

Overview

- What is ICAT?
- Uses and Benefits
- Installation
- Data Integration
- Interfaces
- About the ICAT project
- Future developments

Science & Technology Facilities Council

What is ICAT?

ICAT is a database, with a well defined API that provides a uniform interface to experimental data and a mechanism to link all aspects of research from proposal through to publication.

What is ICAT?

Example ISIS Proposal

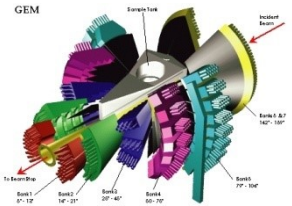

GEM – High intensity, high resolution neutron diffractometer

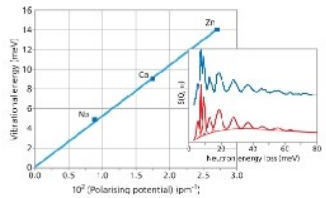

H2-(zeolite) vibrational frequencies vs polarising potential of cations

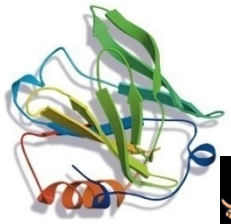

B-lactoglobulin protein interfacial structure

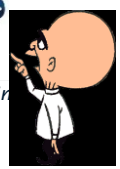

ICAT

Proposals

Once awarded beamtime at ISIS, an entry will be created in ICAT that describes your proposed experiment.

Experiment

Data collected from your experiment will be indexed by ICAT (with additional experimental conditions) and made available to your experimental team

Analysed Data

You will have the capability to upload any desired analysed data and associate it with your experiments.

Publication

Using ICAT you will also be able to associate publications to your experiment and even reference data from your publications.

Science & Technology Facilities Council

Benefits

- Links proposal to data to analysis to publication
- Makes data searchable
- Remote data access
- Assign DOIs to data
- Increased opportunities for data sharing and re-use
- Implements a data policy
- Provenance (Creation, ownership, history)
- Integration with applications

Where is it used?

- STFC (main developer)
  - ISIS pulsed neutron and muon facility
  - CLF (UK national laser facility)
- Diamond
  - UK national synchrotron
- ILL
  - European reactor neutron source
- Oak Ridge National Laboratory (USA)
  - Spallation Neutron Source
- Currently being rolled out at: ESRF, ALBA, SOLEIL, other PaNdata partners

Science & Technology Facilities Council

Technical

Installation

- Relational database - Oracle, MySQL
- Java application server - Glassfish
- Support available

Components

- ICAT is a modular system
  - Authentication
    - Plain text simple database – testing and development
    - LDAP/Active Directory
    - Custom user office connections
  - Data Server
    - Defines an interface – can be implemented as appropriate
    - ‘Disk only’ implementation (ISIS)
    - Tape/disk cache (Diamond)

Flexibility

- Access to data is defined through a ‘rules’ system
  - E.g. Grant read access to data when a user is a Co-Investigator
  - Grant full control to data when a user is the Principal Investigator
  - Grant full control to all data on an instrument when the user is the instrument scientist for that instrument.
- Facility specific: Instruments (beamlines) , Instrument Scientists, Datafile types, parameters
  - E.g. wavelength, total_proton_count, etc

Interfaces

- ICAT exposes a web service (SOAP) API
  - Available for client applications
  - Enables integration with data analysis applications such as DAWN and Mantid.
  - RESTful interface planned
- Default web interface – ‘TopCAT’
  - Allows basic data browsing
  - Able to search many ICAT’s in parallel

Interfaces - TopCAT

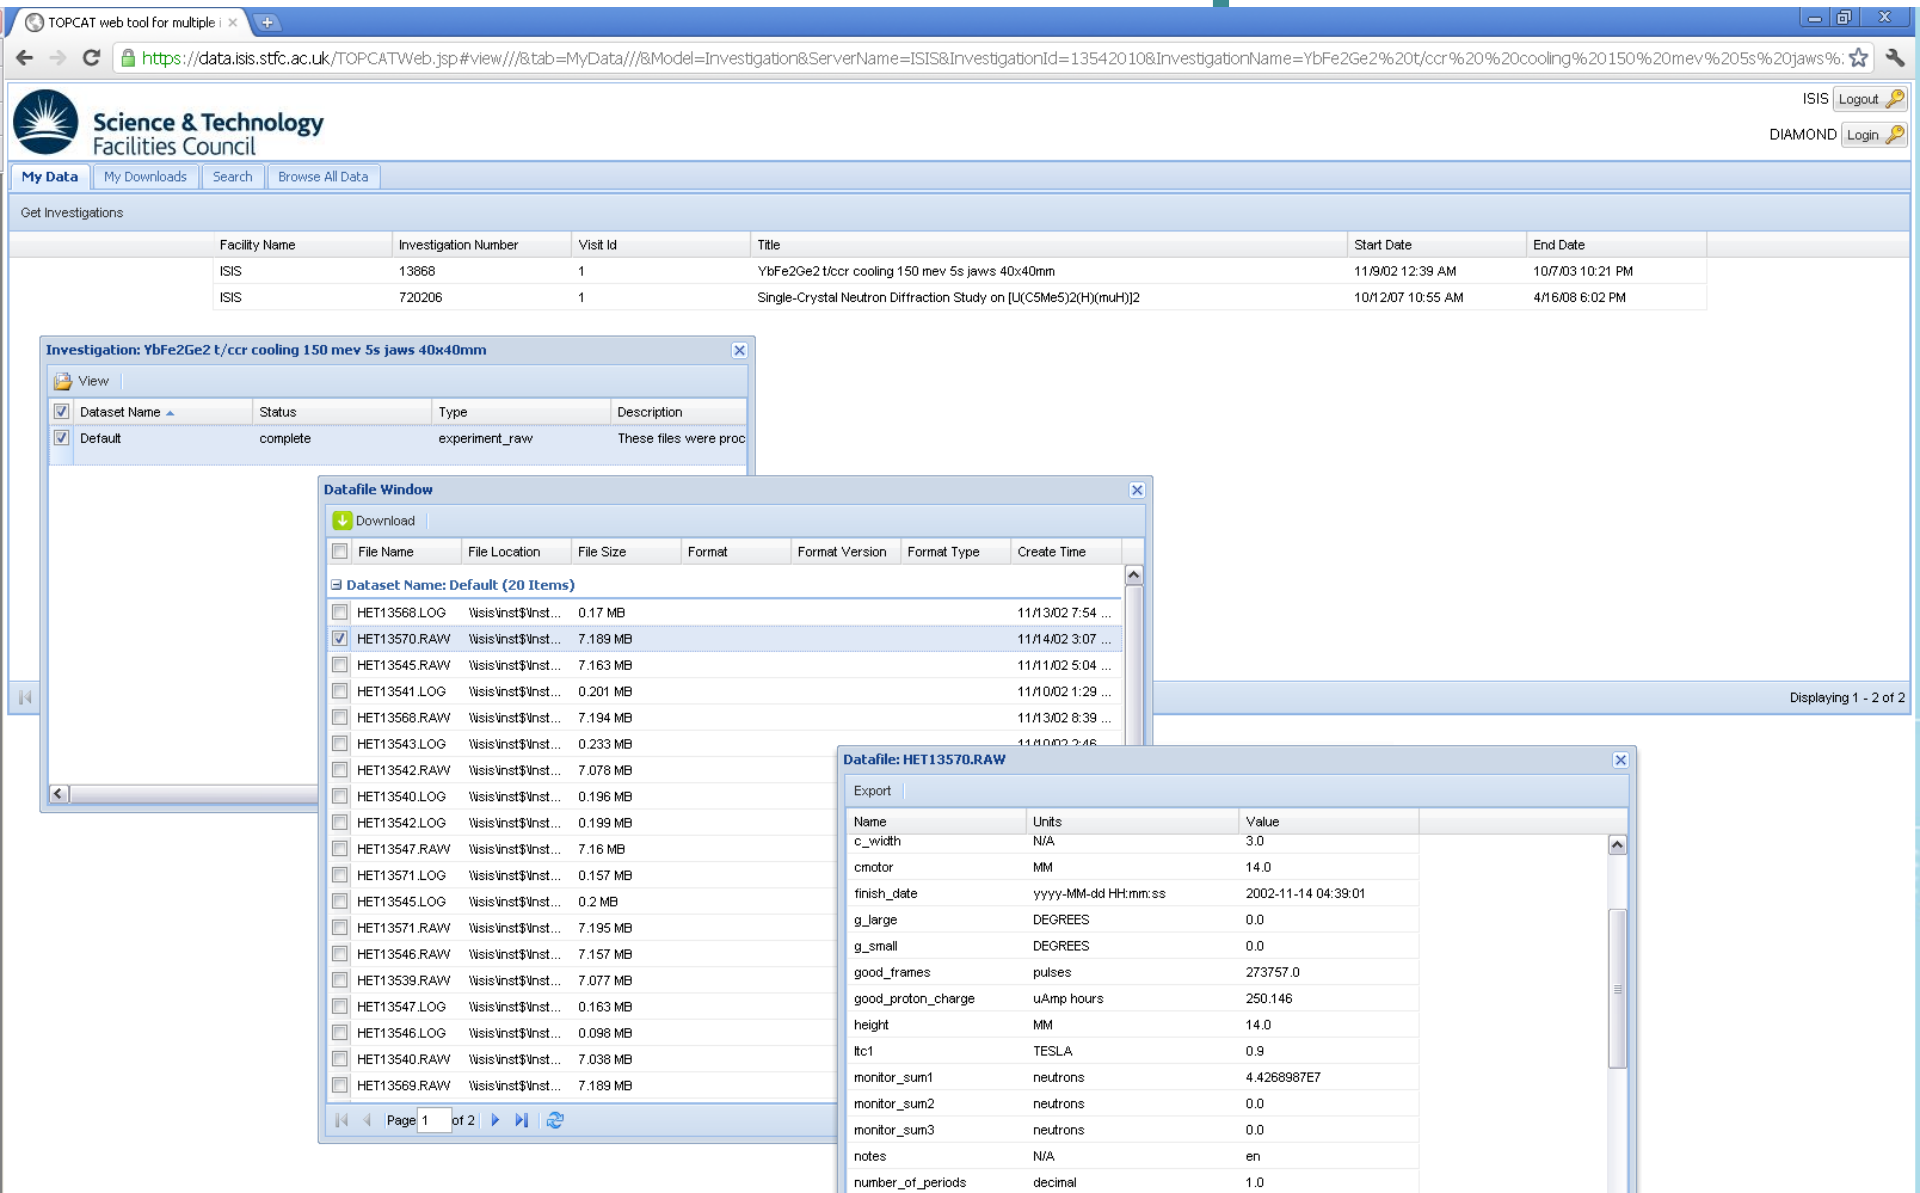

Interfaces - TopCAT

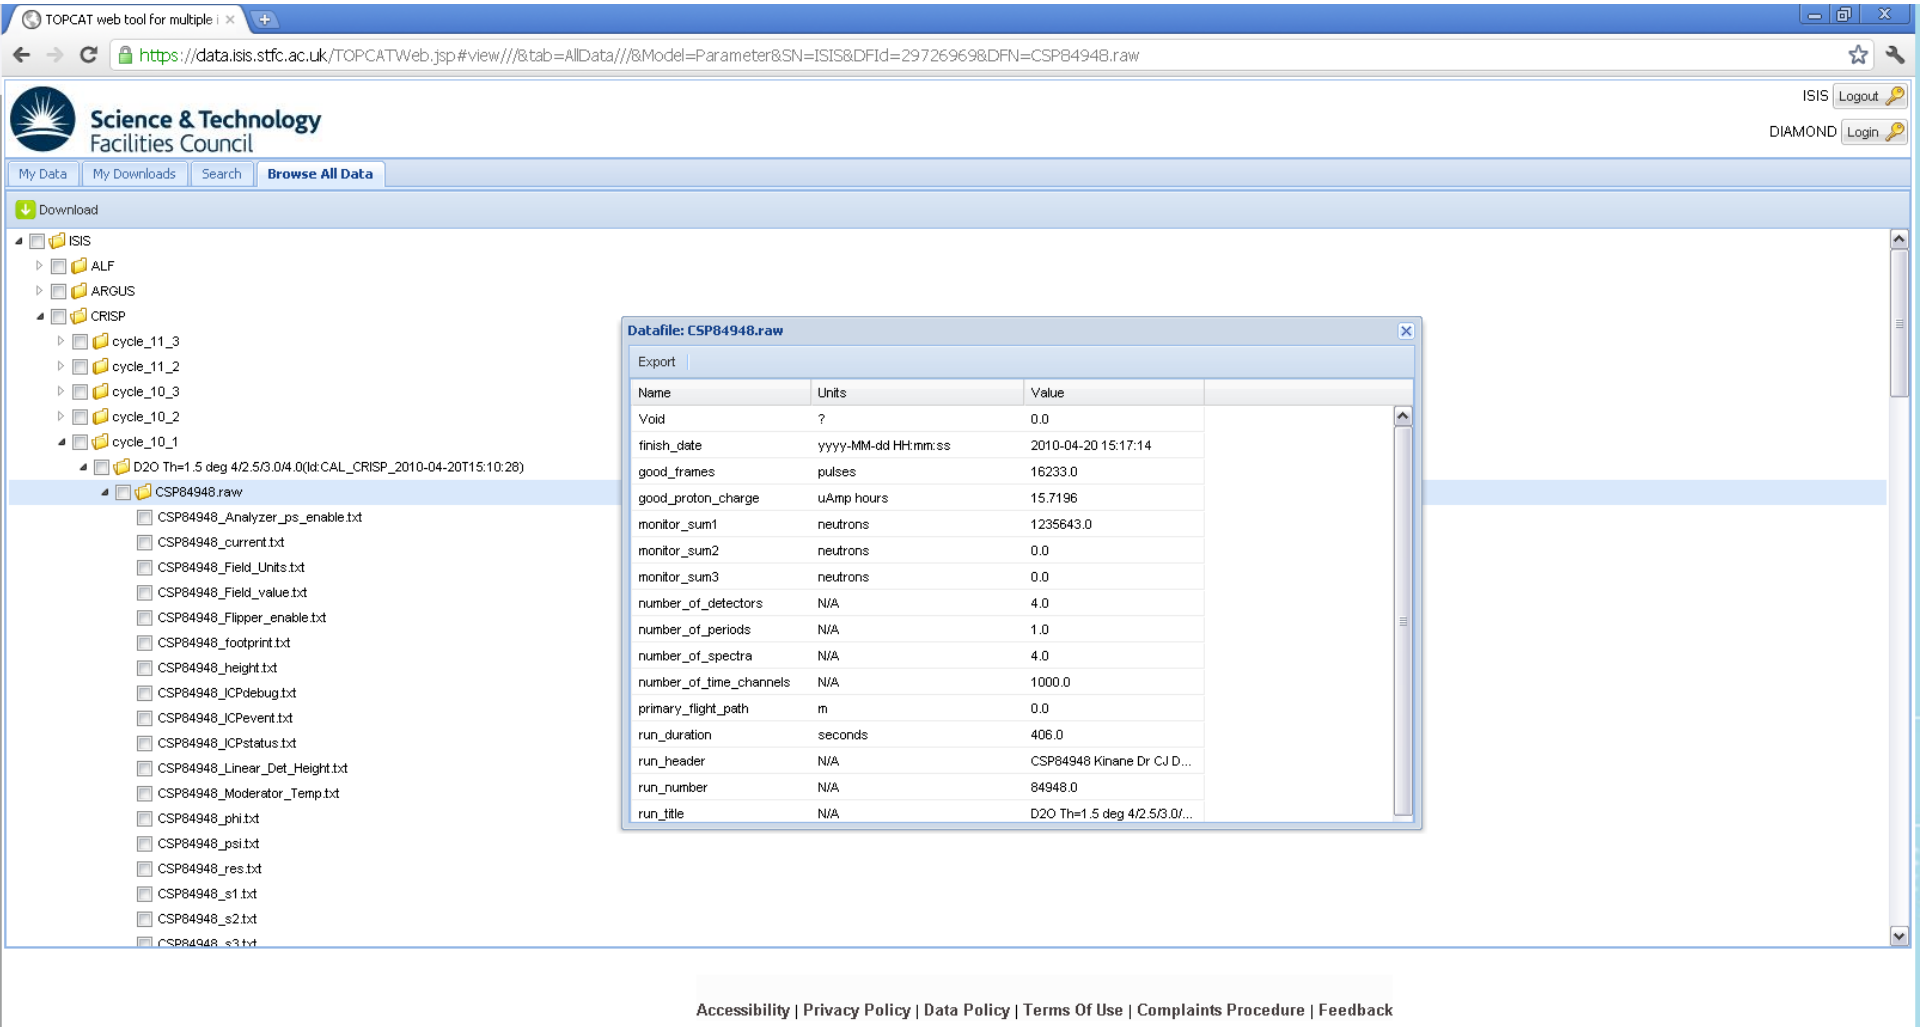

Interfaces - Mantid

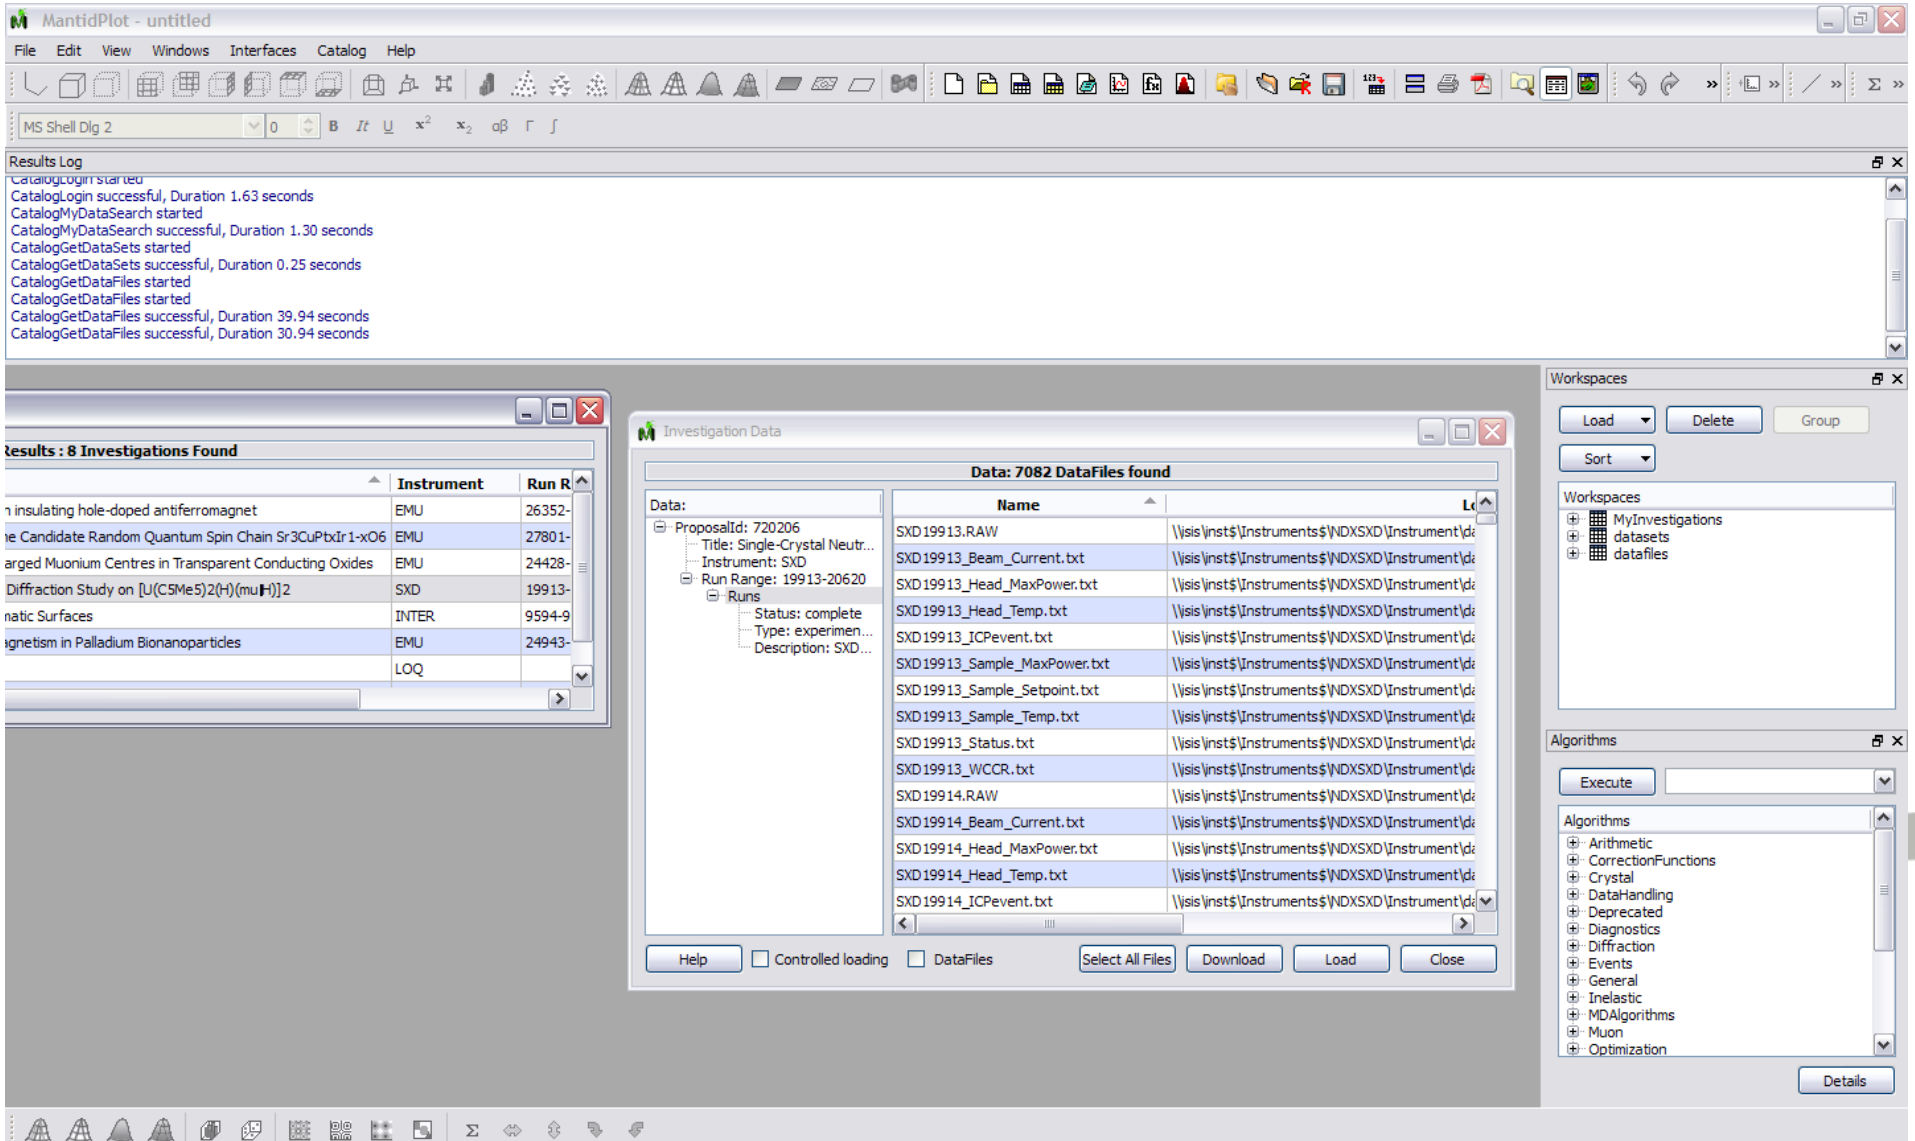

Interfaces - Mantid

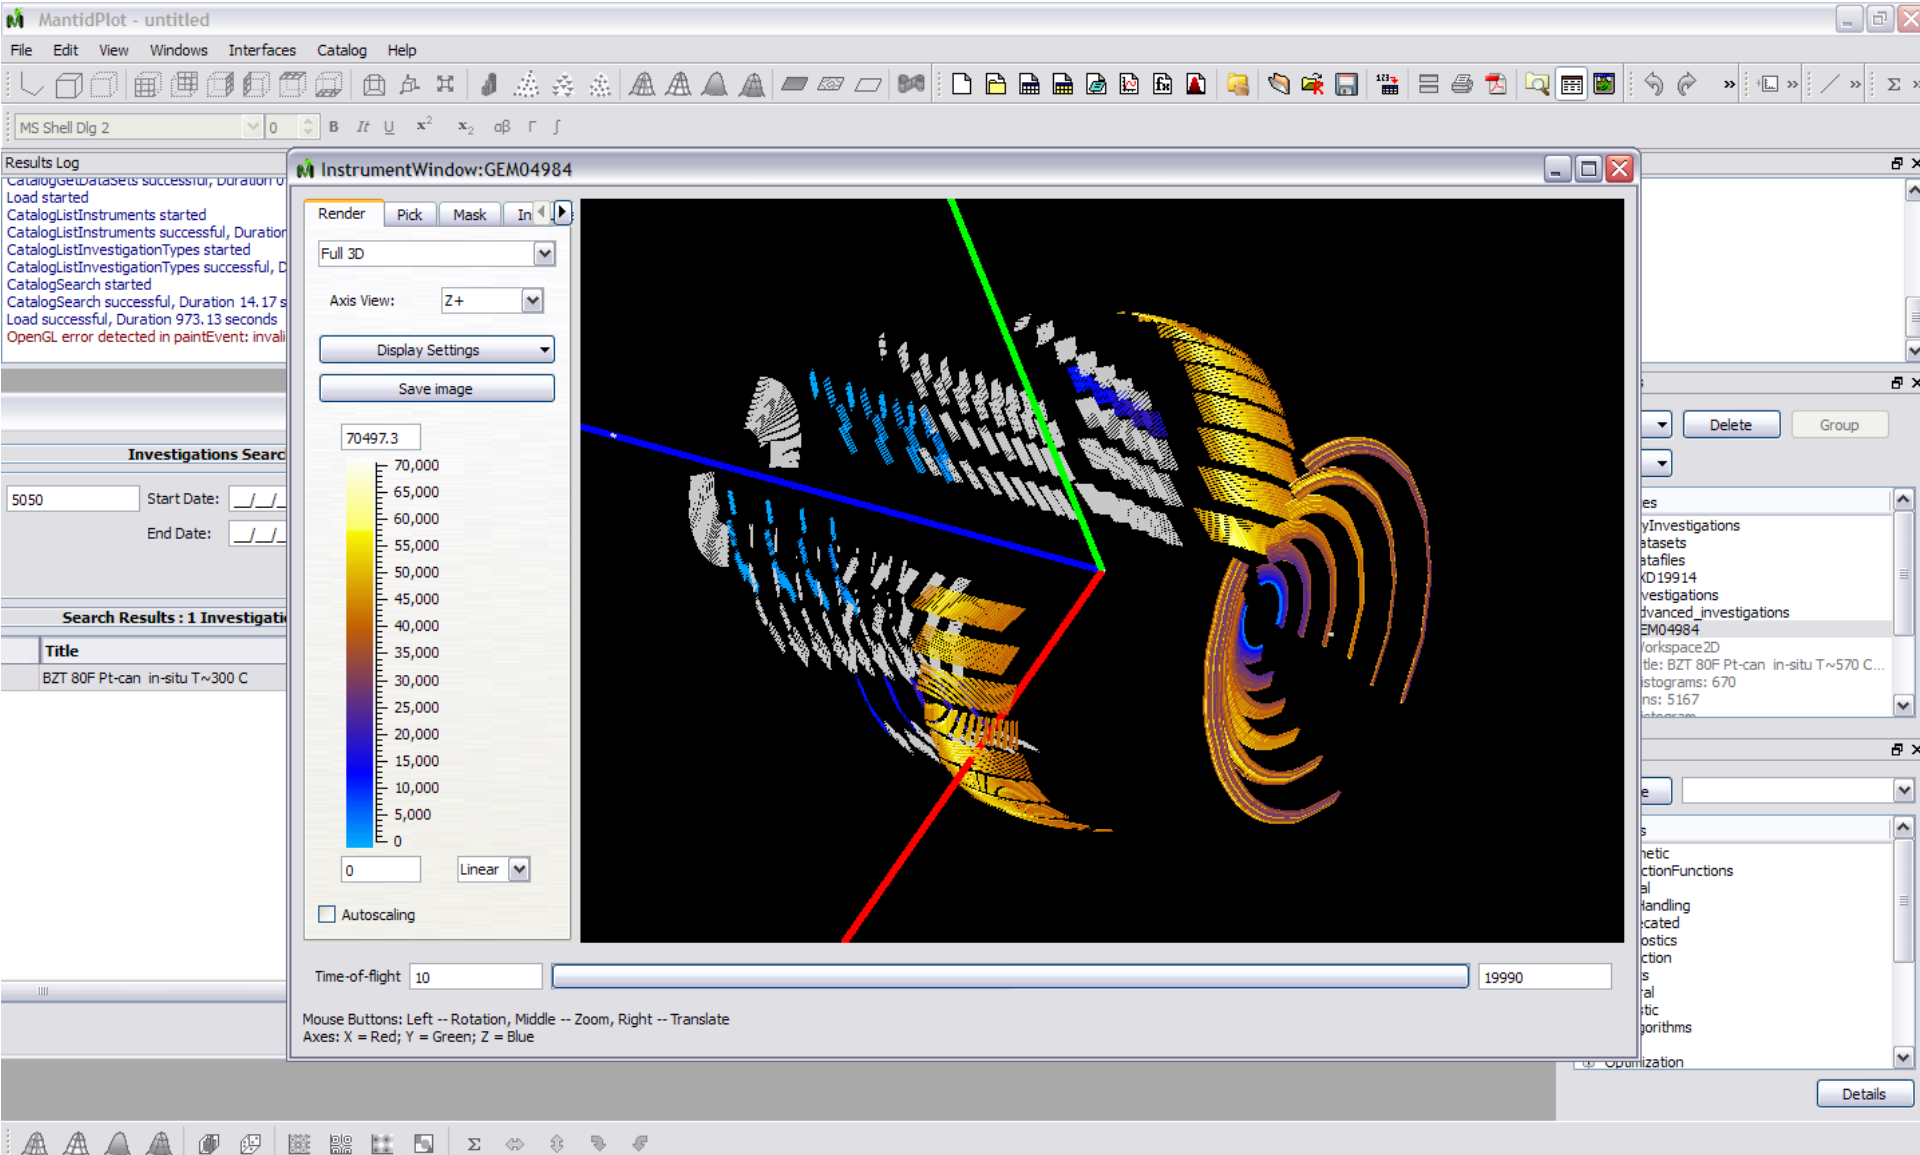

Interfaces - DAWN

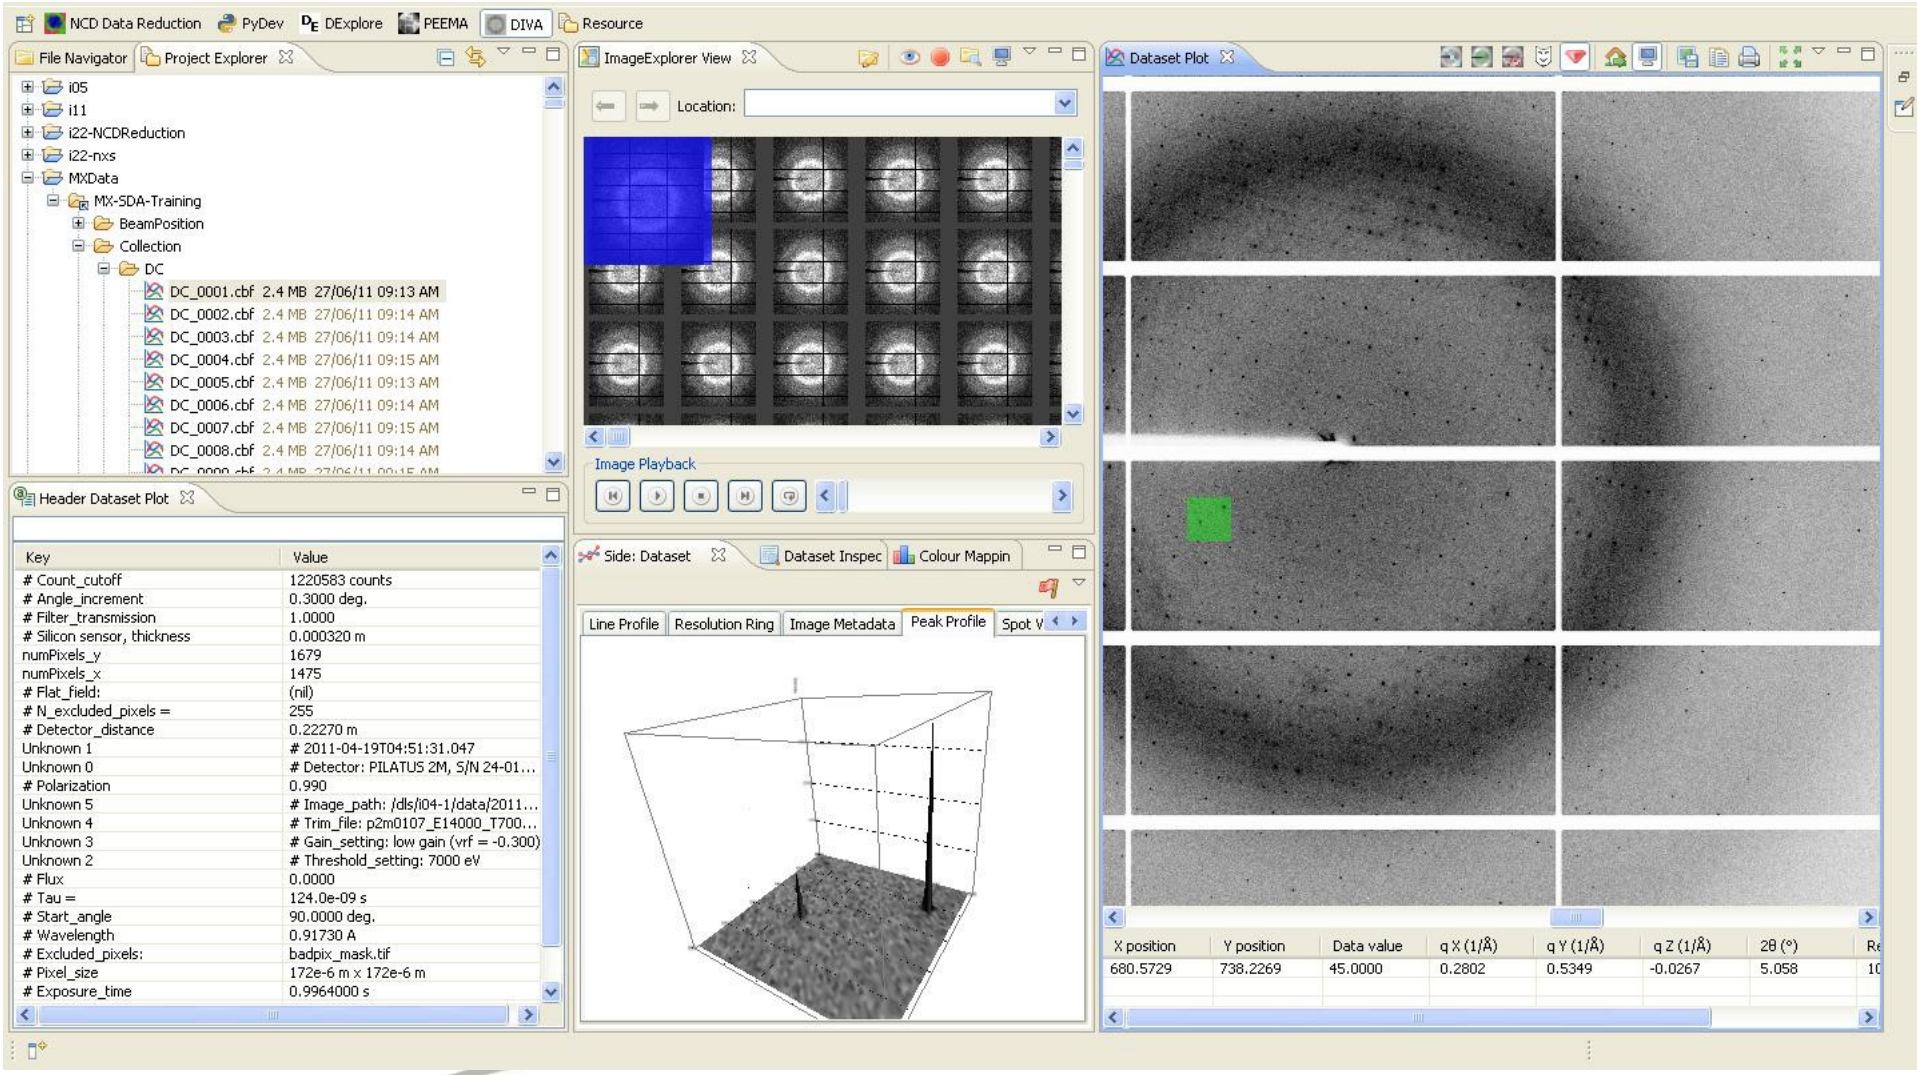

Interfaces - DAWN

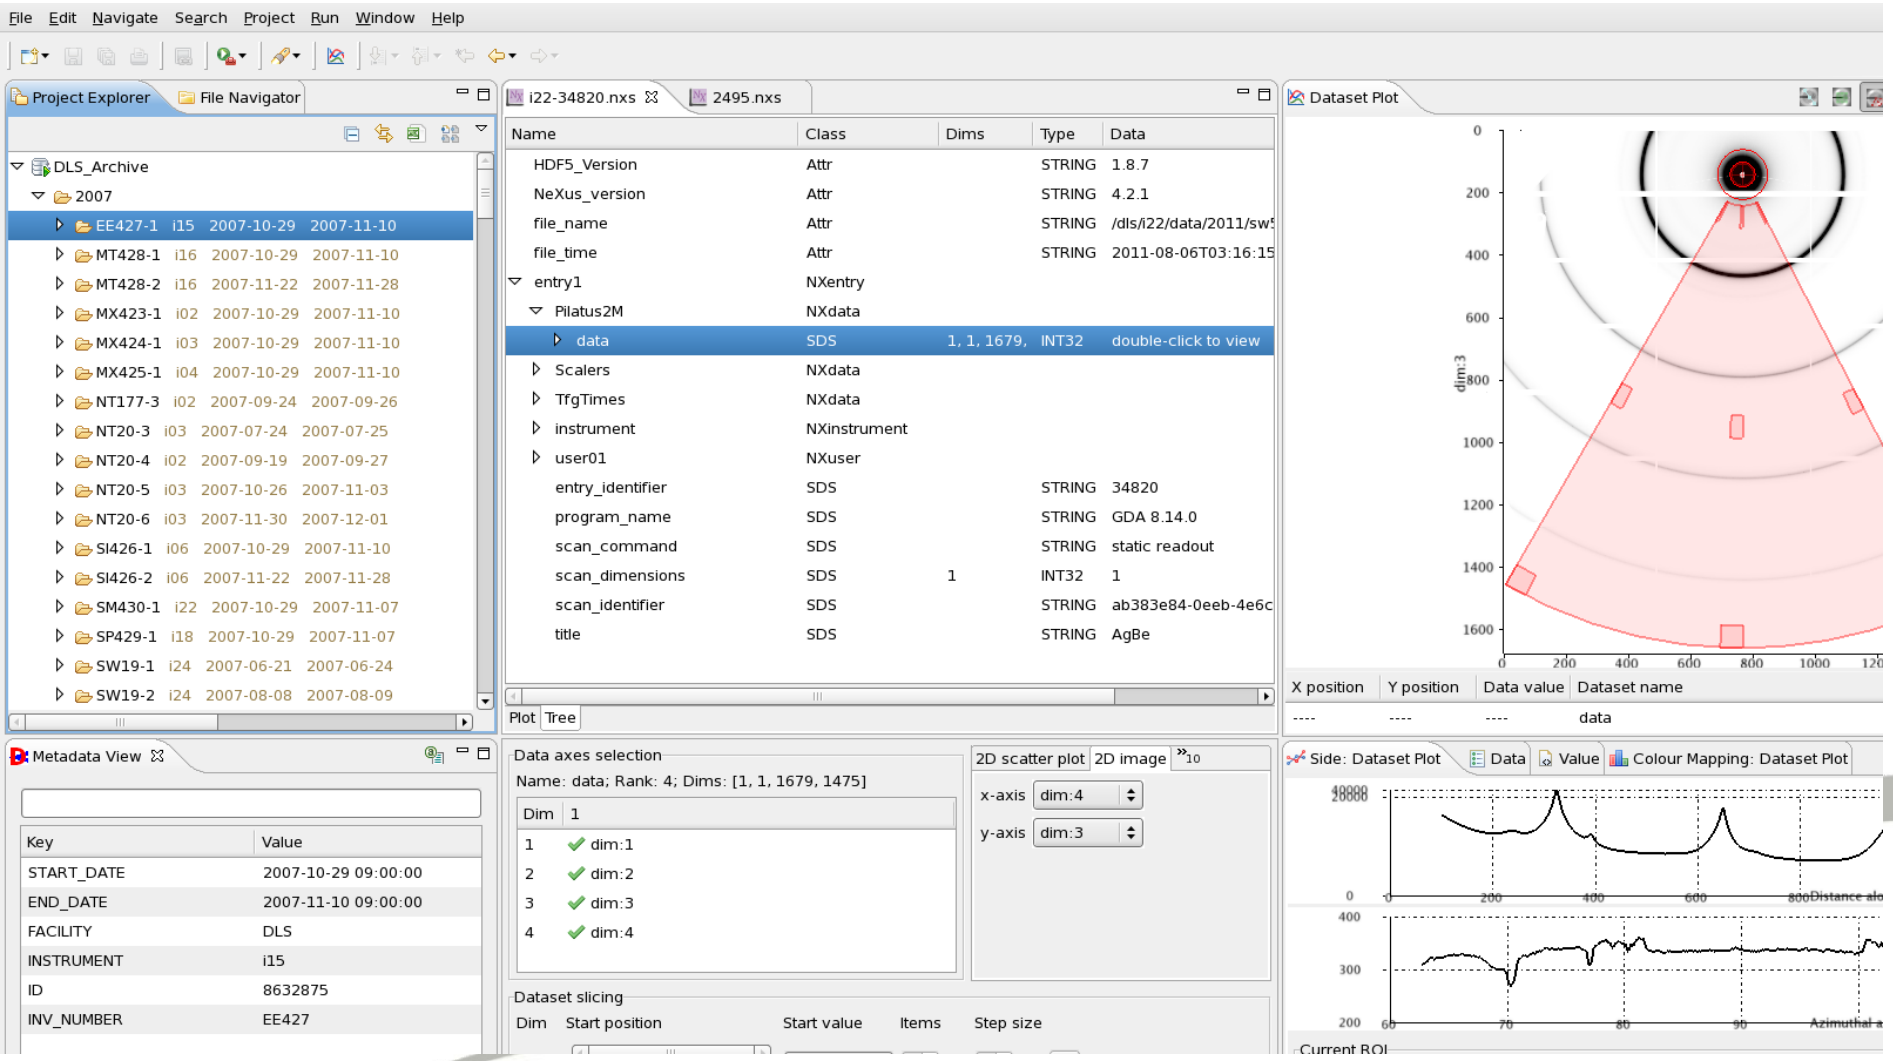

Populating ICAT with Data

- ICAT has a SOAP API so data can be pushed in from most languages
- Typically metadata is imported from two sources
  - User office
  - Experiment data files
- User office link code tends to be bespoke
- Can be simple (experiment title, people) or more complex (abstracts, samples, links between related work etc)

Populating ICAT with Data

- Data files
  - Tools exist to assist ingesting nexus files
  - Ingest custom/specific data formats will require code to extract the metadata from them
- Can be done using the API or ‘XML Ingest’ layer
- XML Ingest defines a simple schema that describes a set of datafiles and inserts them into an ICAT
- Prototype available, but still work-in-progress
- Will simplify data ingestion.

Science & Technology
Facilities Council

The ICAT Project

The ICAT Project

- ICAT is an open source project
- Currently managed by STFC
- Released under BSD and Apache licenses (permissive)
- www.icatproject.org
- http://code.google.com/p/icatproject
- Monthly collaboration meetings (telephone)
- Annual developer meetings (face-to-face)
- Steering committee

Science & Technology Facilities Council

Future Developments

- Under active development
- Releases every 6 months
- Current roadmap
  - Clustered deployment
  - Improved documentation
  - Data import/export/migration
  - Non-relational databases (hybrid)
  - Interface improvements to TopCAT

  http://icatproject.org/releases/road-map/
- Tell us what you want to see…..

Science & Technology Facilities Council

Summary

- ICAT is a mature solution for facility metadata management
- Enables remote data access and linking between proposals, data, analysis and publications
- Flexible architecture enables integration with differing facility systems and requirements
- API enables integration with other software (analysis)
- Open source project with an active collaboration

Questions?

Thanks to: Steve Fisher, Brian Matthews, Alistair Mills, Alun Ashton, Antony Wilson

MONASH University

Store.Synchrotron.org.au

Steve Androulakis

Store.Synchrotron.org.au

• Store.Synchrotron is a system that captures all macromolecular beamline data, available online to all non-commercial Australian Synchrotron users. It was developed by Monash University in a strategic, ongoing partnership.

• Data is **immediately shareable** by the researcher on the web and able to be **published**.

• The service operates on the Australian NeCTAR Research compute cloud in a scalable setup able to withstand load. (http://nectar.org.au/)

• We're actively opening access to raw data behind high-impact research publications under CC BY licenses. Six institutions have opened data so far.

• Built on MyTardis – an open source, Australian made data management platform used all over Australia in proteomics, genomics, electron microscopy, medical imaging, astrophysics, quantum physics and more.

• Visit store.synchrotron.org.au for access / contact steve.Androulakis@monash.edu

MONASH University Australian Synchrotron nectar

Store.Synchrotron Home About My Data Public Data Stats Help Log In

Experiment

Derivatives for structure solution of the peripheral stalk from T.thermophilus A-ATPase

Alastair Stewart, Daniela Stock

4 575 10.1 GB 18th October 2013 Public

Toggle Full Description

Rotary ATPases couple ATP hydrolysis/synthesis with proton translocation across biological membranes and so are central components of the biological energy conversion machinery. Their peripheral stalks are essential components

Description Metadata

Institution Victor Chang Cardiac Research Institute, Australian Synchrotron

Licensing This experiment data is licensed under Creative Commons Attribution 3.0 Australia (CC BY 3.0).

Administrators Steve Androulakis

Download All TAR

4 Datasets

Download Selected

Just start typing to filter datasets based on descriptions

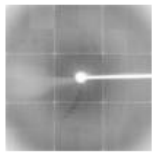

2K3e_5
Synchrotron MX Data
94
1.7 GB

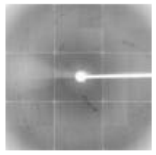

2K3d_4
Synchrotron MX Data
180
3.2 GB

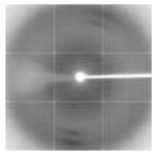

2K7c_3
Synchrotron MX Data
120
2.1 GB

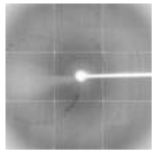

2K3c_3
Synchrotron MX Data
181
3.2 GB

Open Data

MONASH University Australian Synchrotron nectar

Real-time instrument data capture

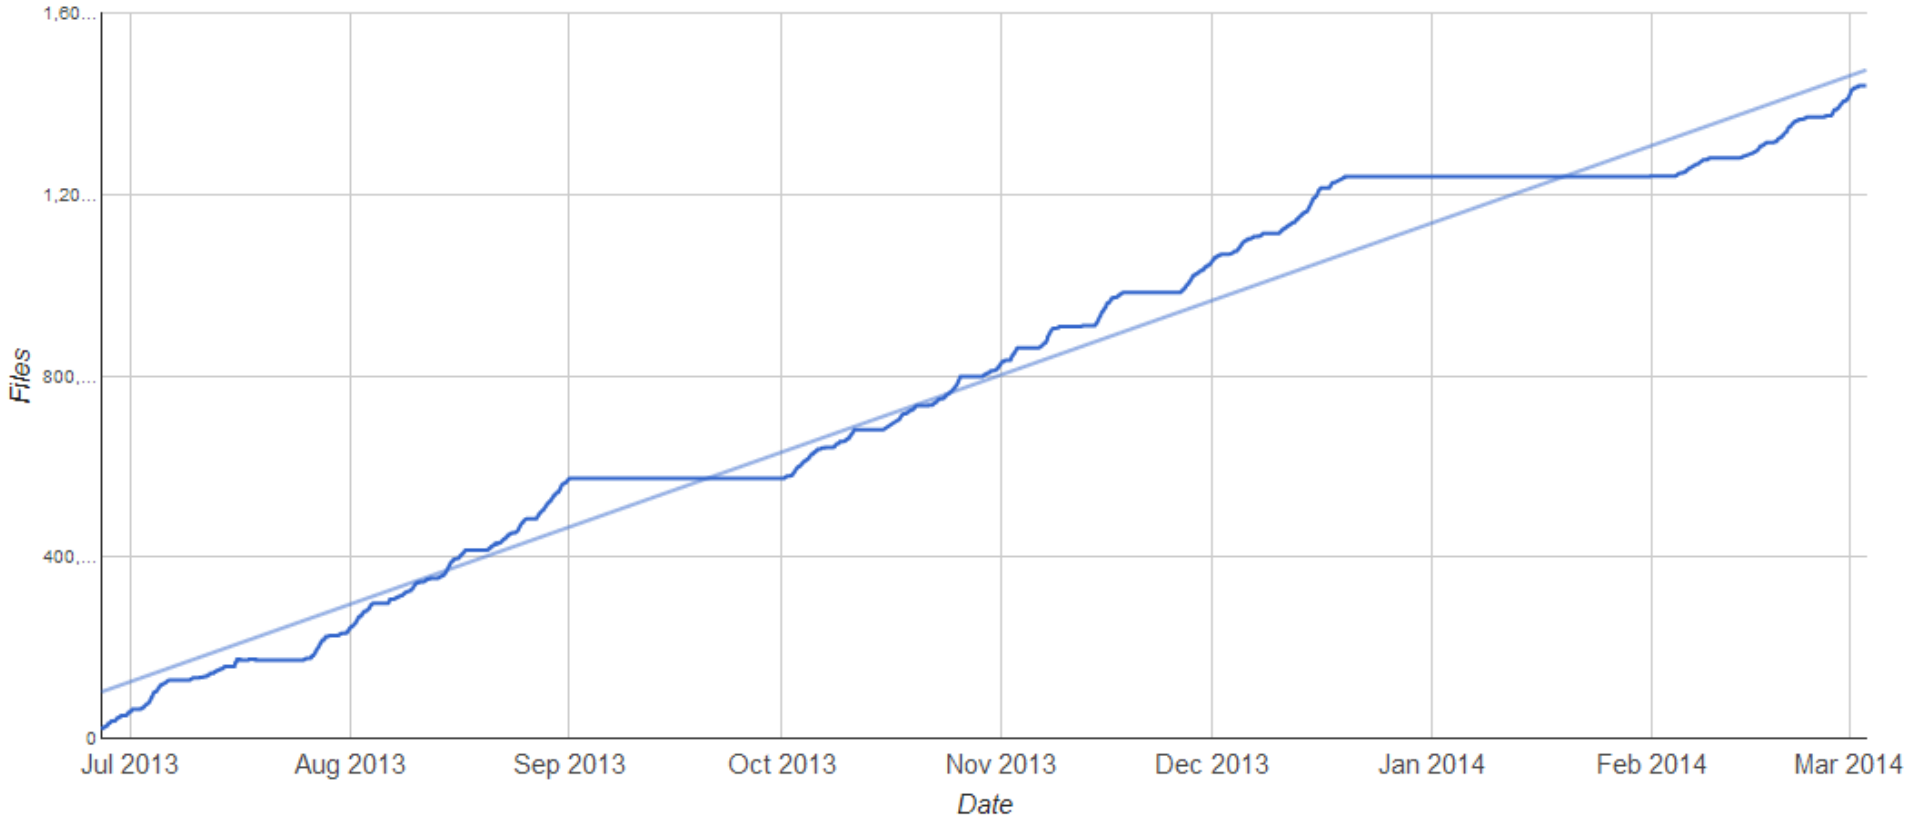

Capture began June 2013. As of March 2013, it has captured over 18 terabytes of data in over 1.5 million raw diffraction images.

Publish your research, then publish your raw data with us.

The Australian Synchrotron is working with researchers to open access to datasets associated with publications. This process is entirely opt-in for researchers and gives permission for the Synchrotron to host your data publicly.

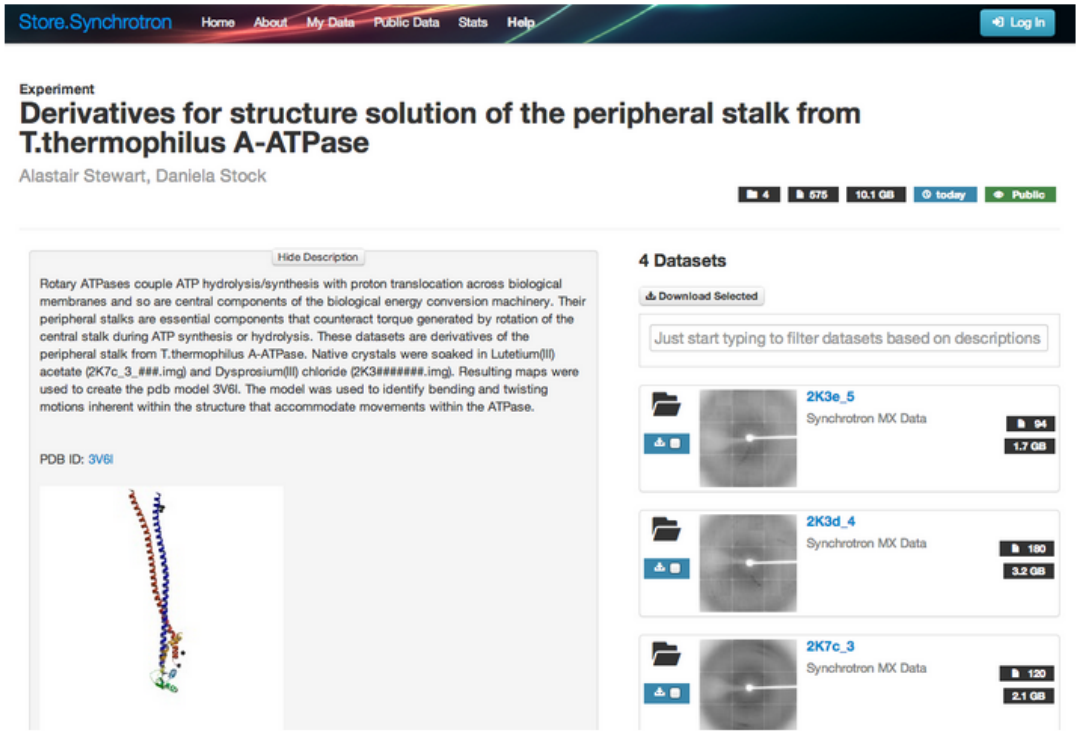

Data is made available via MyTardis at Store.Synchrotron.org.au

What do you need to do?

Contact
mxtickets@synchrotron.org.au

You'll be emailed a short form to fill in basic information about open access to your datasets associated with PDBs and publications.

Once data is made available via our service, you will receive a DOI to this data for citation.

Note: The data you're collecting today is also archived privately and securely via this system.

Log in to https://store.synchrotron.org.au with your Australian Synchrotron credentials to browse and download your data.

Presented on the Synchrotron's Control PCs in a slideshow.

MONASH University
Australian Synchrotron
nectar

How can we contribute to PaNSig?

• Store.Synchrotron.org.au is a scalable cloud service that was created to withstand network bandwidth, CPU, disk IO. We have considerable engineering experience in this area and can provide advice and guidance.

• We have worked with Synchrotron users in order to have their raw data released publicly (http://tardis.edu.au/syncpublish/) and have a basic model for this process.

• The Australian National Data Service (ANDS.org.au) is funding an 18 month project to refine the Store.Synchrotron.org.au process of collecting quality metadata at time of collection, and time of publication. Many of its outputs would be relevant to PaNSig.

• Monash University are in the process of increasing its 'instrument integration' of this kind using the same underlying system as Store.Synchrotron. This includes expansion to more beamlines at the Australian Synchrotron and likely the Australian neutron source ANSTO. Experiences here with different research communities, policies, practices, technology and culture should be broadly useful to this group.

Developing a plan of activities for RDA PaNSig

RDA Plenary, Dublin, March 2014

Co-chairs

Currently:

- Amber Boehnlein , SLAC, US
- Frank Schluenzen, DESY, DE
- Brian Matthews, STFC, UK

Topics for Discussion

- Purpose
- Technical activities: RDA WGs etc
- Interaction within RDA
- Future Meetings
- Interactions between meetings
- Members

Purpose

- Data related issues of science applications associated with large scale source facilities
- Source facilities share a number of issues in their data handling.
- These could include:
  - Scalability of data volumes and data access rates
  - Standardization of (meta-)data and vocabularies
  - Data, cataloguing, publishing, discovery, sharing, transfer and access, policies
  - Data analysis tools and frameworks supporting workflows and provenance
  - Interaction with the data handling practices and standards within different communities.

Technical Activities

- Spinning out RDA technical working groups
  - Specific
  - Time-bound (c. 18 months)
  - Clear outcome (document, format, code etc)
  - Need commitment
  - **Topics** ?
- Any documents as a group ?
  - capturing community views.
- Future projects?

Interaction with other RDA groups

- We will be requested to comment on draft recommendations from other RDA working groups
  - Will notify group and request if people want to review documents
- Are there specific groups we want to interact with?
  - Domain specific groups: Structural Biology, Materials Data
  - Technical infrastructure: Persistent IDs, Big data Analytics, cloud computing, metadata, data publication, …..

Future meetings

- RDA Plenary
  - 4th RDA Plenary: Amsterdam, 22-24 September, 2014
  - 5th RDA Plenary: USA, March 2015
- NoBugs 2014 ?
- Other events?
- Opportunities to establish a PaN Computing event similar to CHEP ??

Interactions in between f2f

- Keeping up momentum is important
  - Website/Wiki (RDA Mediated)
  - mailing list (RDA Mediated)
  - phone/video meetings?
- Subgroups on particular topics?

Other members

- Who should we be aiming at ?
- PaNS outside US/EU/AU ?

Possible immediate activities?

- Online technical forum
- Counting users
- Collecting views data usage
- Project/software “market”
- Common terms
- Other RDA Groups: Structural Biology, Materials, Persistent ID …
- Meetings : RDA the Hague

Other discussion?

DataCite –
Persistent links to scientific data

DataCite

DOI names for citations

URLs are not persistent

- (e.g. Wren JD: **URL decay in MEDLINE- a 4-year follow-up study**. Bioinformatics. 2008, Jun 1;24(11):1381-5).

Digital Object Identifiers (DOI names) offer a solution

- Mostly widely used identifier for scientific articles
- Researchers, authors, publishers know how to use them
- Put datasets on the same playing field as articles

Dataset
Yancheva et al (2007). Analyses on sediment of Lake Maar. PANGAEA.
doi:10.1594/PANGAEA.587840

DataCite

DataCite members

1. Technische Informationsbibliothek (TIB)
2. Canada Institute for Scientific and Technical Information (CISTI),
3. California Digital Library, USA
4. Purdue University, USA
5. Office of Scientific and Technical Information (OSTI), USA
6. Library of TU Delft, The Netherlands
7. Technical Information Center of Denmark
8. The British Library
9. ZB Med, Germany
10. ZBW, Germany
11. Gesis, Germany
12. Library of ETH Zürich
13. L'Institut de l'Information Scientifique et Technique (INIST), France
14. Swedish National Data Service (SND)
15. Australian National Data Service (ANDS)
16. Conferenza dei Rettori delle Università Italiane (CRUI)
17. National Research Council of Thailand (NRCT)
18. The Hungarian Academy of Sciences
19. University of Tartu, Estonia
20. Japan Link Center (JaLC)
21. South African Environmental Observation Network (SAEON)
22. European Organisation for Nuclear Research (CERN)

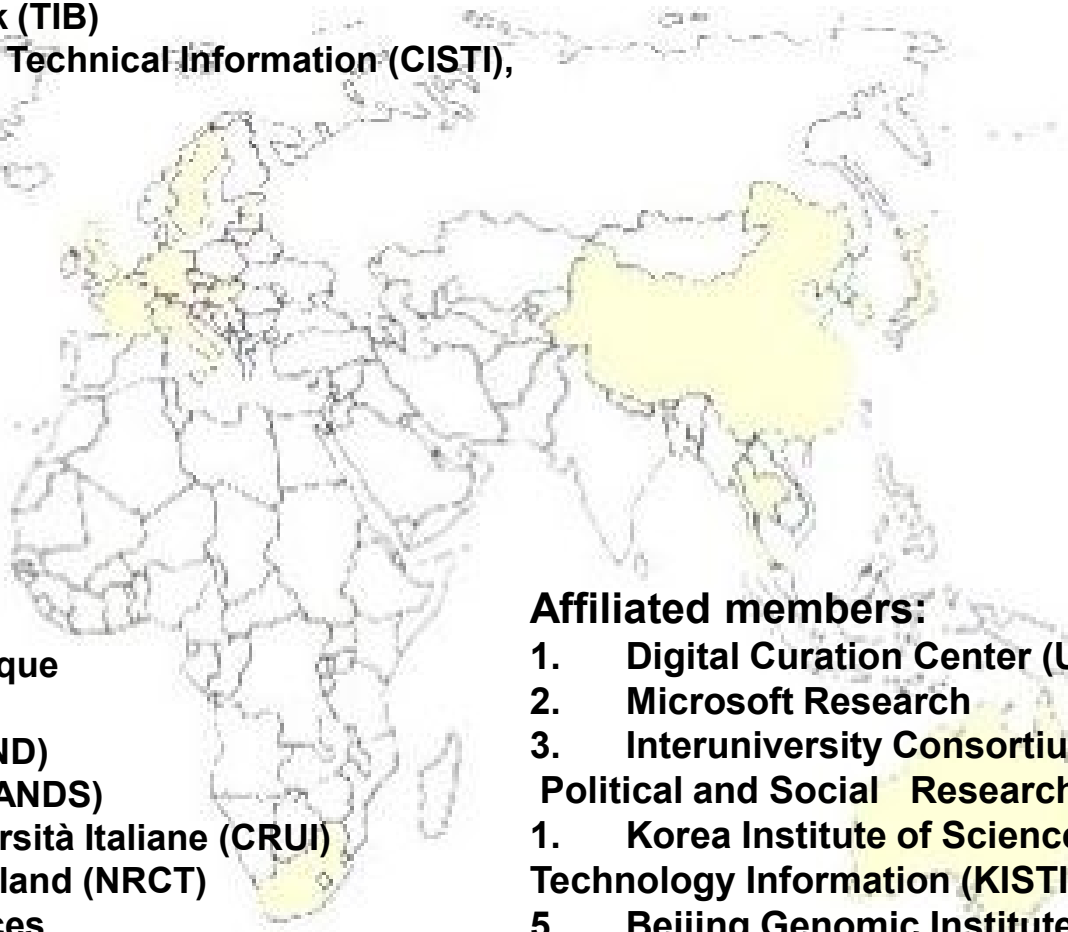

Affiliated members:

1. Digital Curation Center (UK)
2. Microsoft Research
3. Interuniversity Consortium for Political and Social Research (ICPS

1. Korea Institute of Science and Technology Information (KISTI)
5. Bejiing Genomic Institute (BGI)
6. IEEE
7. Harvard University Library
8. World Data System (WDS)
9. GWDG

DataCite structure

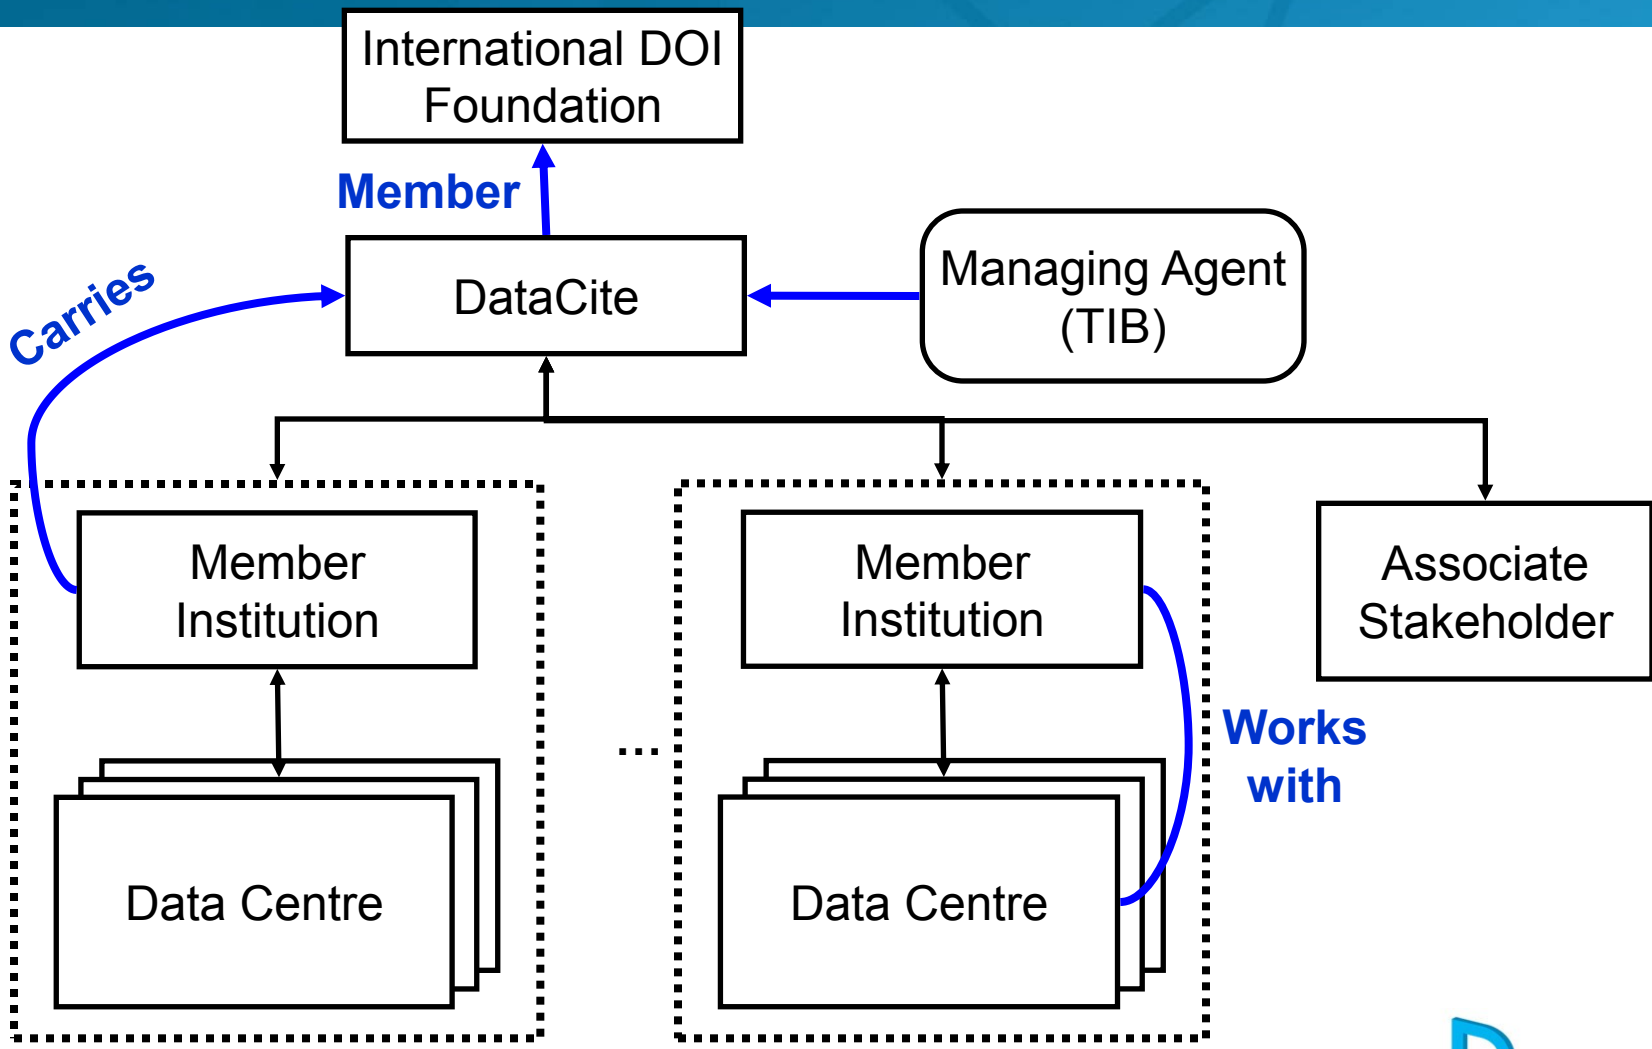

DataCite

DataCite in 2014

Over 3,200,000 DOI names registered so far.

290 data centers.

8,000,000 resolutions in 2013.

DataCite Metadata schema published (in cooperation with all members) http://schema.datacite.org

DataCite MetadataStore

http://search.datacite.org

DataCite

OAI and Statistics

OAI Harvester

http://oai.datacite.org

DataCite statistics (resolution and registration)

http://stats.datacite.org

DataCite

ODIN project

ORCID and DataCite interoperability network. Funded under FP7

http://www.odin-project.eu

Claim your DataCite DOI with your ORCID profile:

http://datacite.labs.orcid-eu.org/

DataCite

2012: STM, CrossRef and DataCite Joint Statement

1. To improve the availability and findability of research data, the signers encourage authors of research papers to **deposit researcher validated data in trustworthy and reliable Data Archives**.

2. The Signers encourage Data Archives to **enable bi-directional linking between datasets and publications** by using established and community endorsed unique persistent identifiers such as database accession codes and DOI's.

3. The Signers encourage publishers and data archives to make visible or increase **visibility of these links** from publications to datasets and vice versa

DataCite

Example

The dataset:

Storz, D et al. (2009):

Planktic foraminiferal flux and faunal composition of sediment trap L1_K276 in the northeastern Atlantic.

http://dx.doi.org/10.1594/PANGAEA.724325

Is supplement to the article:

Storz, David; Schulz, Hartmut; Waniek, Joanna J; Schulz-Bull, Detlef; Kucera, Michal (2009): *Seasonal and interannual variability of the planktic foraminiferal flux in the vicinity of the Azores Current.*

Deep-Sea Research Part I-Oceanographic Research Papers, **56(1),** 107-124,

http://dx.doi.org/10.1016/j.dsr.2008.08.009

DataCite

Cooperation

MoU with ORCID

Agreement with Re3Data and DataBib to include their service in 2016

MoU with RDA to become organisational affiliate

Joint Declaration of Data Citation Principles

https://www.force11.org/datacitation

DataCite

2014 Annual conference

DataCite Annual Conference 2014

25-26 August 2014, Inist-CNRS, Nancy, France

(Just after the IFLA World Library and Information Congress in Lyon)

Giving value to data:
advocacy, guidance, services

Coming to Nancy:

- *by train:*
1h30 from Paris
4 hours from Lyon

- *by plane:*
direct link from Paris Airport

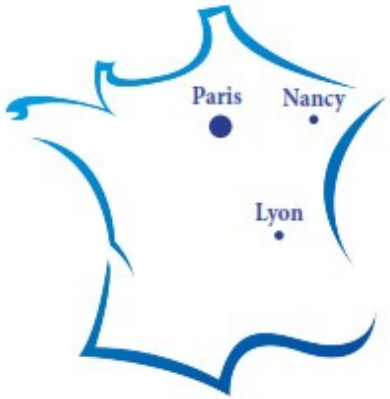

DataCite

Thank you!


Knowledge for Life
ZB MED


gesis
Leibniz Institute for the Social Sciences


ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics


TIB
TECHNISCHE INFORMATIONSBIBLIOTHEK


สำนักงานคณะกรรมการวิจัยแห่งชาติ
วช.
NRCT
National Research Council of Thailand


CERN


NRC·CNRC
CISTI
ICIST


inist


TU Delft


VE RI TAS


BRITISH LIBRARY


SAEON
South African Environmental Observation Network


IEEE


MAGYAR TUDOMÁNYOS AKADÉMIA
1825


TARTU ÜLIKOOL
1632
UNIVERSITAS TARTUENSIS


DTU


GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen


KiSTi
www.kisti.re.kr


SND Swedish National Data Service


華大基因
BGI


ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich


Japan Link Center
Welcome to Japan Link Center, JaLC!

University of California
Calif


ICSU
WORLD DATA SYSTEM


osti.gov


PURDUE
UNIVERSITY
LIBRARIES


DCC


ands
AUSTRALIAN NATIONAL DATA SERVICE


ICPSR
A PARTNER IN SOCIAL SCIENCE RESEARCH
INTER-UNIVERSITY CONSORTIUM FOR POLITICAL AND SOCIAL RESEARCH


Microsoft
Research


CRUI
Conferenza dei Rettori delle Università Italiane

Data analysis issues and frameworks

Alun Ashton

Group Leader, Data Analysis Software

diamond

Diamond Light Source

- EM Facility
- SFX XFEL UK Hub

diamond

Data Analysis – Experimental Challenges
Diversity: 33 beamlines > 120 experiment techniques

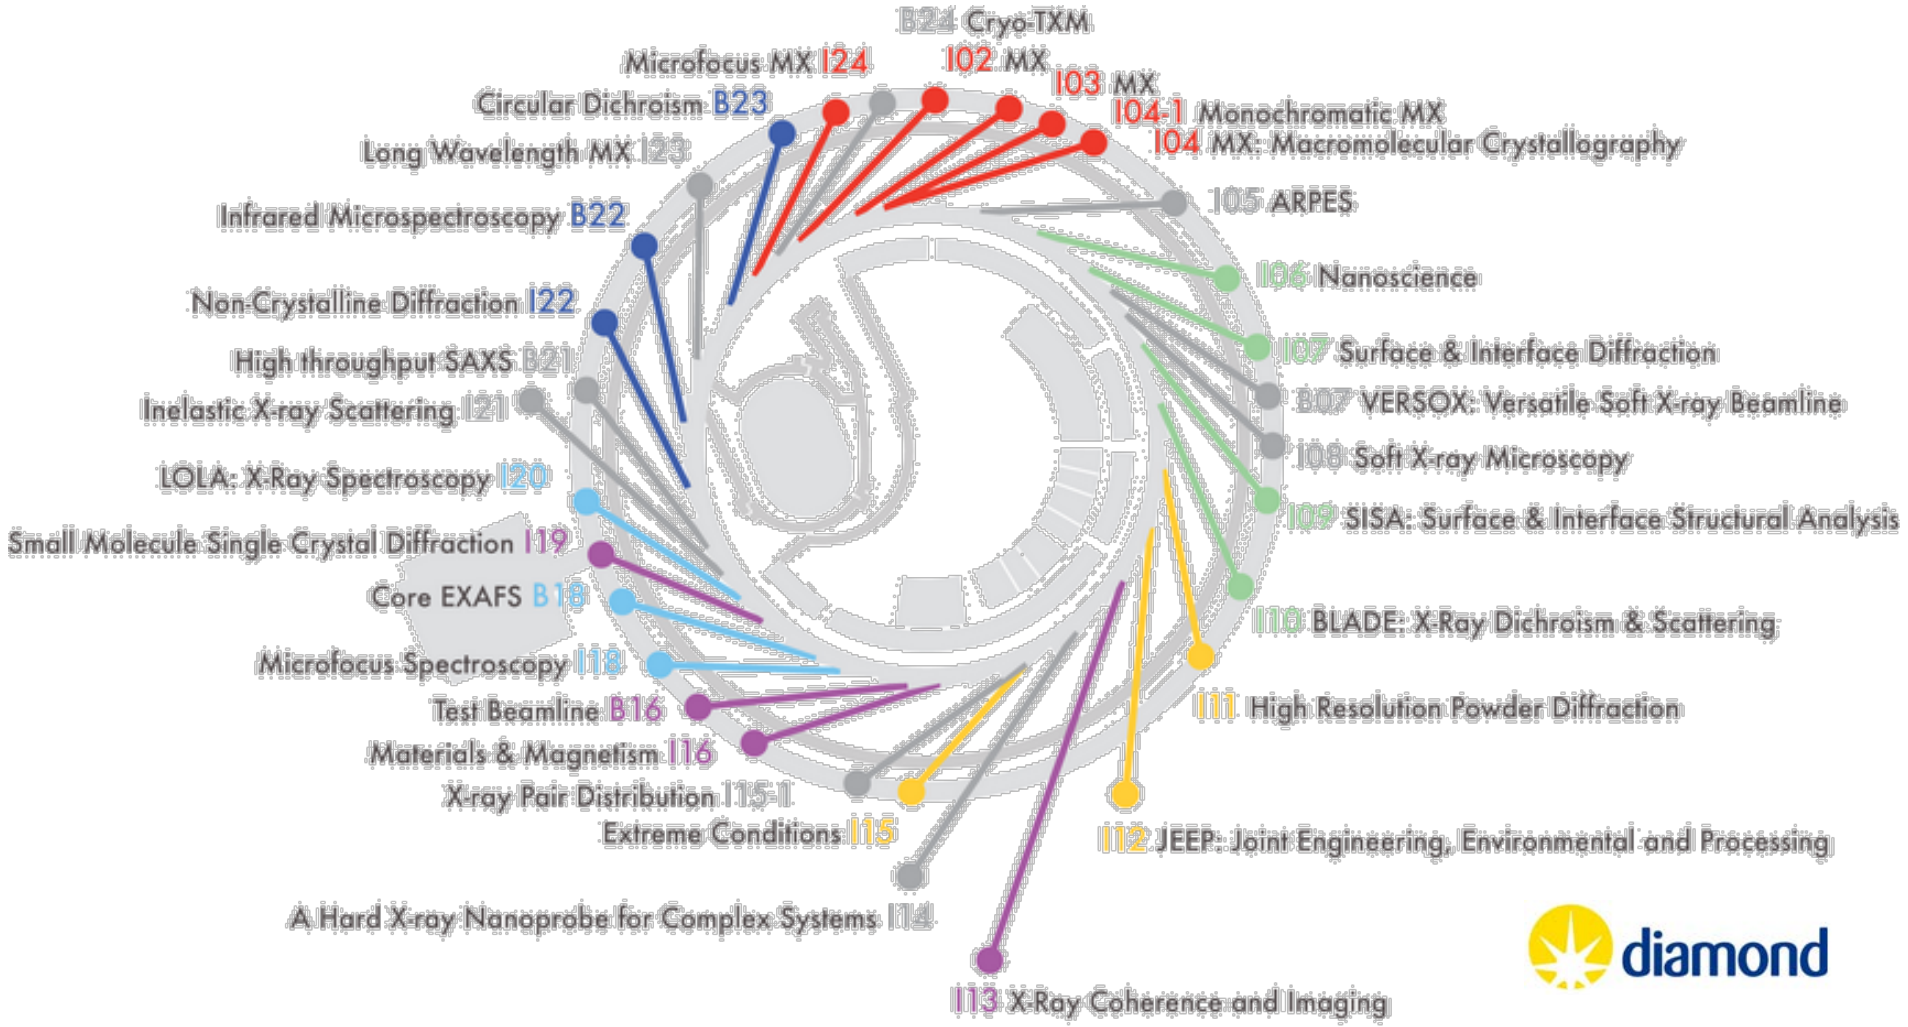

diamond

Data Analysis – Experimental Challenges

BEFORE (FROM DLS/USER'S INSTITUTION)	**IMMEDIATE** (DURING EXPERIMENTS)	**SHORT TERM** (BEFORE THE USER GOES HOME)	**LONG TERM** (FROM DLS/USER'S INSTITUTION)
Simulations Processing of older datasets	"Real time" data processing, analysis and visualisation – to make experimental decisions	Data reduction and processing – Users go home with clean data free of instrument artefacts. Preliminary data analysis – helpful, but may require significant processing power and know-how	Detailed analysis – from data to information. Incorporating results from other techniques. Experiments: ➢Provide parameters for a model. ➢Test/verify a model or theory. ➢Show where a new theory or model is required.

Data Analysis – Experimental Challenges

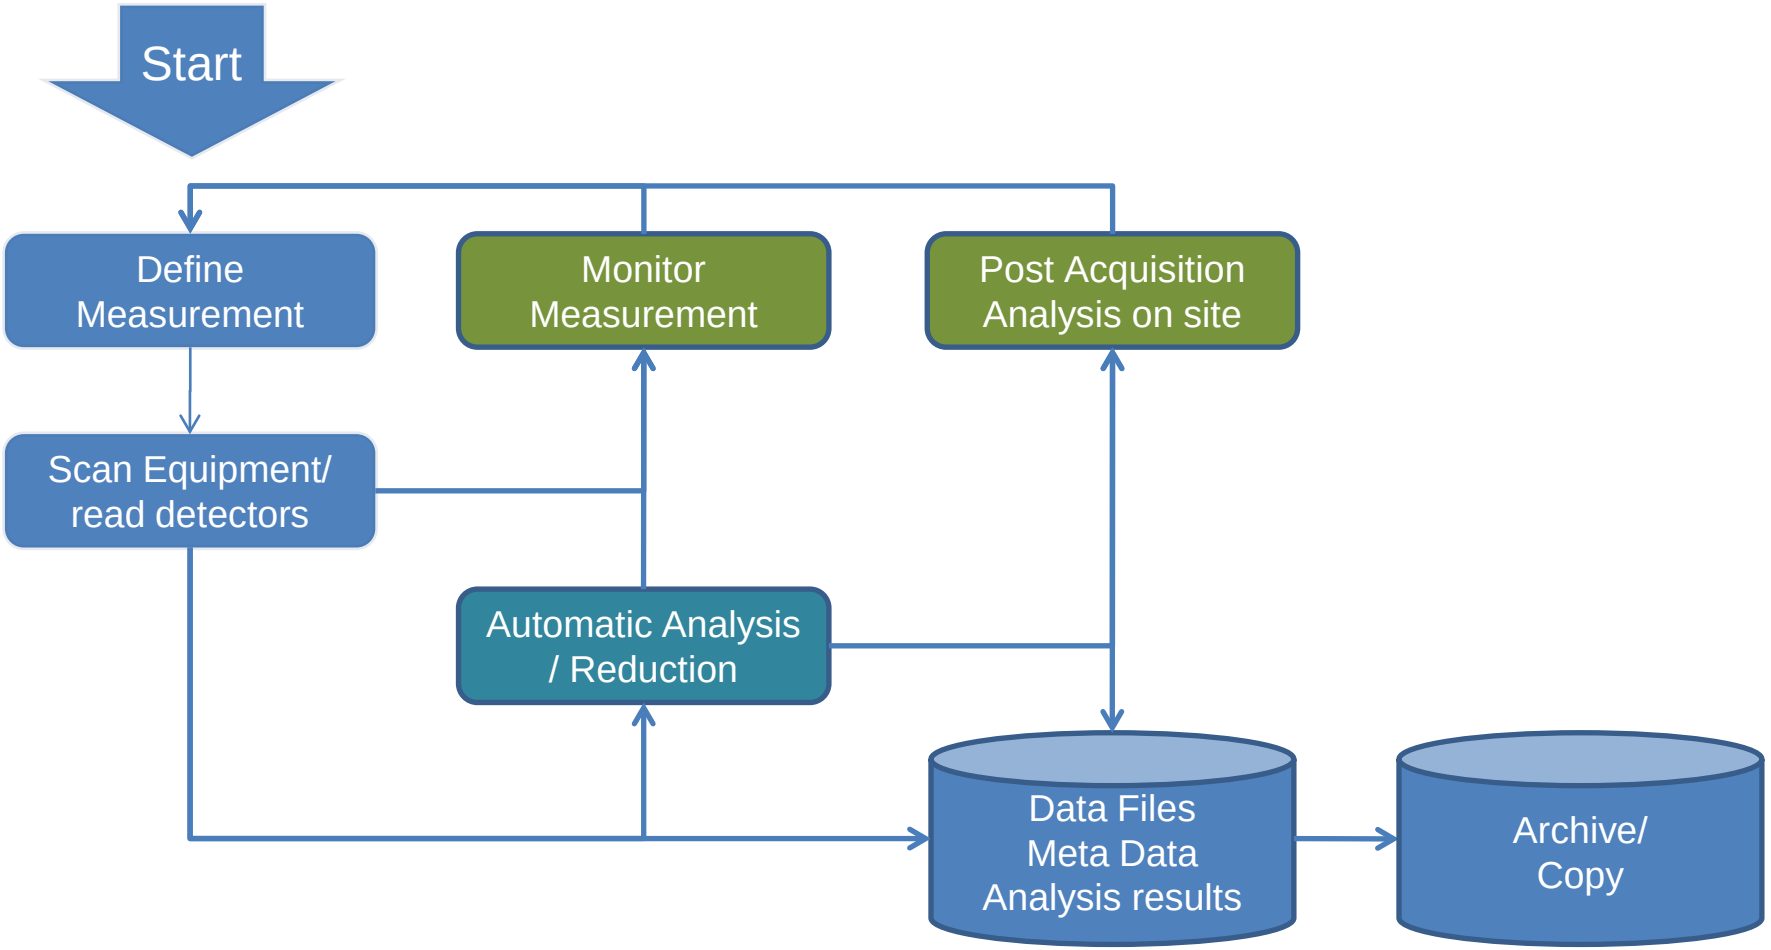

diamond

The ‘Customer’

Data Analysis – Experimental Challenges

How facility users want to analyse their data (sample):

- Command line, interactive analysis
- ‘scripts’
- Black box
- Wizards
- GUI
- Automatically
- Someone else do it
- Don't you look at it

- Quickly
- Estimates
- Properly
- Publication quality
- The way I used to do it
- Use the new stuff
- On all the data
- On some of the data

diamond

Data Analysis – Experimental Challenges

How facility users want to analyse their data (sample):

- Only when I am at the facility
- When I am at home
- On a web page
- On my laptop
- On your computers
- Yesterday
- Next year
- By the student
- By the expert
- Securely
- Shared data
- ASCII not binary...

Other Factors

- Existing codes and expertise
- Plethora of data structures/file formats/data rates
- Available computing resources
  - Enough computing
  - Quick access to Data
- Information management
- Collaborations

diamond

How we are dealing with the problem in Diamond

examples

Beamline Computing/Software Support

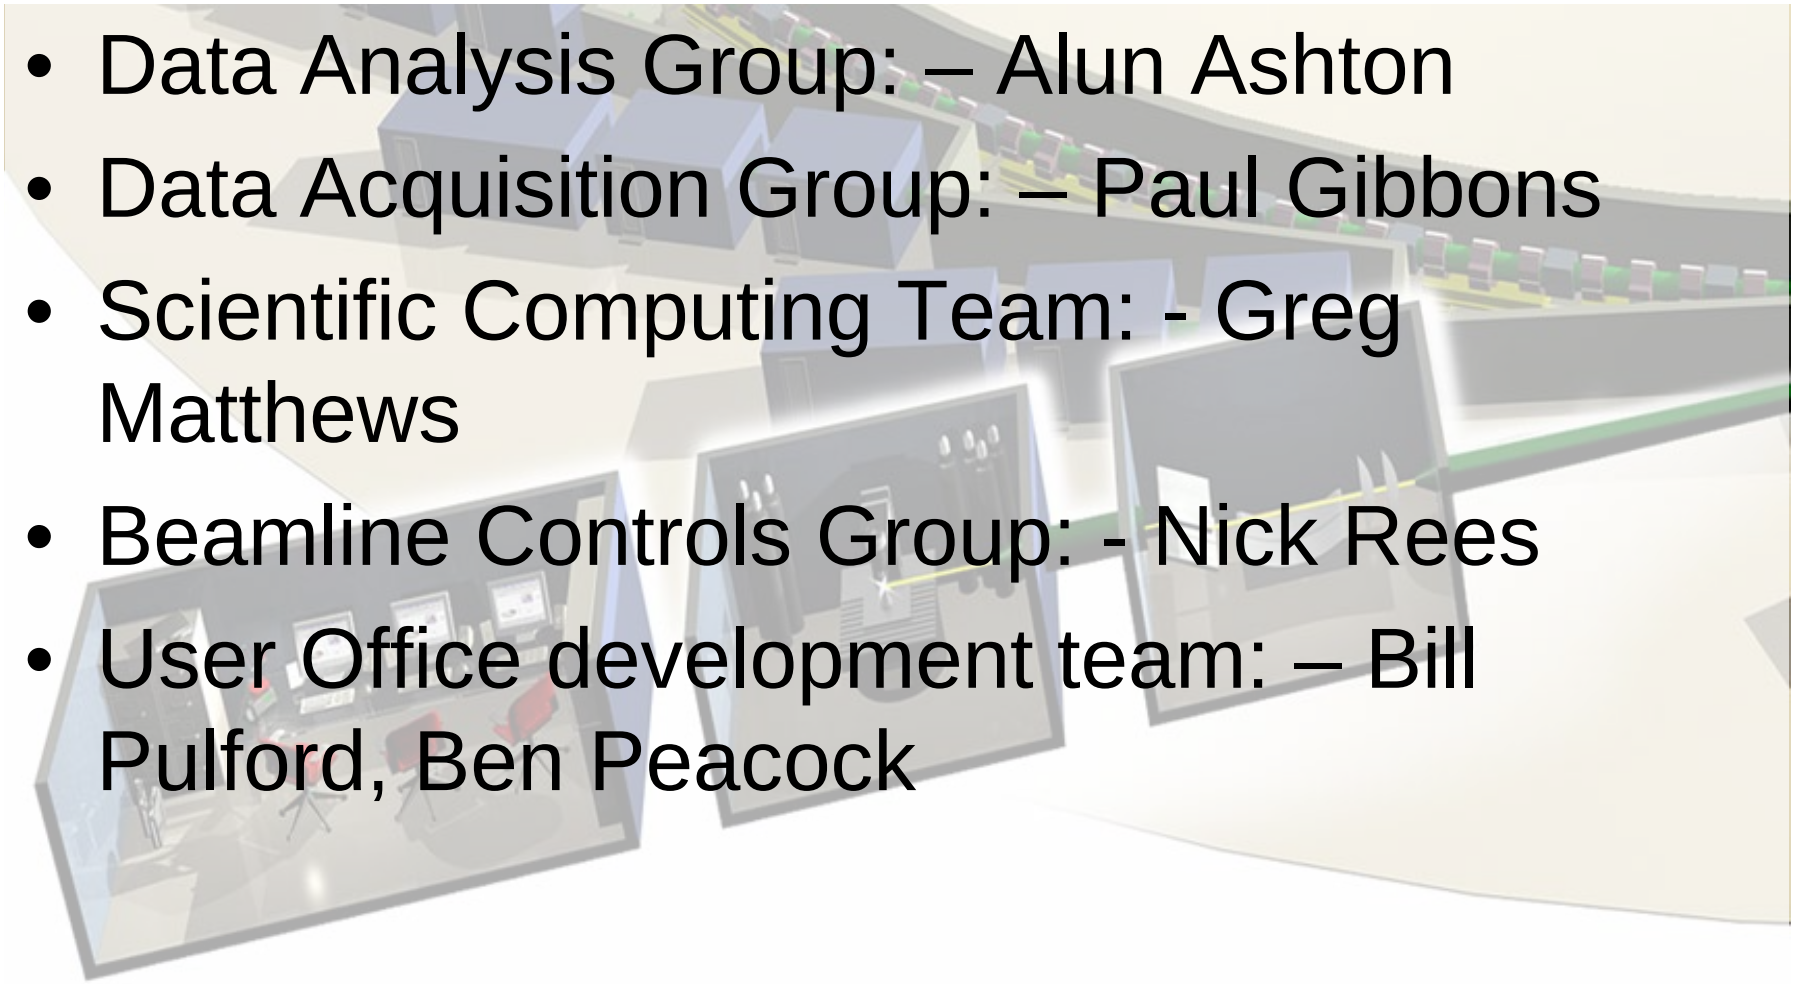

- Data Analysis Group: – Alun Ashton
- Data Acquisition Group: – Paul Gibbons
- Scientific Computing Team: - Greg Matthews
- Beamline Controls Group: - Nick Rees
- User Office development team: – Bill Pulford, Ben Peacock

Acquisition and Analysis in the Eclipse Framework (www.eclipse.org)

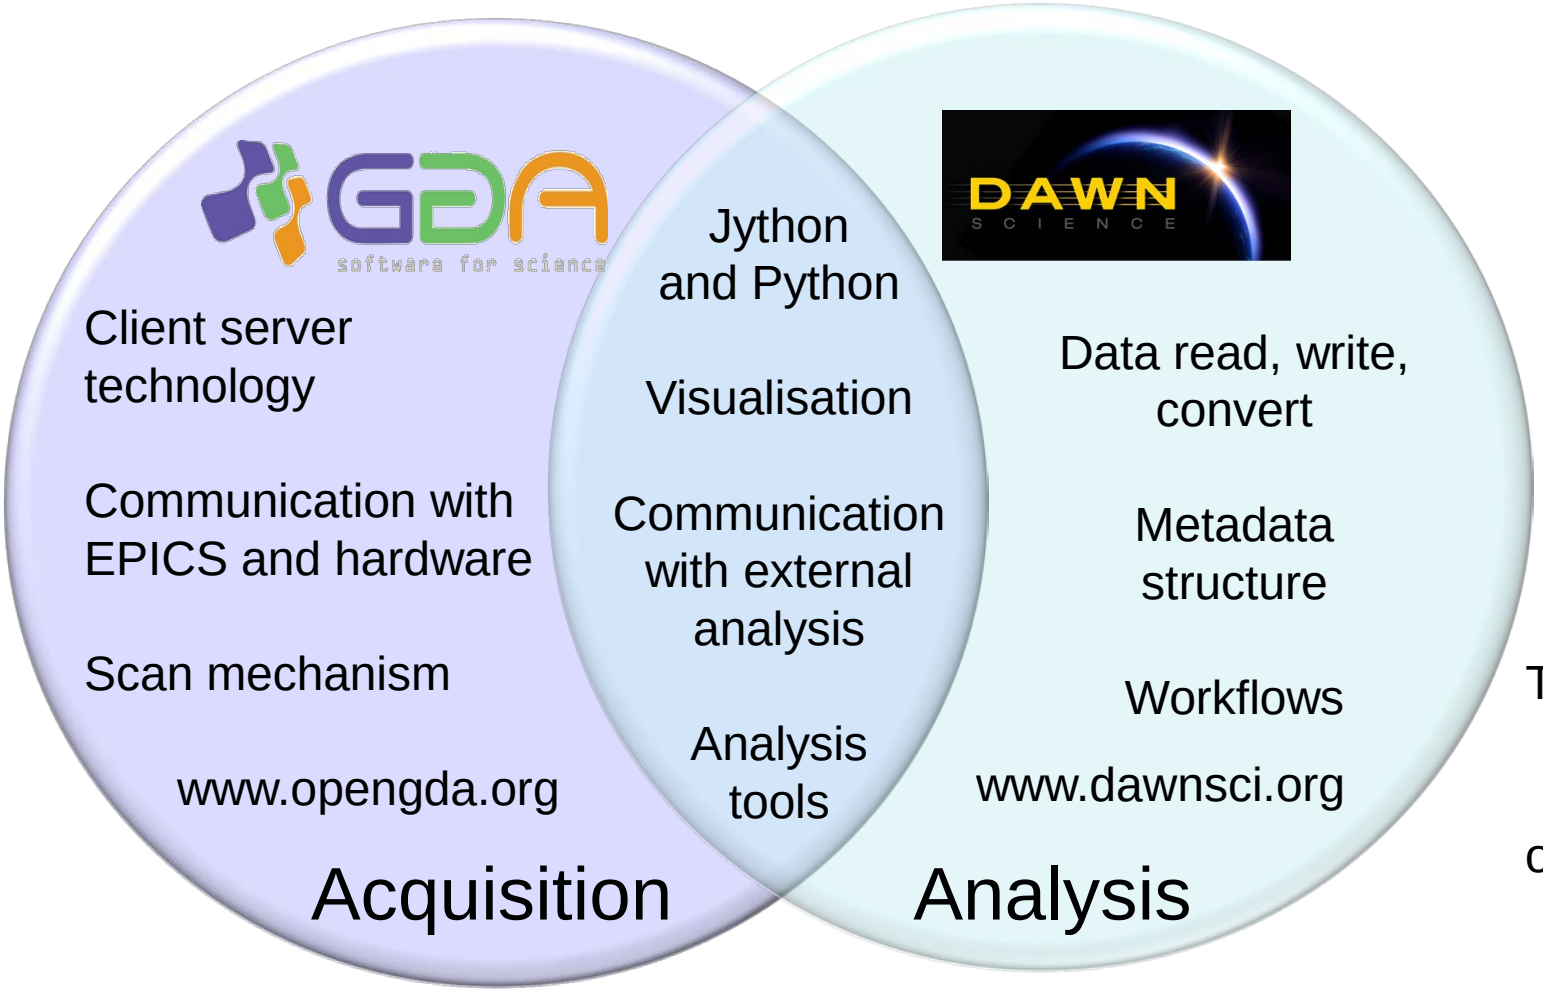

DAWN is a collection of generic and bespoke ‘views’ collated into ‘perspectives’.

The perspectives and views can be used in part or whole in either the GDA or DAWN.

All core technologies open source

diamond

An open source collaboration with a core of generic tools with local extensions and implementations.

- Visualisation of scientific data
- Workflow tool for the treatment/manipulation of scientific data
- Integrated python environment

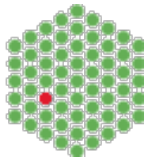

DAWN SCIENCE

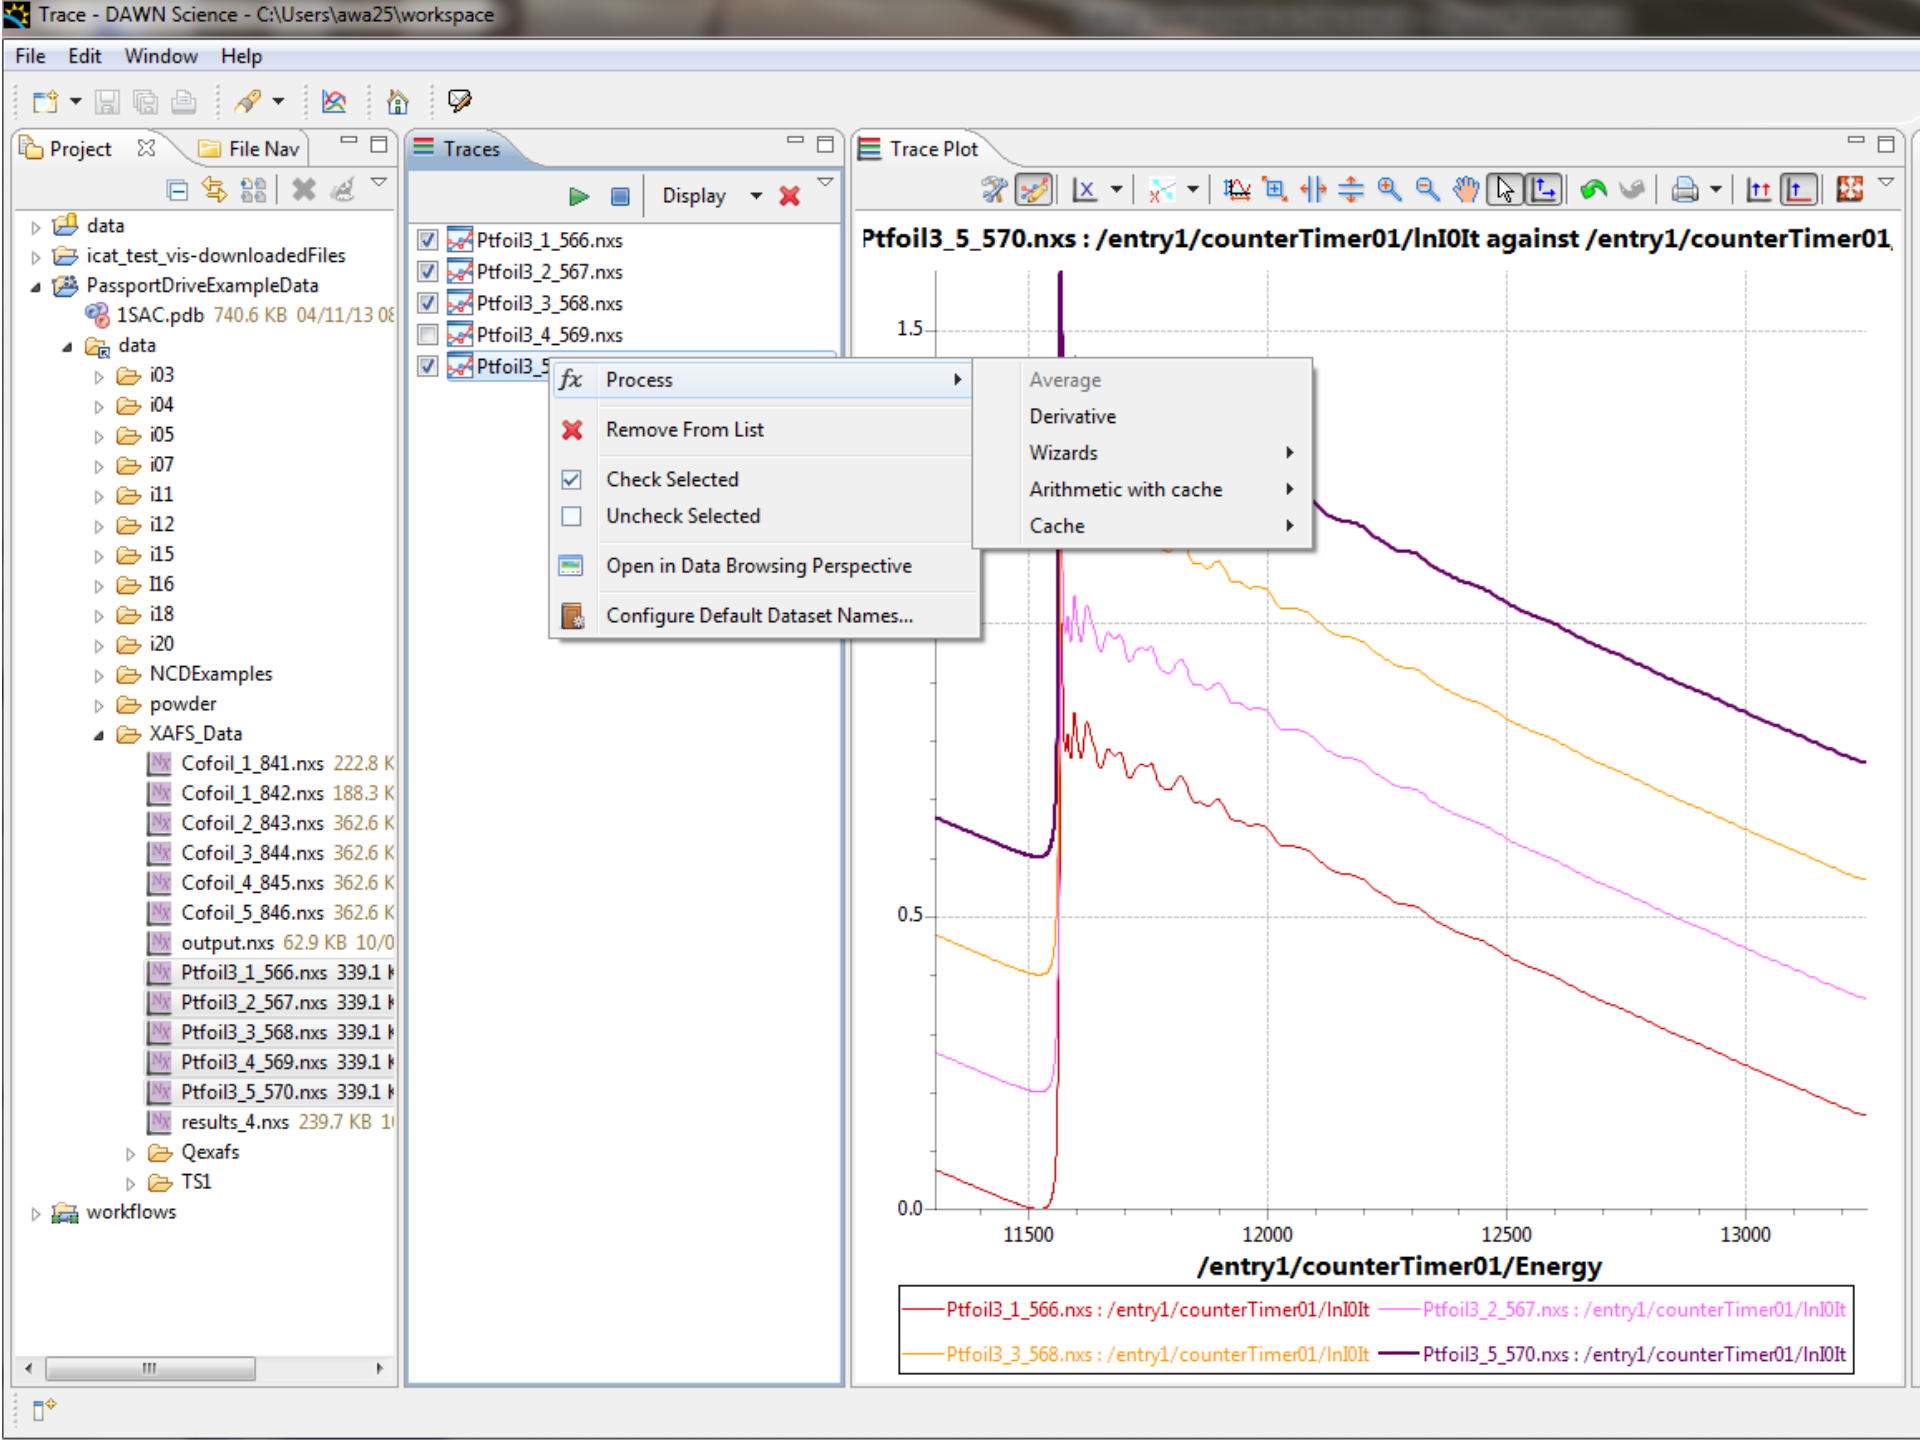
Trace - DAWN Science - C:\Users\awa25\workspace
File Edit Window Help
Project
File Nav
data
icat_test_vis-downloadedFiles
PassportDriveExampleData
1SAC.pdb 740.6 KB 04/11/13 08
data
i03
i04
i05
i07
i11
i12
i15
I16
i18
i20
NCDExamples
powder
XAFS_Data
Cofoil_1_841.nxs 222.8 K
Cofoil_1_842.nxs 188.3 K
Cofoil_2_843.nxs 362.6 K
Cofoil_3_844.nxs 362.6 K
Cofoil_4_845.nxs 362.6 K
Cofoil_5_846.nxs 362.6 K
output.nxs 62.9 KB 10/0
Ptfoil3_1_566.nxs 339.1 K
Ptfoil3_2_567.nxs 339.1 K
Ptfoil3_3_568.nxs 339.1 K
Ptfoil3_4_569.nxs 339.1 K
Ptfoil3_5_570.nxs 339.1 K
results_4.nxs 239.7 KB 1
Qexafs
TS1
workflows
Traces
Display
Ptfoil3_1_566.nxs
Ptfoil3_2_567.nxs
Ptfoil3_3_568.nxs
Ptfoil3_4_569.nxs
Ptfoil3_5
Process
Remove From List
Check Selected
Uncheck Selected
Open in Data Browsing Perspective
Configure Default Dataset Names...
Average
Derivative
Wizards
Arithmetic with cache
Cache
Trace Plot
Ptfoil3_5_570.nxs : /entry1/counterTimer01/lnI0It against /entry1/counterTimer01
1.5
0.5
0.0
11500
12000
12500
13000
/entry1/counterTimer01/Energy
Ptfoil3_1_566.nxs : /entry1/counterTimer01/lnI0It
Ptfoil3_2_567.nxs : /entry1/counterTimer01/lnI0It
Ptfoil3_3_568.nxs : /entry1/counterTimer01/lnI0It
Ptfoil3_5_570.nxs : /entry1/counterTimer01/lnI0It

Image stack analysis

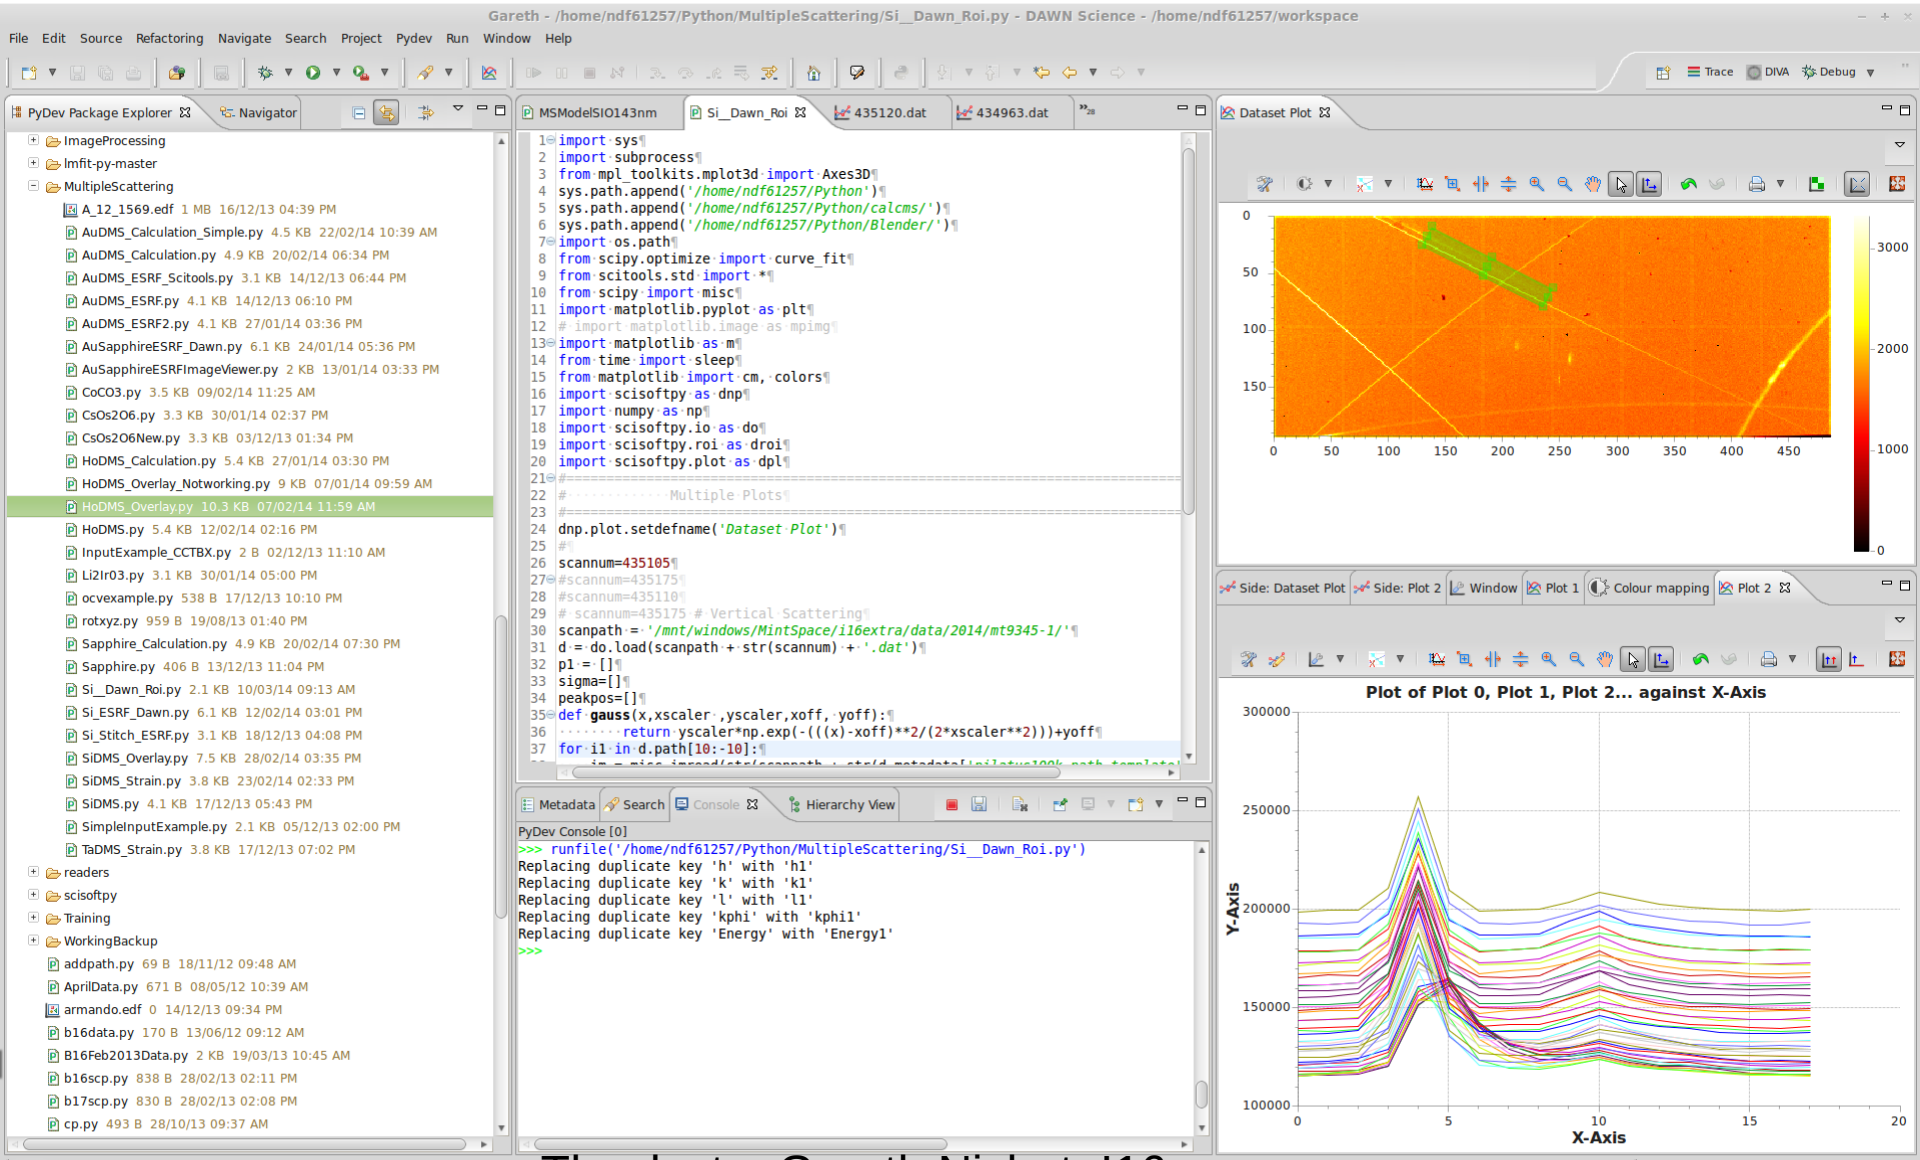

Thanks to: Gareth Nisbet, I16

2D Diffraction Processing

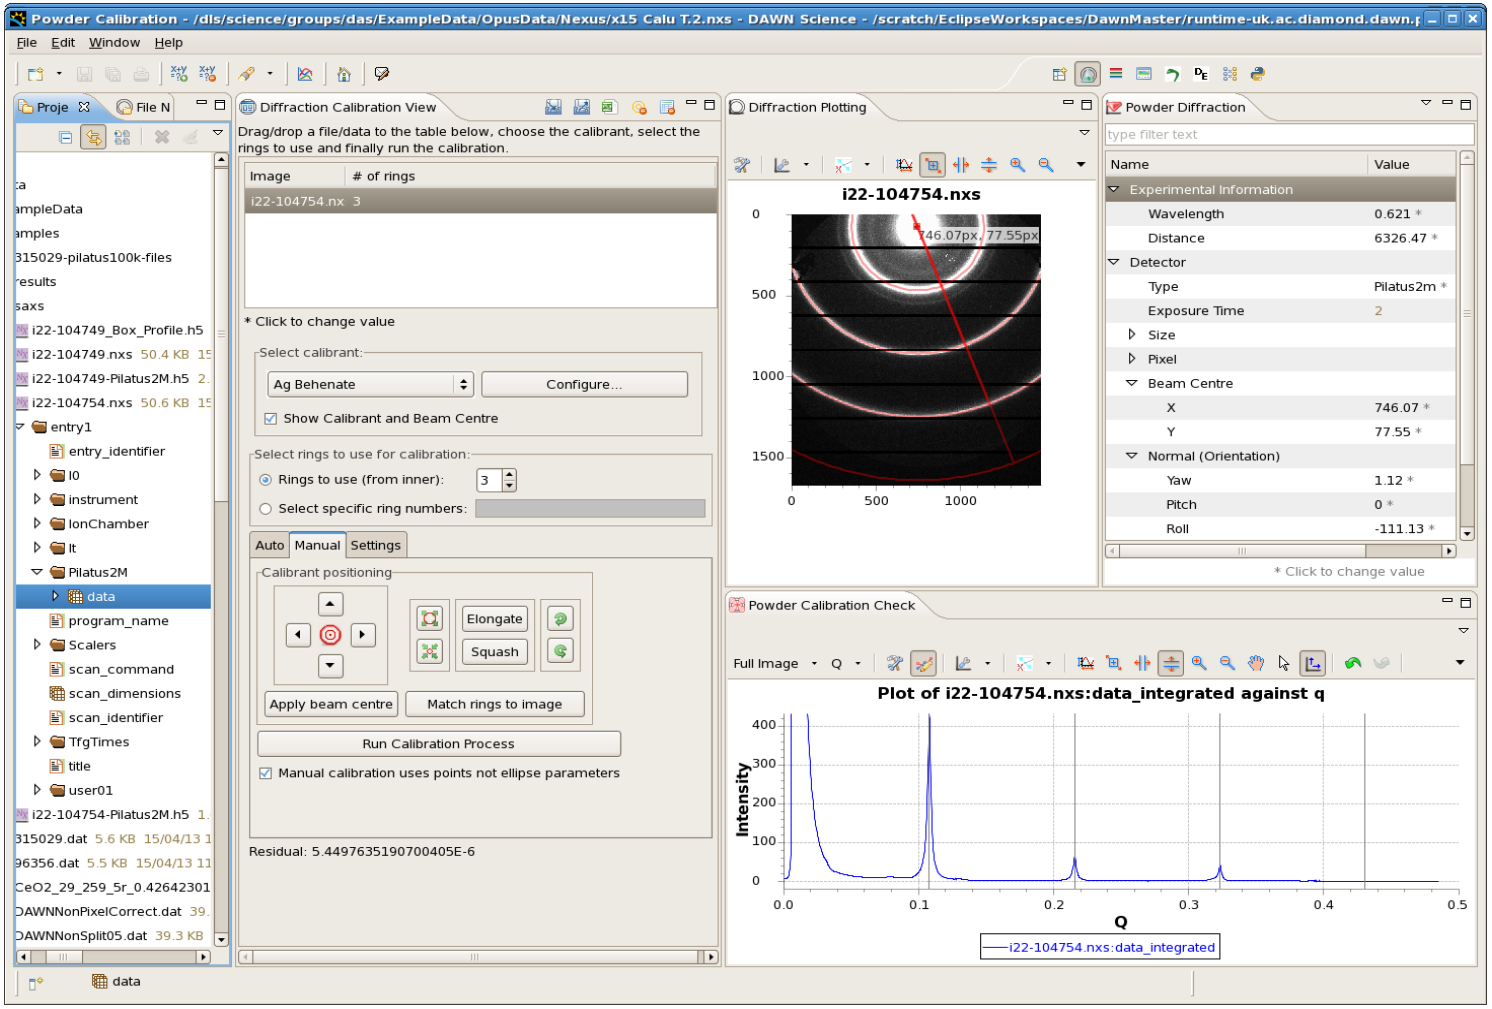

BioSAXS experiment

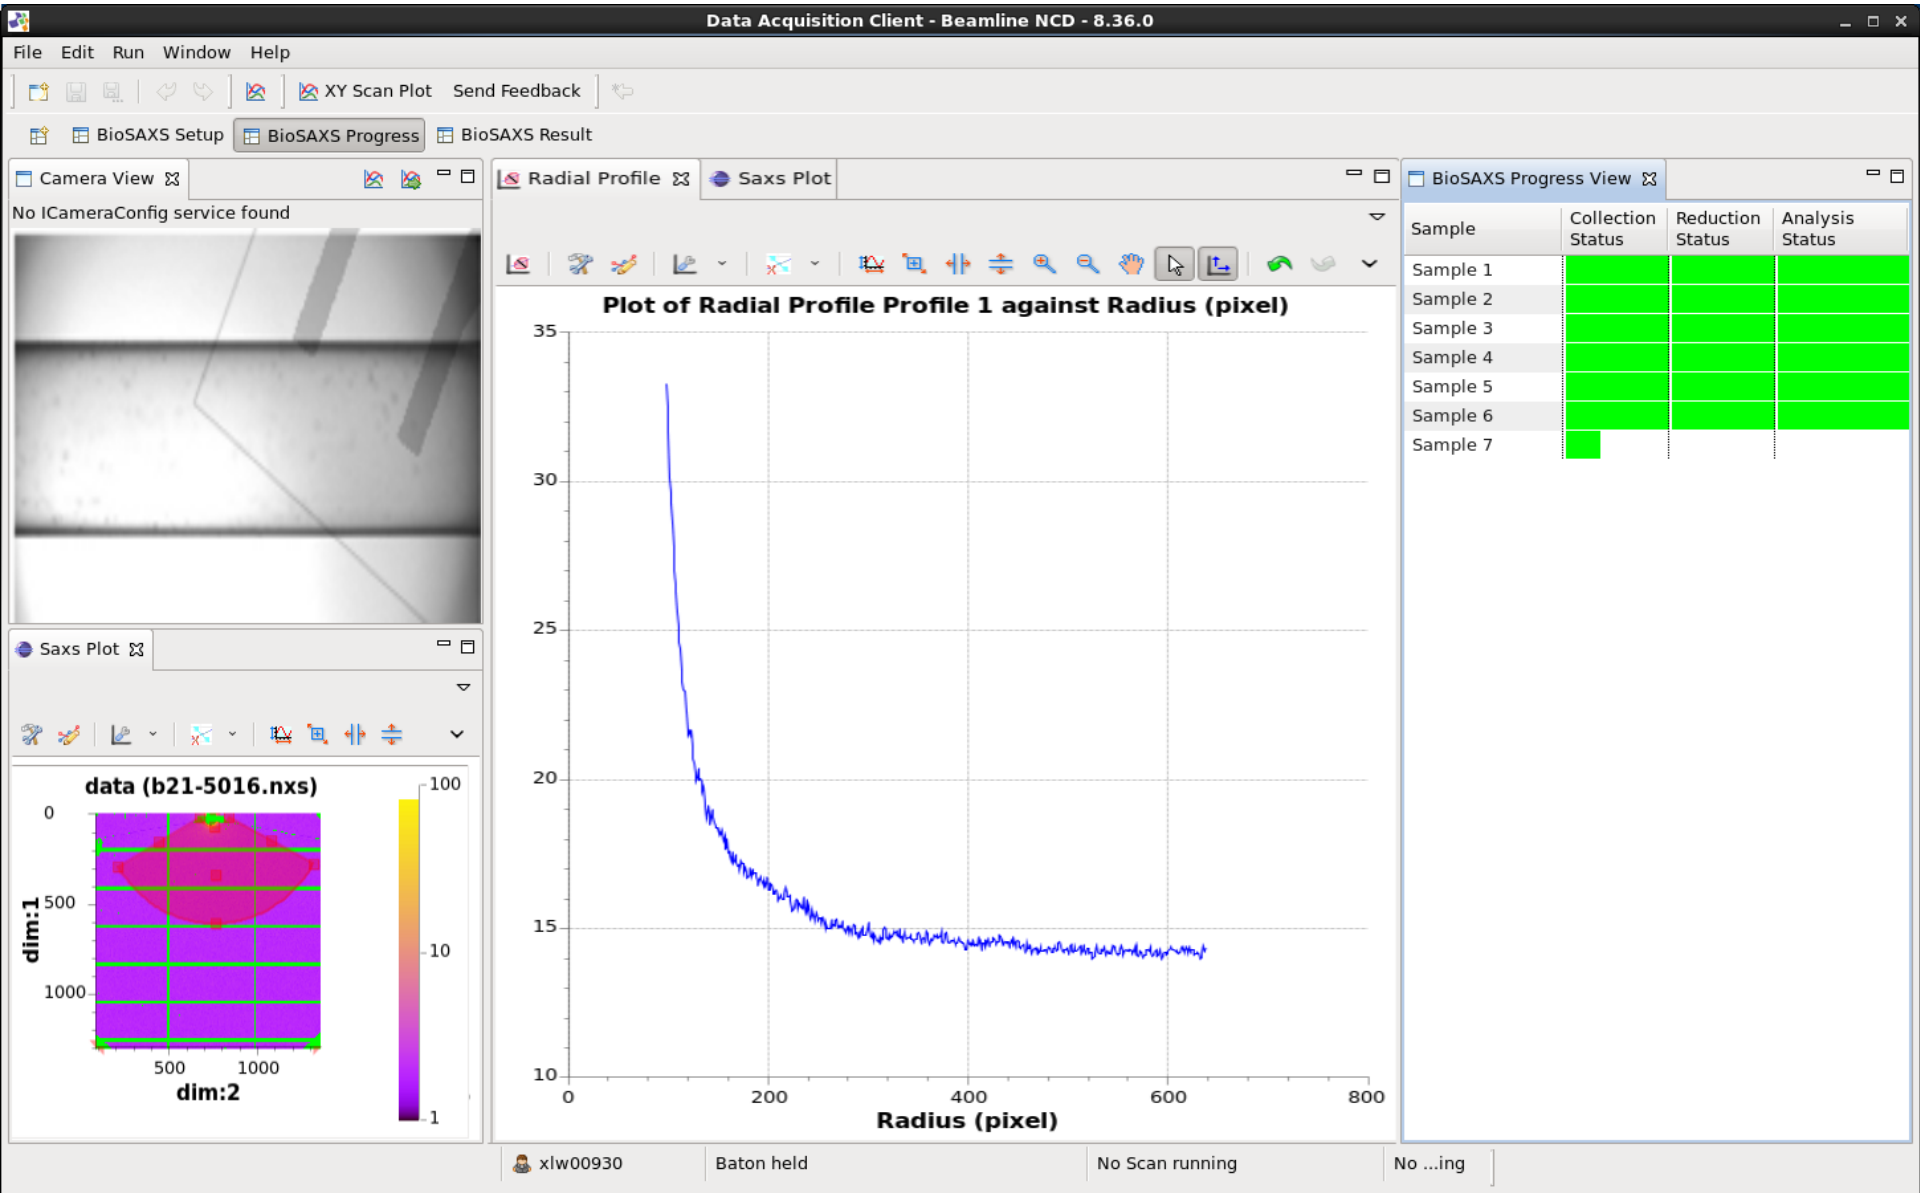

SAXS Data Reduction

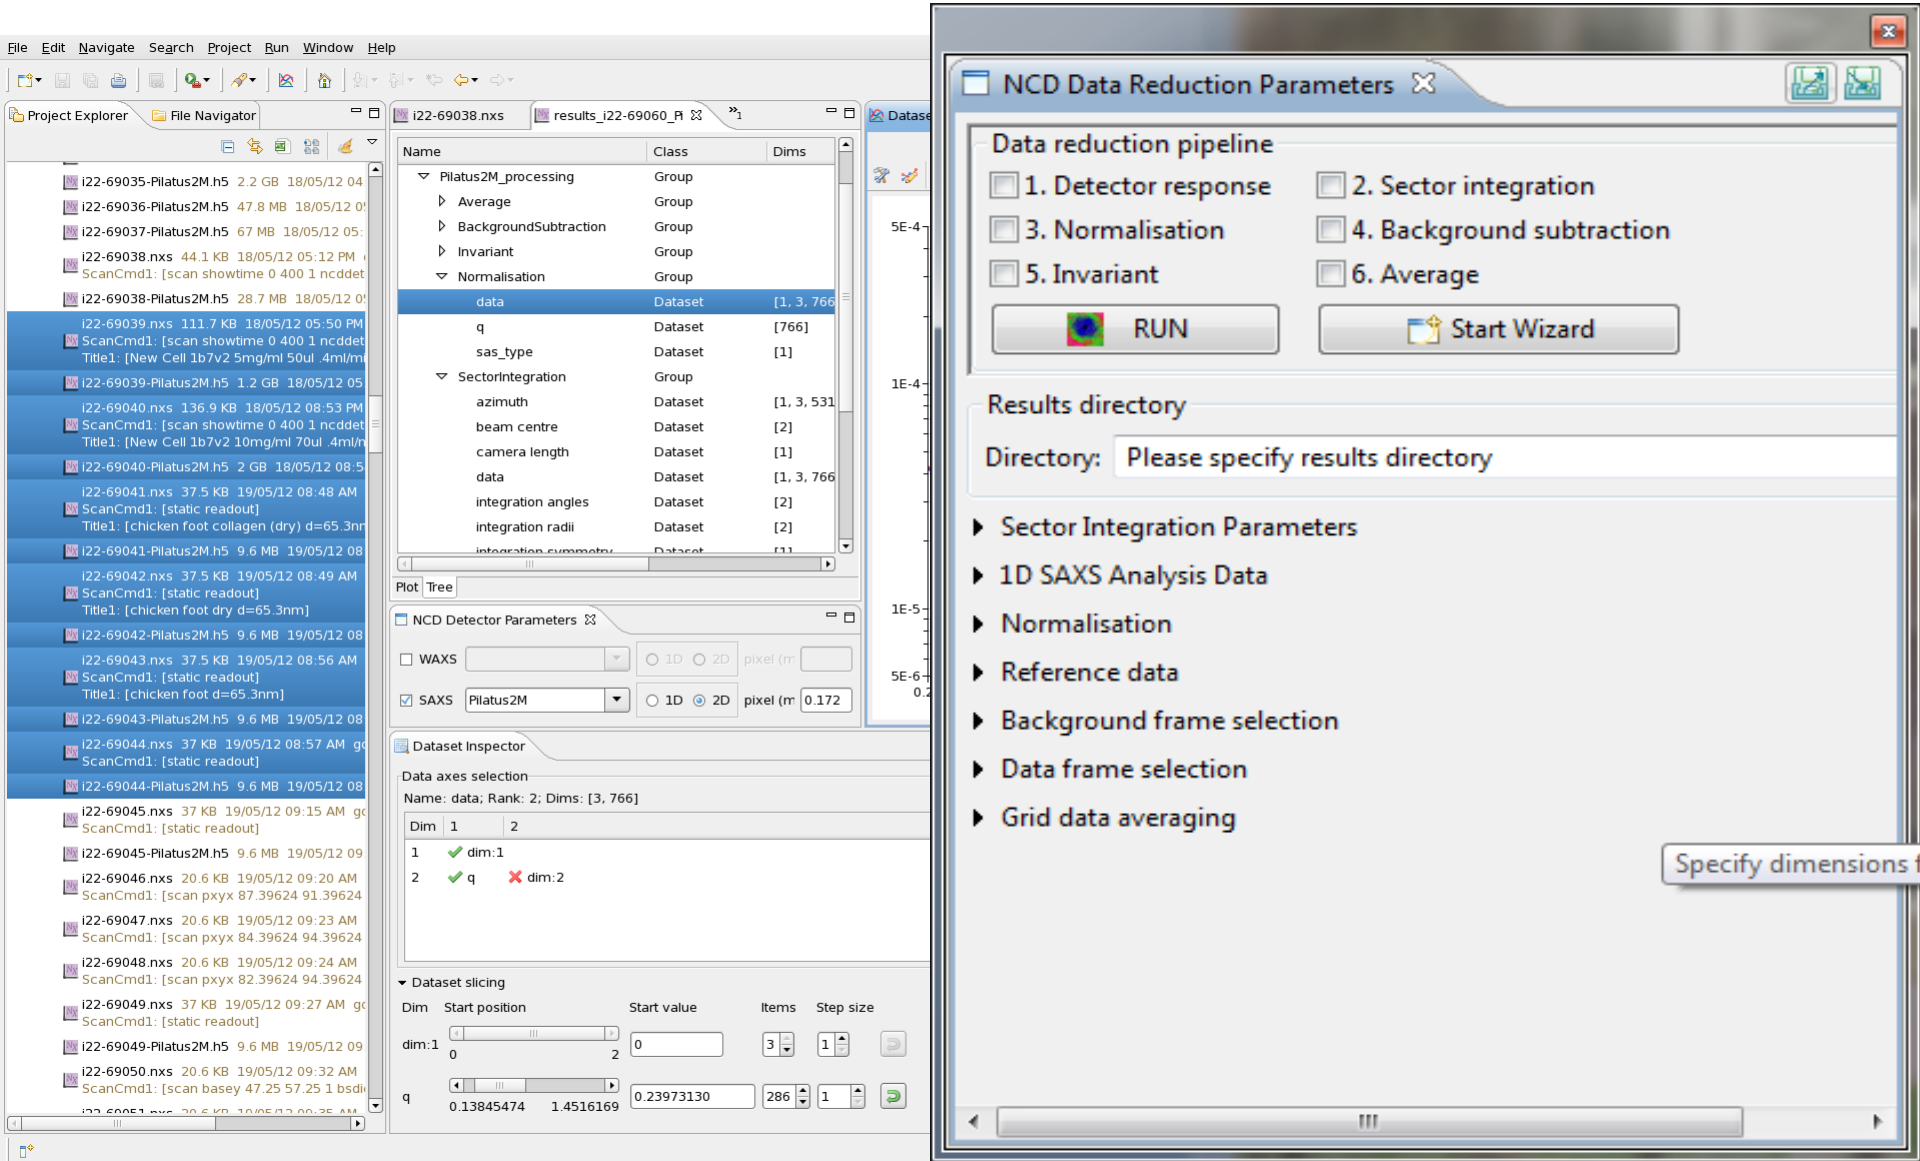

Automated Data Reduction

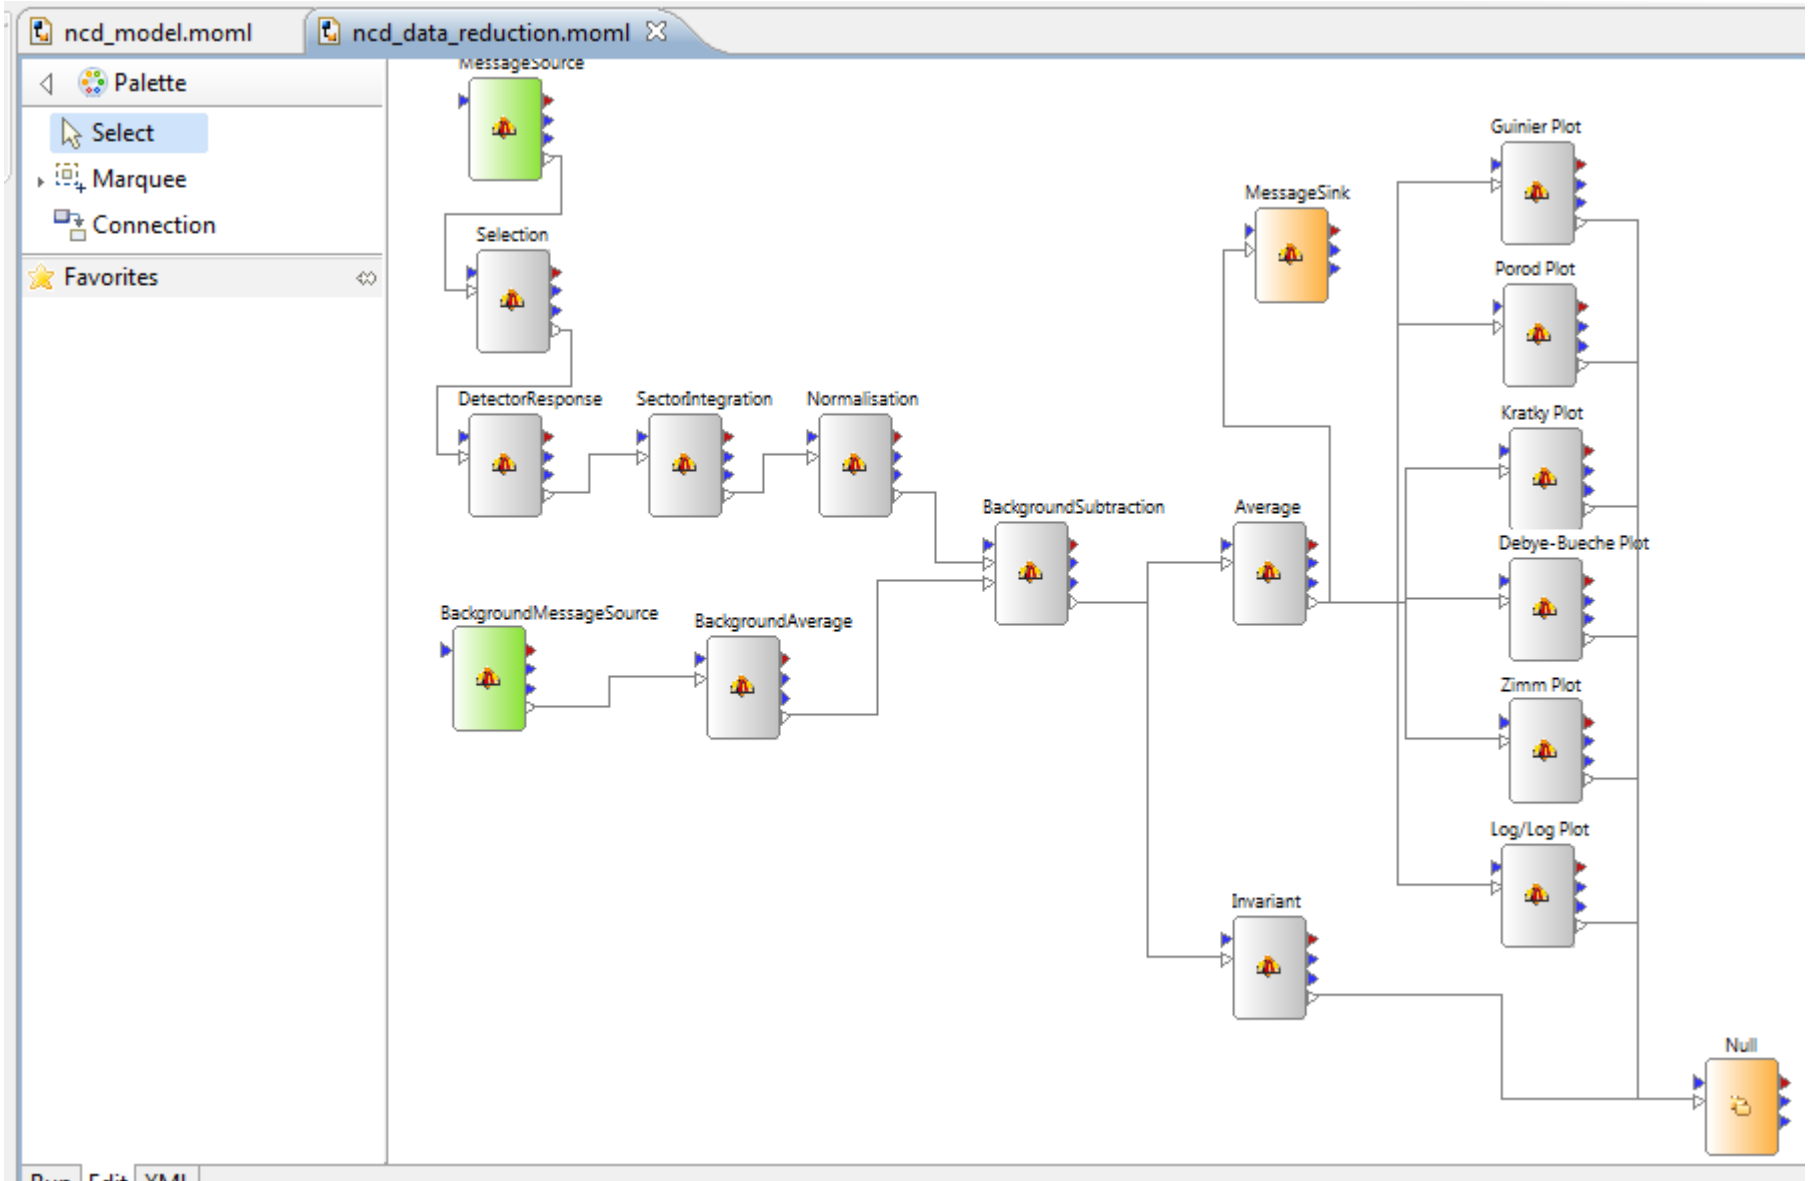

Tomography Reconstruction GUI

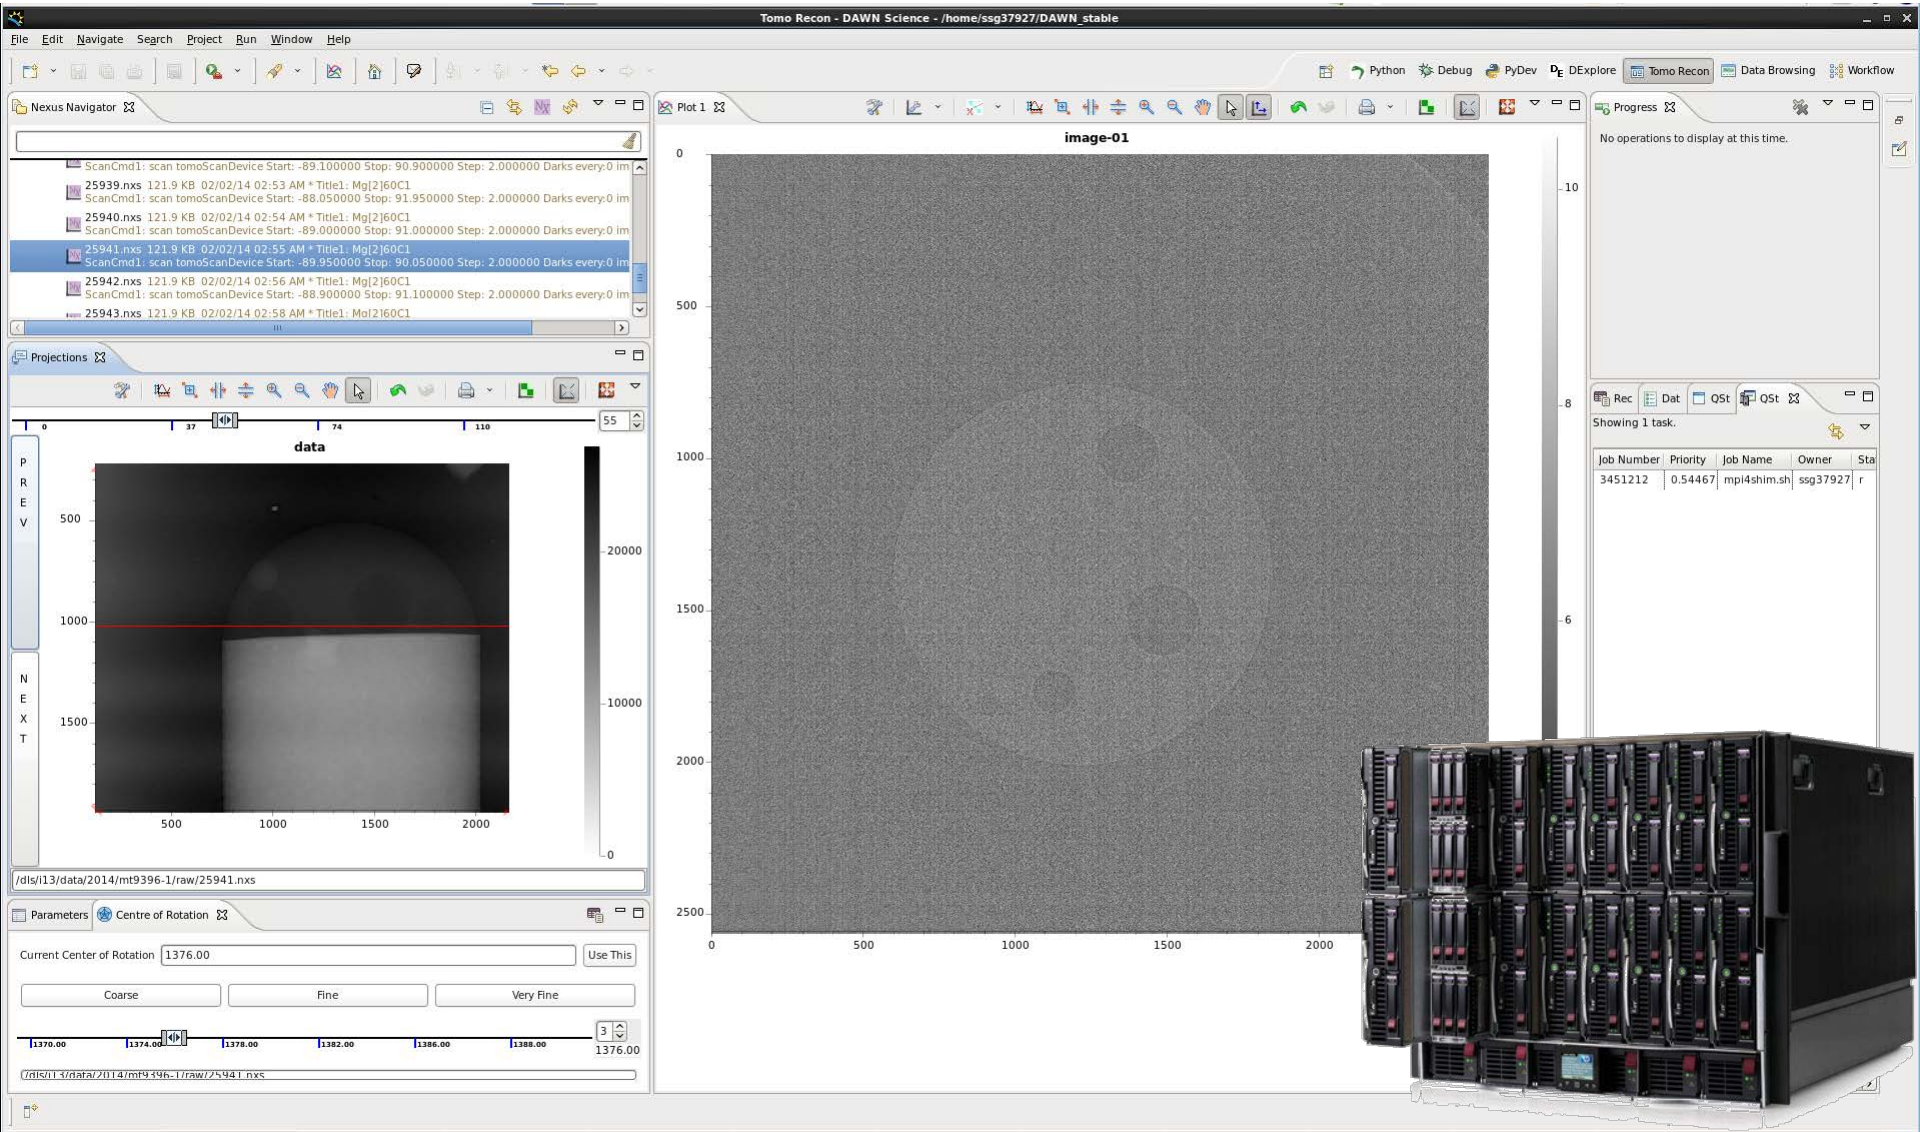

ARPES Angle-Resolved Photo-Emission Spectroscopy

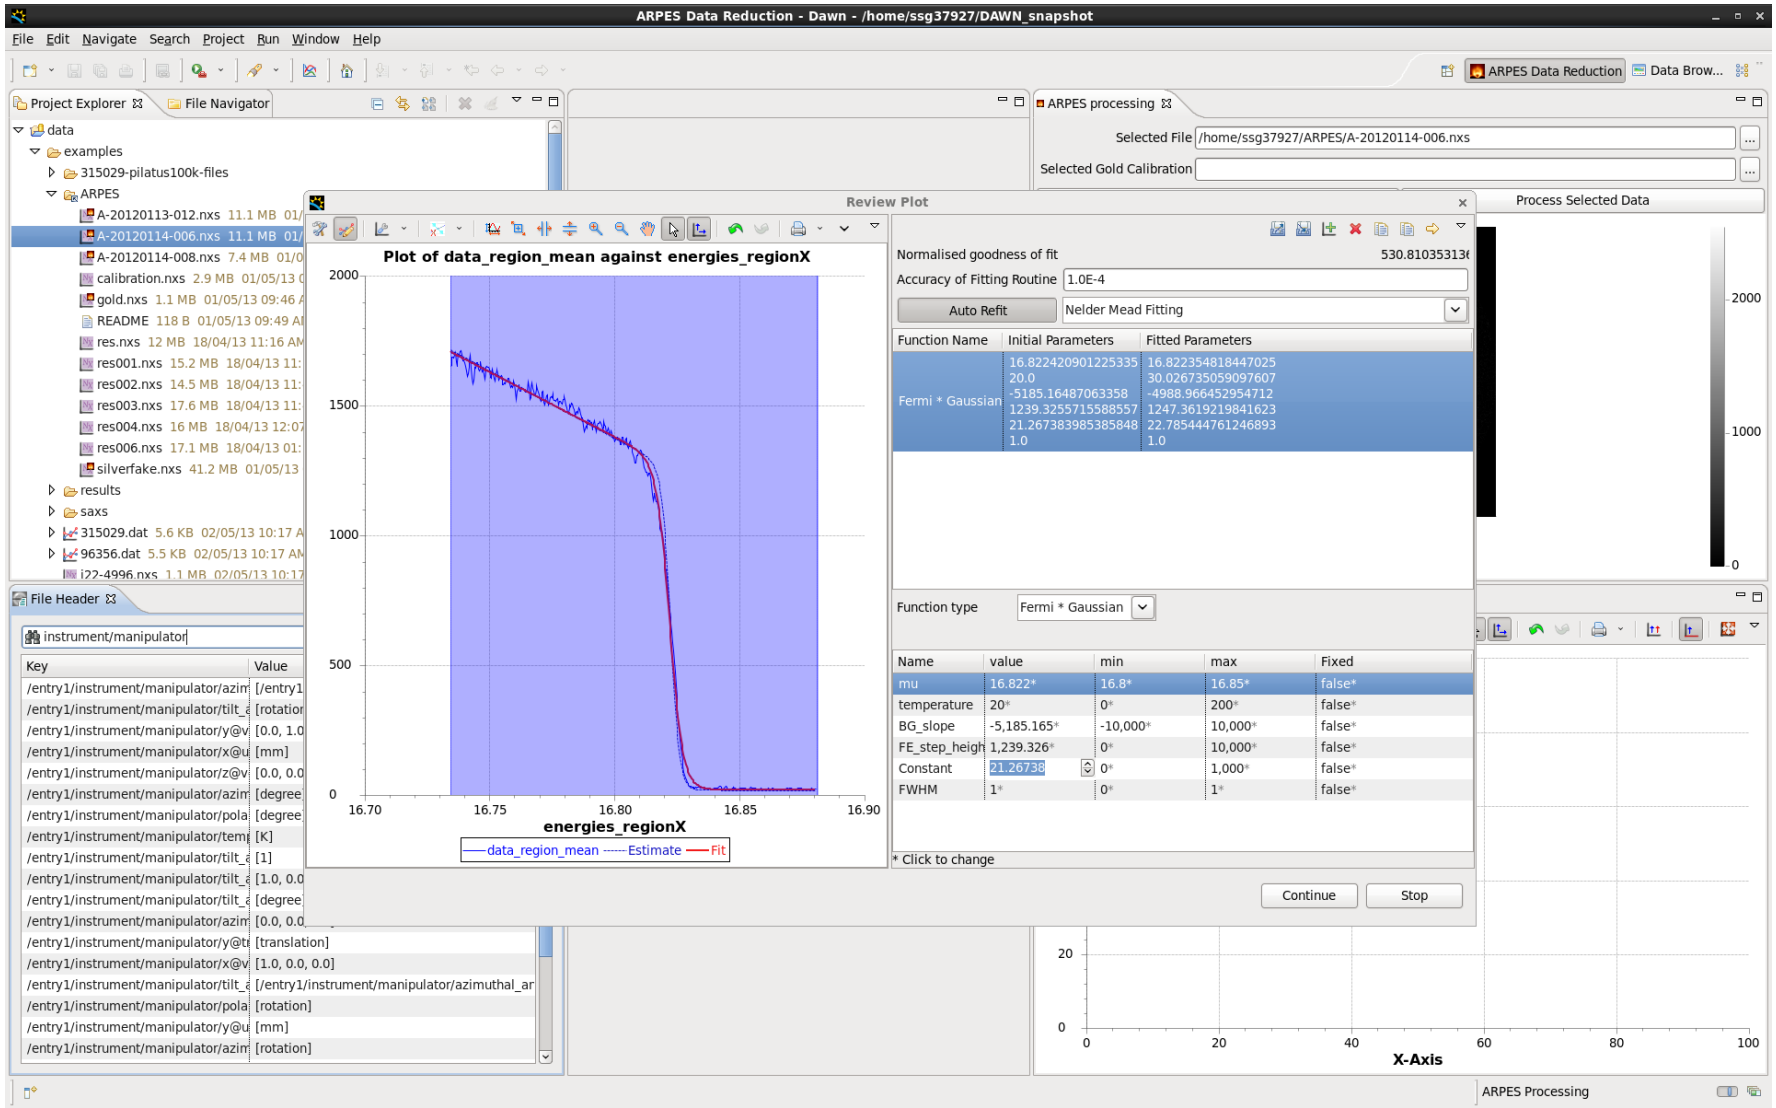

ARPES

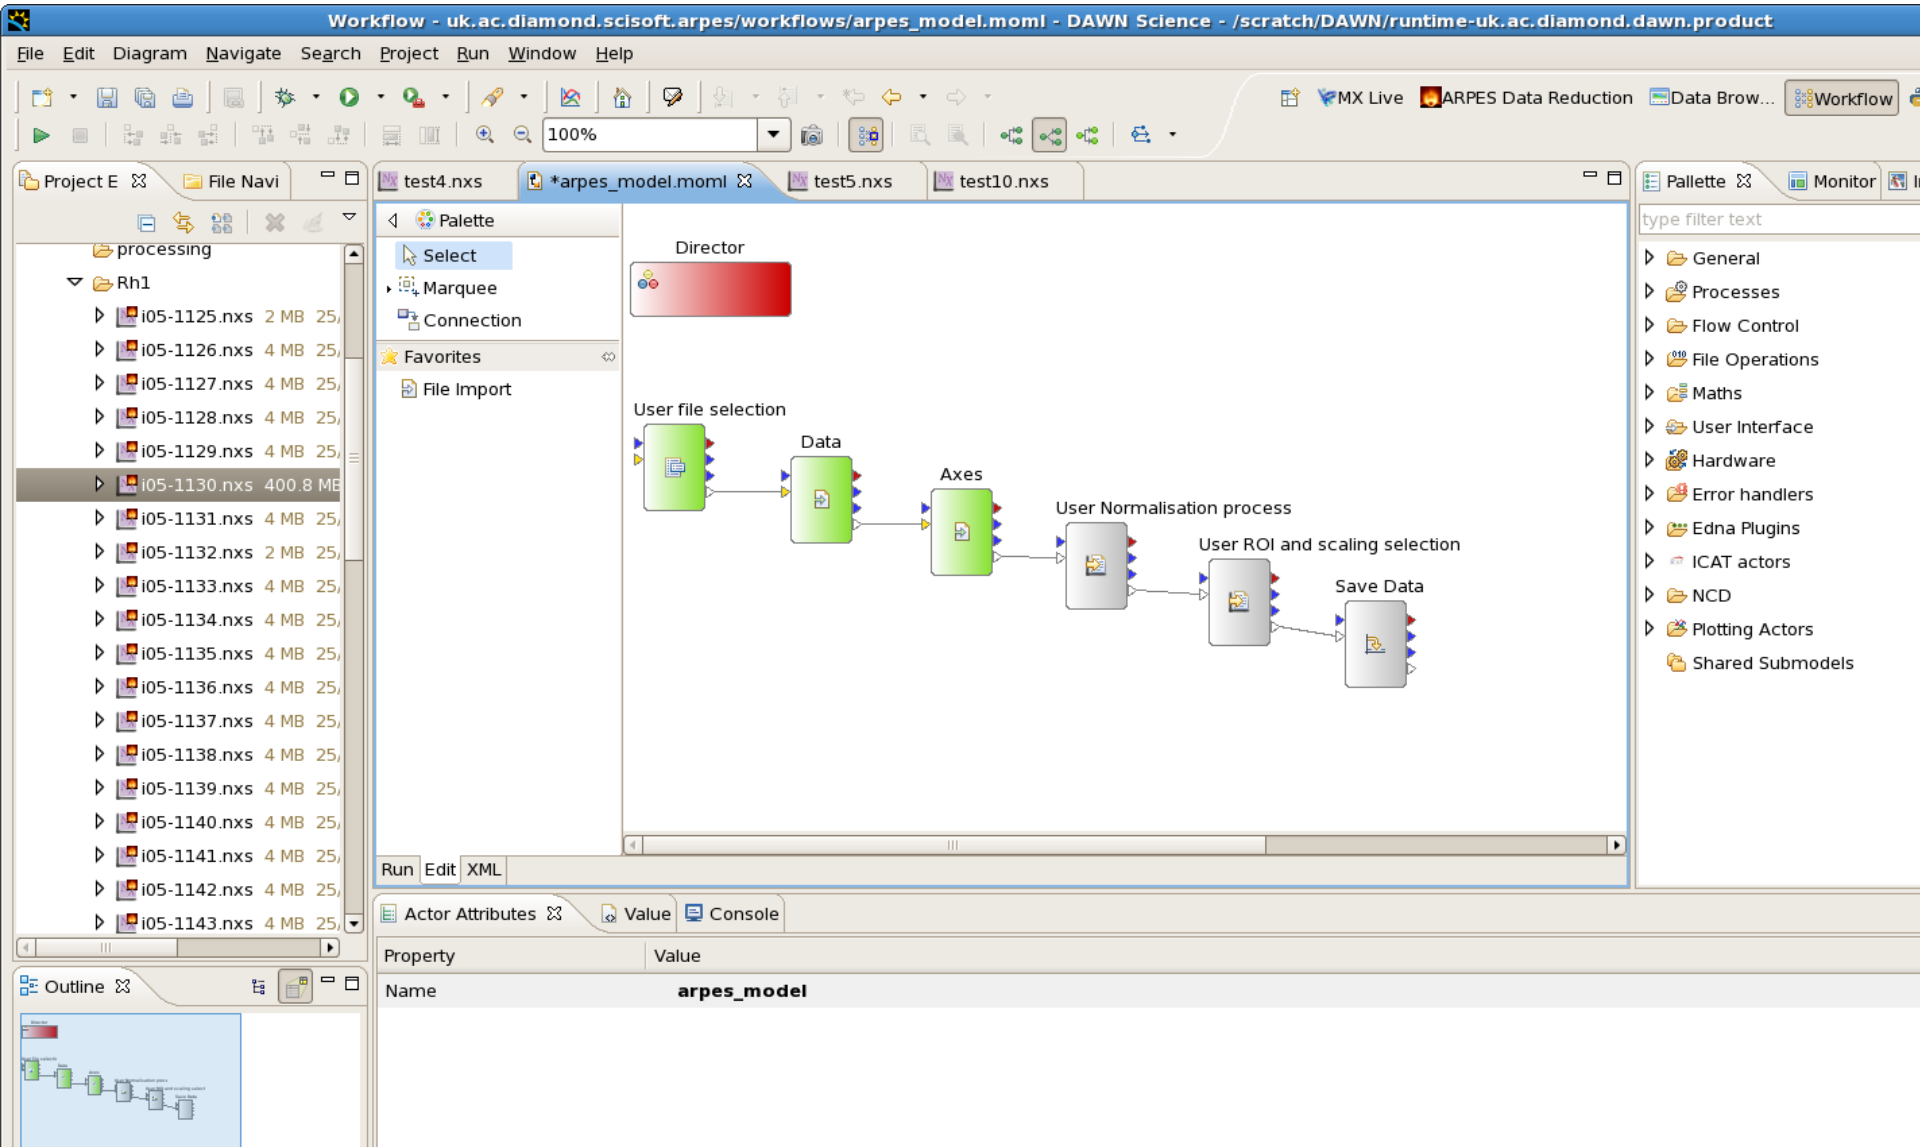
Workflow - uk.ac.diamond.scisoft.arpes/workflows/arpes_model.moml - DAWN Science - /scratch/DAWN/runtime-uk.ac.diamond.dawn.product
File Edit Diagram Navigate Search Project Run Window Help
MX Live
ARPES Data Reduction
Data Brow...
Workflow
100%
Project E
File Navi
processing
Rh1
i05-1125.nxs 2 MB
i05-1126.nxs 4 MB
i05-1127.nxs 4 MB
i05-1128.nxs 4 MB
i05-1129.nxs 4 MB
i05-1130.nxs 400.8 MB
i05-1131.nxs 4 MB
i05-1132.nxs 2 MB
i05-1133.nxs 4 MB
i05-1134.nxs 4 MB
i05-1135.nxs 4 MB
i05-1136.nxs 4 MB
i05-1137.nxs 4 MB
i05-1138.nxs 4 MB
i05-1139.nxs 4 MB
i05-1140.nxs 4 MB
i05-1141.nxs 4 MB
i05-1142.nxs 4 MB
i05-1143.nxs 4 MB
Outline
test4.nxs
*arpes_model.moml
test5.nxs
test10.nxs
Palette
Select
Marquee
Connection
Favorites
File Import
Director
User file selection
Data
Axes
User Normalisation process
User ROI and scaling selection
Save Data
Run
Edit
XML
Pallette
Monitor
type filter text
General
Processes
Flow Control
File Operations
Maths
User Interface
Hardware
Error handlers
Edna Plugins
ICAT actors
NCD
Plotting Actors
Shared Submodels
Actor Attributes
Value
Console
Property
Value
Name
arpes_model

ARPES

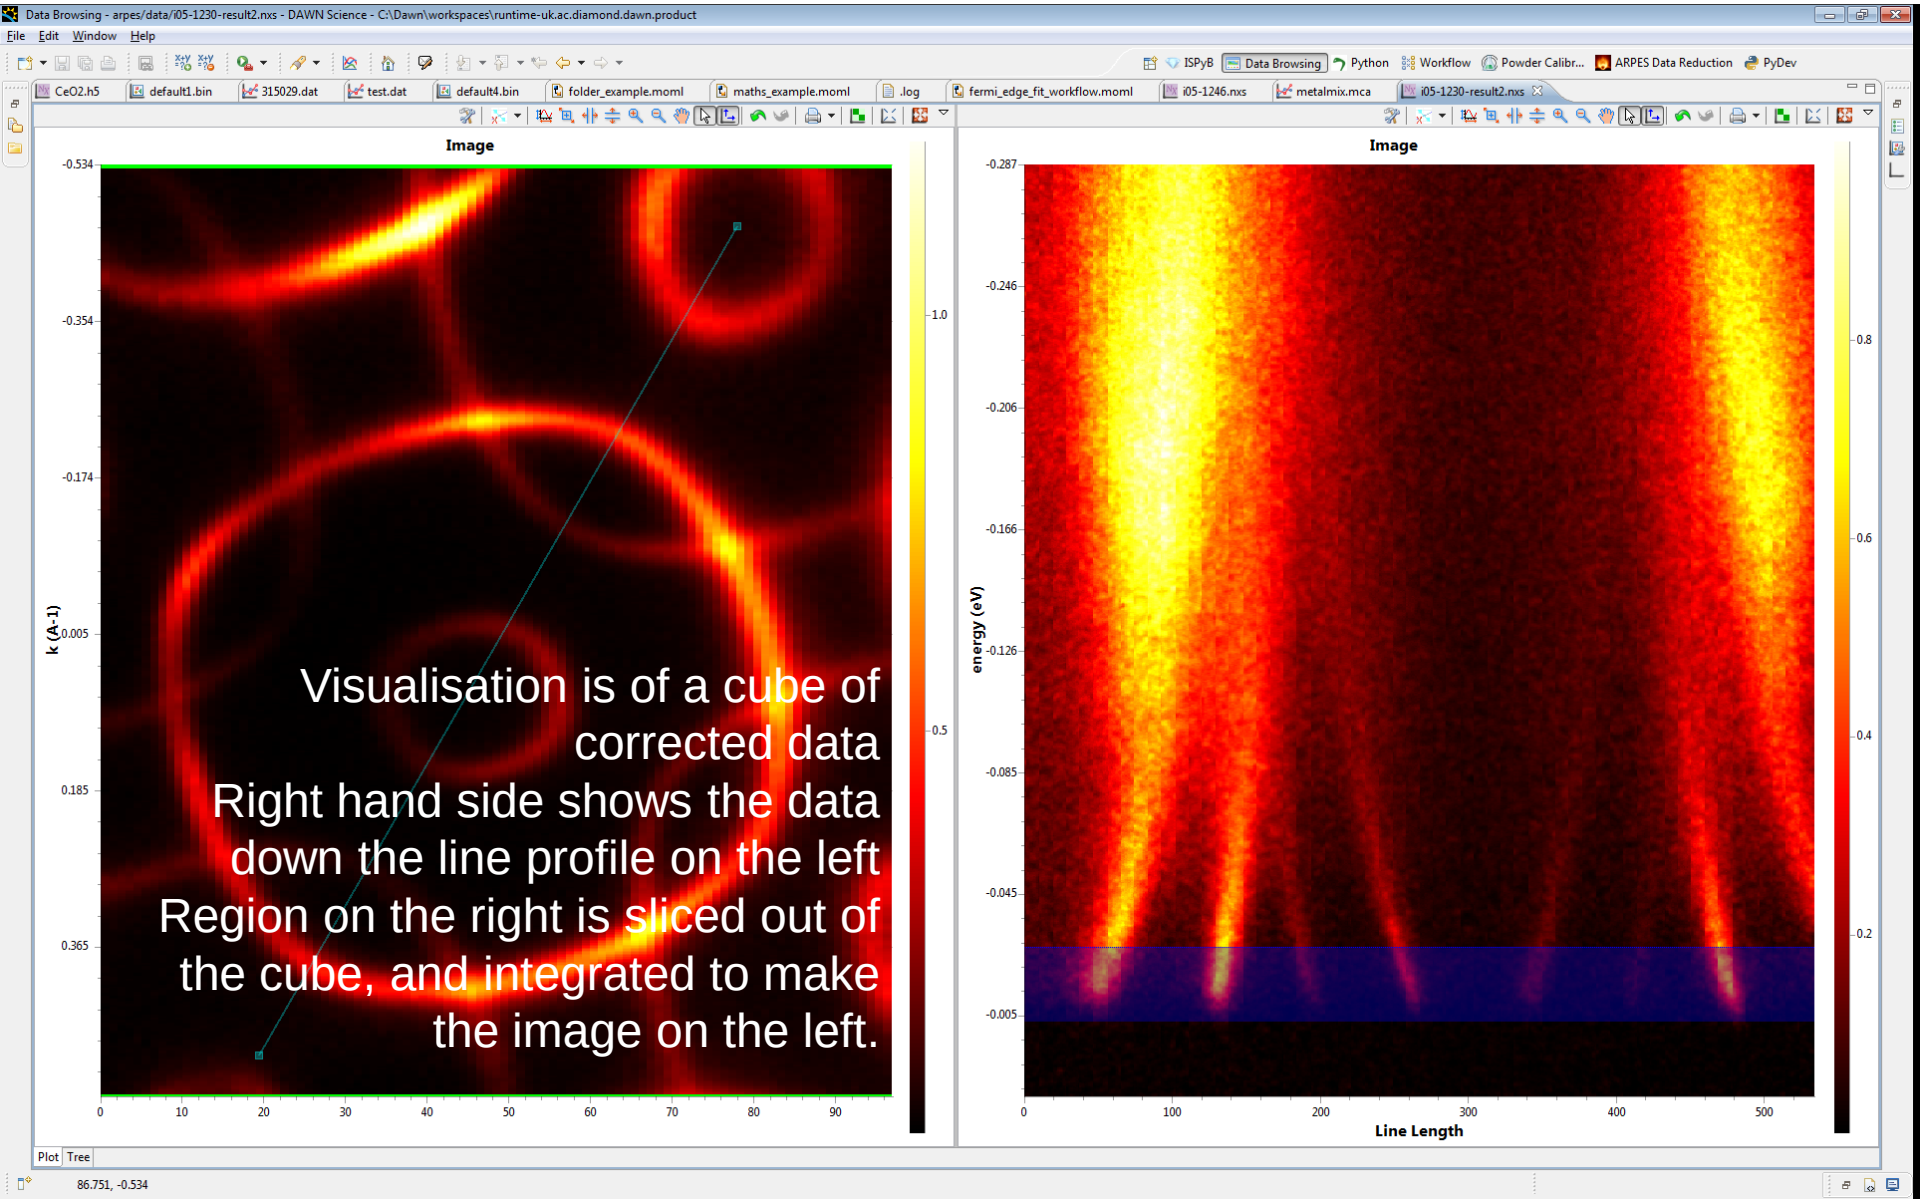
Data Browsing - arpes/data/i05-1230-result2.nxs - DAWN Science - C:\Dawn\workspaces\runtime-uk.ac.diamond.dawn.product
File Edit Window Help
ISPyB Data Browsing Python Workflow Powder Calibr... ARPES Data Reduction PyDev
CeO2.h5 default1.bin 315029.dat test.dat default4.bin folder_example.moml maths_example.moml .log fermi_edge_fit_workflow.moml i05-1246.nxs metalmix.mca i05-1230-result2.nxs
Image
-0.534
-0.354
-0.174
0.005
0.185
0.365
k (A-1)
0 10 20 30 40 50 60 70 80 90
1.0
0.5
Visualisation is of a cube of corrected data
Right hand side shows the data down the line profile on the left
Region on the right is sliced out of the cube, and integrated to make the image on the left.
Image
-0.287
-0.246
-0.206
-0.166
-0.126
-0.085
-0.045
-0.005
energy (eV)
0 100 200 300 400 500
Line Length
0.8
0.6
0.4
0.2
Plot Tree
86.751, -0.534

File Formats x n

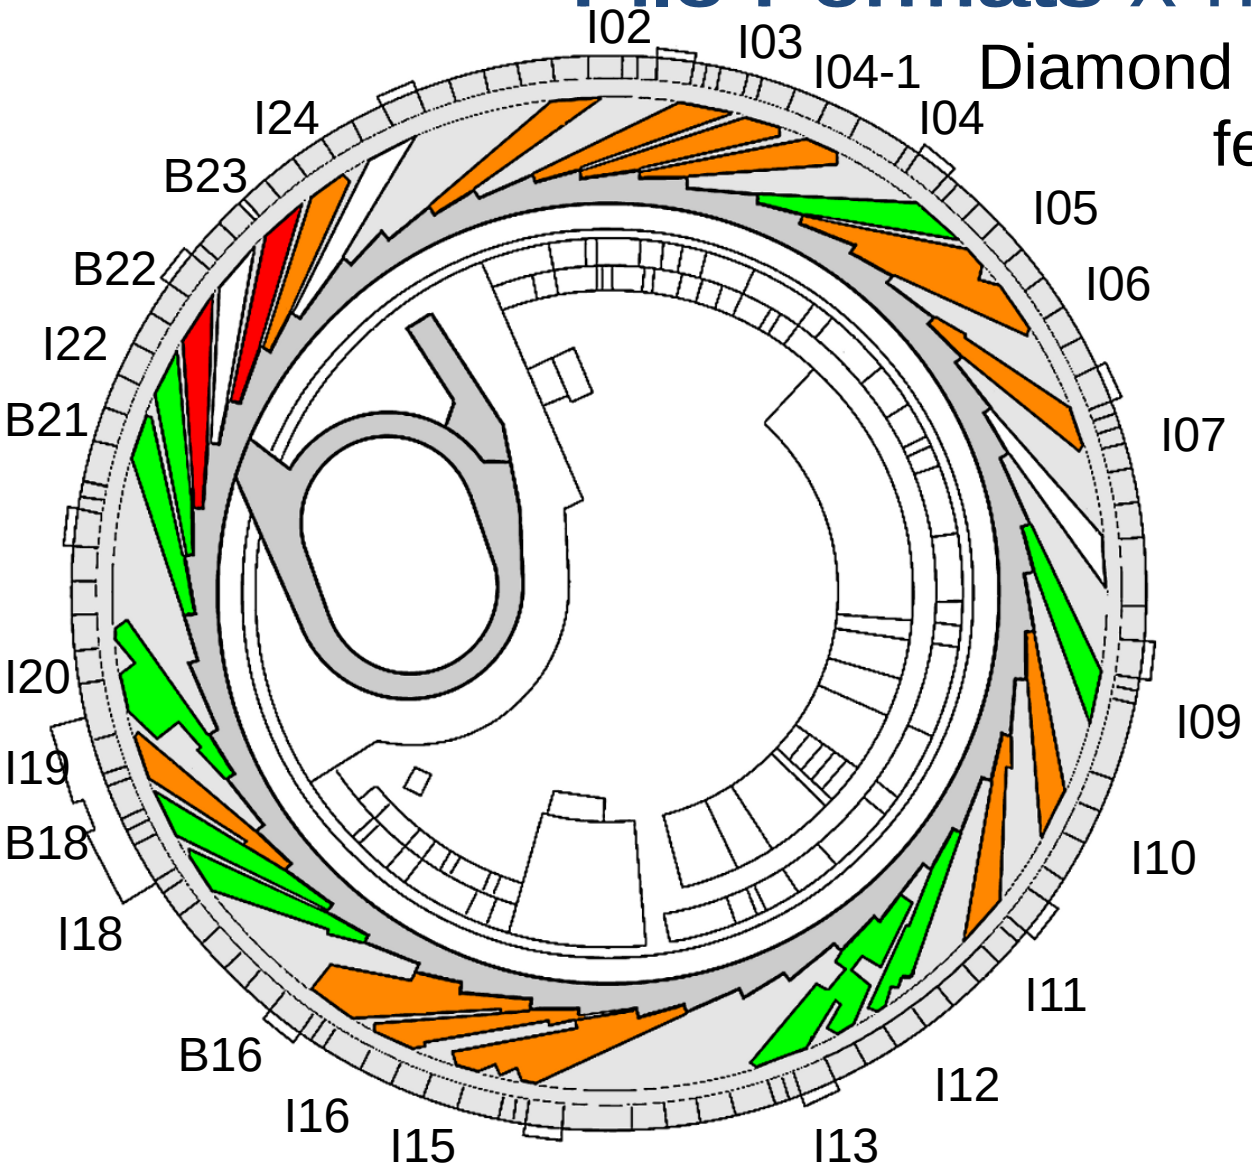

Diamond has a policy of, where feasible, to standardise on file formats, the choice being NeXus/HDF5

Green: predominantly using NeXus.

Orange: Mixed NeXus and other formats or considering NeXus in the next 12 months.

Files can be generated by Detector, EPICS or Data Acquisition

ISPyB information Flow: Display of Results

Data tracking and status checking of data acquired:

- Automatic plot generation of Scattering curves, Kratky, Guinier and P(r) plots

Macromolecule	Concentration	Scattering	Kratky	Guinier	P(r)	Frames (Averaged/Total)	Guinier				Gnom		Porod	
							Rg (nm)	Points	Quality (%)	I(0)	Rg (nm)	Dmax (nm)	Volume (nm³)	MM (kD) Vol. est.
bsa	5.80 mg/ml					bbsa (10 of 10) bsa (10 of 10) bbsa (10 of 10)	3.13 nm	44 - 87 (43)	78.37	68.25 ± 4.964017e-2	3.20 nm	10.94	118.64	59.3 - 79.1

- Display of 1D curves, averaged curves and subtracted curves

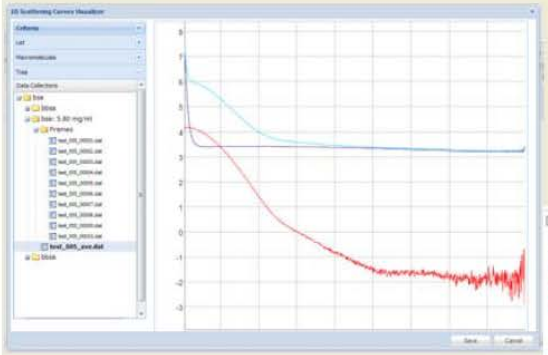

- Concentration Effects

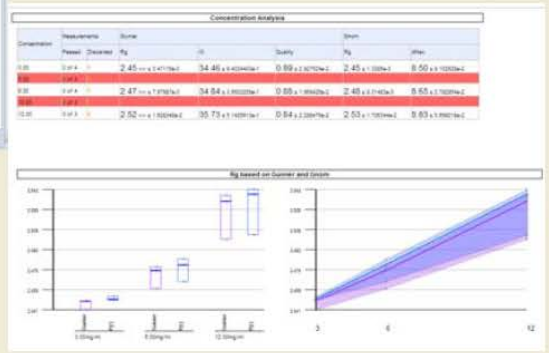

- 3D Models displayed using WebGL

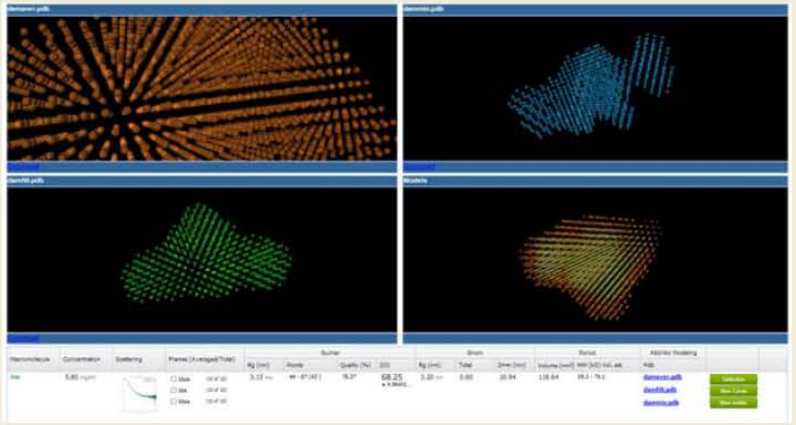

A prototype version of ISPyB for BioSAXS (ISPyBB) database is now available at BM29@ESRF and is being ported to P12@EMBL

Visits | Experiments (i02) | Cancel

lysozyme

LysoCT200_99bcryo 1.66 Å
120511_hc1_work/lysozymeCT200/
AM P1 40 x 77 x 78 Å

LysoCT200_99bcryo 1.66 Å
AM RSymm = 0.115 I/Sigma = 20.2

LysoCT200_80cryo 1.66 Å
120511_hc1_work/lysozymeCT200/
AM RSymm = 0.070 I/Sigma = 24.3

LysoCT200_80cryo 2.20 Å
120511_hc1_work/lysozymeCT200/
AM P4 76 x 76 x 37 Å

LysoCT200_80cryo 2.20 Å
120511_hc1_work/lysozymeCT200/
AM P4 77 x 77 x 37 Å

LysoCT200_80cryo 2.20 Å
120511_hc1_work/lysozymeCT200/
AM P4 78 x 78 x 38 Å

Visits | Status | Settings

Visits | Experiment details (nt5814-2, i02)

Omega end	300.0°
Rotation per image	0.30°
Exposure time	1.000 s
Beamsize X	50.0 µm
Beamsize Y	25.0 µm
Transmission	20.0%
Resolution	1.7 Å

Autoprocessing results

Autoprocessing successful

1 - fast_dp results
/dls_sw/apps/fast_dp_refactor/svn/src/fast_dp.py -j 0 -J 36 /dls/i02/data/2012/nt5814-2/120…

2 - xia2 results
xia2 -min_images 3 -3dii -atom s -blend -project nt5814v2 -crystal LysoCT20099bcryo1 -is…

Comments

(2247,9,164) EDNAStrategy1: subWedge:1Aperture: Large

Associated images

ISPyB - workflows/examples/ncd_data_reduction.moml - DAWN Science - C:\Users\awa25\workspace (Not Responding)

File Edit Diagram Navigate Search Project Run Window Help

Data Collections | Edge Scans | MCA Spectra

awa25 ▸ Beamlines ▸ I03 ▸ 2014 ▸ cm4950-1 ▸

nimages > 100 Execute Query (Syntax)

Name	Date	#Images	Protein Acronym	Crystal Type	Sample Name	Completeness [%]	I/Sigma
test_M6S1_3	2014-03-11 10:15:31	2400	-	-	-	98.9	51.9
test_M6S1_1	2014-03-11 10:01:56	2400	-	-	-	92.4	60.5
th_8_1	2014-02-12 08:54:08	540	-	-	-	99.4	13.4
collect_1	2014-02-06 08:54:05	3600	-	-	-	91.4	50.9

Wedge Pa | Header T | Experime

Experimental Details

test_M6S1_1

Images Collected	2,400
Wavelength	2 Å
Omega Start	0 °
Omega End	0 °
Rotation per Image	0.3 °
Exposure Time	0.04 s
Beamsize X	80 µm
Beamsize Y	20 µm
Transmission	0.499%

Crystal | Actor Attributes

1 2 3 4

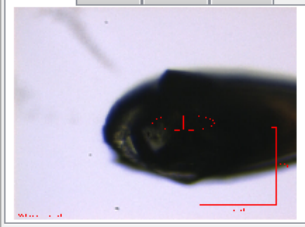

Autoprocessing Results

Name	Resolution (Å)	Type	Completeness (
test_M6S1_1			
xia2 (3daii)	55.21 - 2.01	overall	92.5
xia2 (3da)	30.94 - 2.01	overall	92.5
fast_dp	27.6 - 2.01	overall	92.4

Spot Summary

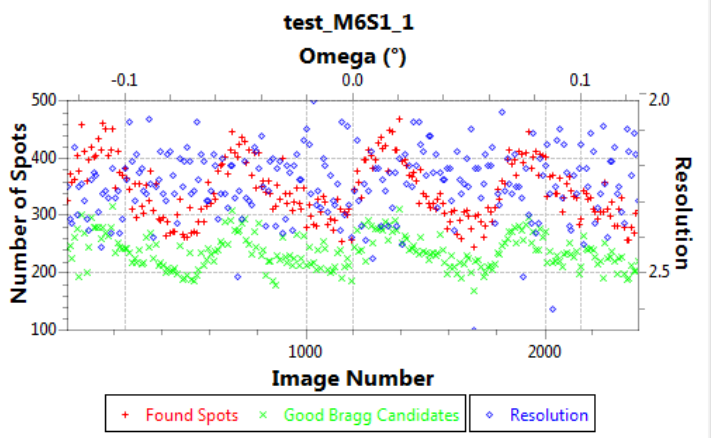

Image preview

12:28
28/03/2014

diamond

Archive and Reprocessing Service

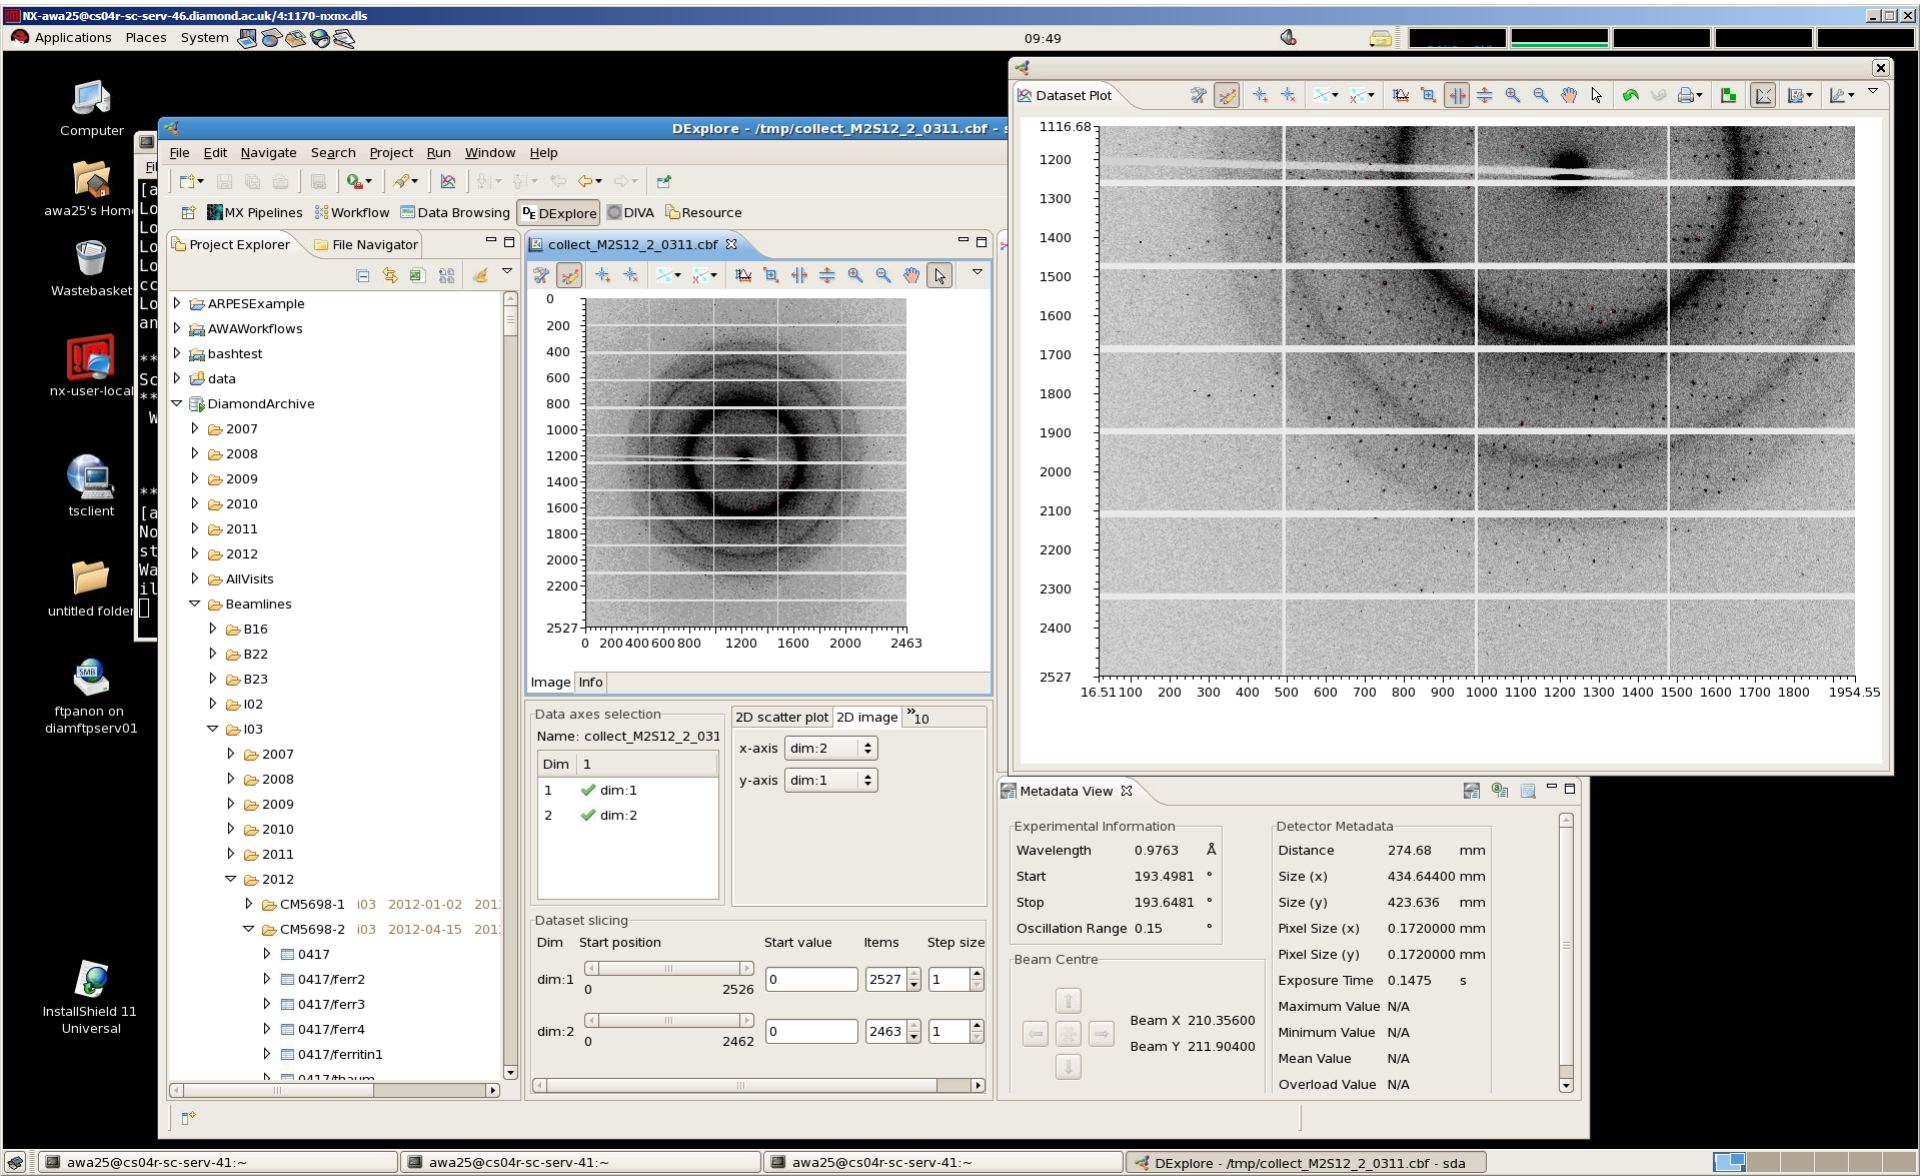

Collaboration experience

- Collaborating Inside the facility is often as challenging as with external groups.
- ‘customer’ involvement essential
- Management enthusiasm and support essential
- The scope of a collaborative project must fit in with the ambitions of your facility.
- Collaborating is expensive and ideally all parties must have adequate and commensurate resources though a centre of mass can produce results quicker, but at the cost of less well resourced partners.
- Although service providers are often happy to collaborate, our customers might not....

Current status/volumes in Diamond
ICAT = 285,198,074 files

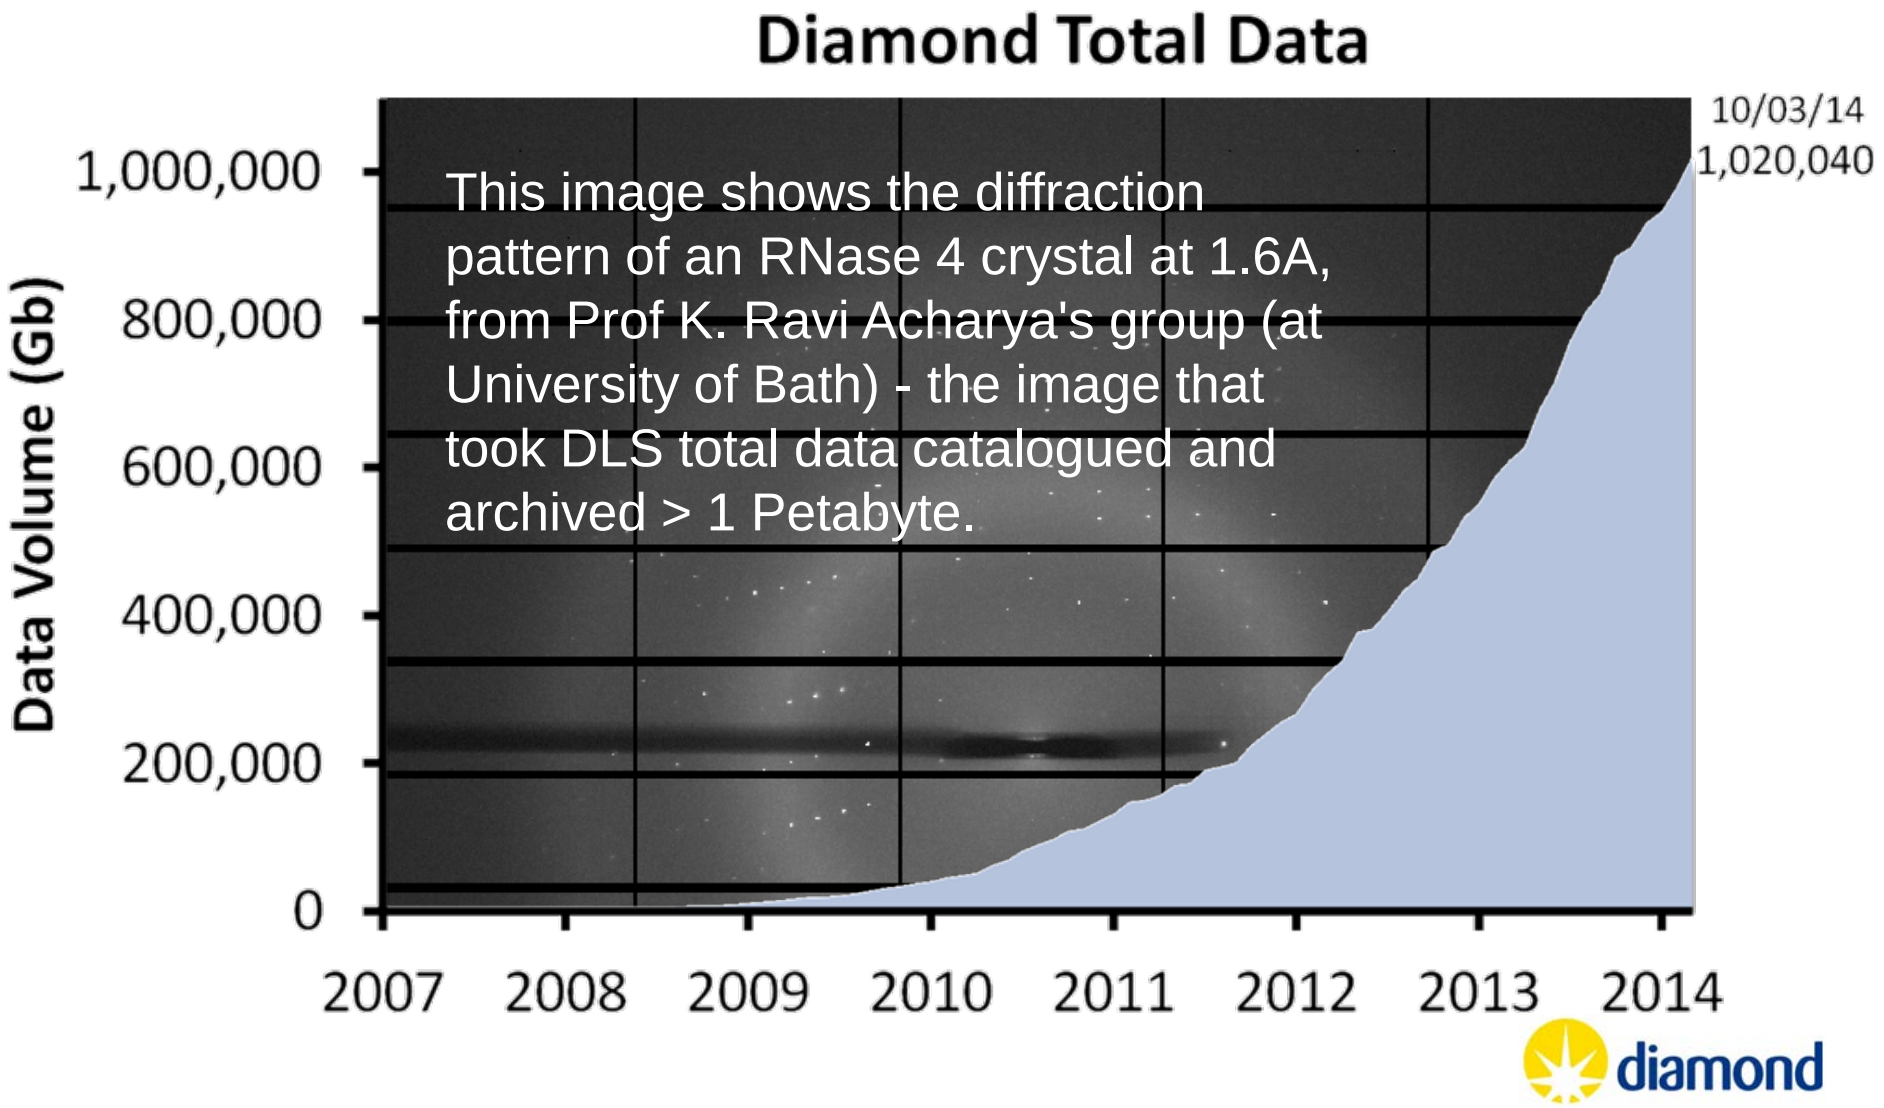

Novel developments in HDF5

An overview on new developments in the HDF5 library

Eugen Wintersberger
Dublin, 28.03.2014

DESY

HDF5 at a glance ...

HDF5 is a binary file format for large datasets

Data organized as trees by means of groups and datasets

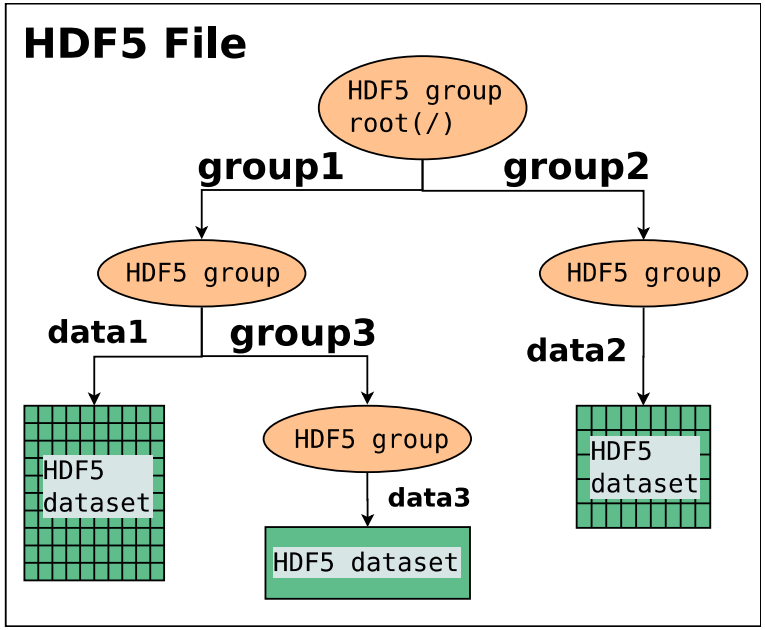

Groups and datasets can be annotated by attributes

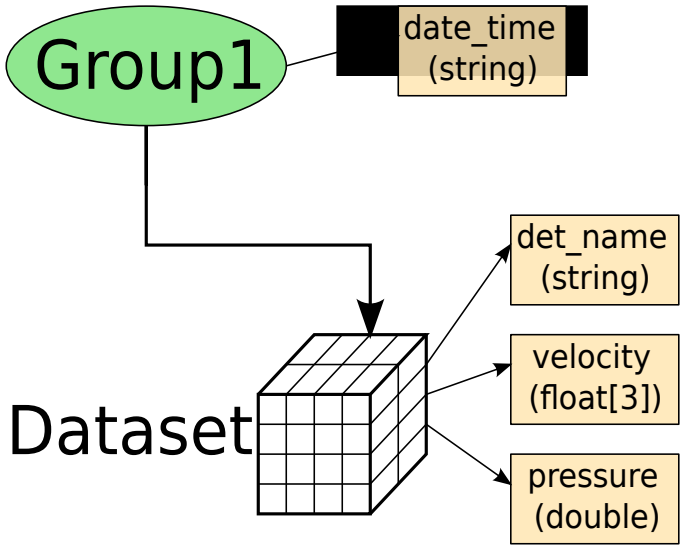

Developed by the HDF Group in Champaign, IL, USA

DESY

HDF5 key features

- Transparent compression of individual datasets
- platform independent (Windows, Linux, AIX, Solaris, OSX, OpenVMS)
- library is implemented in C (easy to interface)
- bindings to many languages: C++, Fortran, Python, Java, Perl, R, Go
- Supported by commercial products like: Matlab, Mathematica, IDL

Easy parsing: data access via library → it is not about how data is stored but rather where (path).

Recently developed features already available in 1.8.12

How compression works in HDF5

Writing compressed data

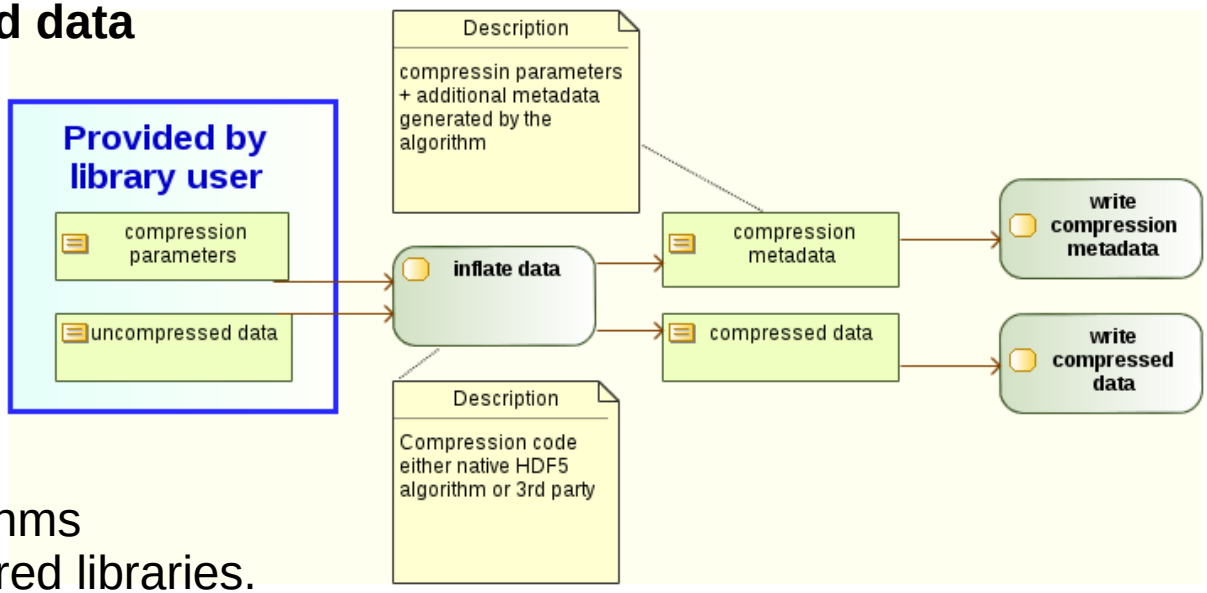

Compression algorithms implemented as shared libraries.

Reading compressed data

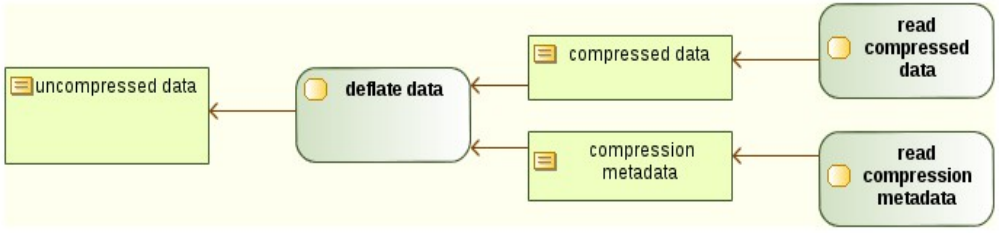

DESY

Writing precompressed data ...

Sometimes compression can be done better in hardware (FPGAs)
→ reduced bandwidth requirements for network and disk I/O!

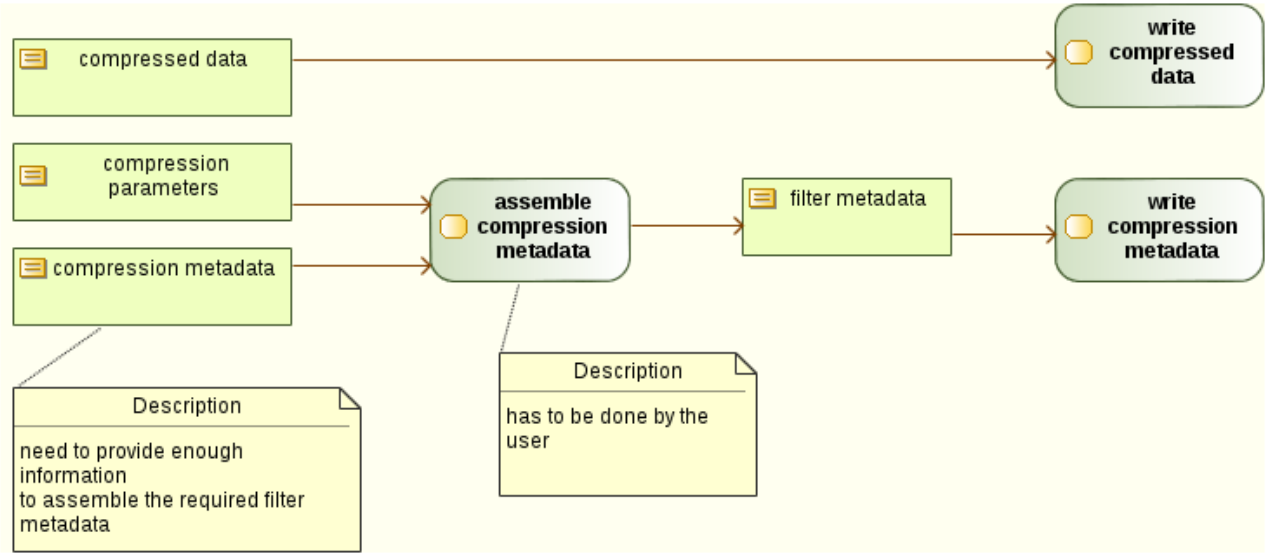

Feature funded by DECTRIS

This features allows bypassing the filter chain and write compressed data directly.

Reading the data is the same as for conventionally written data!

External filter API

Detector vendors or facilities may want to use custom compression algorithms.

Possible with all HDF5 1.8.X versions but: reading software needs to be recompiled => commercial applications do not have access to the data.

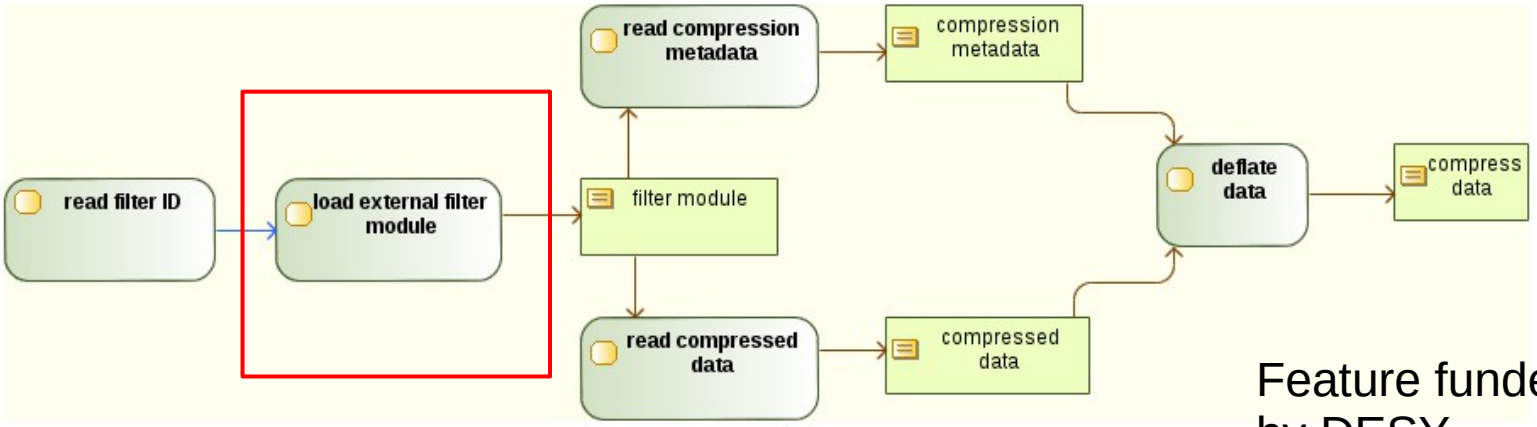

Feature funded by DESY

External filters can be loaded at runtime on demand.

Requirement:
Detector vendors and facilities provide libraries for all target platforms.

New features addressing concurrent data access

Single Writer – Multiple Reader = SWMR

Current situation: cannot read data from a file while it is written by two different processes!

SWMR would make this possible!

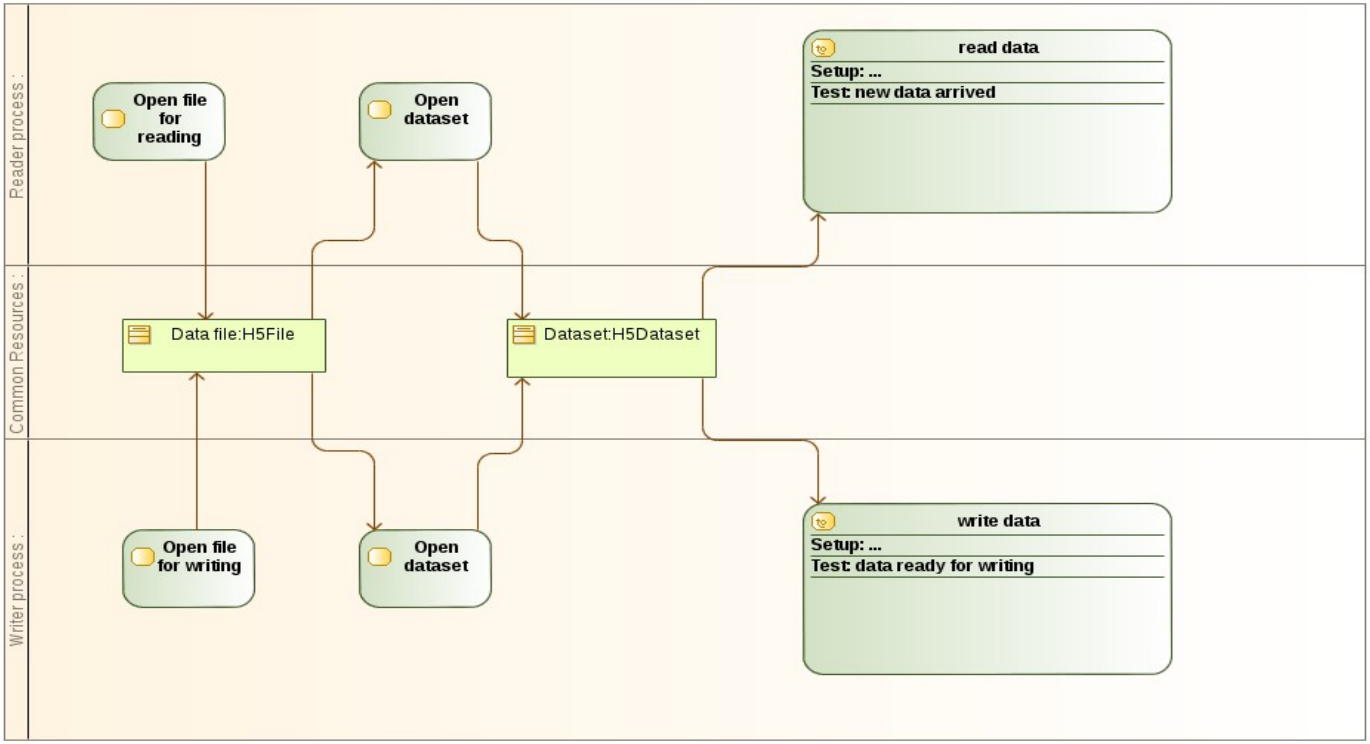

Status: A prototype is available – missing funding for library integration.

DESY

Multithreading support

Option 1: Make HDF5 fully threadsafe!

- would give best performance
- major rewrite of 300K lines of C-Code
- costs of 4-6 FTEs
- significant future maintenance efforts

Very unlikely to happen without extensive funding!

Option 2: Use threads within the library!

- use multiple threads for compression
- asynchronous (non-blocking) I/O
- costs ~1.5 FTEs
- lower maintenance effort
- could improve continuously

Remote access to data stored in an HDF5 file

Open Source Project for a Network Data Access Protocol

Advantages

- ➔ Already existing and actively developed protocol
- ➔ Data access via HTTP requests where URLs encode the information of what shall be retrieved from a data source
- ➔ Independent of the dataformat (other formats can easily be supported)
- ➔ Data can be accessed with any software that can dereference URLS (webbrowser, Excel, Libreoffice,...)
- ➔ Matlab 2012a provides native support for OPeNDAP
- ➔ C++, Java, and Python libraries available

www.opendap.org

HDF5 data server

Serialization of HDF5 library calls rather than using HTTP and URLs

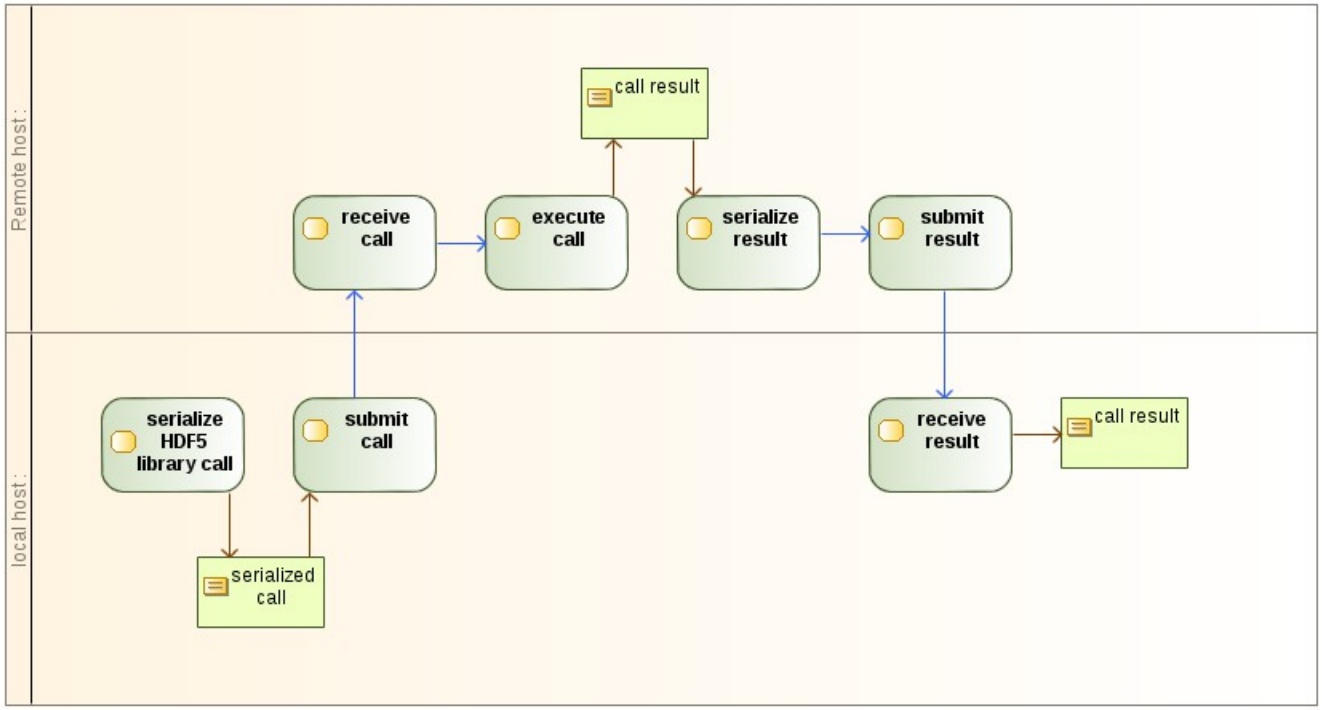

Status: only as a design draft – no implementation will be done without funding.

Licenses and funding issues

Currently the HDF5 group is entirely project funded!

… they are looking for new funding (licensing) models

➔ the core library will always remain open source

➔ dual licensing would be an alternative

➔ some features may become commercial products in future

As a user community we should take funding issues of the HDF5 group serious and think about how we can improve their financial backup!

Conclusions: what the RDA could do

- Establish an HDF5 special interest group as a hub for HDF5 related work
- collect feature requests from various user communities
- organize funding activities if a new feature should be implemented
- manage filters (compression algorithms) which should be available for HDF5 and provide hosting resources
- Provided solutions based on HDF5 for certain selected use-cases
- Organize workshops around HDF5 (maybe as satellite events around meetings)

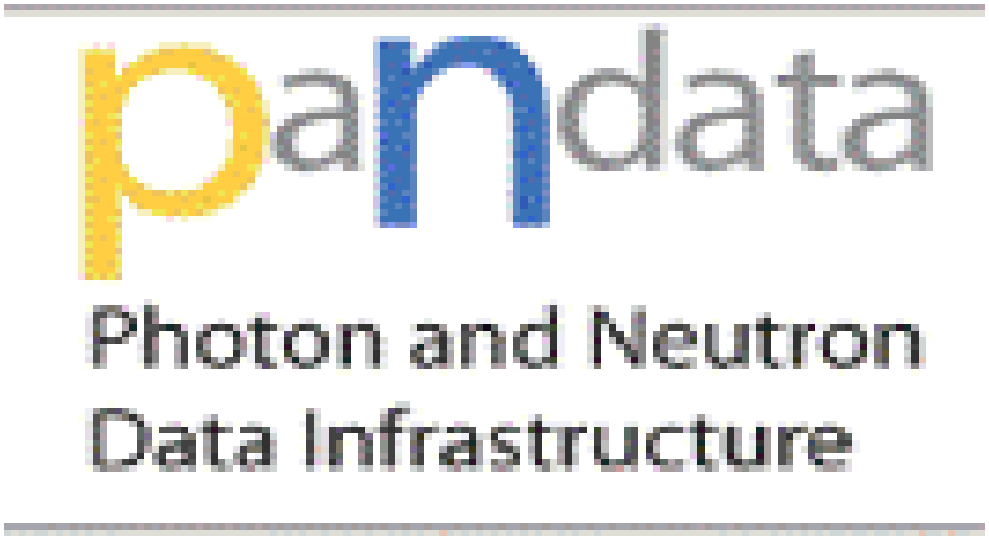

Brian Matthews

STFC

PaNdata

- Photon and Neutron Data Infrastructure
- Established in 2007 with 4 facilities
  - now standing at 13
  - With “friends” around the world
- Combined Number of Unique Users more than 35000 in 2011
- Combines Scientific and IT staff from the collaborating facilities
- European Framework 7 Projects
  - PaNdata-Europe: SA, 2009-11
  - PaNdata-Open Data Infrastructure, IP, 2011-14

–Guestimates
- Investment > €4.000.000.000*
- Running costs > €500.000.000/yr*
- Publications > 10.000/yr*
- RCosts/Publication ~ €50.000*%
- Data volume >> 10PB/yr*

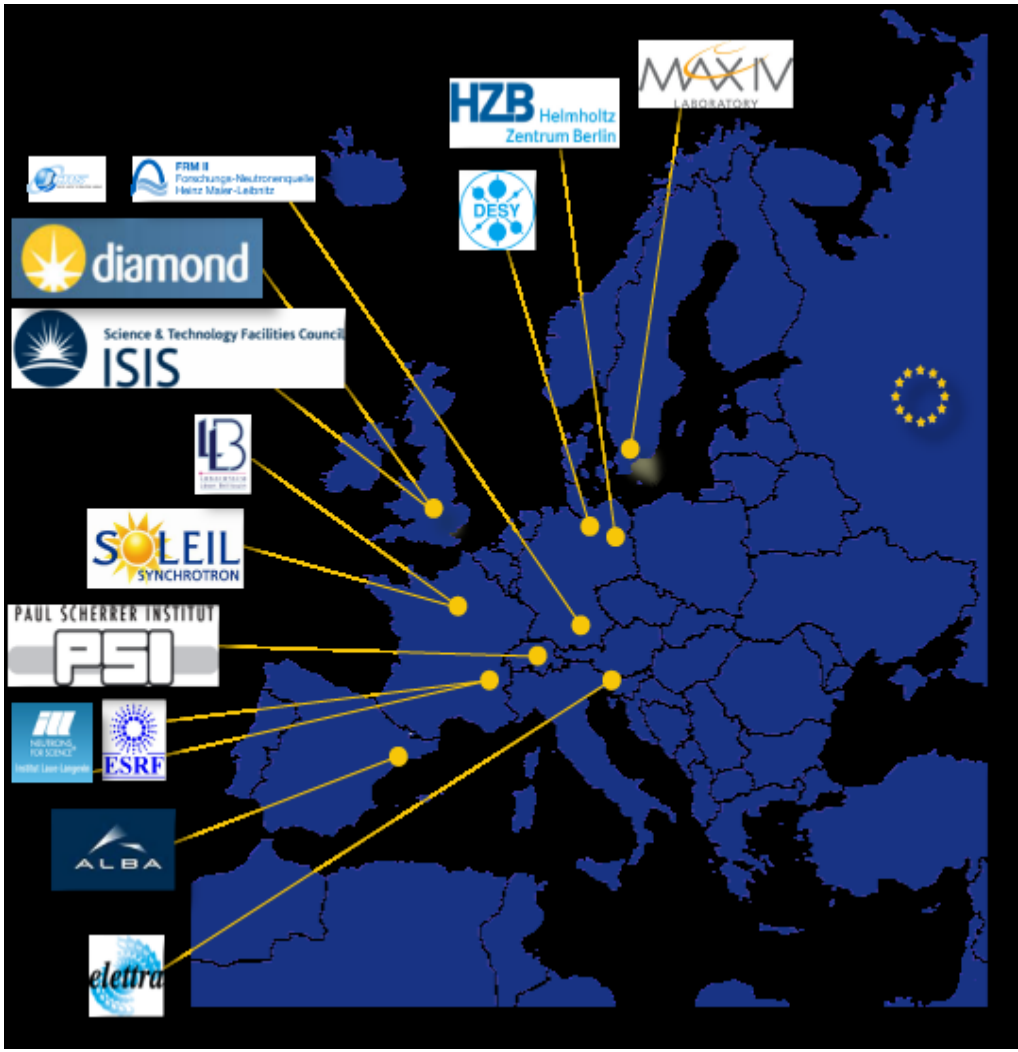

Counting Users

Number of Users shared between facilities																
	ALBA	BER II	DESY	DLS	ELETTRA	ESRF	FRM-II	ILL	ISIS	LLB	SINQ	SLS	SOLEIL	neutron	photon	all
ALBA	773	7	61	58	51	281	2	51	13	5	10	77	105	69	400	773
BER II	7	1563	115	46	27	179	157	383	198	98	191	62	36	580	329	1563
DESY	61	115	4197	137	222	851	116	255	113	62	95	315	188	469	1294	4197
DLS	58	46	137	4407	102	810	30	267	399	33	52	229	192	546	1130	4407
ELETTRA	51	27	222	102	3167	433	11	77	35	20	18	179	367	141	900	3167
ESRF	281	179	851	810	433	10287	139	900	369	190	174	963	1286	1313	3586	10287
FRM-II	2	157	116	30	11	139	1095	347	137	89	161	33	29	509	259	1095
ILL	51	383	255	267	77	900	347	4649	731	301	395	156	222	1518	1347	4649
ISIS	13	198	113	399	35	369	137	731	2880	89	233	94	56	936	745	2880
LLB	5	98	62	33	20	190	89	301	89	1235	74	39	151	391	323	1235
SINQ	10	191	95	52	18	174	161	395	233	74	1219	224	31	590	415	1219
SLS	77	62	315	229	179	963	33	156	94	39	224	3827	399	371	1470	3827
SOLEIL	105	36	188	192	367	1286	29	222	56	151	31	399	4568	394	1817	4568
neutron	69	1563	469	546	141	1313	1095	4649	2880	1235	1219	371	394	10023	2334	10023
photon	773	329	4197	4407	3167	10287	259	1347	745	323	415	3827	4568	2334	25336	25336
all	773	1563	4197	4407	3167	10287	1095	4649	2880	1235	1219	3827	4568	10023	25336	33025

http://pan-data.eu/Users2012-Results

PaN-Data Integration

Common data environment, common user experience

Shared Data Policy Framework

Federated User Authentication — umbrella

Federated Data Catalogue — ICAT

Common Data Format — *NeXus*





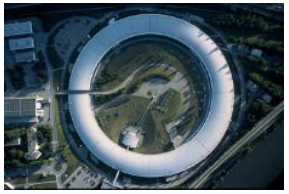

Towards the Future

- Provenance
  - Integrating context, analysis and publication into the record
- Preservatiom
  - Long-term need for archiving and curating data
  - Persistence Identifiers, itegtrity, context,
  - Costs and Benefits of data preservation
- Scalability
  - Managing high data rates and volumes
  - Parallel file stores

Infrastructure for managing data flows

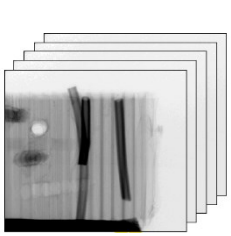

Scan

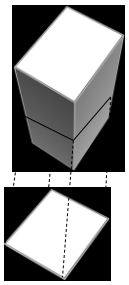

Reconstruct

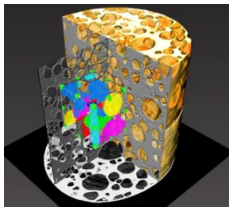

Segment
+ Quantify

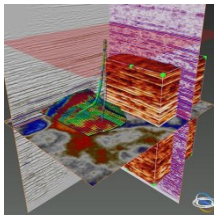

3D mesh +
Image based Modelling

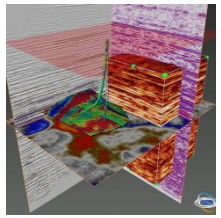

Predict + Compare

ICAT

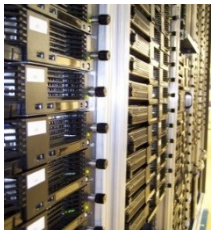

Data
Catalogue

Petabyte
Data storage

Parallel
File system

HPC
CPU+GPU

Visualisation

Infrastructure + Software + Expertise!

- **Tomography**: Dealing with high data volumes – 200Gb/scan, ~5 TB/day (one experiment)
- **MX**: high data volumes, smaller files, but a lot more experiments
- Hard to move the data – needs to be handled at the facility?

Some mage credit: Avizo, Visualization Sciences Group (VSG)