BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

Real-Time Data Pipeline and analysis using SPOT and HIPGISAXS

Alexander Hexemer and Craig E. Tull
Lawrence Berkeley National Lab

OLCF Users' Meeting
June 24, 2015 @ OLCF

Data Challenge

up to 2 Petabytes /year of raw data.

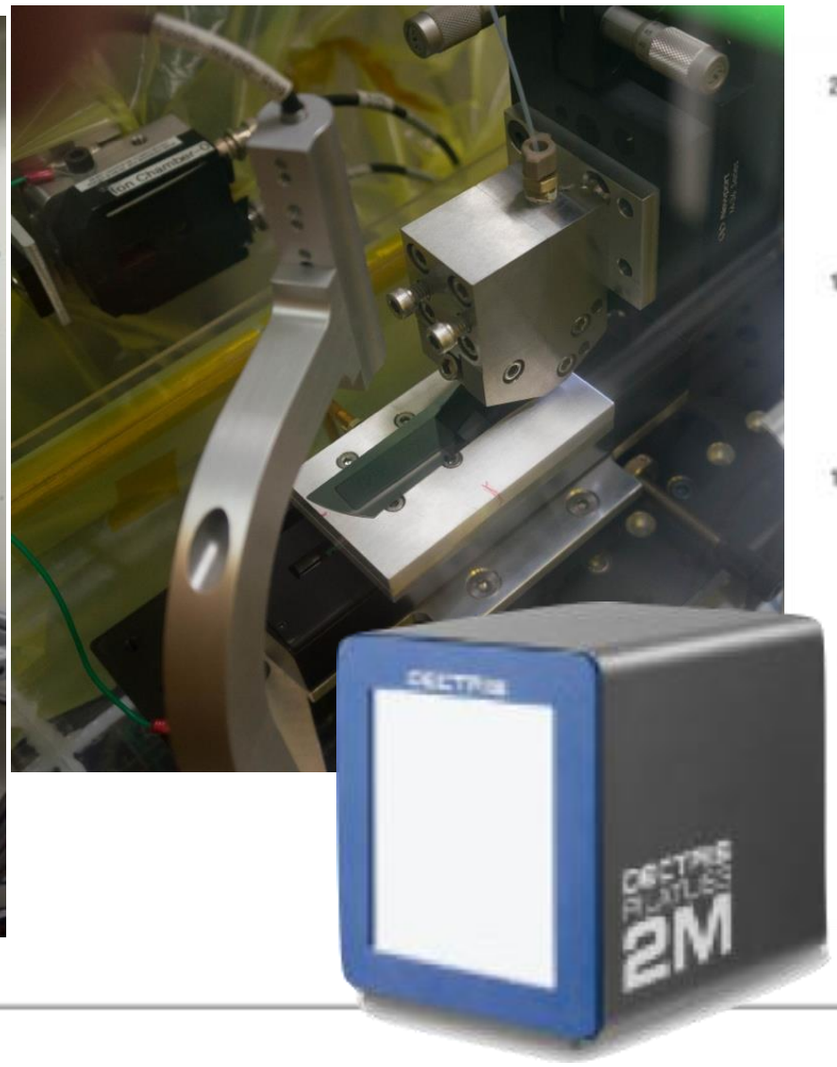

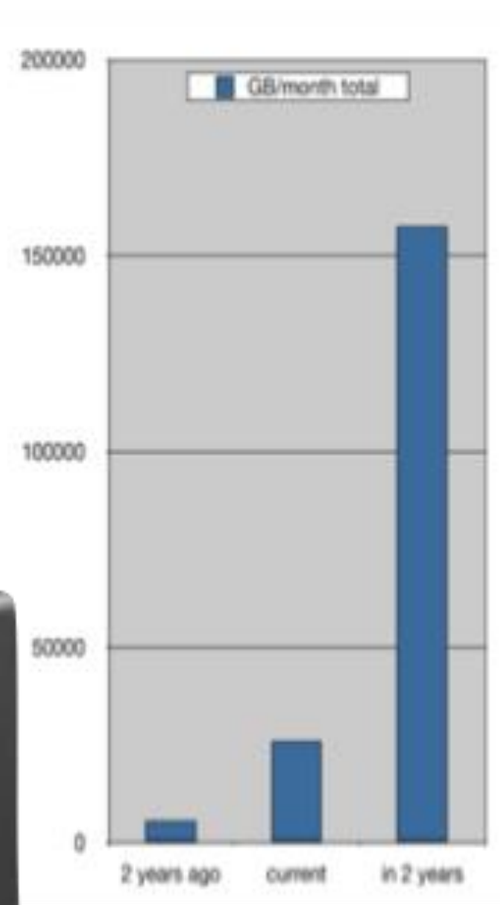

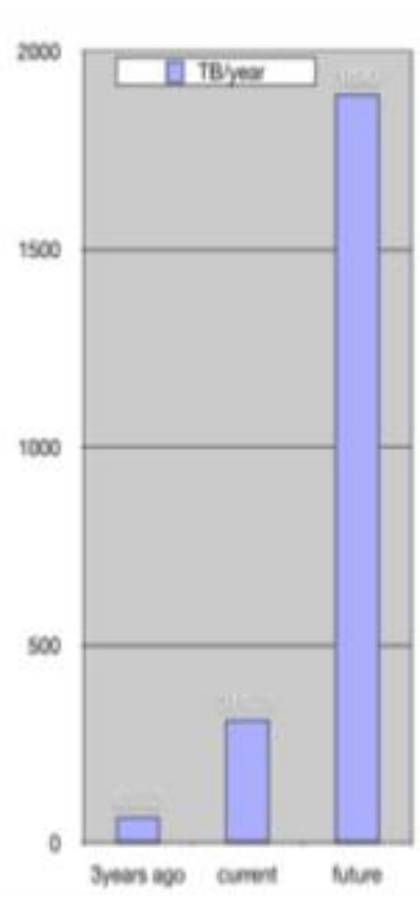

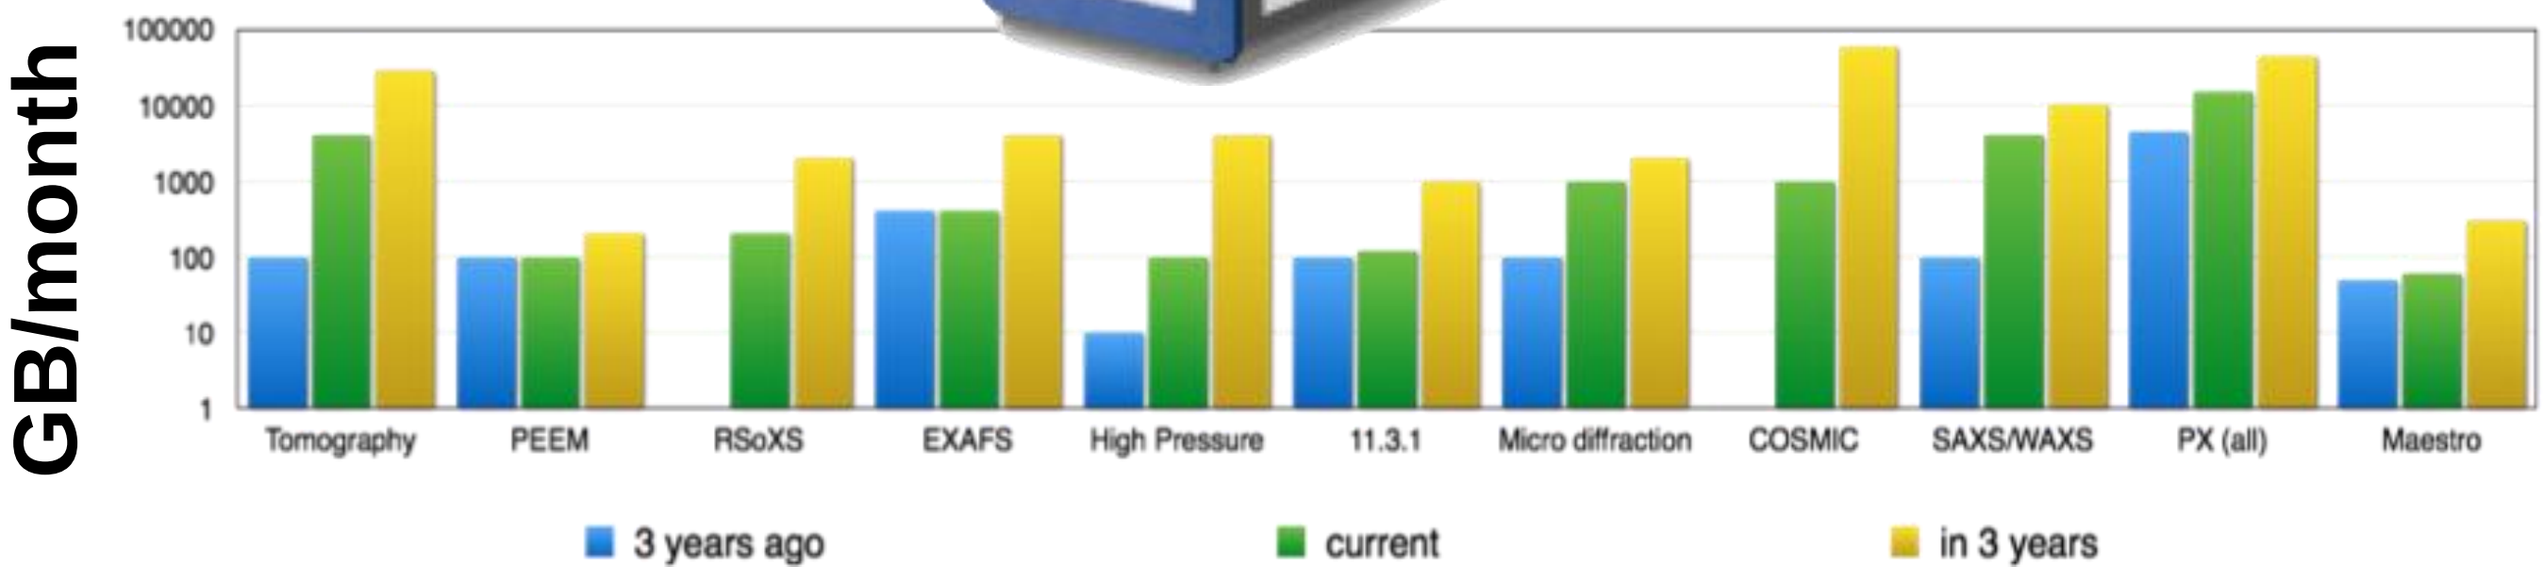

Hard X-ray Scattering

Recover morphology from scattering

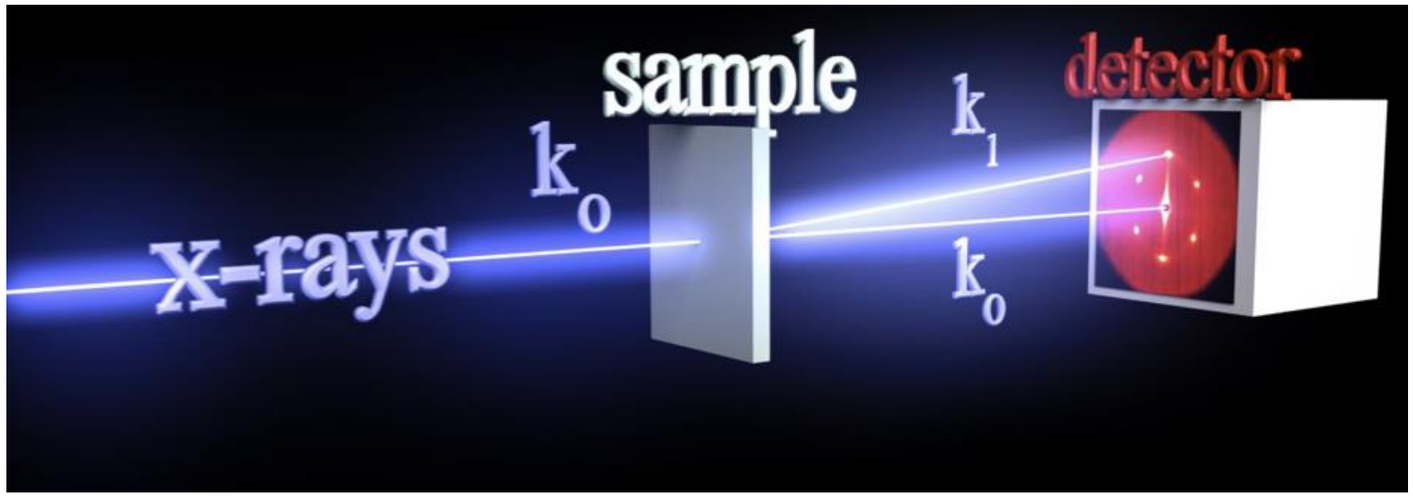

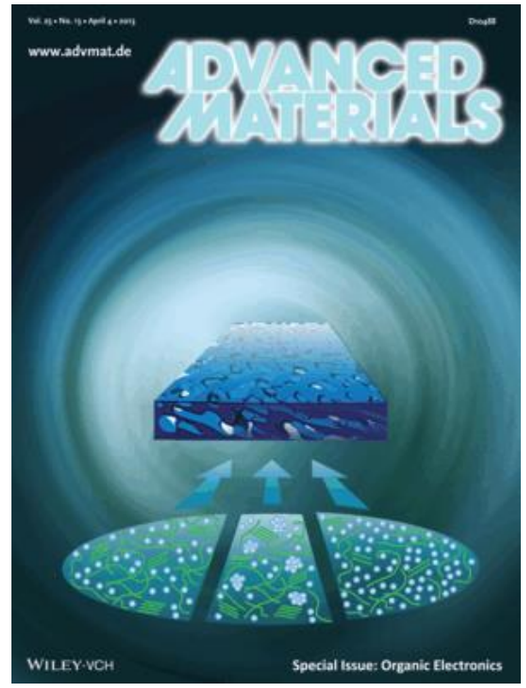

Contrast from electron density difference

Hard X-ray energy: 10keV

BL 7.3.3

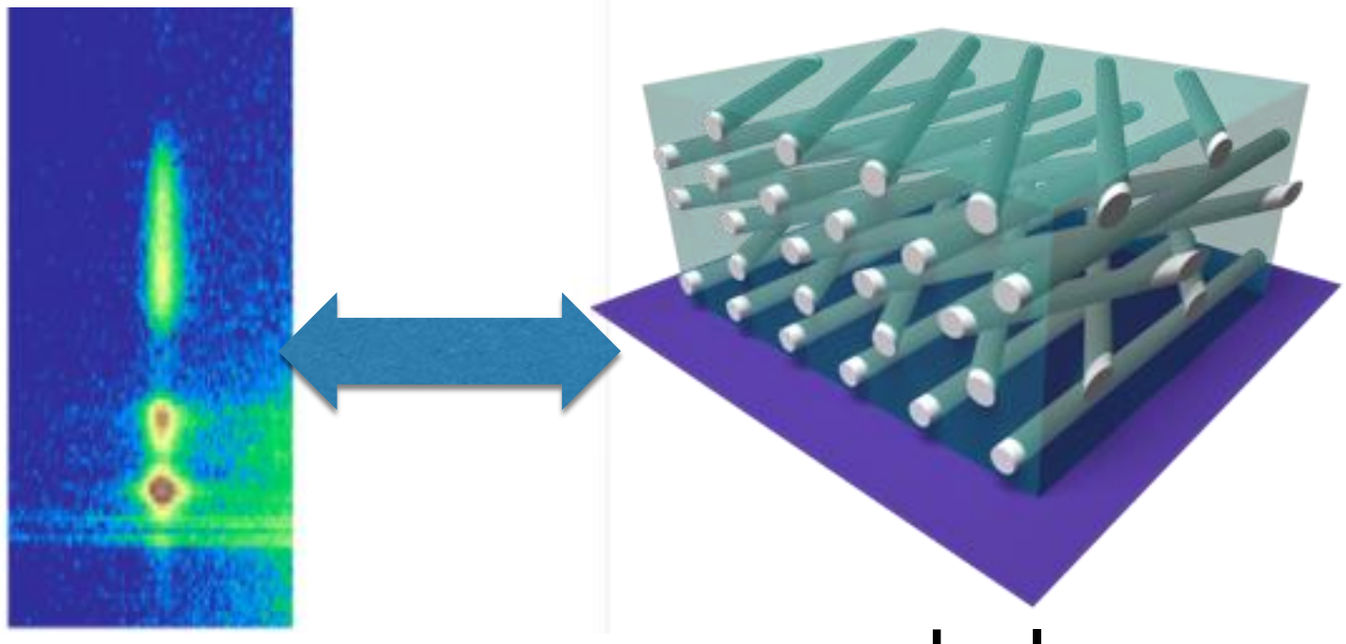

data

morphology

X-ray Scattering

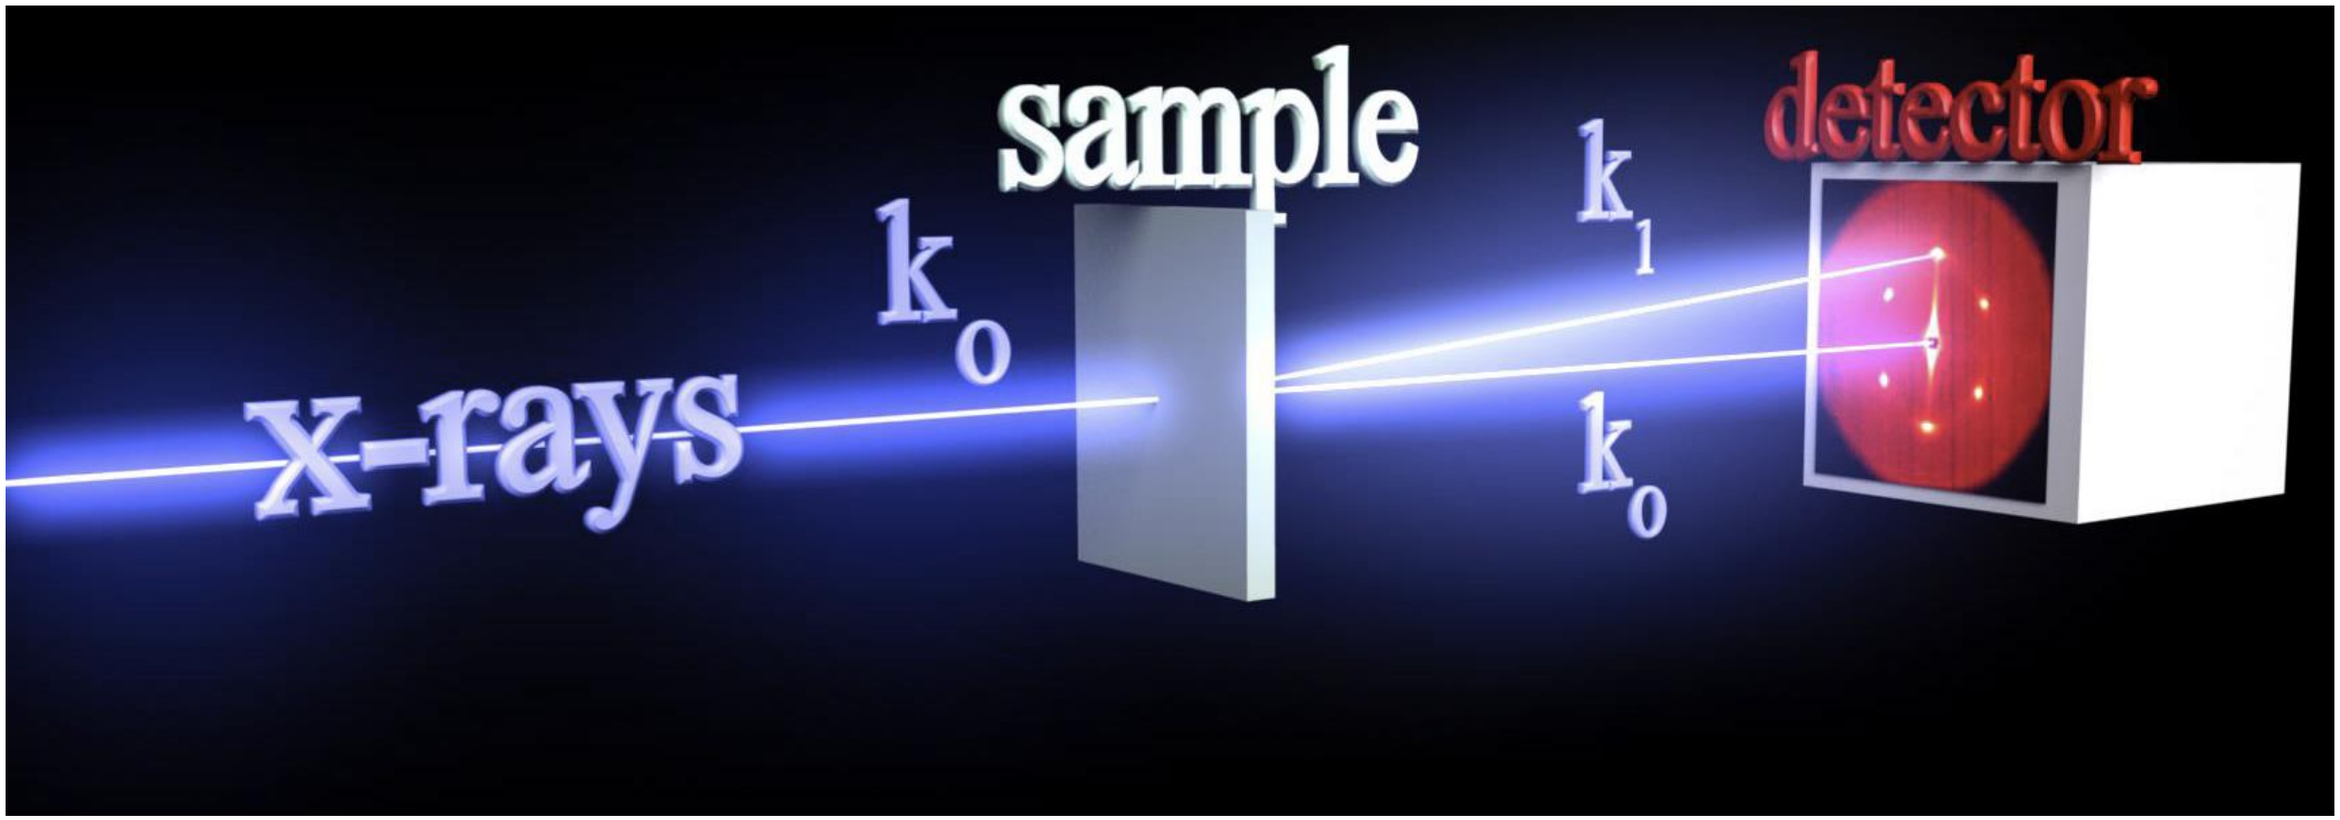

$$I(q) = \left| \int \rho(r) e^{-i\, q \cdot r} \right|^2$$

Measure structure size from 1A to 100's on nm

Fast Detectors + in situ experiments

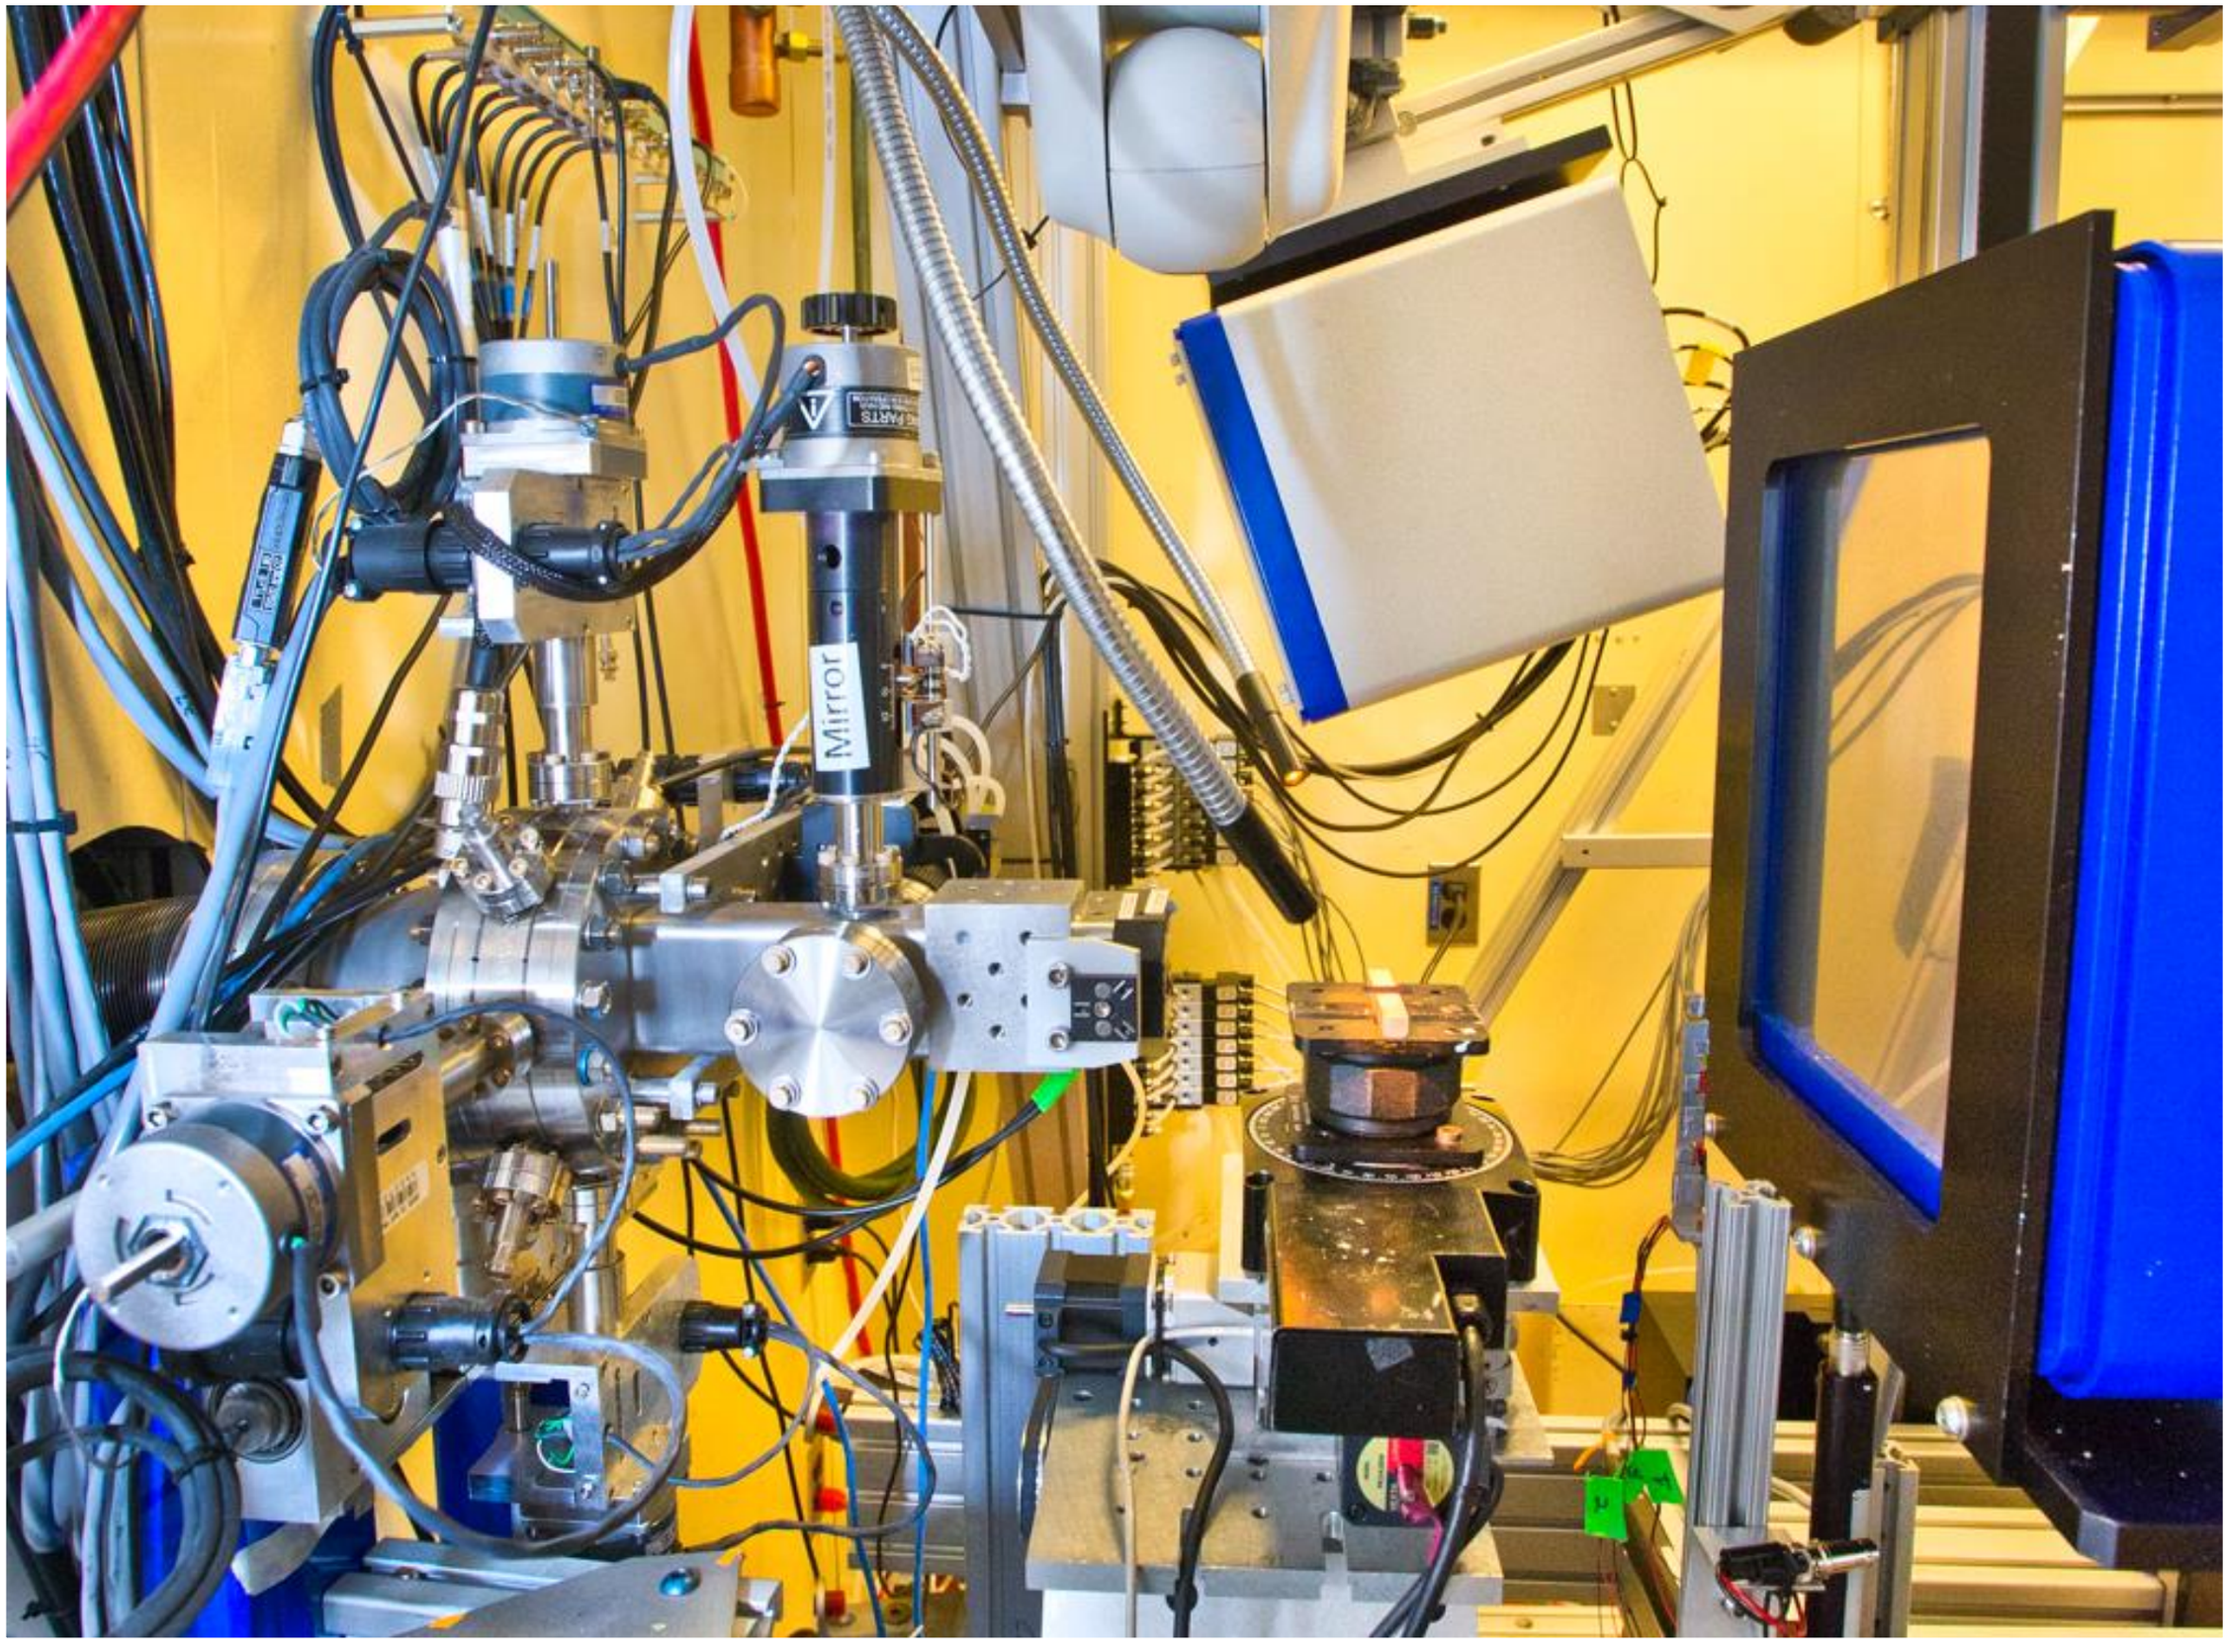

Data Collection: Robot

Epson G6 SCARA robot

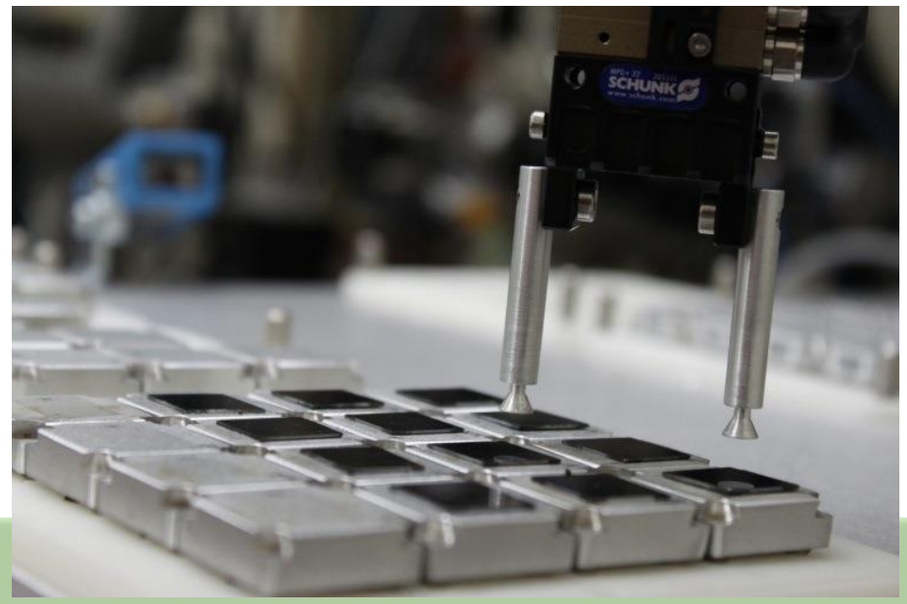

Each puck is a kinematic mount for repeatable, secure transfers

Barcode reader

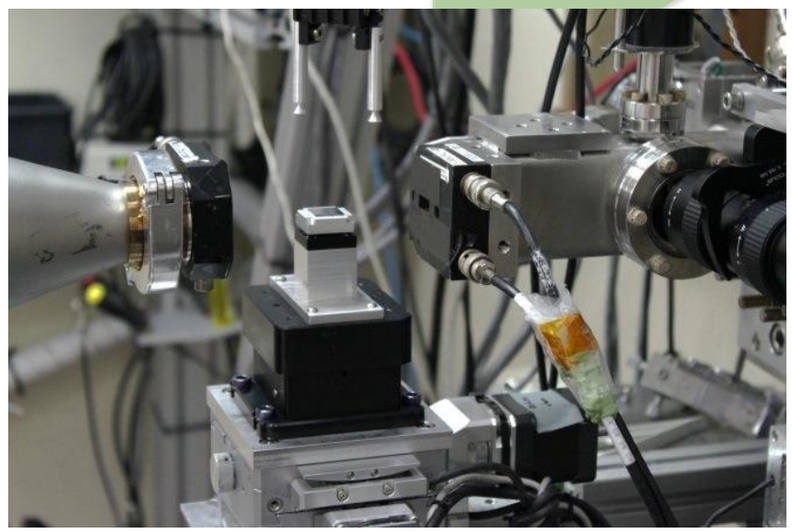

Measurement stage

Data Flow

Transfer to NERSC

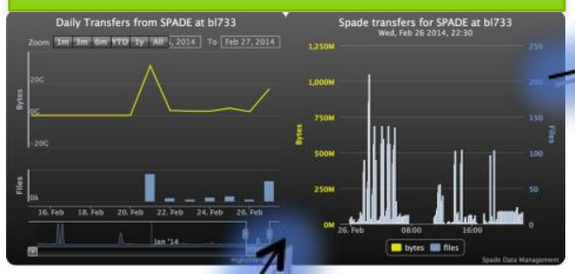

On-the-fly calibration, processing

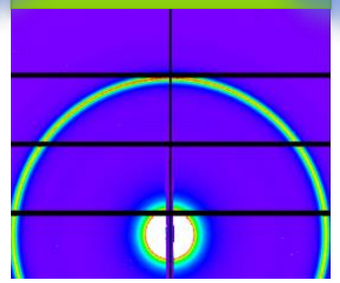

Real-time access via web portal

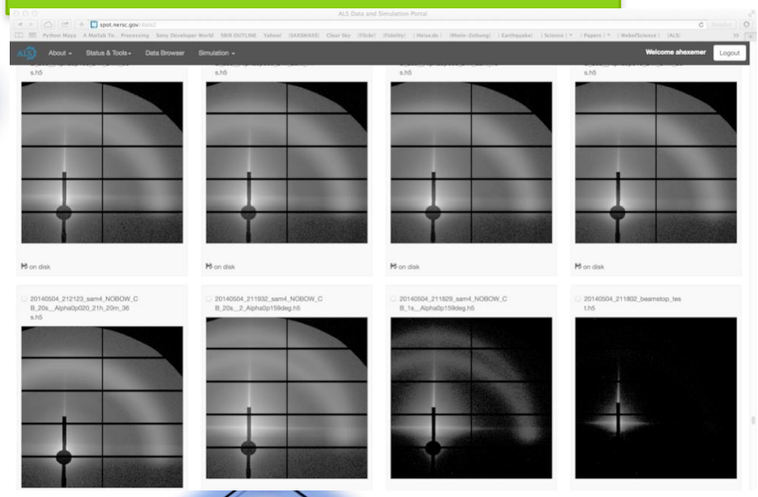

Combining:
GIXSGUI, dpdak + …

Data collection

Analysis and modeling on NERSC supercomputers:

HipGISAXS simulation

HipRMC fitting

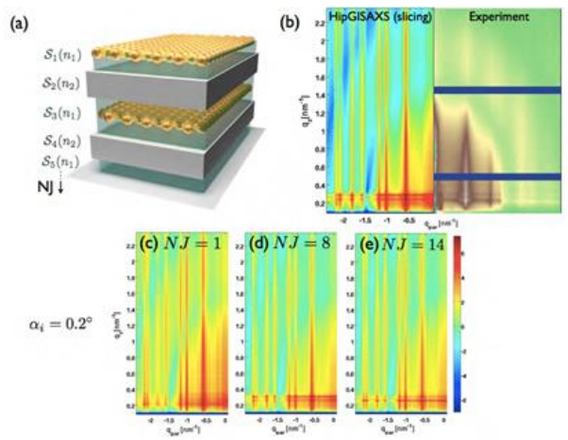

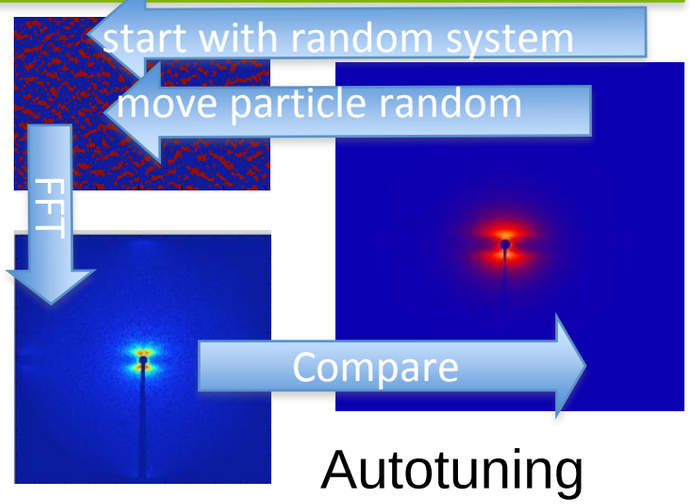

Autotuning

HiPIES: High Performance Interactive Environment for Scattering

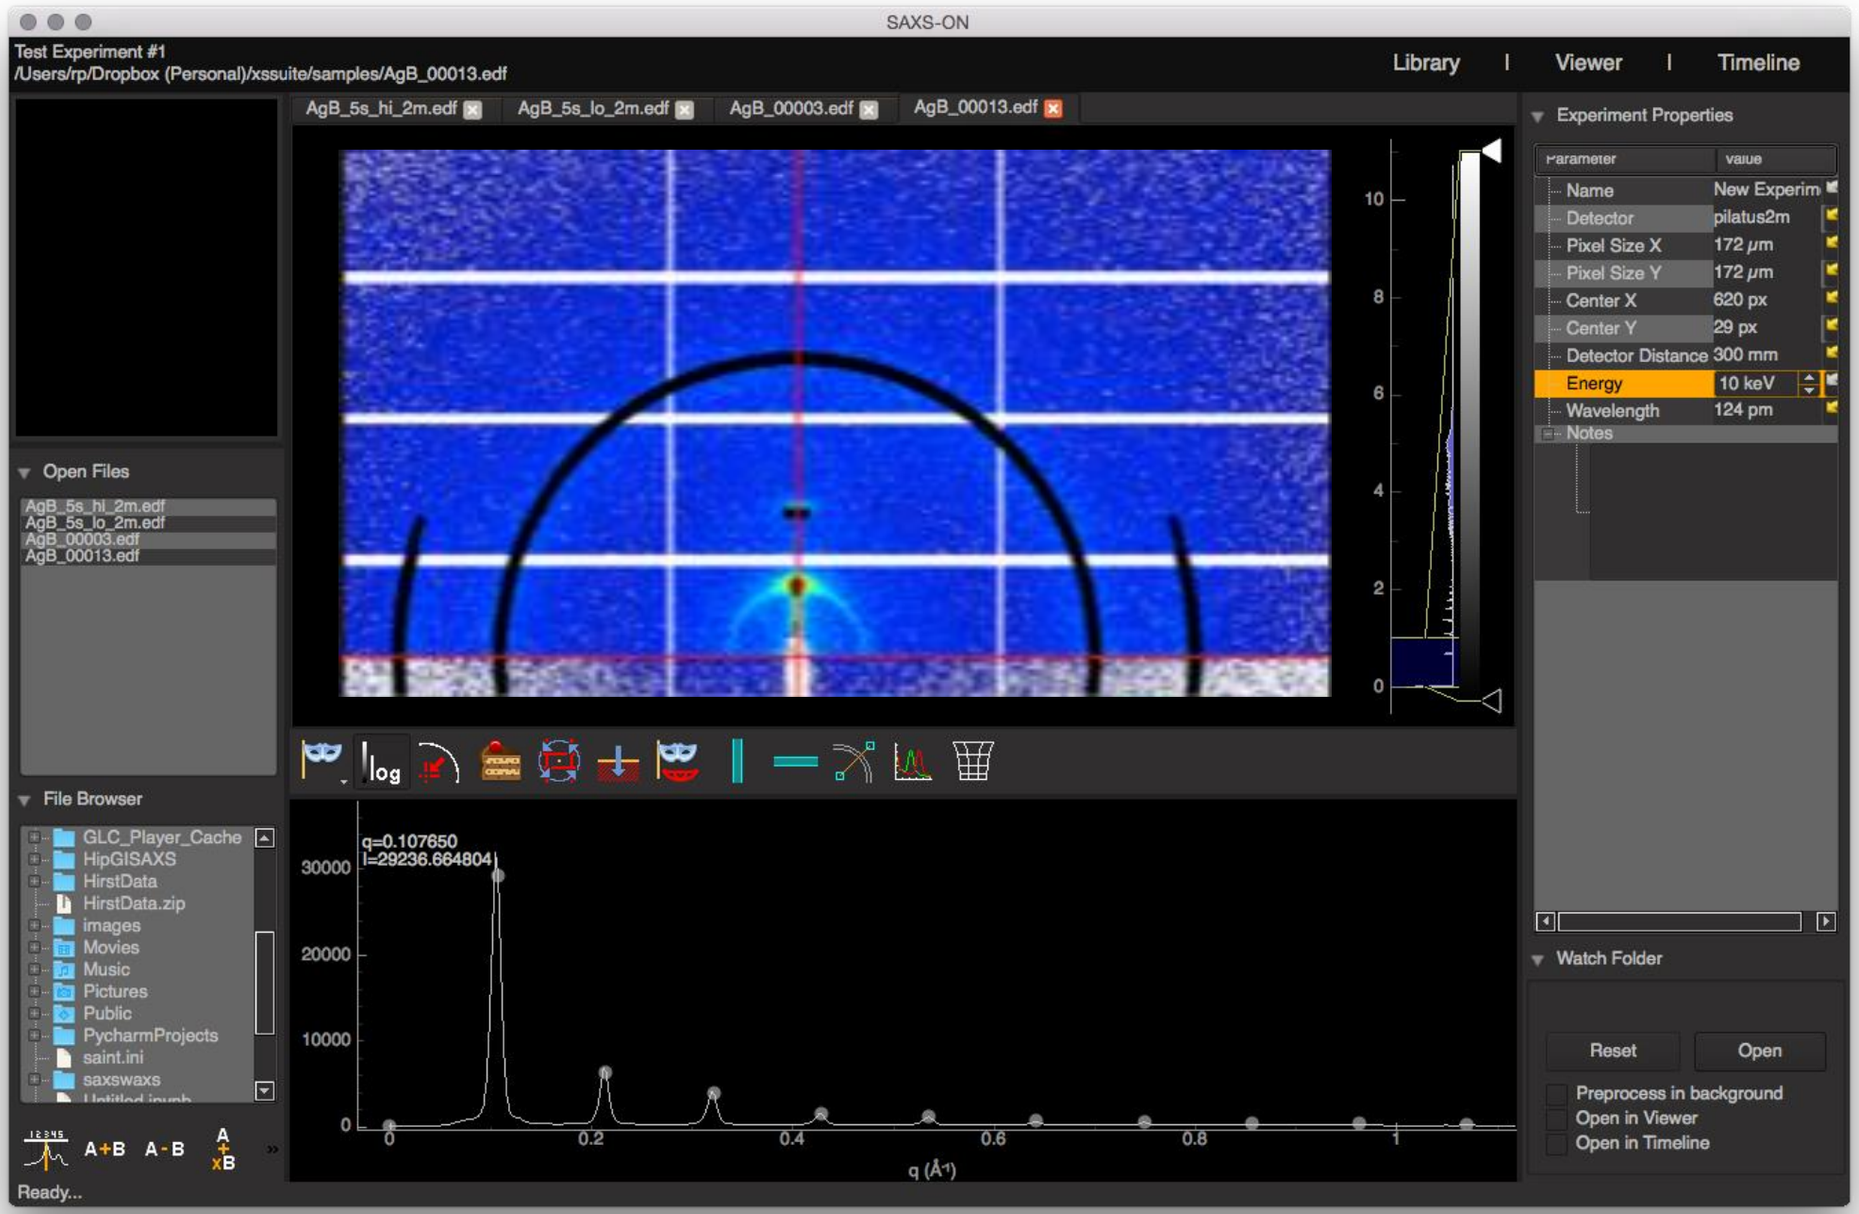

Same code for client and supercomputer

On demand file loading: large data sets local or cloud

Automation:

Data reduction, peak finding, arc finding, background subtraction, time-line compression etc.

• Printing Organic Photovoltaics & Realtime Analysis

- **Show if on-the-fly data analysis of complex systems is possible by combining:**
    - **a state of the art materials science questions (OPV or Nafion)**
    - **state the art X-ray detectors and instrumentation**
    - **advanced mathematical algorithms and software**
    - **fast data movement and visualization**
    - **run on some of the fastest computer in the world**

A. Hexemer (LBNL/CAMERA), C.E.Tull (LBNL), J. Deslippe (NERSC), R.S. Canon (NERSC), E. Dart (ESnet), I.Foster (ANL), J.A. Sethian (LBNL/CAMERA), G. Shipman (ORNL), J. Wells (ORNL), K. Kleese van Dam (PNNL), T.P. Russell (UMass), E. Gomez (PennState)

Facilities: ALS (BES), NERSC (ASCR), ANL(ASCR), OLCF (ASCR), ESnet (ASCR), CAMERA (ASCR)

Nafion

Ionomer
Polymer based proton conductive membranes
used as fuel cell membranes

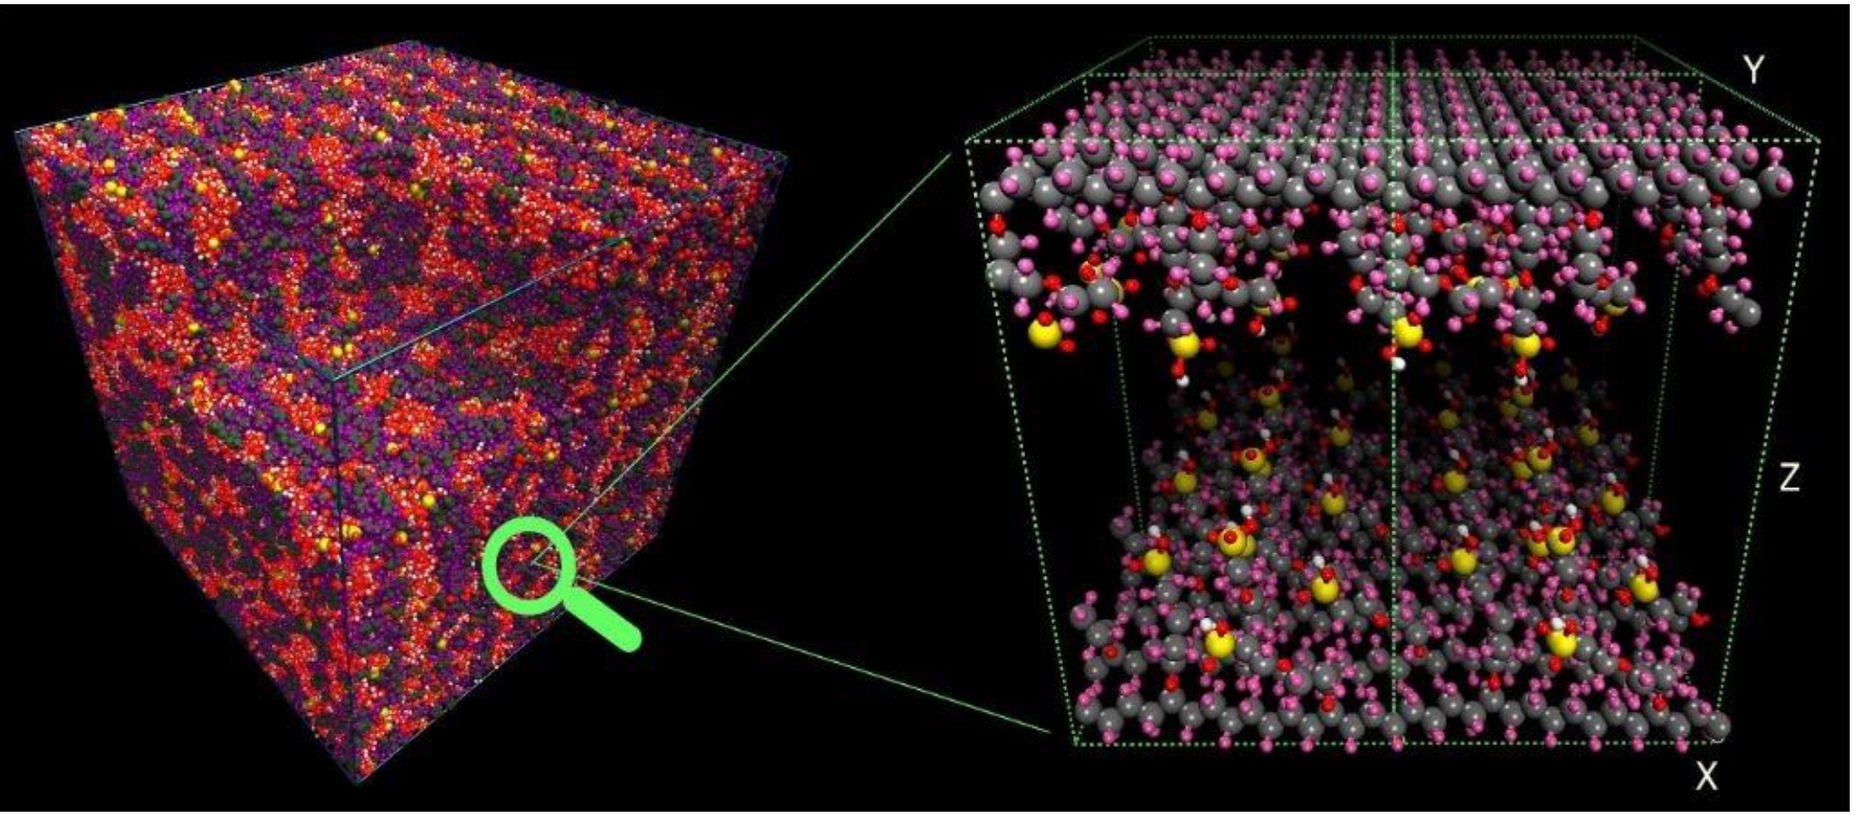

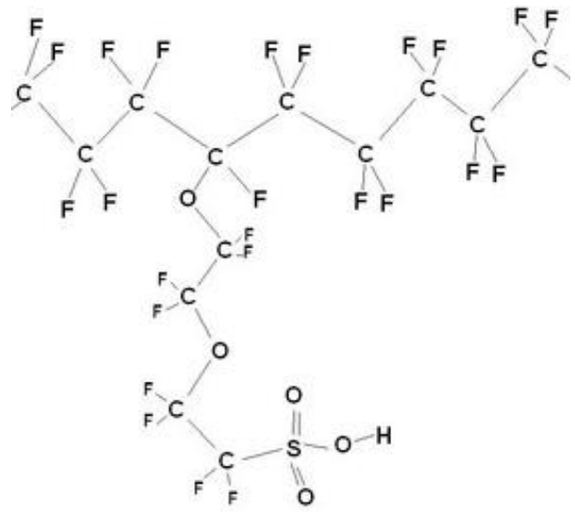

“Teflon” plus side groups

Pavel V. Komarov doi:10.3762/bjnano.4.65

Tensile Test

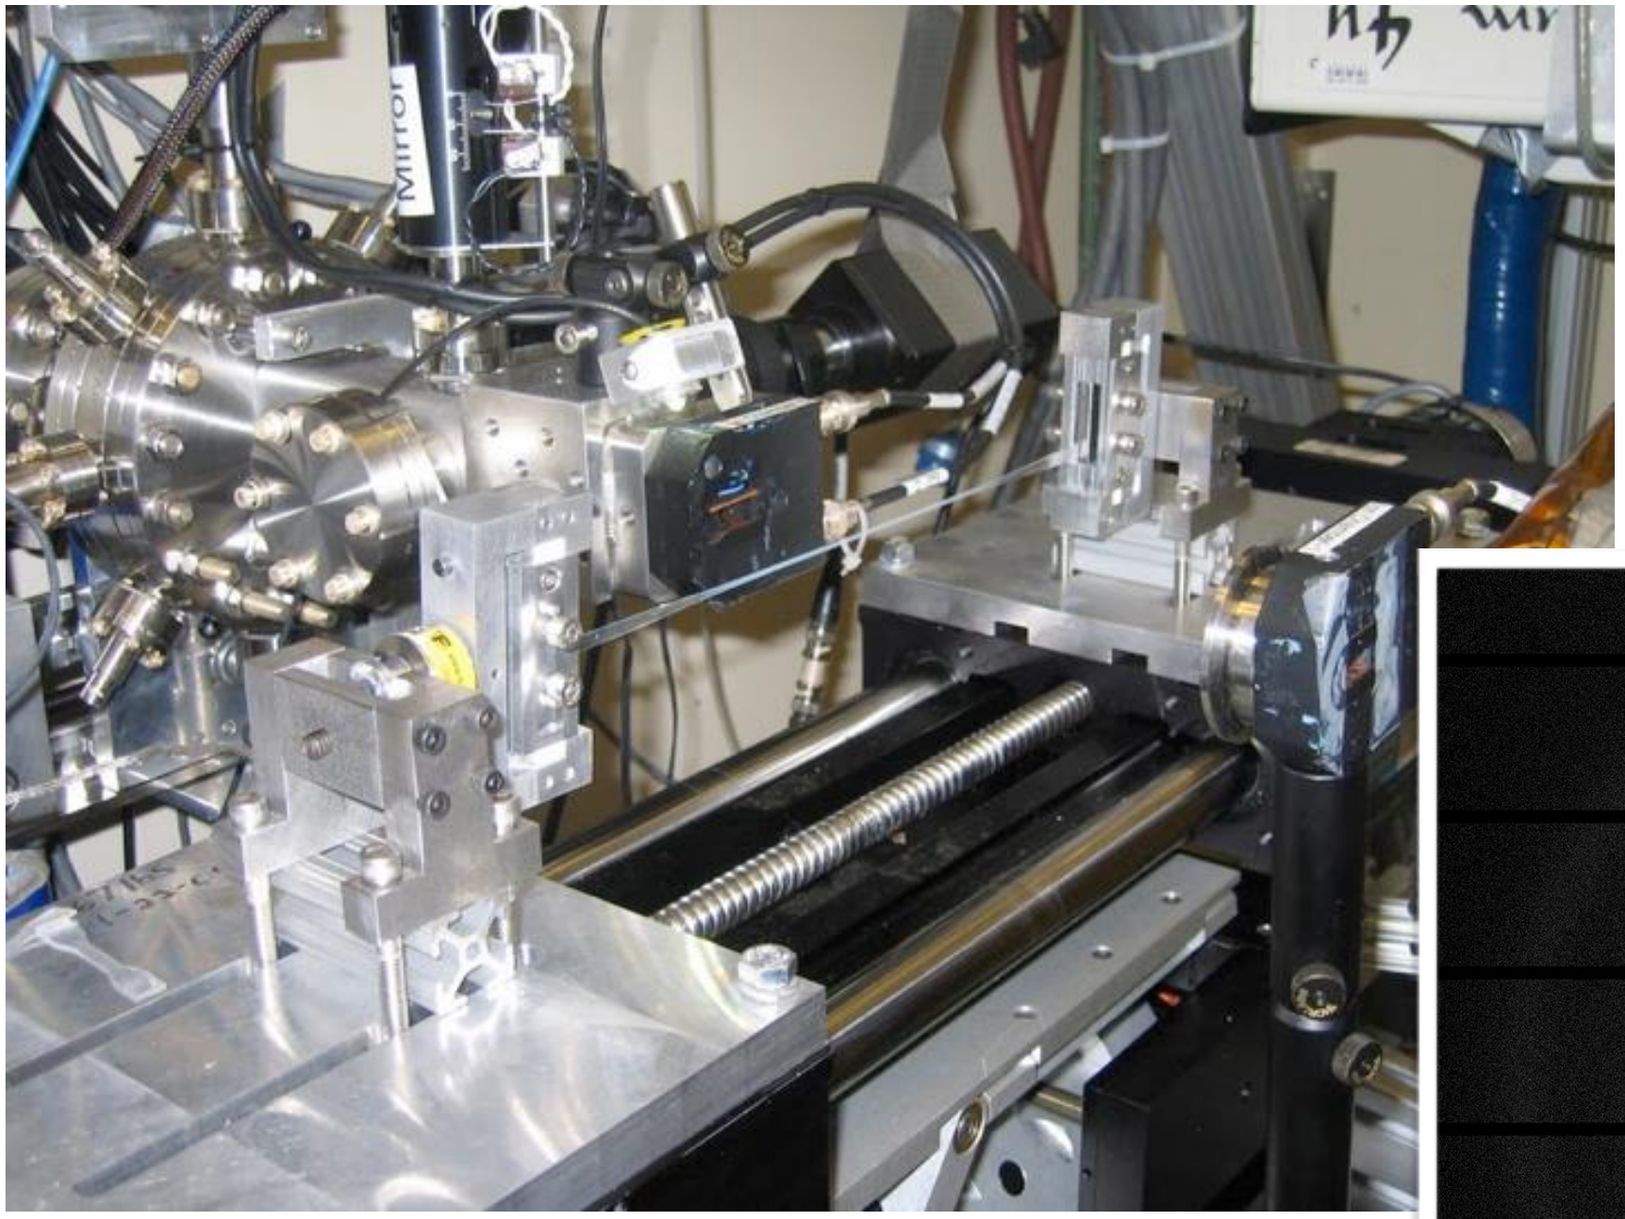

Nafion
Polymeric fuel cell membrane

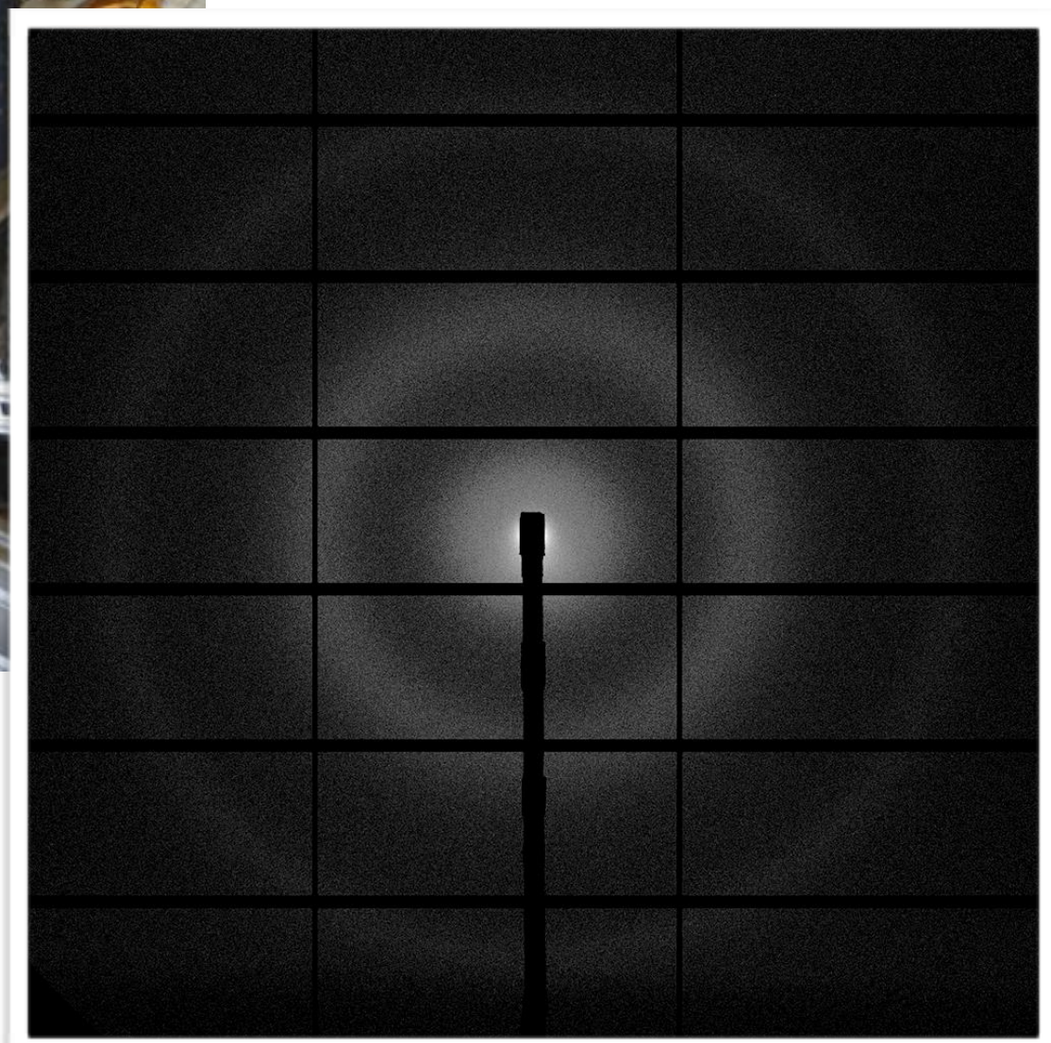

FFT + mask

olymer system with filler particles

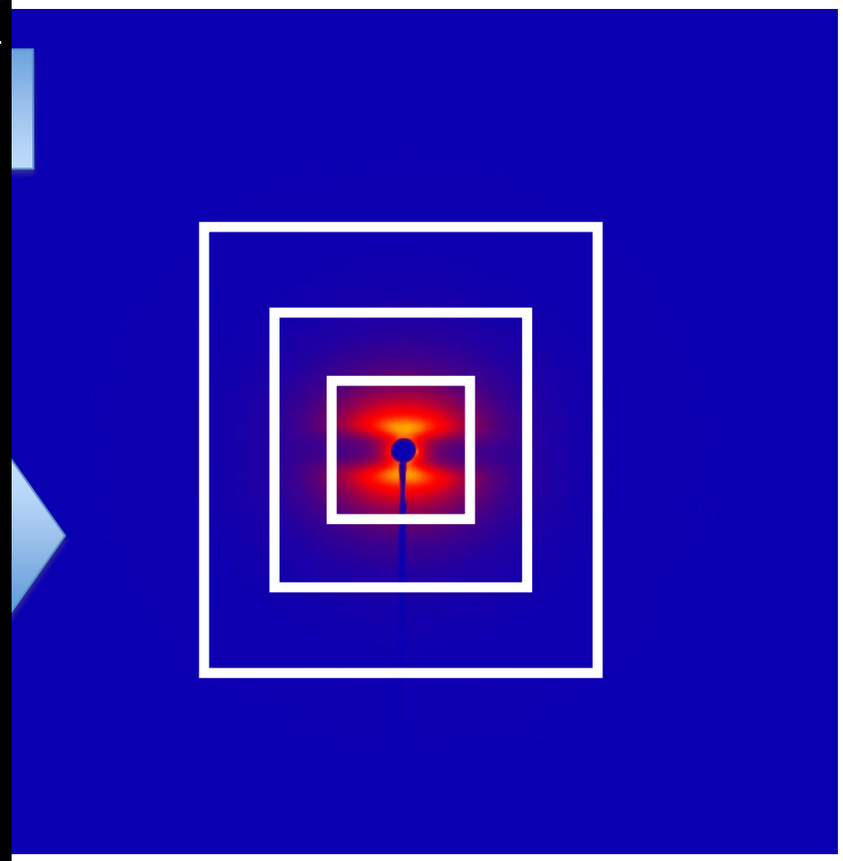

data from experiment

Reverse Monte Carlo Results

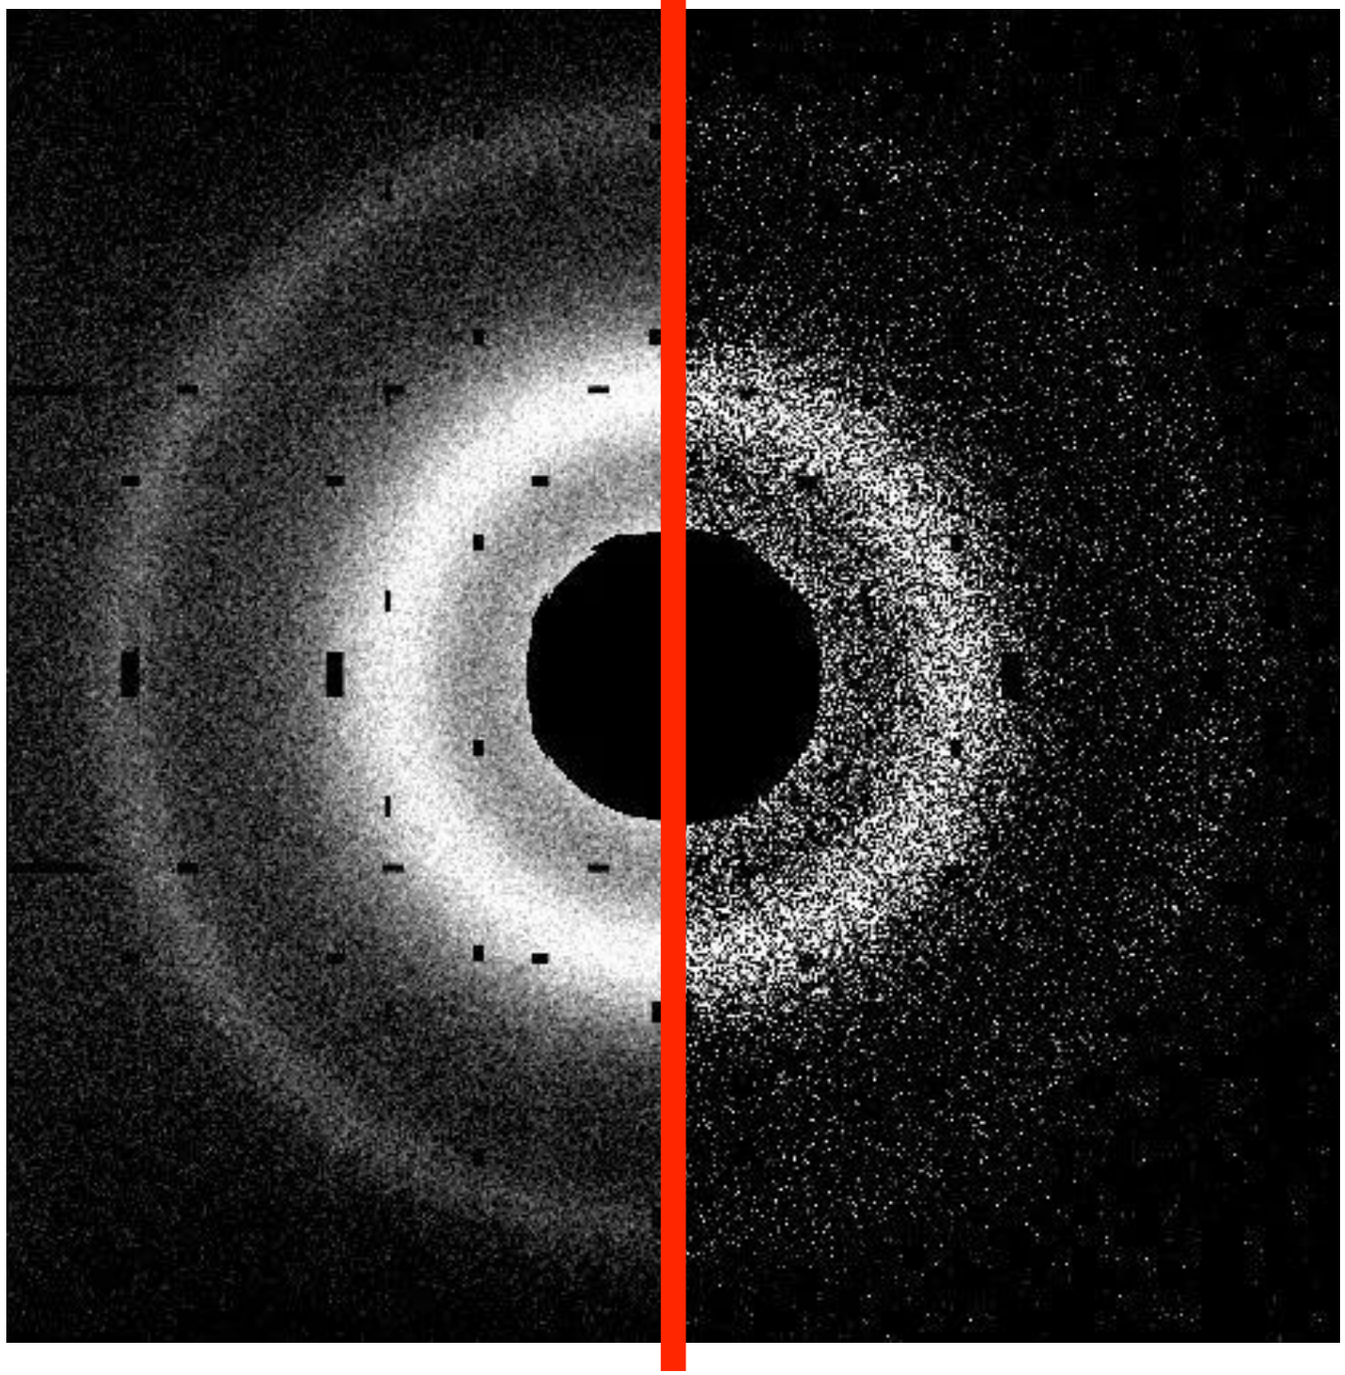

data fit

Typical structure from fitting

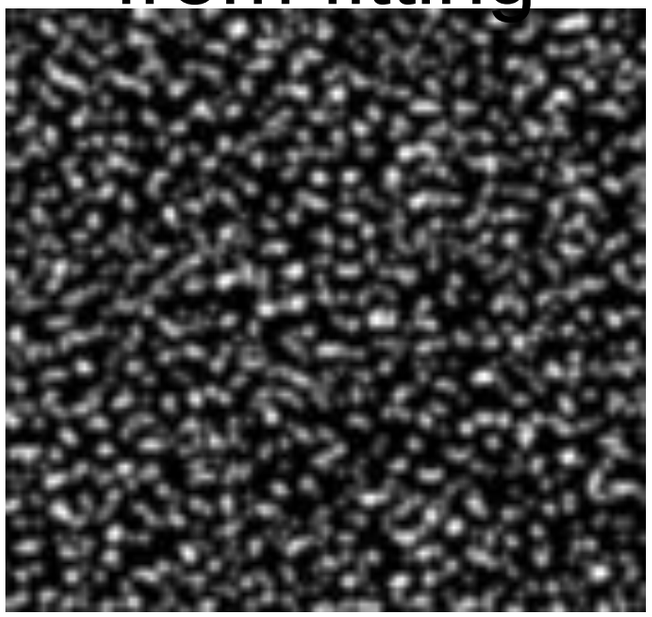

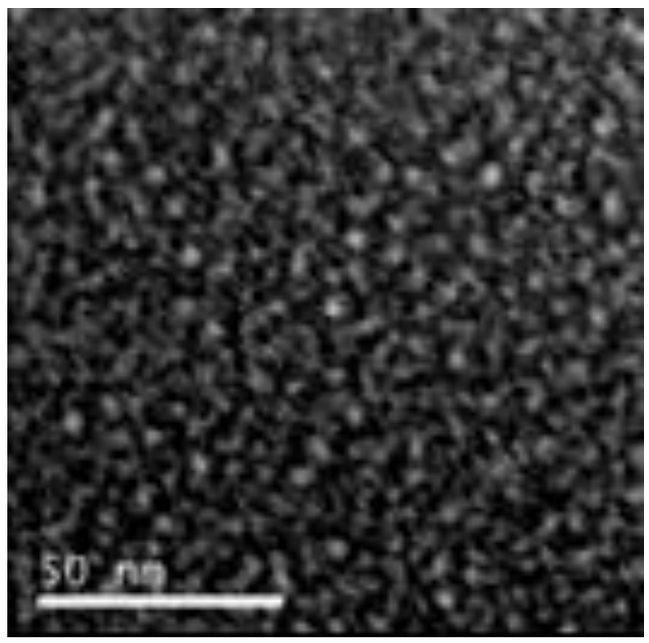

TEM (stained)

How about Real-time?

Experiment:	1 frame per second 600 frames total/sample + 15 min sample change total of 25 min
Time to fit single frame per node (actually 20x same data frame with different initial conditions, since we need statistics) (TITAN and EDISON)	12-20 min

Possible !!!

OPV

OLED

OTFT

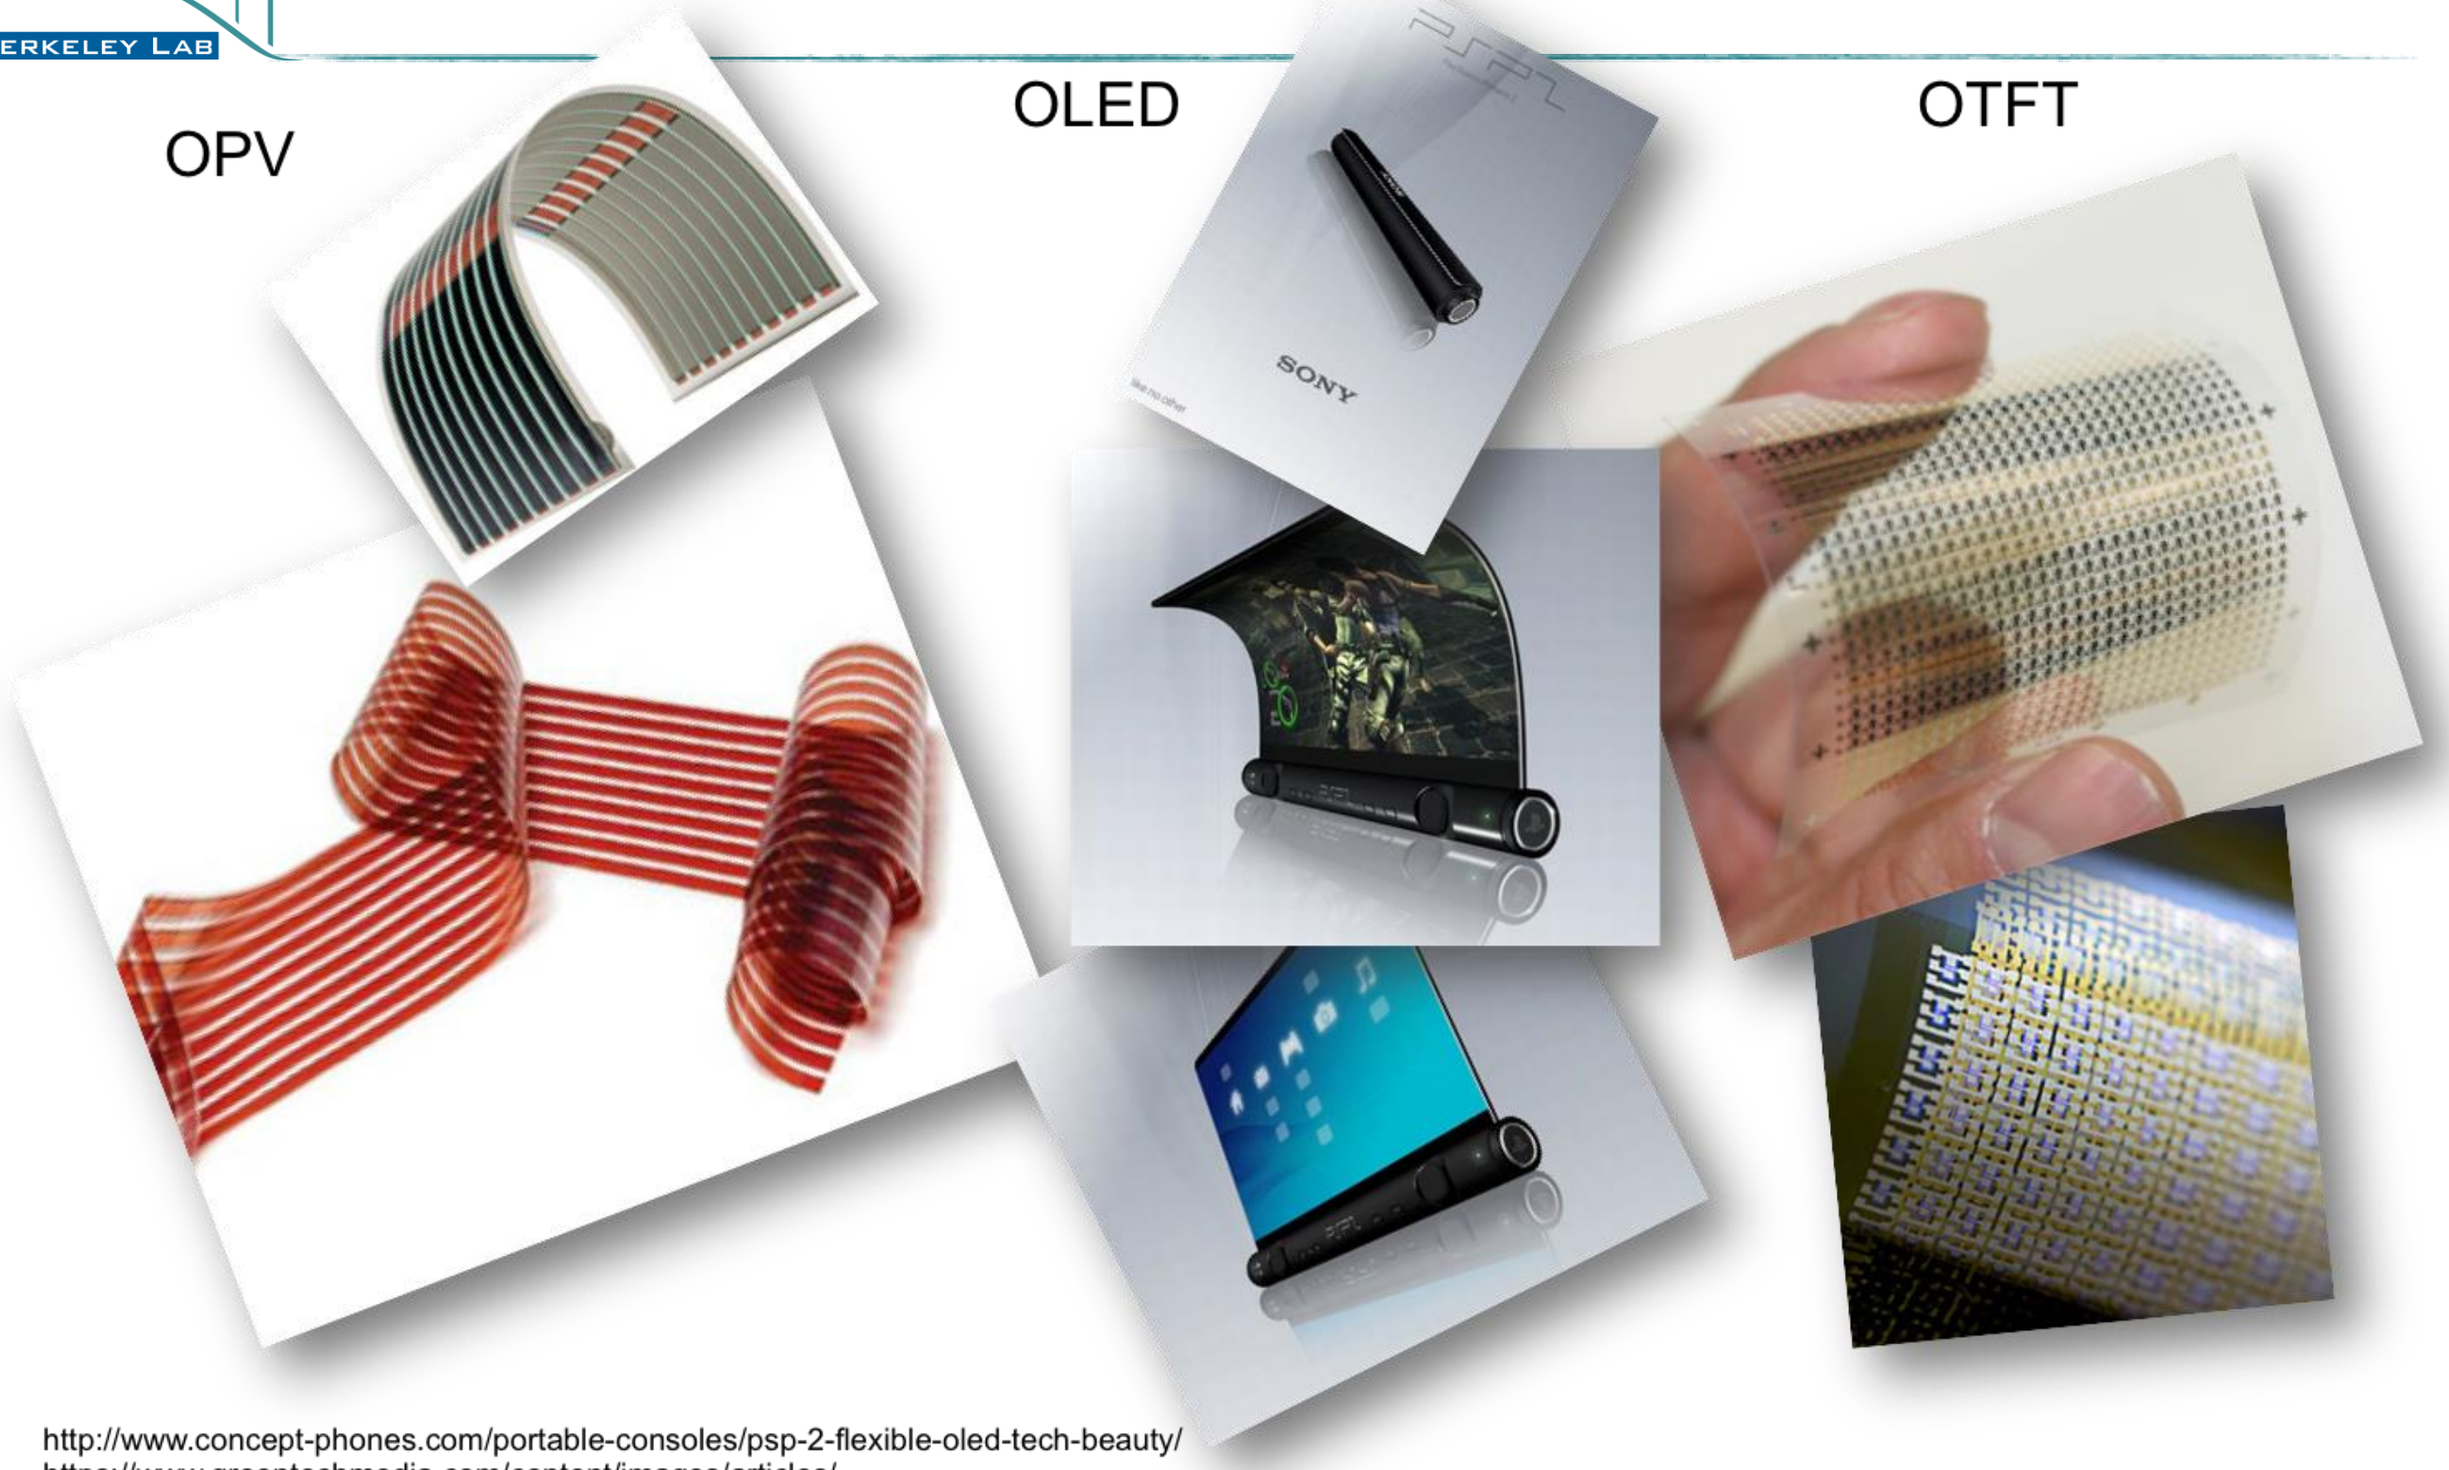

http://www.concept-phones.com/portable-consoles/psp-2-flexible-oled-tech-beauty/
https://www.greentechmedia.com/content/images/articles/KonarkaGets45MFromTotal_medium_image1_1560.jpg
http://www.4engr.com/research/catalog/13/index.html

R2R OPV Manufacturing (UMass Amherst)

Goal: Low-Cost Materials, Low-Cost Manufacturing, Low-Cost Installation

Slot die printer

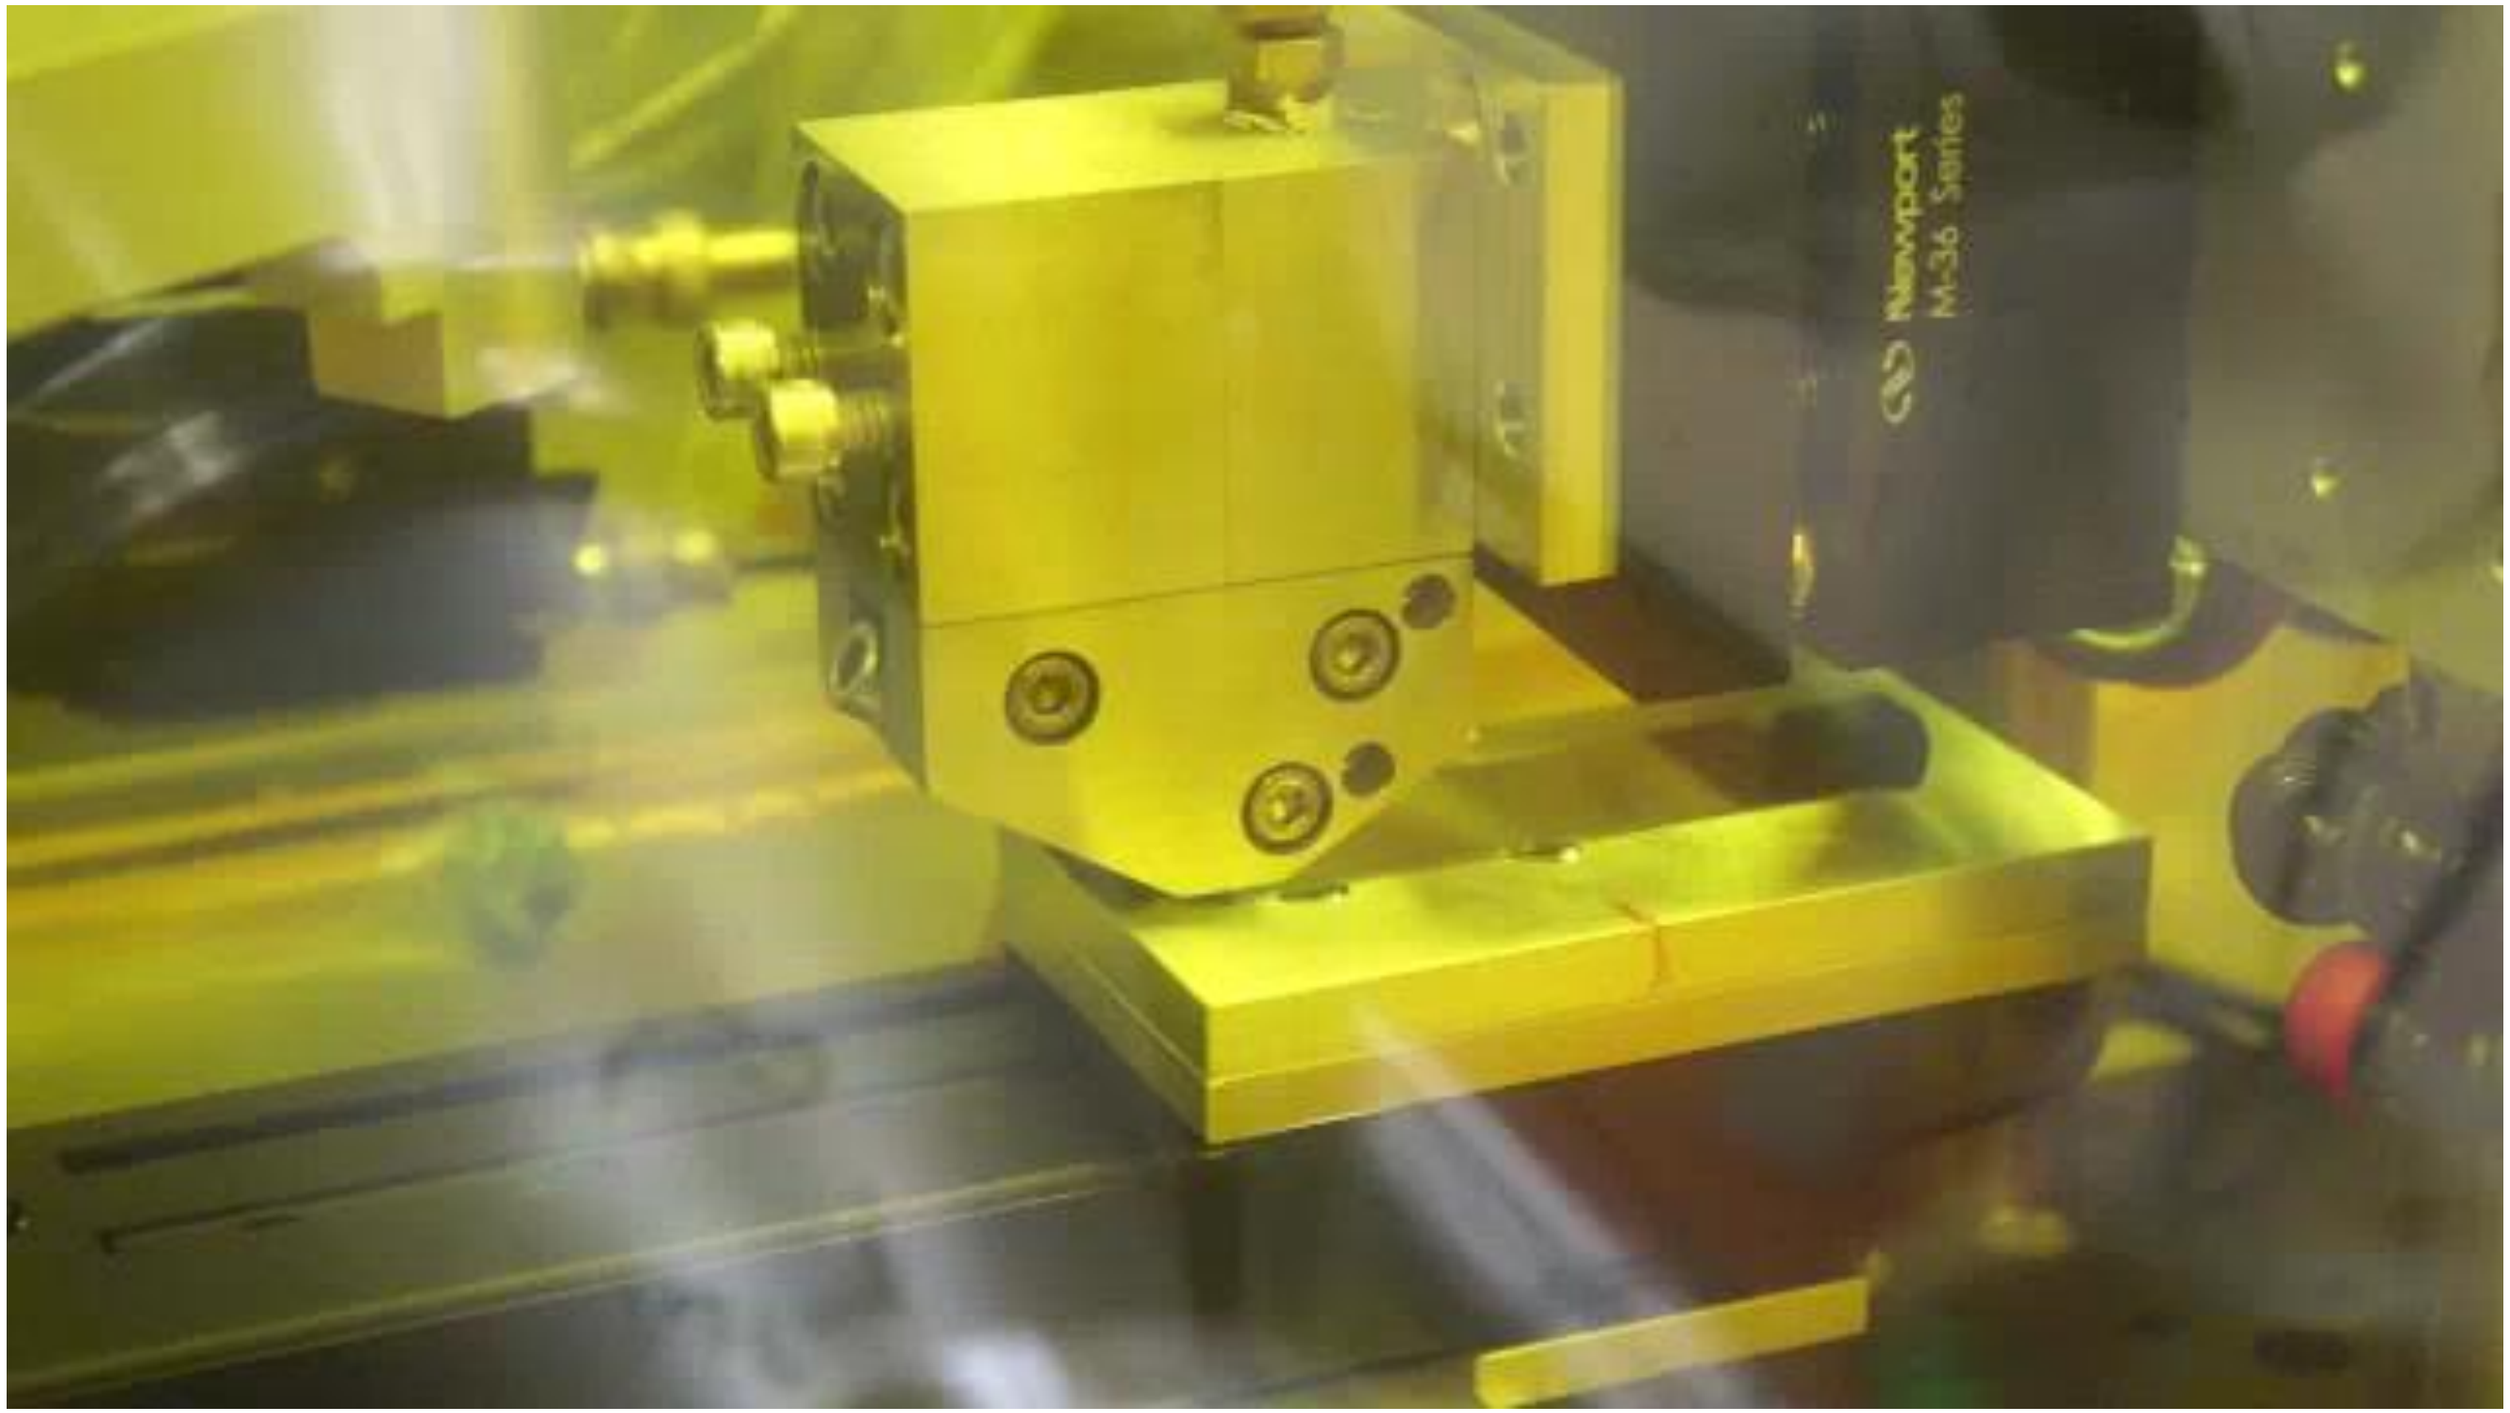

In situ slot die printer for OPVs

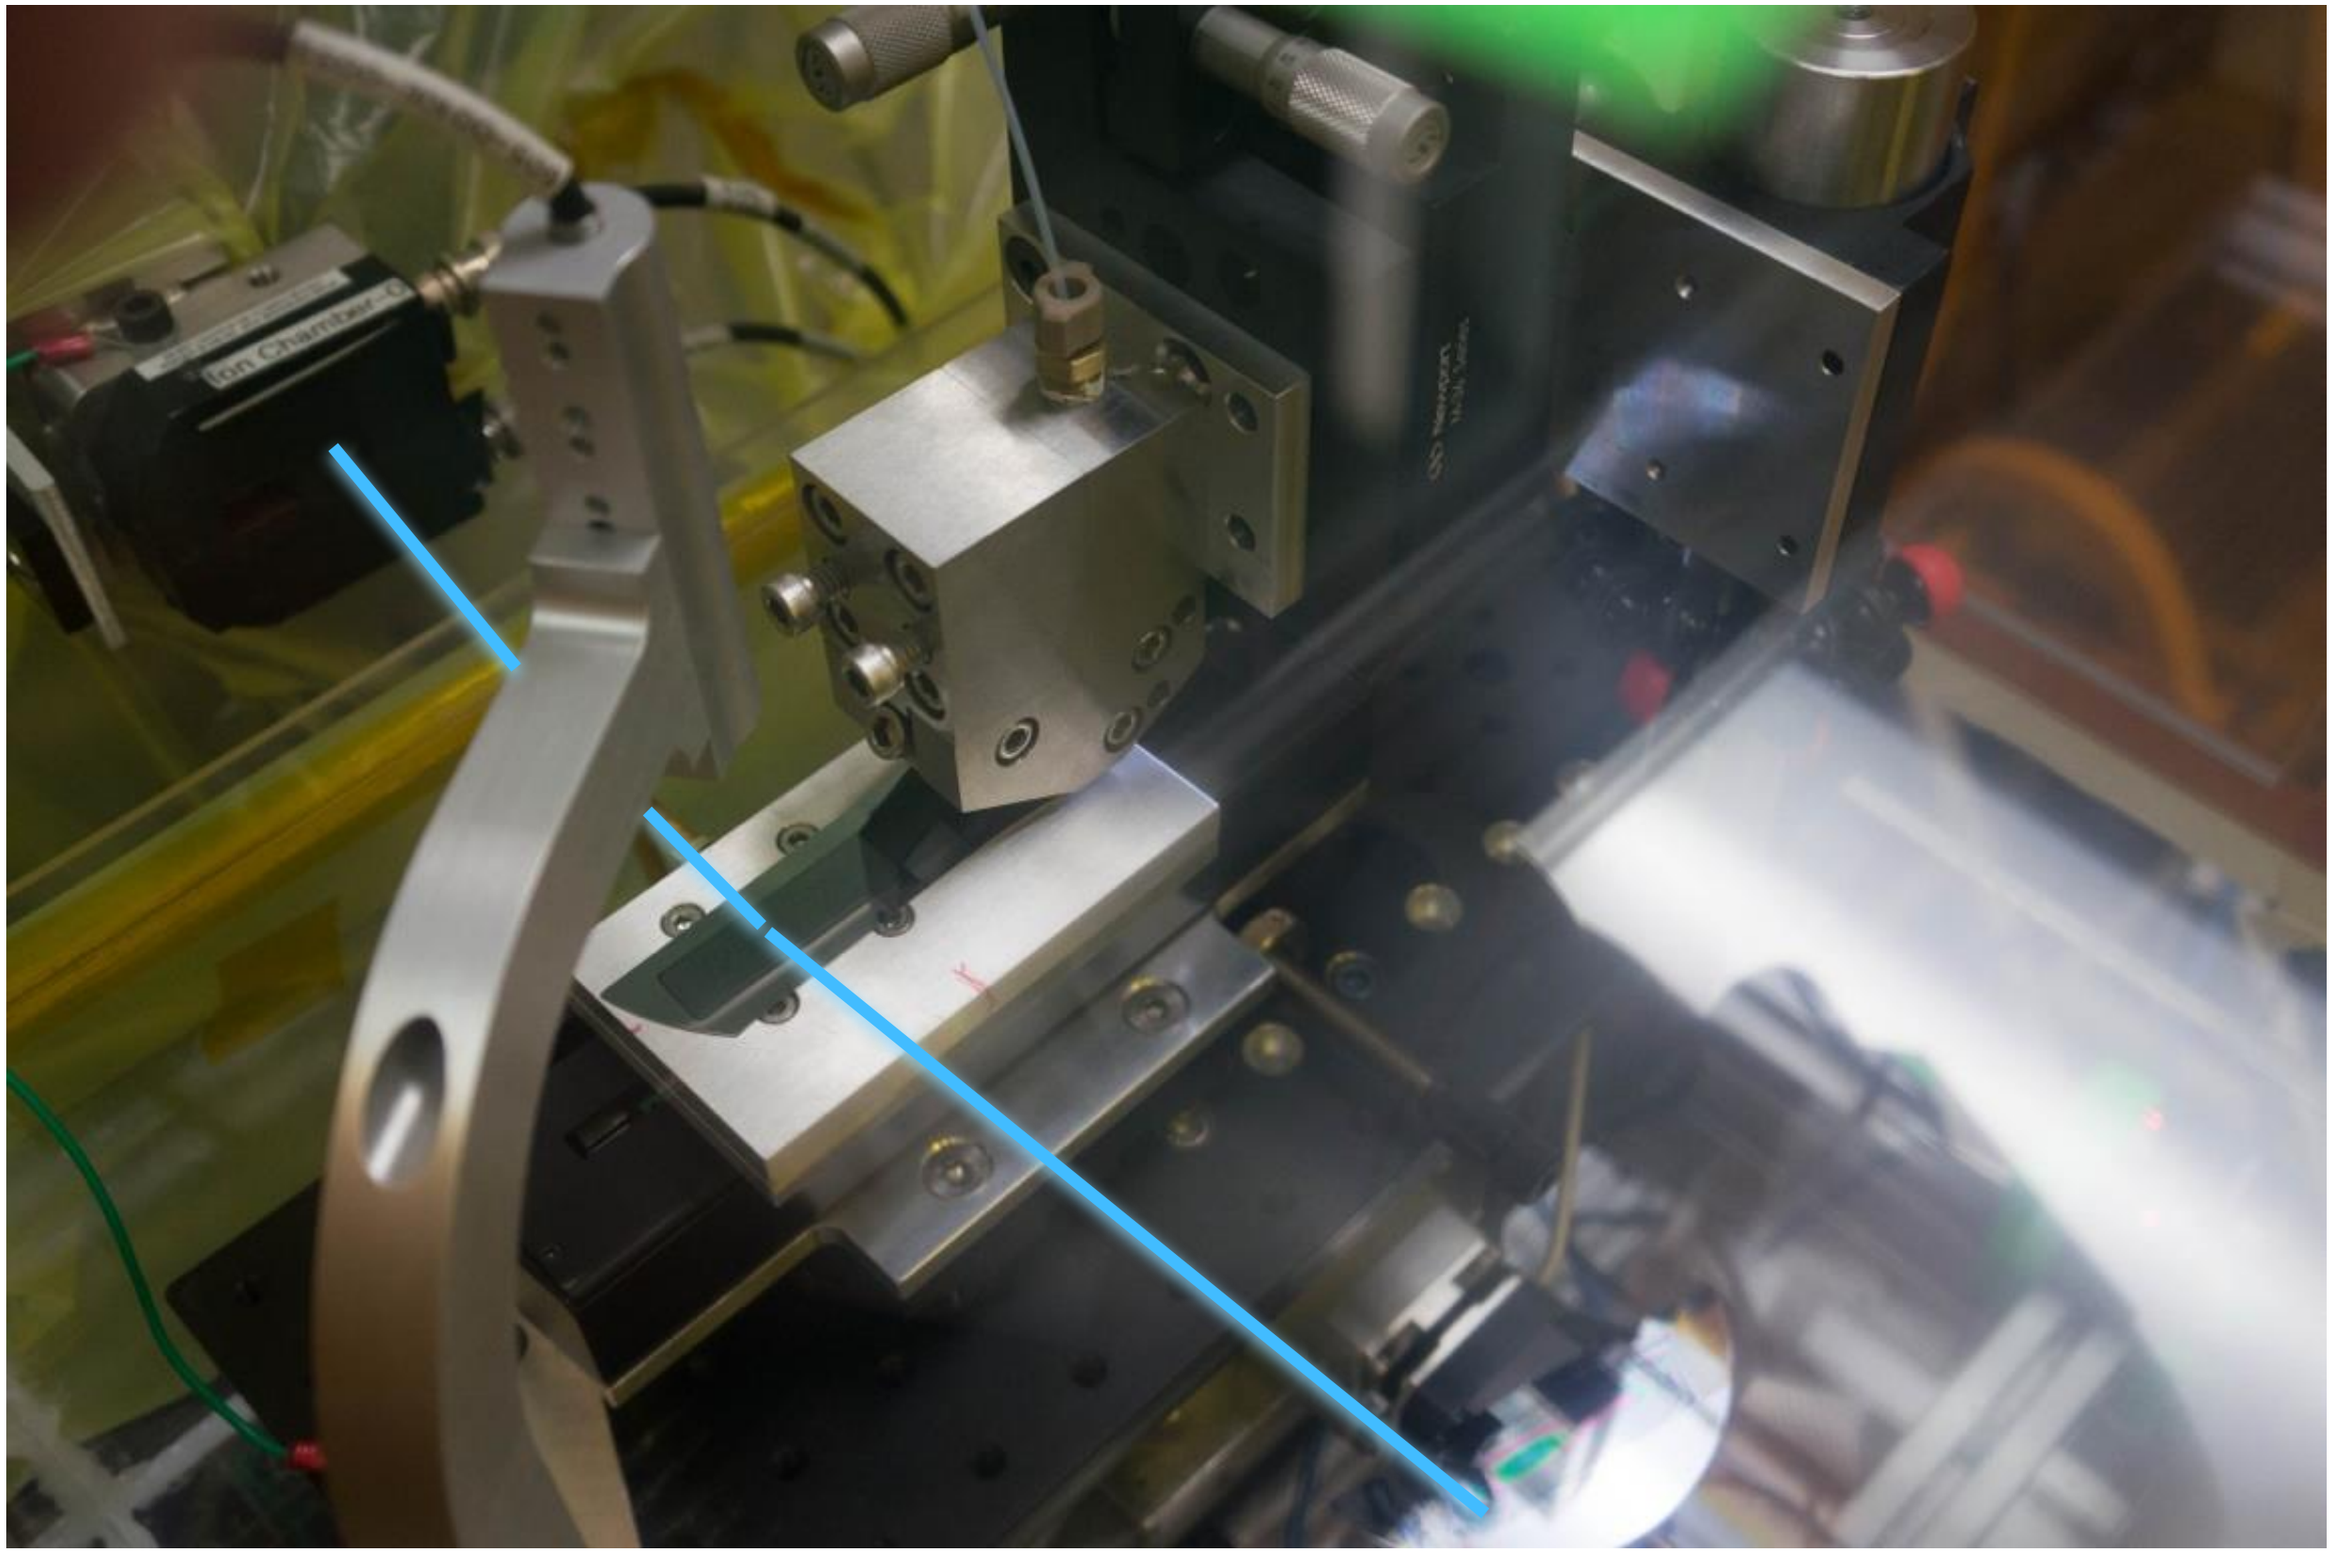

Organic Photovoltaic Processing using Roll-to-Roll

One of the revolutionary appeals of roll-to-roll manufacturing of organic photovoltaics (OPV) is the potential to achieve energy recovery times as low as 10 days. R. Sondergaard et al.: materialstoday Volume 15, Issues 1–2, January–February 2012, Pages 36–49

Roll to Roll printer placed inside the beamline

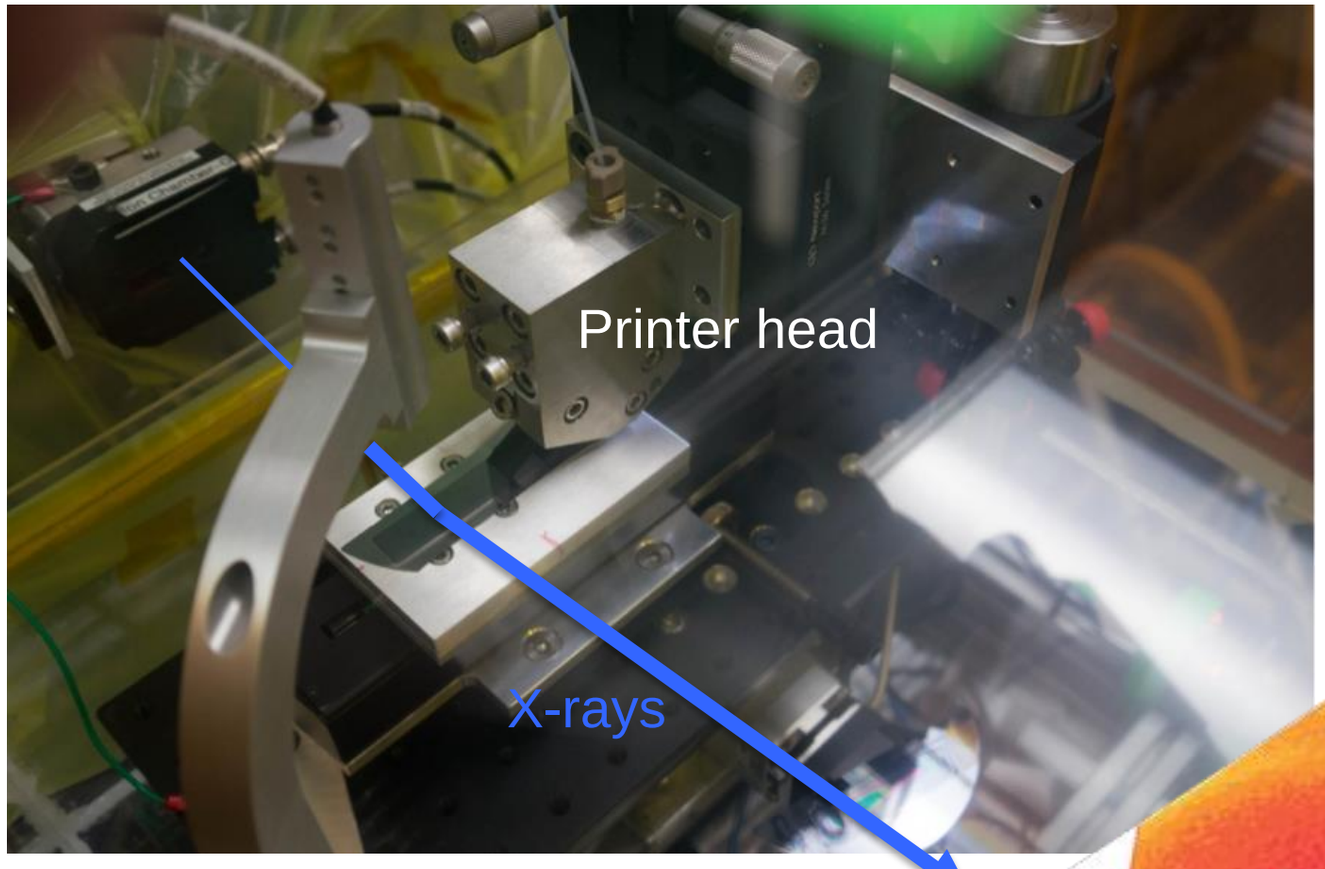

Conjugated polymer and small molecule

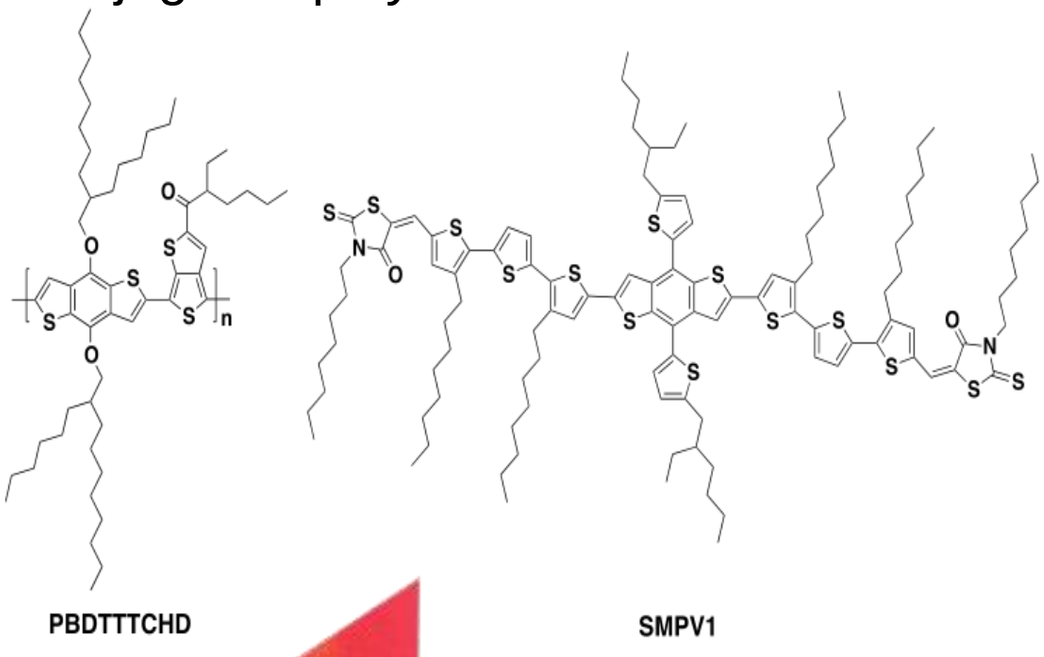

GIWAXS data we collected during printing

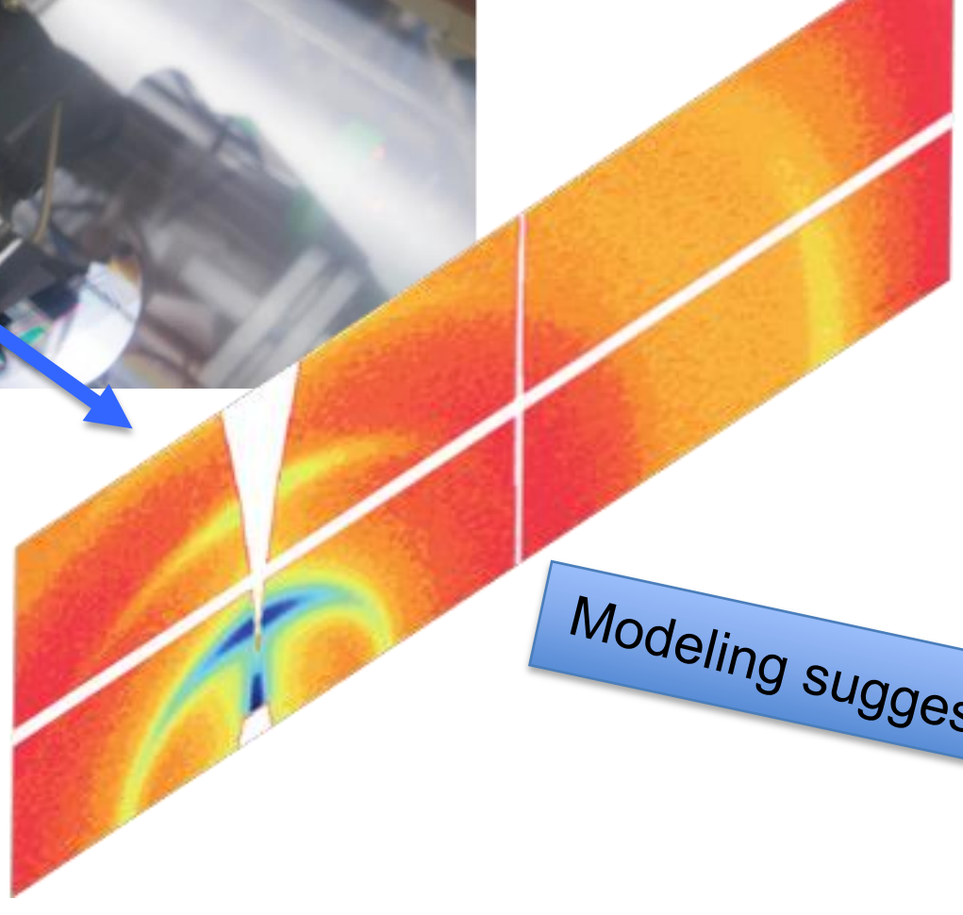

Modeling suggests

Edge-on packing

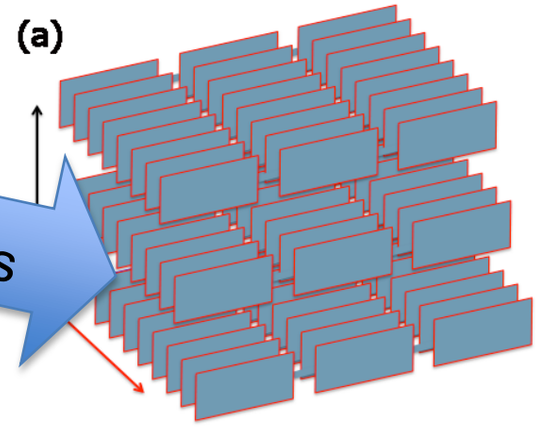

High Performance GISAXS (part of CAMERA now)

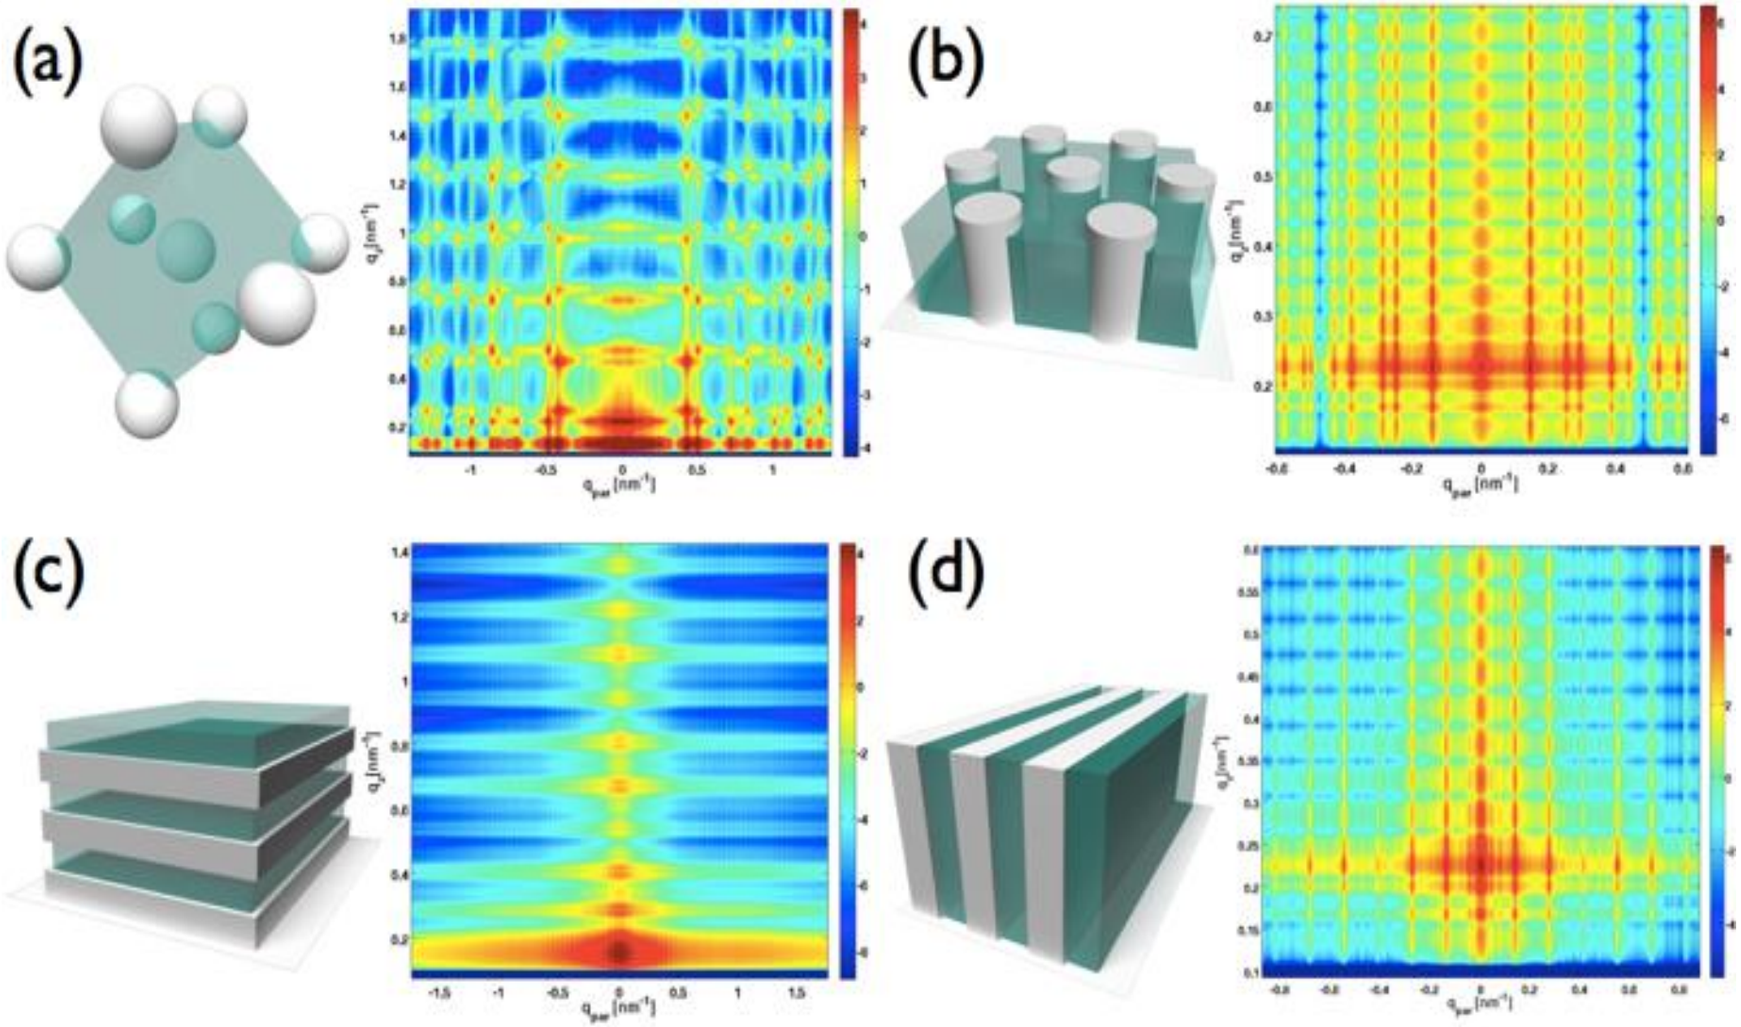

Scattering simulation examples

Include latest theory take advantage of latest architecture: Multi CPU/GPU
Scientist don't care were they do the calculation!!!

Jamie Sethian
Head of CAMERA

www.camera.lbl.gov

Time evolution of GIWAXS patterns

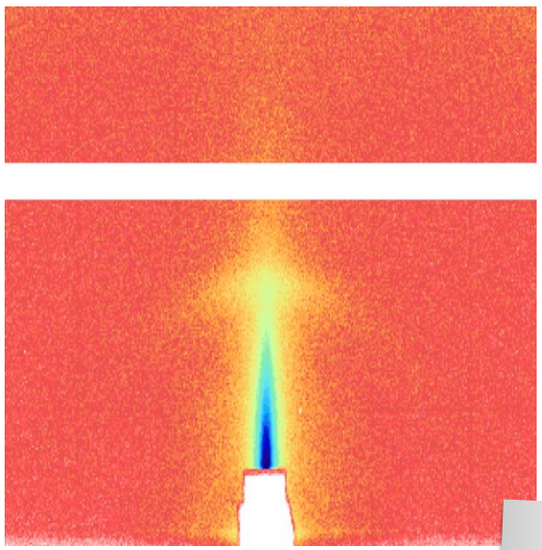

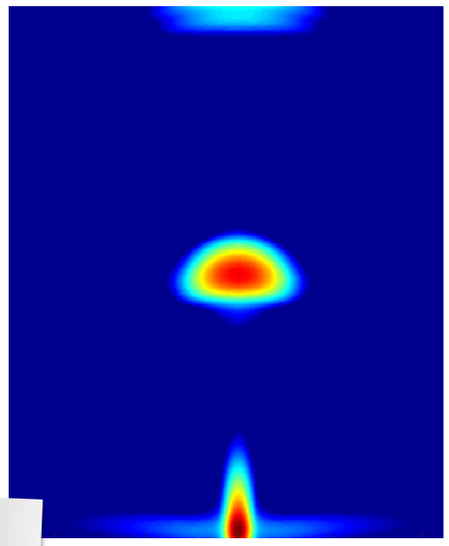

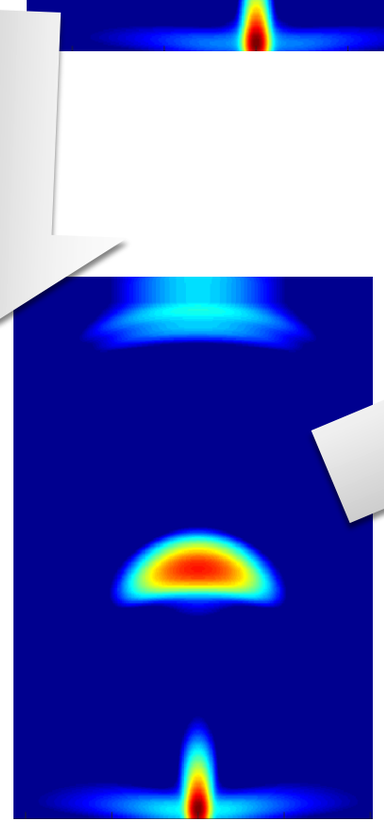

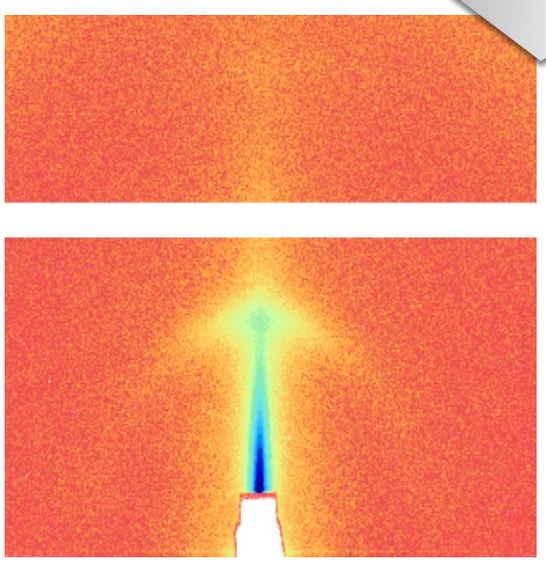

During the drying

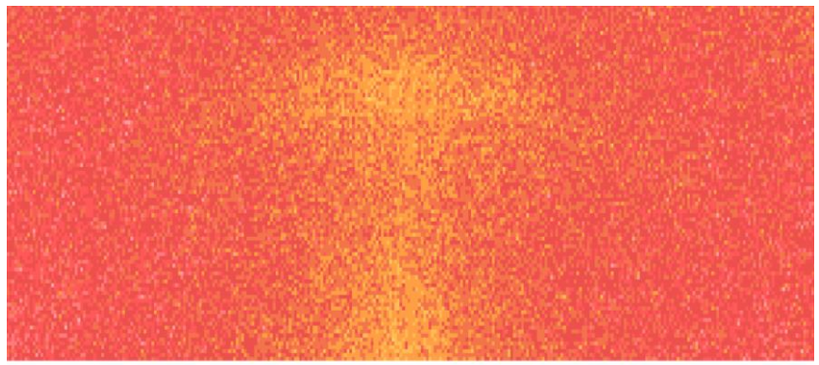

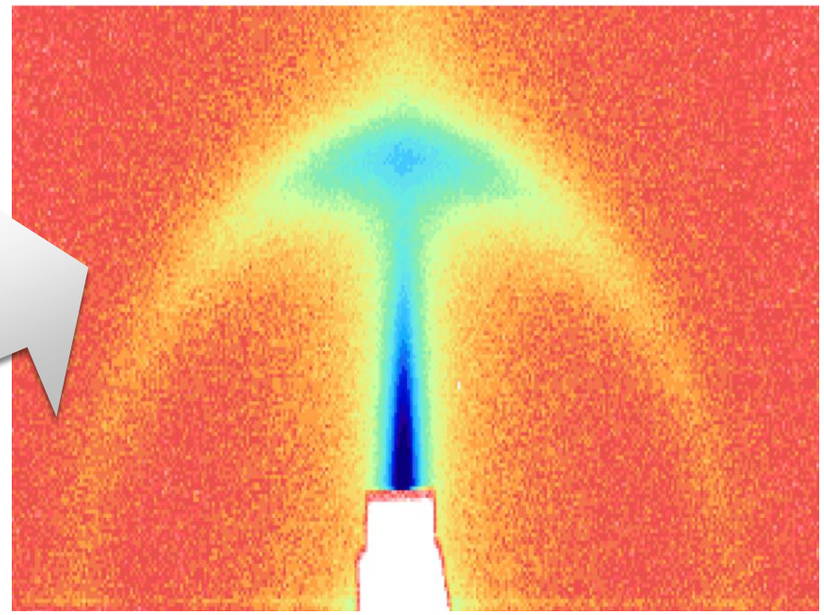

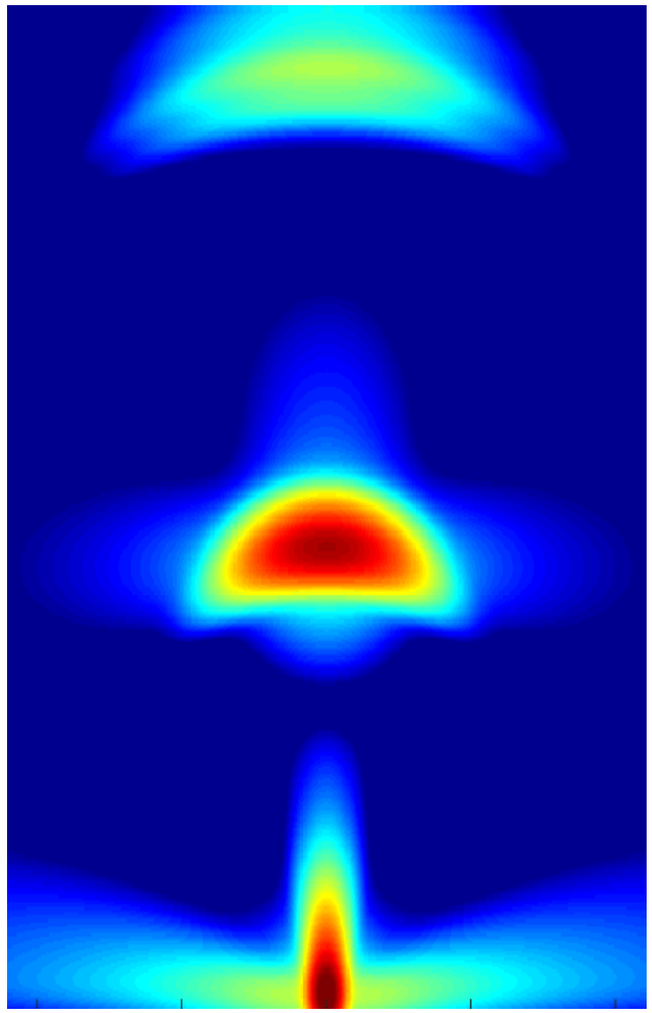

Later stages after the drying

Real time?

Titan: 1 K20X GPU/node

 # nodes = 500, # agents = 20, # generations = 20:
Total time = 3110.00 sec [avg. generation time = 155.50 sec]

 # nodes = 2000, # agents = 50, # generations = 20:
Total time = 2071.60 sec [avg. generation time = 103.58 sec]

 # nodes = 8000, # agents = 80, # generations = 20:
Total time = 865.60 sec [avg. generation time = 43.28 sec]

Printing demo experiments created 36,000 frames in 3 days (1/2 year on TITAN)

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

Real-Time Data Pipeline and analysis using SPOT and HIPGISAXS

Alexander Hexemer and Craig E. Tull
Lawrence Berkeley National Lab

OLCF Users' Meeting
June 24, 2015 @ OLCF

Value, Volume, Velocity, Variety

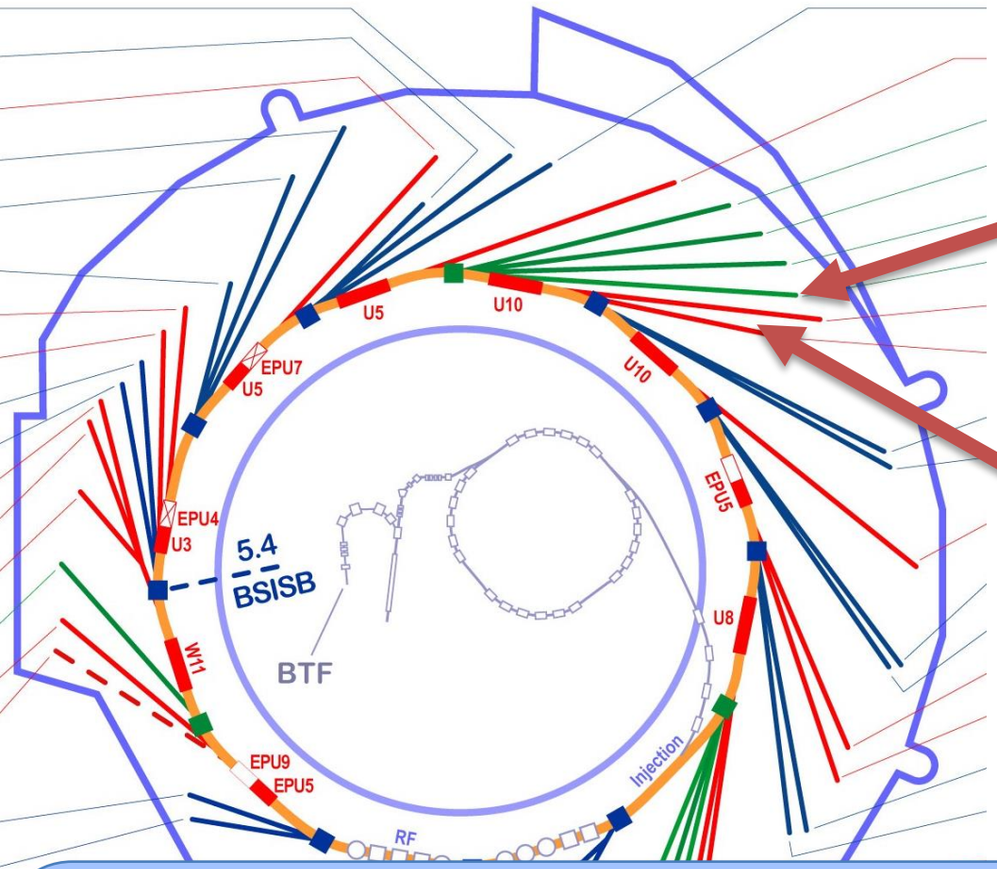

Reciprocal Space
(Scattering)

100,000's images per day

HipGISAXS/HipMC
parallel Scattering

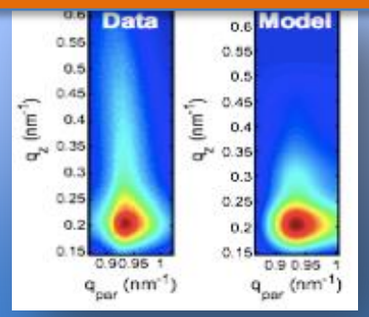

Real Space
(Tomography)

100's of TB per year

Arec3d, QuantCT,
ImgRec, GridRec

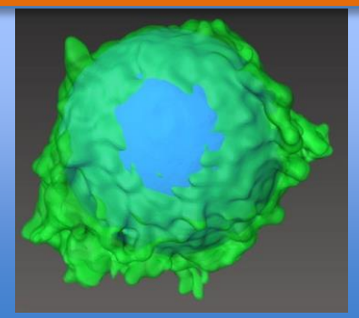

Spectroscopy
(MicroXas)

100's of TFlops of GPU/CPU

ShirleyXAS (MSD)

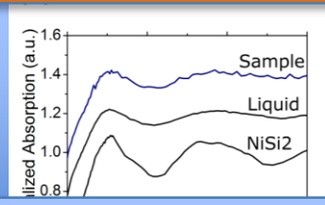

"Past limitations of detector technology have been largely solved through recent DOE investment and that the present bottleneck for research throughput is the lack of availability of appropriate analysis software and modeling tools."[1]

[1]Advanced Data Analysis and Modeling Tools for Scattering Methods Workshop - Sept, 2010

Light source scientists need accessible…

- scalable software systems.
- HPC/HTC/network resources.
- advanced algorithms & analysis.
- advanced simulation & theory.
- realtime feedback.
- advanced visualizations.

The definition of facility must expand to provide scientific knowledge, not files.

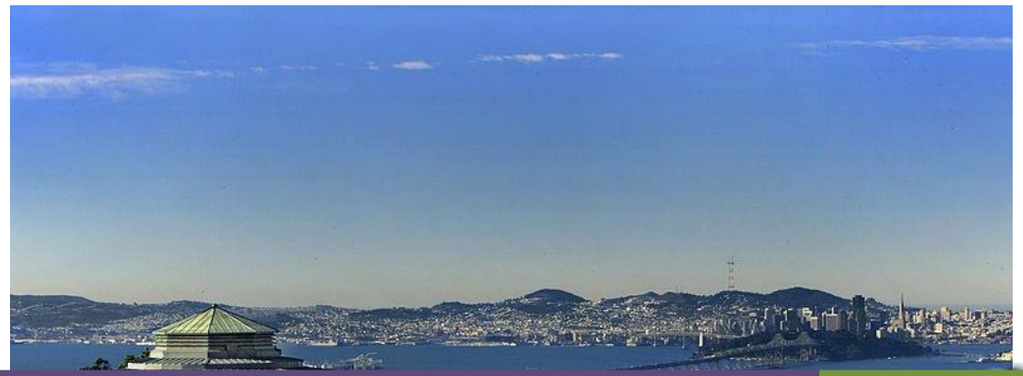

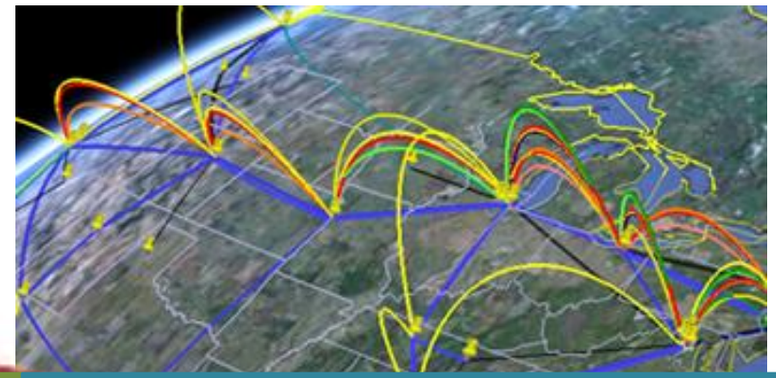

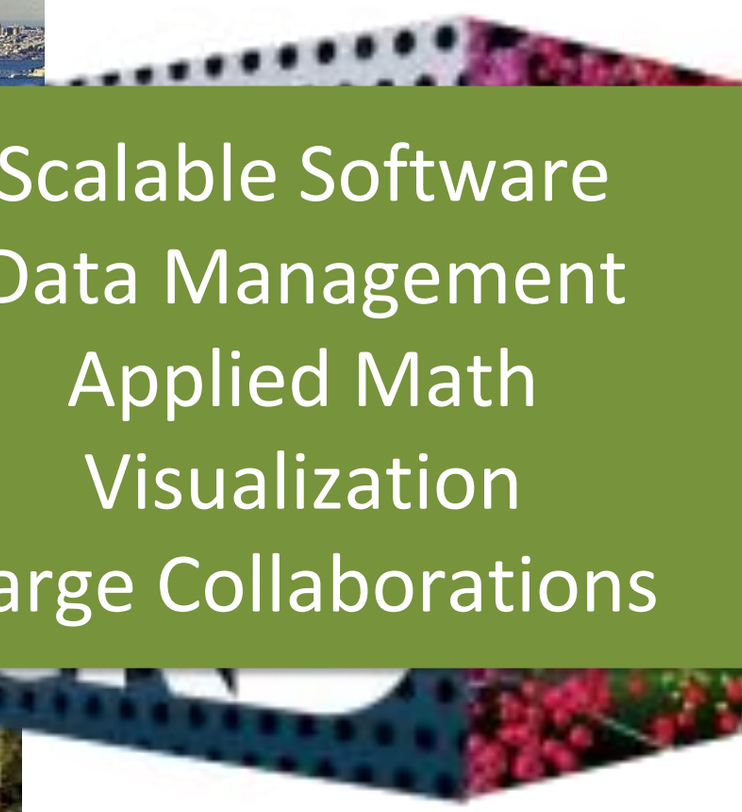

Experimental Data	Scalable Software	Advanced Network
Science Expertise	Data Management	HPC & HTC Facility
Simulations	Applied Math	Data Systems
Theory	Visualization	Science Gateways
Algorithms	Large Collaborations	User Support

SPOKE

CETull@lbl.gov - 24 June 2015

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

SPOT Suite: Integration of ALS, ESnet, and NERSC into a proto-super-facility.

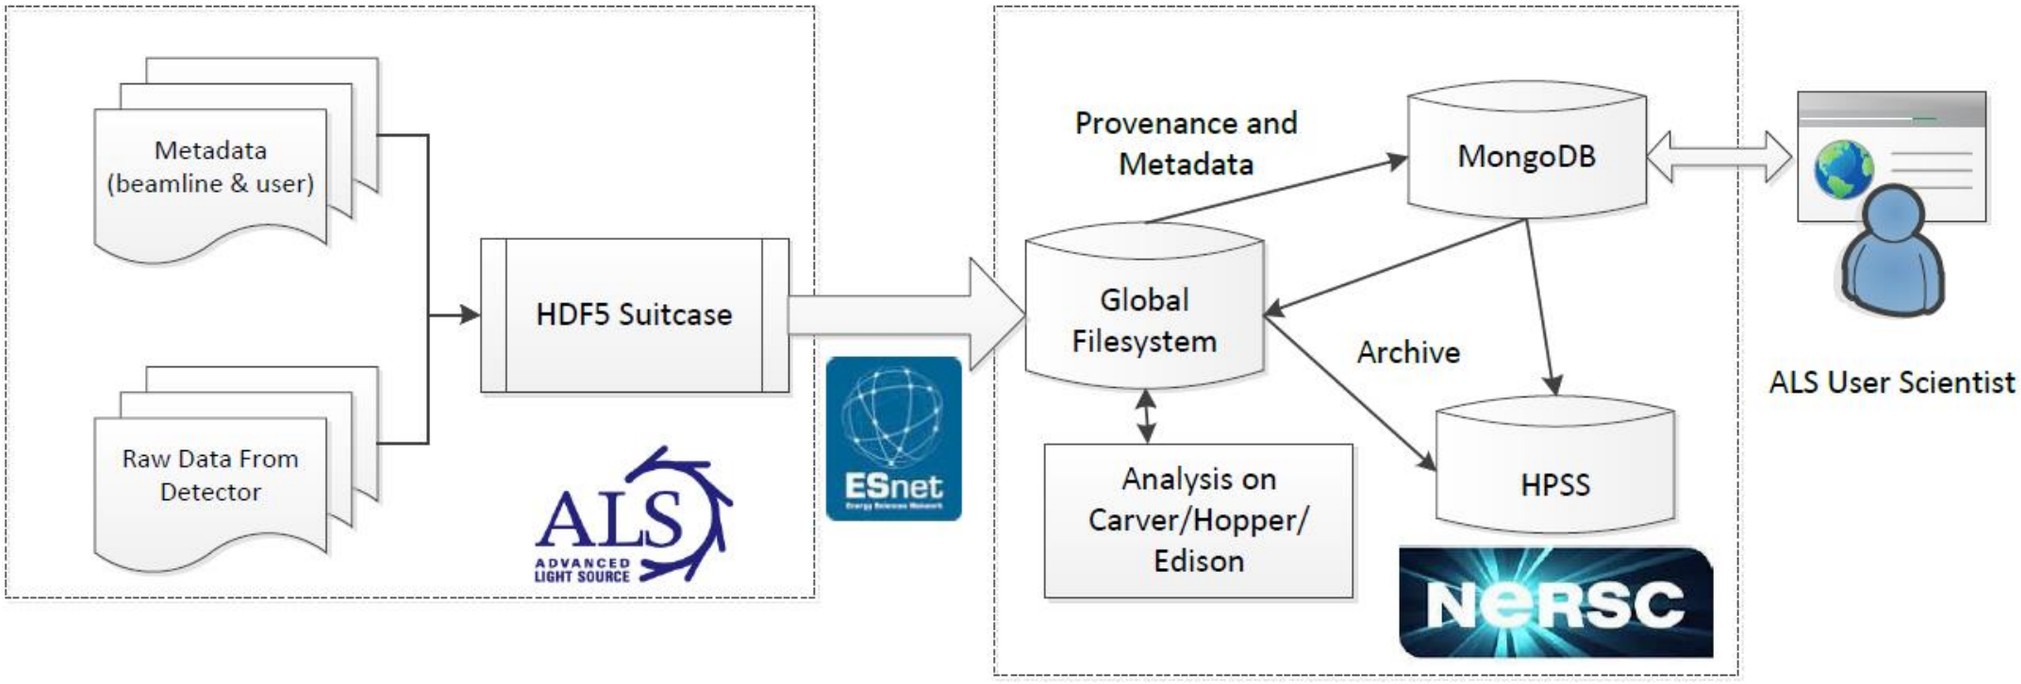

- **Computing Research Div., Advanced Light Source, Material Science Div., ESnet, NERSC**
- CRD brings large-scale data science experience & advanced computational science expertise.
- ALS represents each data "theme", provides algorithm expertise, and allows at-scale tests using real data, real executables, and critical science.
- CRD, MSD, NERSC supply HPC-ready simulation, advanced algorithms (CAMERA), and viz (VISiT)
- NERSC, ESNet provide raw capacity and experience to solve problems at scale. (eg. Science DMZ)

SPOT Suite LDRD: Multi-division team from CRD, ALS, ESnet, MSD, NERSC.

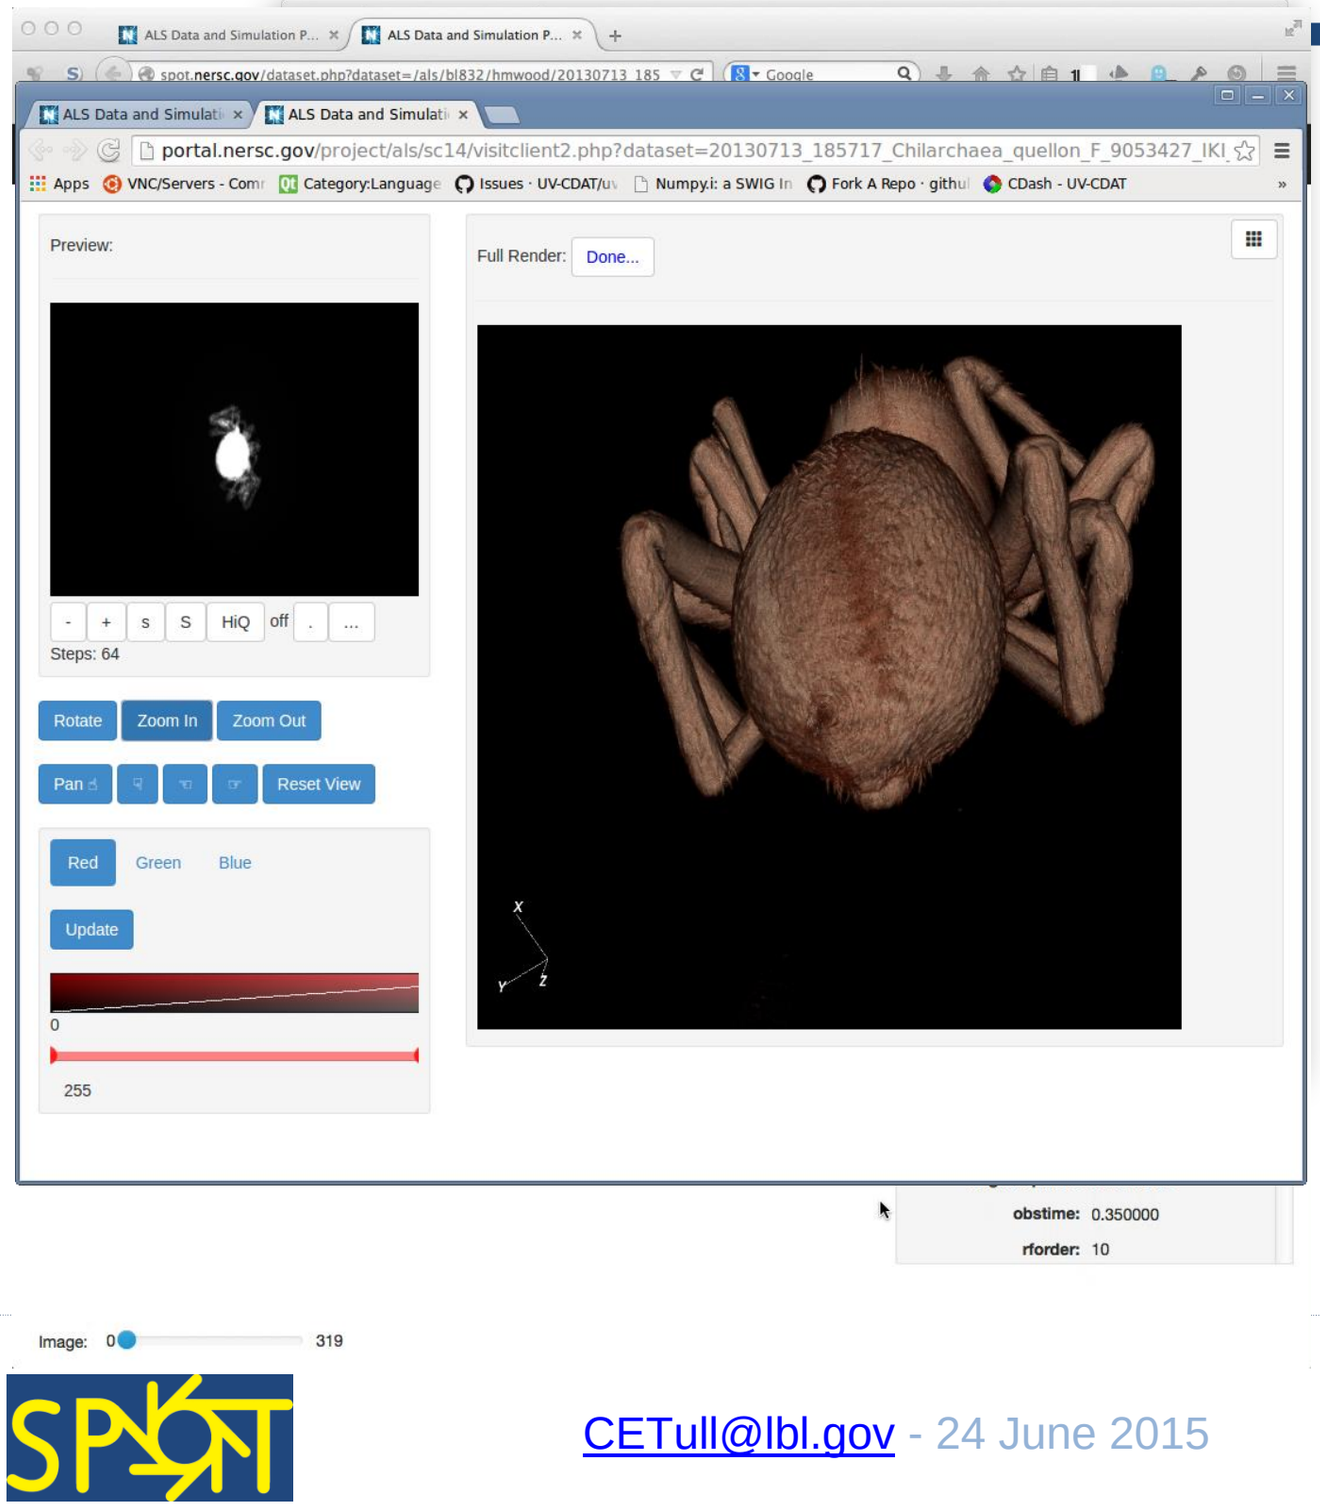

Pseudo-production running at ALS beamlines:

- **24/7 Operation** **(26Apr2015)**
- **180,758 Datasets**
- **157 Beamline Users**
- **1,257 TB Data Stored**
- **>2.8 million Jobs at NERSC**

Portal interfaces for:

- Data monitoring & control
- Job monitoring
- Job inspection
- Data browsing and searching
- Browser-based remote viz
- … metadata, provenance, algorithms, theory-simulation, etc.

Remote experiments now a reality.

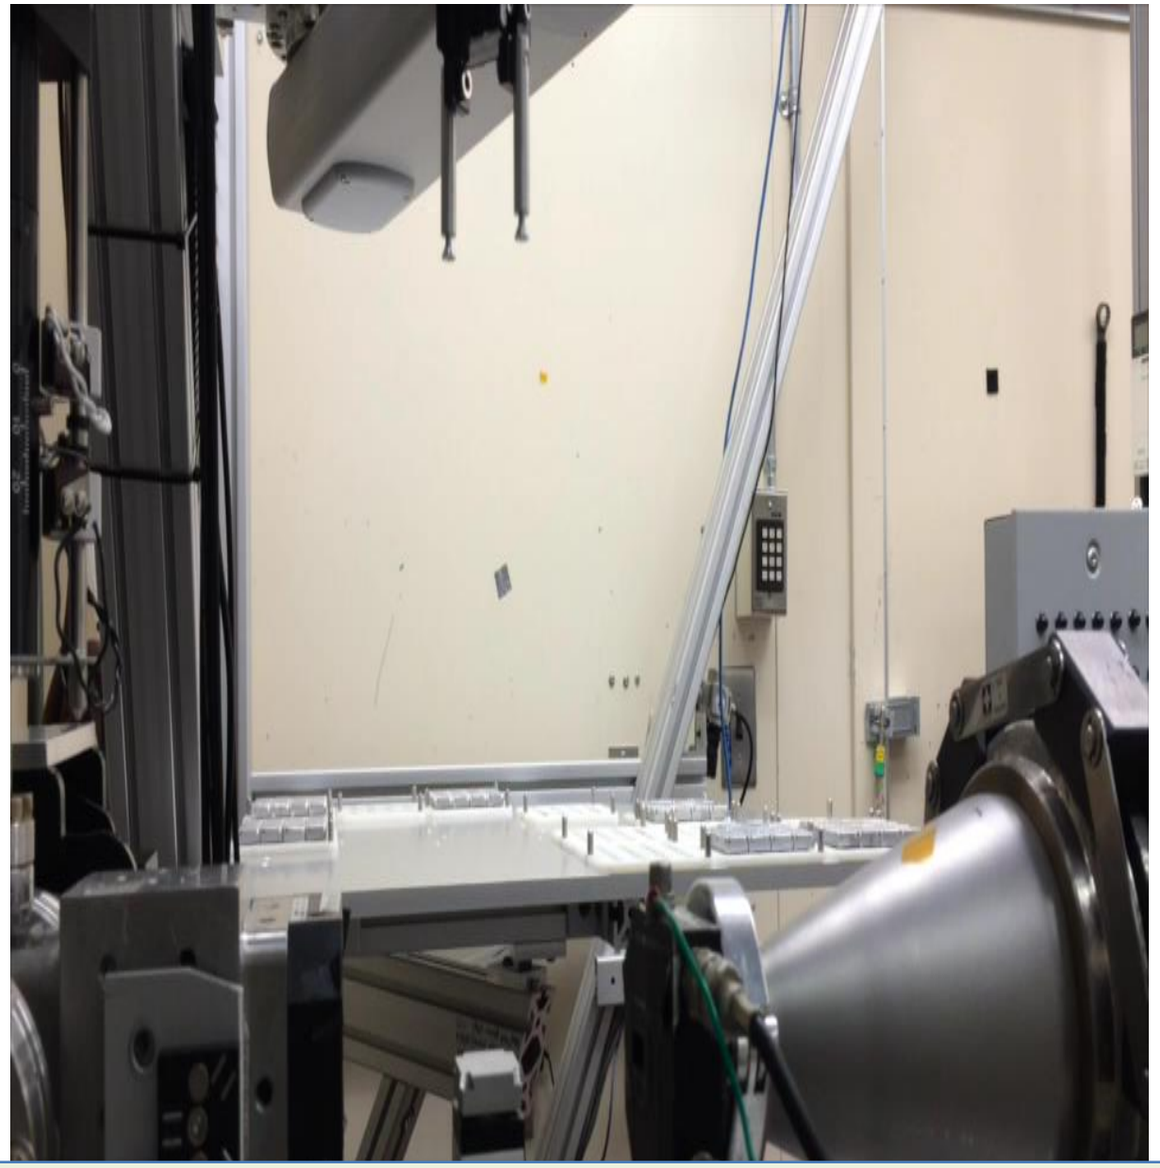

From: Alessandro Sepe as2237@cam.ac.uk -- **Actually, I did not feel any difference between a standard beamtime and this NERSC remotely accessed beamtime, which is quite an extraordinary result.**

Data Demos: Can this model be extended to illustrate the Super-Facility?

- **Printing BCP Organic PhotoVoltaics & Realtime Analysis**
  - **Advanced mathematical algorithms for data analysis and large-scale simulations coupled to data analysis are becoming critical to scientific understanding of BES lightsource data.**
  - **Many Reverse Monte Carlo (RMC) and simulation codes are HPC-scale, even Exa-scale, and heavily optimized for (eg. GPU) architectures.**
  - **We want to eliminate boundaries between the Scientist and the world's best Algorithms running on the best architecture for that code.**
- **Multi-Facility Imaging and Analysis Pipelines for Large X-ray Data**
  - **Replication at new lightsource facilities:**
    - **APS (tomography), LCLS (crystalography), NSLS (3D chemical map)**
  - **Evaluate common file formats & SW for Tomography**
  - **Integration of VisIt high quality, remote visualization**
  - **Integration of QuantCT filtering and segmentation**
  - **Integration of post-segmentation analysis results**

APS-NERSC Demo: Tomography, TomoPy, and Exchange HDF5

- Semiphore for data readiness; Spade(gridftp) in push mode
  - No suitcasing – Data Exchange format
- Wrap TomoPy; Digest DE format; Image/Volume displays unchanged
- Separate APS/ALS Disk & CPU allocations; Spade to bridge
- SPOT workflows for a) TomoPy & b) harvesting img/metadata
- I/O profiling & burst buffer performance tests

"Data sets from different user groups dealing with completely different topics - covering all possible data collection schemes and sizes."

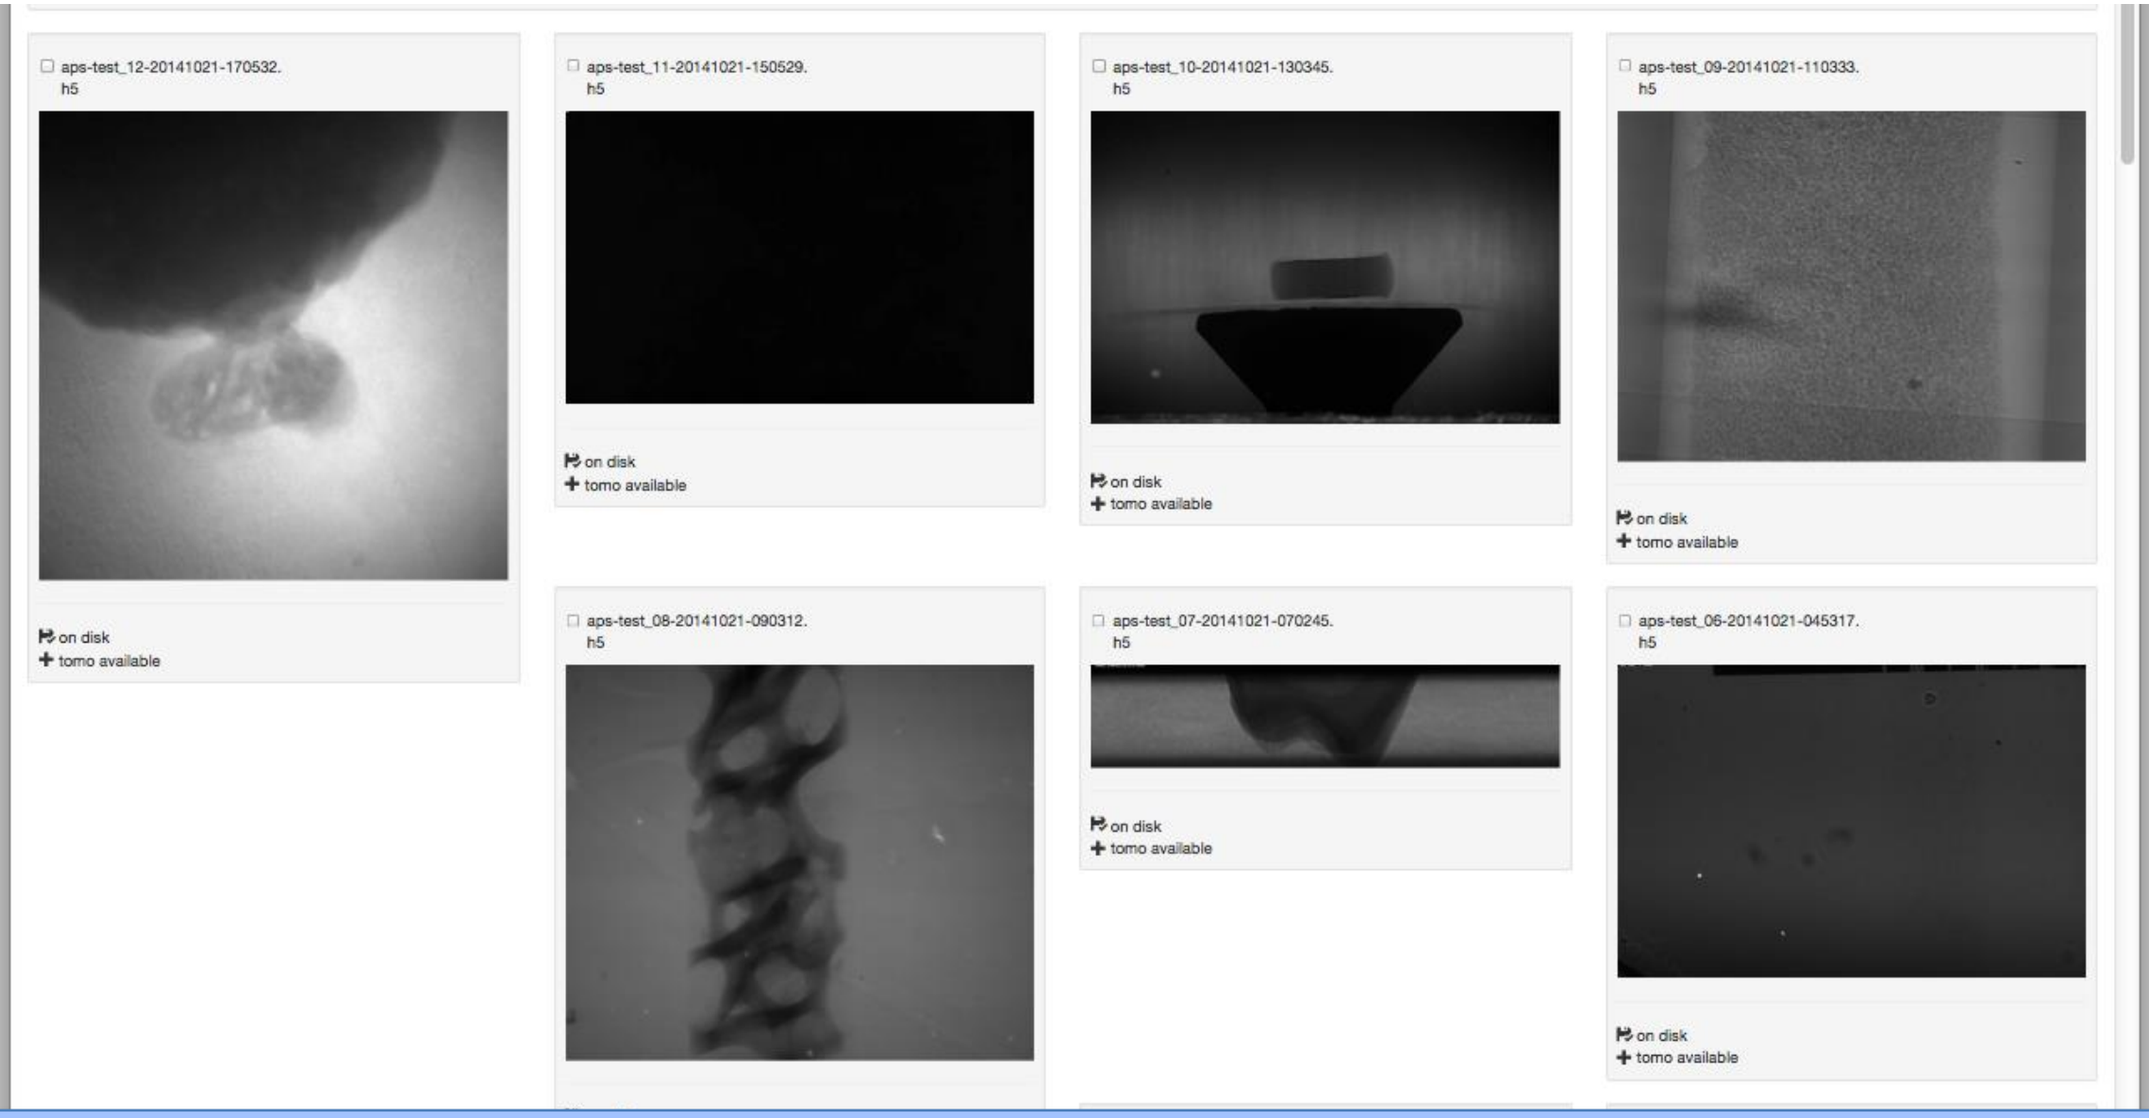

- "Co-designing Big Iron for Big Data" – Nick Wright, Chris Daley
  - Profiling TomoPy and testing effect of Burst Buffers on Alva
- 'Science DMZ as a Service - Creating Science Super-Facilities with GENI'

NSLS-NERSC Demo: 3-D Chemical Mapping using Transmission X-Rays

- Semiphore for data readiness; Spade(gridftp) in push mode
  - No suitcasing – Data Exchange format
- Port/Wrap 3DXANES; Adapt Image/Volume displays to 3 species
- Separate NSLS/ALS Disk & CPU allocations; Spade to bridge
- SPOT workflows for a) 3DXANES & b) harvesting img/metadata
- Ran **live** at SC14

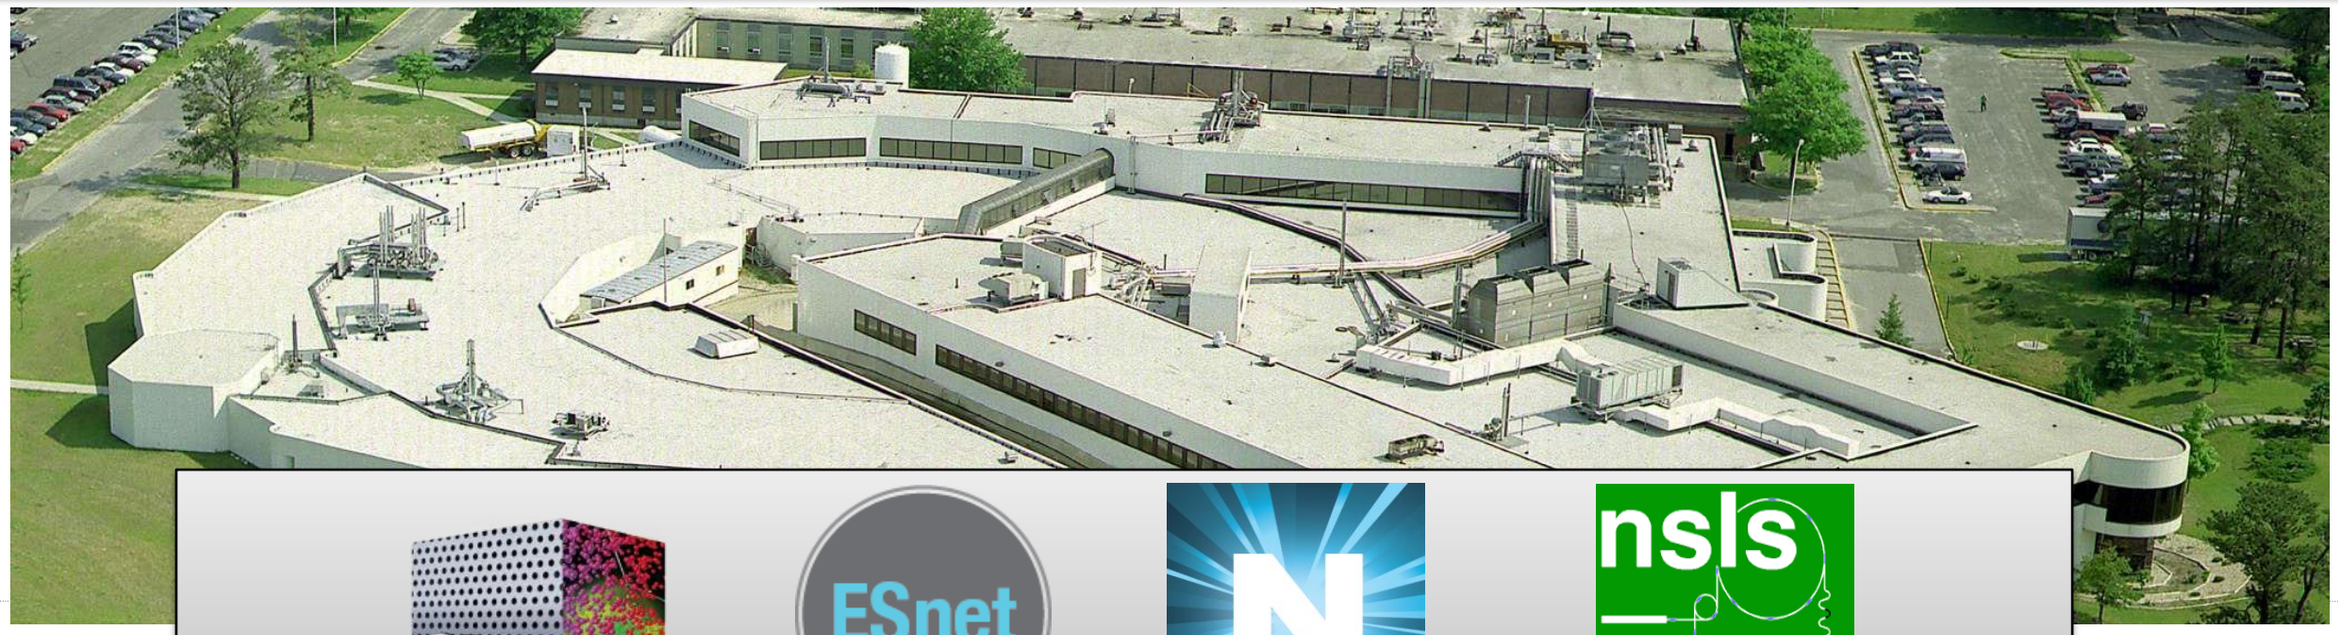

Running 3DXANES at NERSC required specialized visualization for 3D chemicals

- A hand-made sample with powders of LiFePO4, FePO4 and Fe
- 3D spectroscopic imaging:
  - A collection of 2D absorption contrast projections at different angles with a series of monochromatic X-ray, thereby generating a 4D data.
  - Focus on determining the chemical composition of each voxel position in a reconstructed 3D sample object.
- Workflow for each individual voxel position of the 3D sample:
  - Central shift correction: align the centers of projection images
  - Volumetric Reconstruction: recover 3D sample for each energy level
  - Sample Fitting: reassemble the 4D data to form XANES spectra over energy levels and perform least square fitting

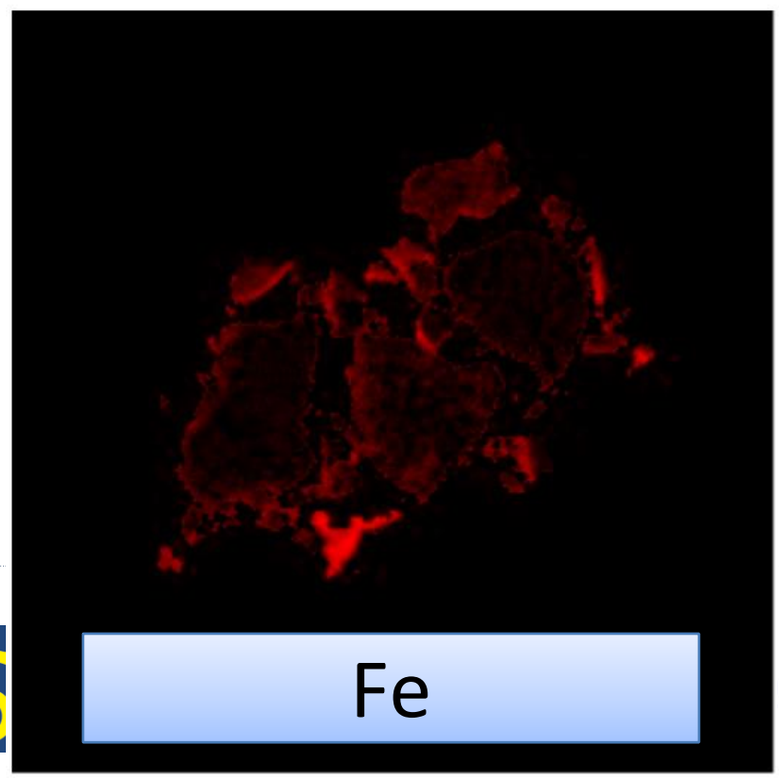

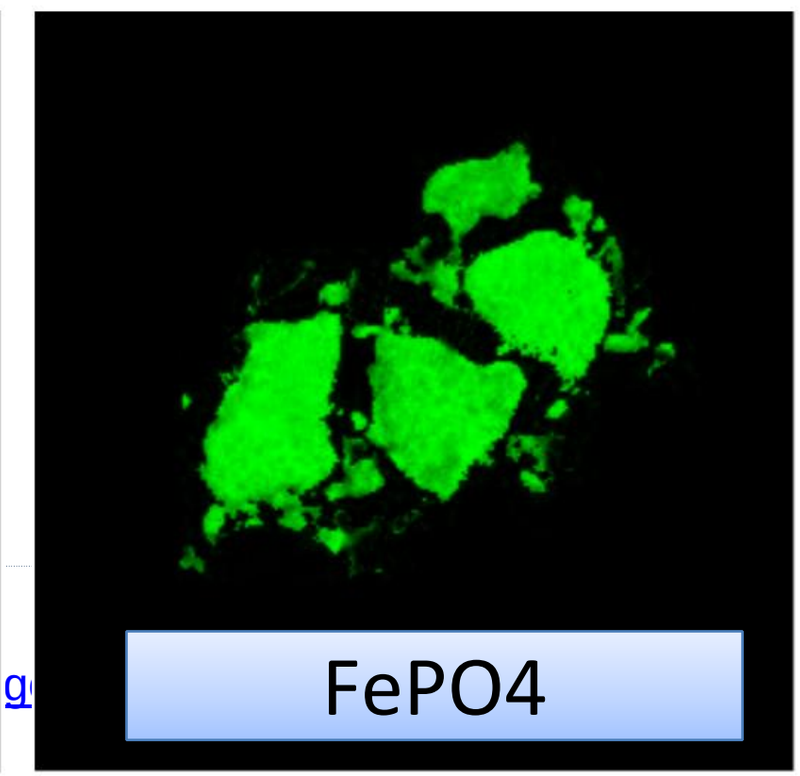

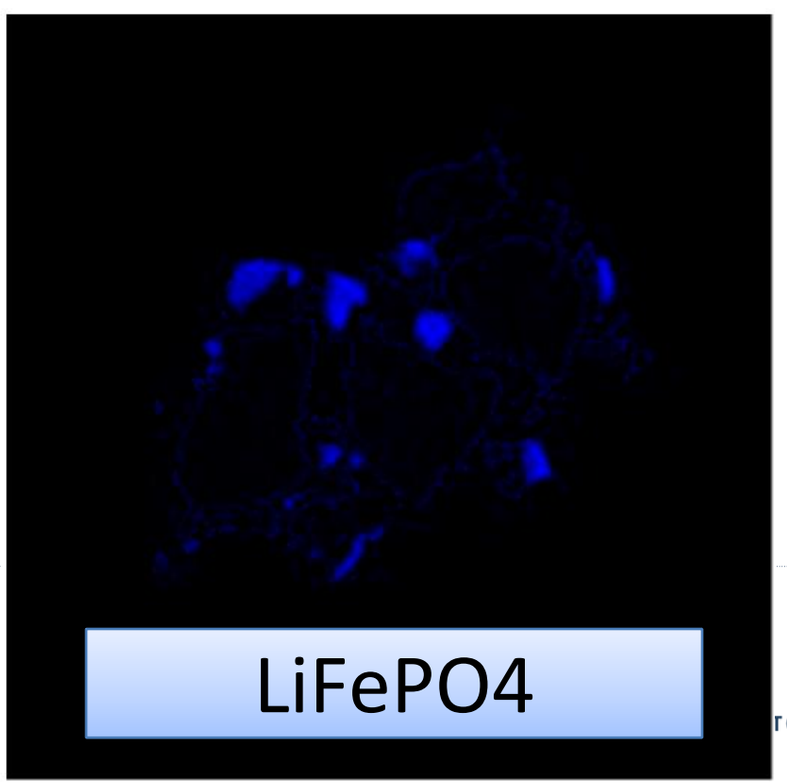

LCLS-NERSC Demo: Photosynthesis and nano-Crystallography

- SLAC/ESnet/NERSC network path; Round-robin DTN load balance
- REST interface for data readiness; Spade(bbcp) in pull mode
  - No suitcasing – raw XTC streams
- Port & wrap psana/cctbx; Digest XTC data; Adapt 2D image displays
- Separate LCLS/ALS Disk & CPU allocations; Spade to bridge
- SPOT workflows for a) Avg/Std/Max & b) cctbx
  - Fast (chunk0) & Full (all chunks)

“SPOT was like an extra pair of hands working in the background.” – N.Sauter

- June: 72 TB data @ CXI
  - Dry run challenges: XTC, psana, network & queuing
- July: 55 T
  - Full sy
    DAQ
    SPOT,

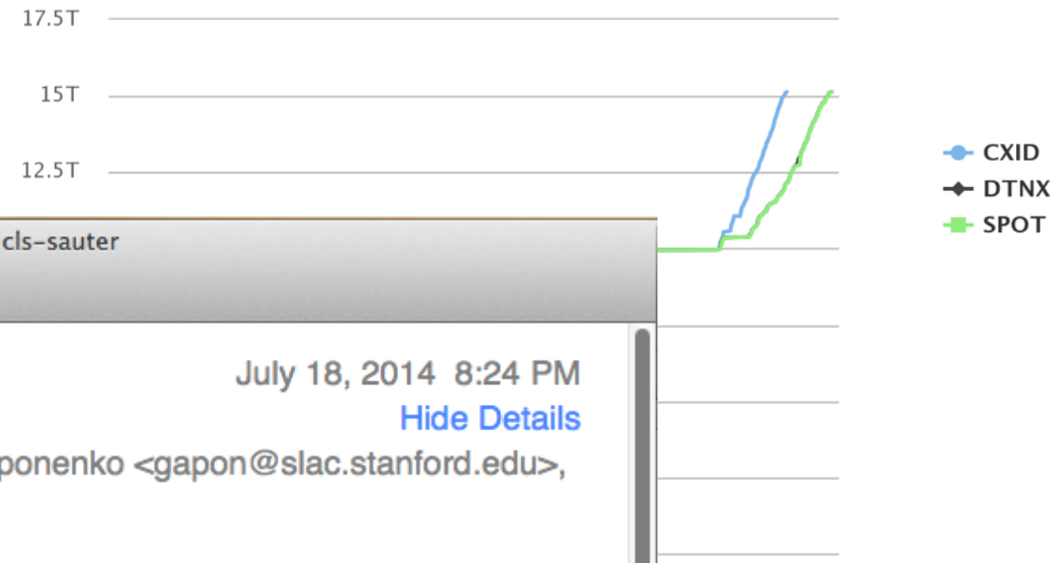

Re: data hiatus — 20140613-lcls-sauter

Aaron Brewster July 18, 2014 8:24 PM

To: Craig E. Tull <cetull@lbl.gov> Hide Details

Cc: Nicholas Sauter, O'Grady, Christopher P., Igor Al Gaponenko <gapon@slac.stanford.edu>, Amedeo Perazzo, spot-dev@lists.lbl.gov

Re: data hiatus

They put the tape on which I think damped down the signal.

Check out run 40. Data!

And success story. the first indication we had diffraction data was when run 40 popped up on SPOT. folks gathered around and pictures were taken on cell phones. yay :)

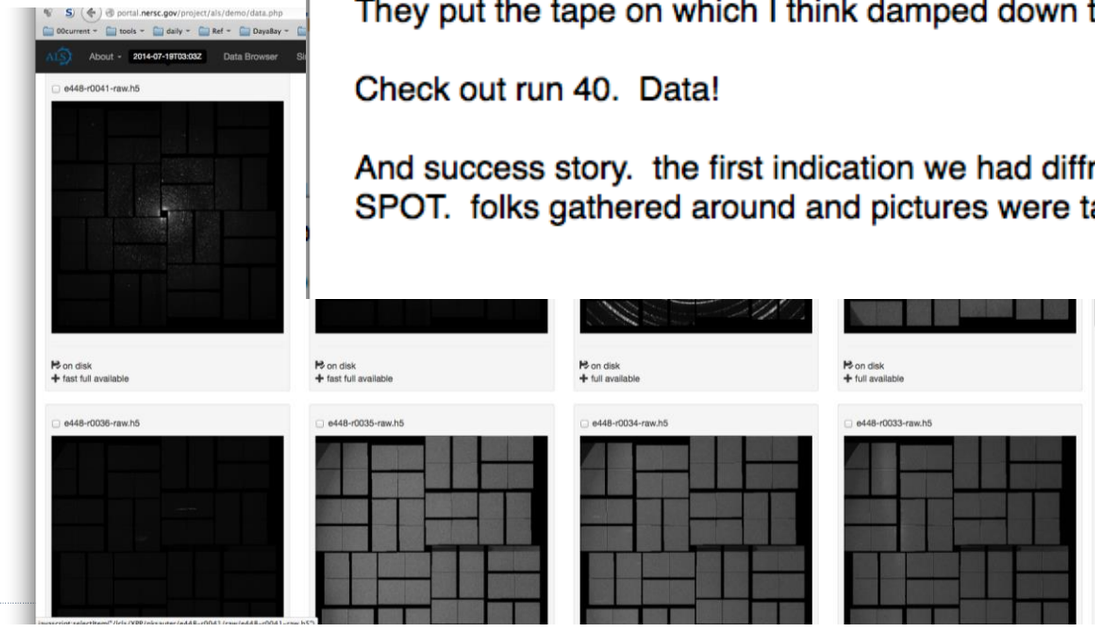

Recent work: LCLS Run – May 2015

- SLAC LCLS run @ XPP beamline:
  - 88 TB data generated in 5 days

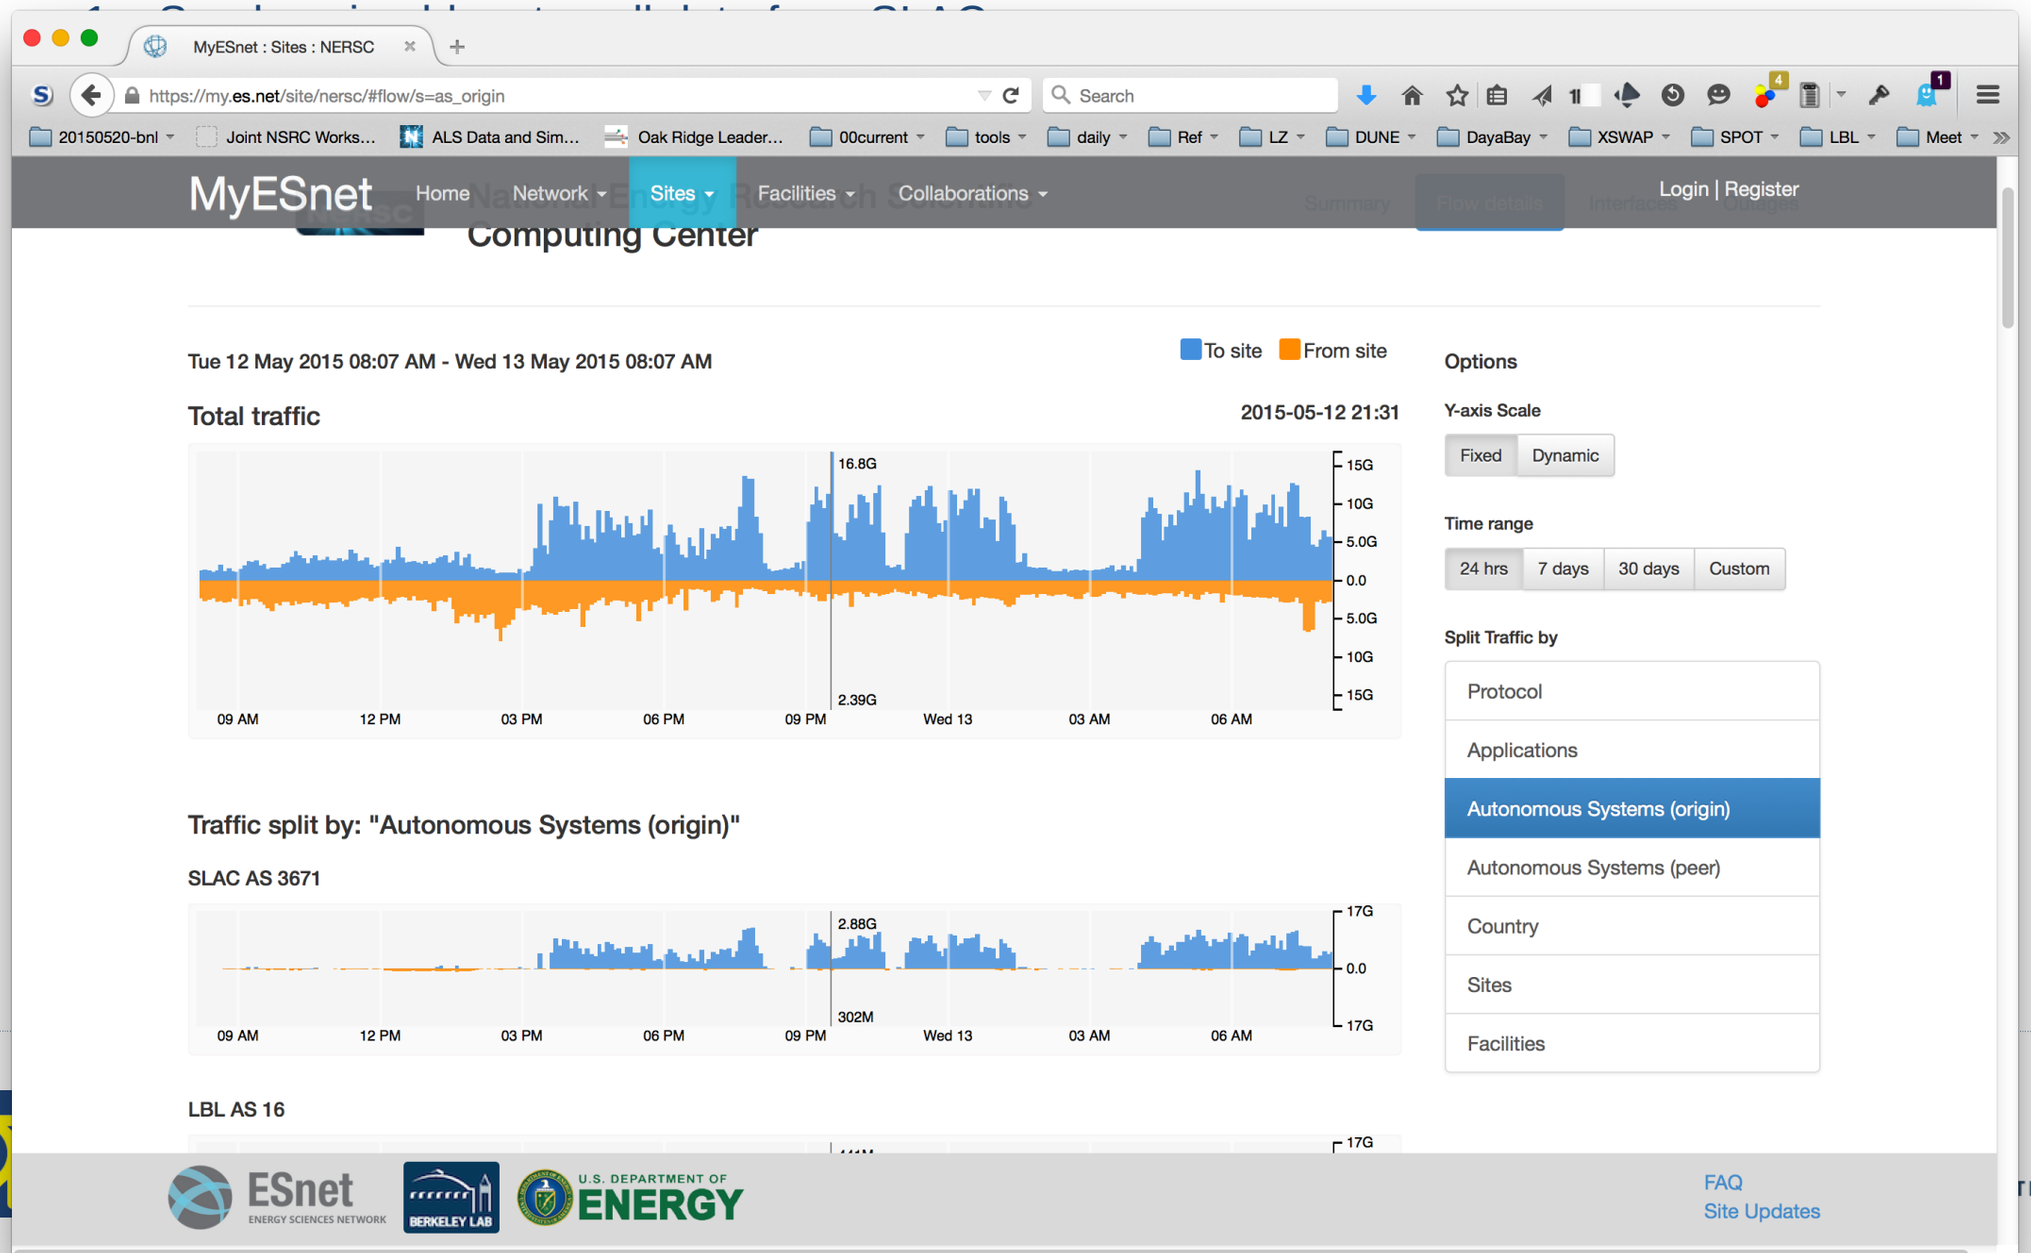

Conclusions of the Data Challenges

- Start with 24/7 “real-time” results for 2 ALS beamlines.
- In 9 months we were able to:
  - establish 3 new facility data pipelines in real-time or pseudo real-time
  - incorporate 3 new facility-developed codes (psana – LCLS, TomoPy – APS, 3DXANES – NSLS)
  - adopt 3 new data file formats
  - integrate Spade and Globus Online
  - interface our workflow to OLCF/Titan
- The Data Demos challenge accelerated our “Goals” by 2-3 years
  - But only in heroic demo mode!

ALS-NERSC-Titan Super-Facility Demo: Real-time analysis for Scattering Data

- Typical ALS chain + dpdak + HipRMC on NERSC reservation
- Co-schedule ALS beamtime & Titan reservation
- Typical ALS chain + dpdak => Spade => Globus Online => ORNL (FW)
- Sentry launches HipGISAXS – 8000 nodes and feeds results back IRT
- HipGISAXS runs **live** during beamtime (see next)
- Integrated Super-Facility Dashboard – **not** production quality

GISAXS Printing Demo Dashboard

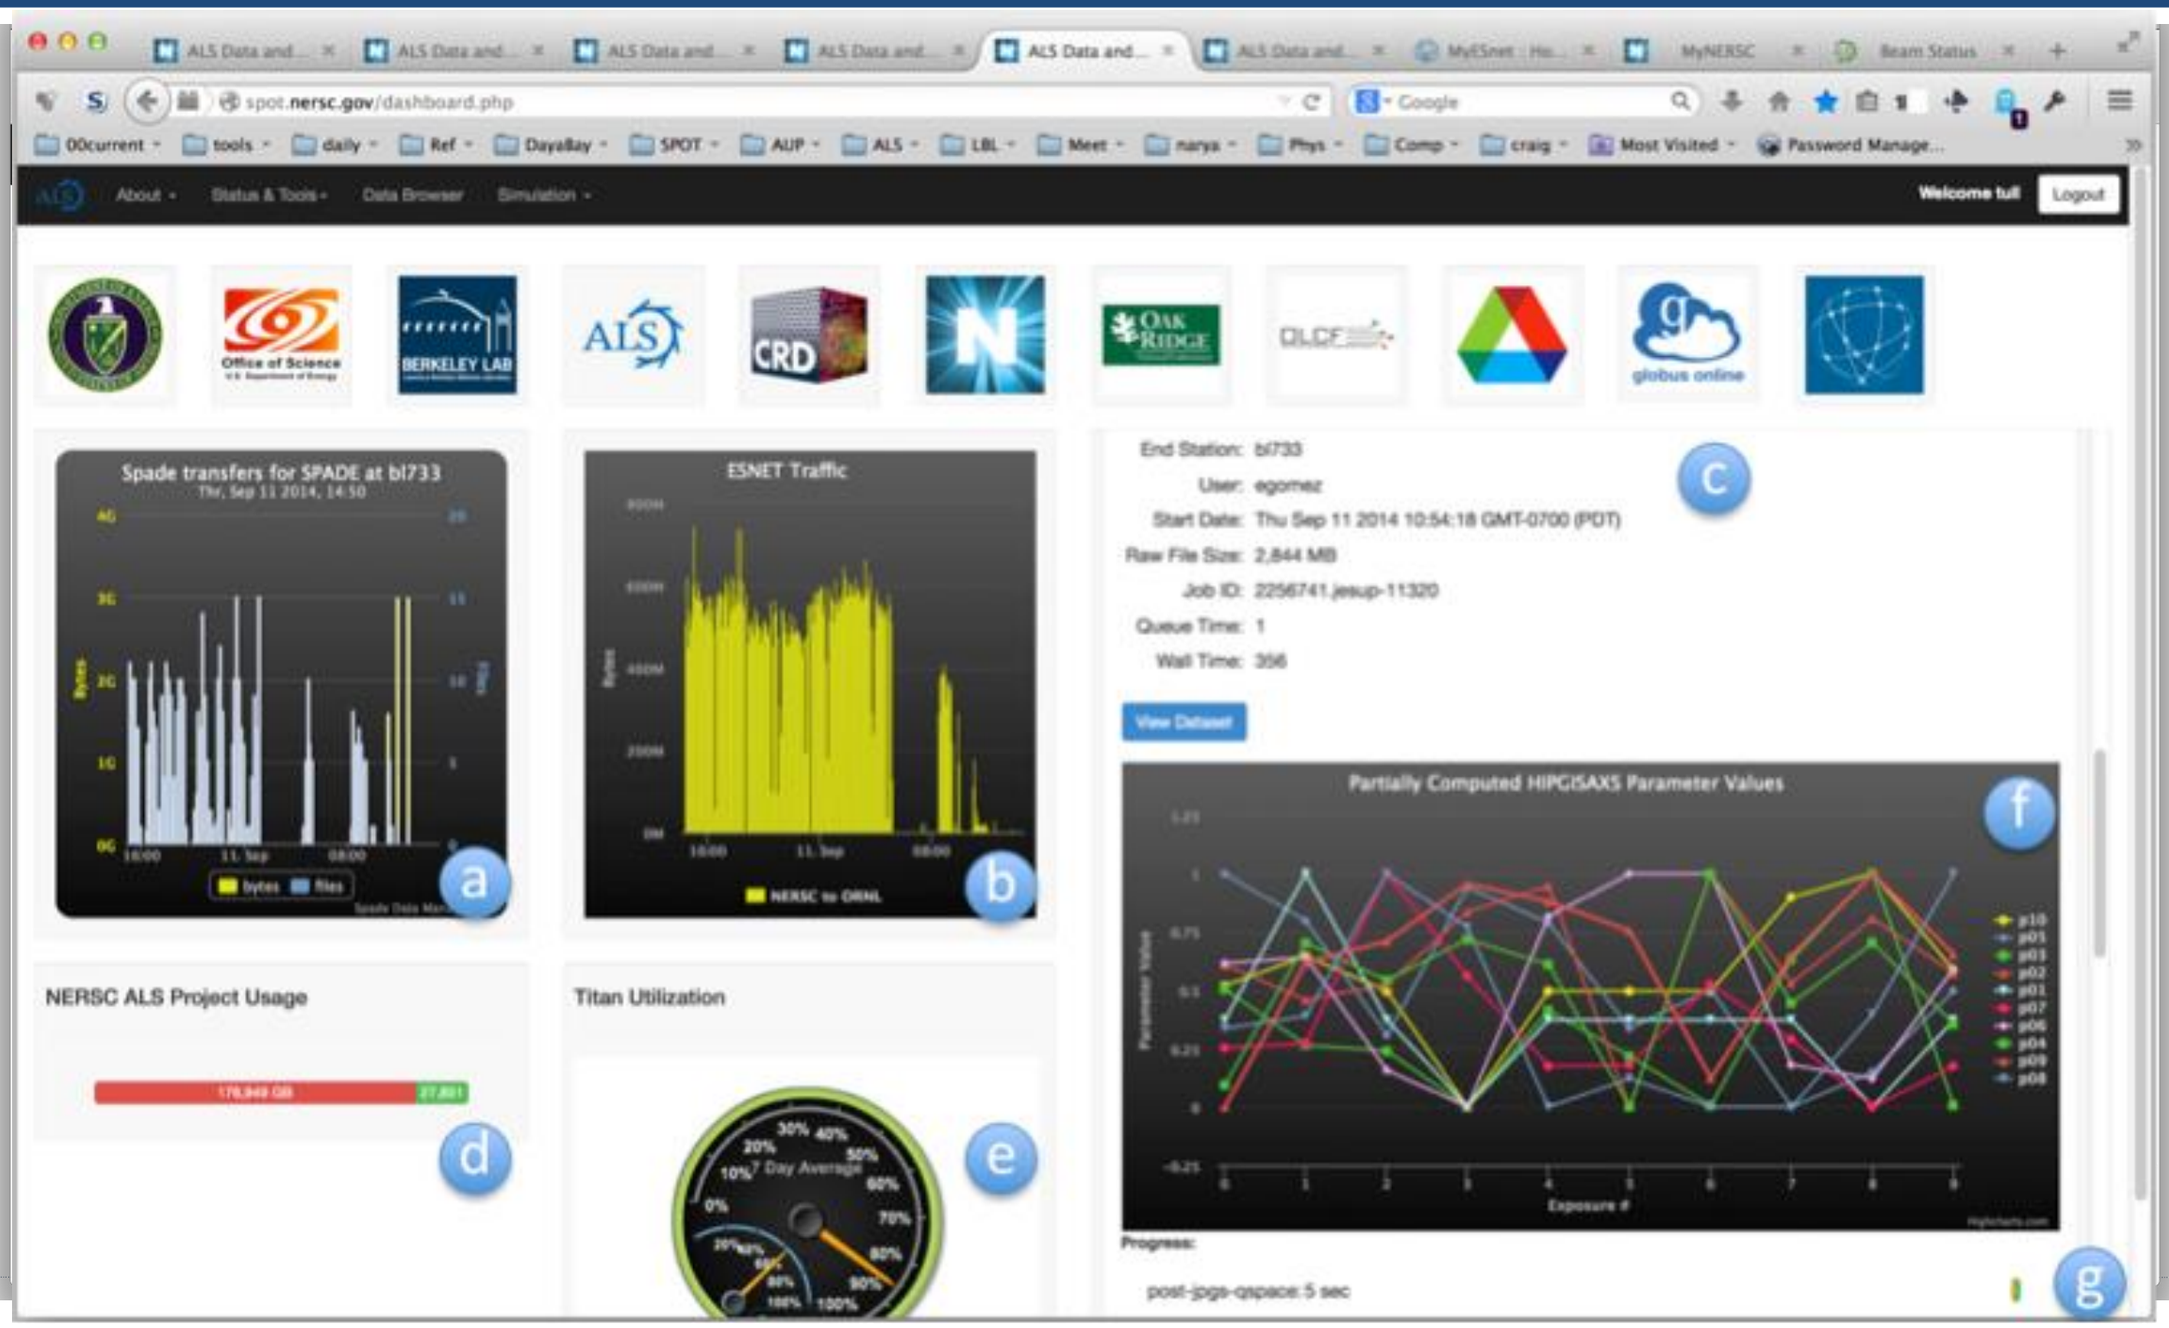

OPV Printing Demo – Sept 11, 2014

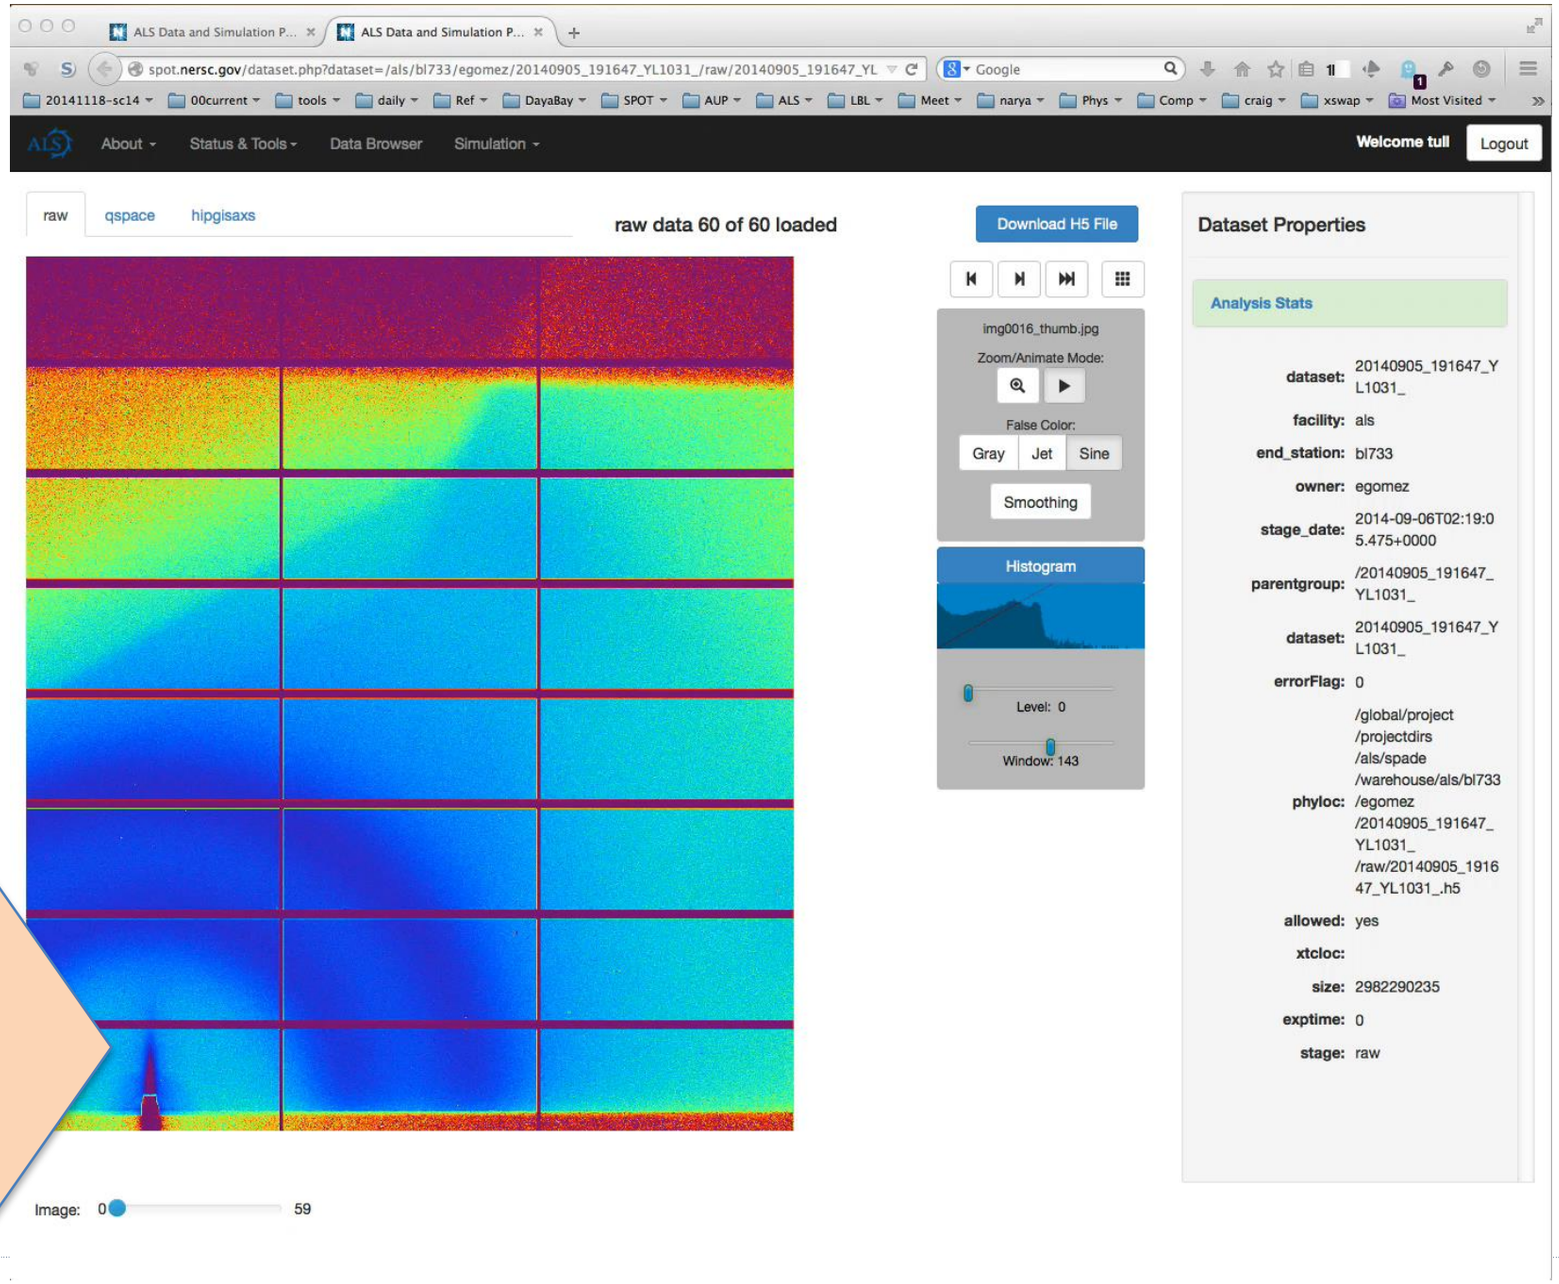

Structures evolve & develop during drying

Super-Facility Demo: Is real-time possible for SAXS/WAXS and GISAXS?

- Reverse Monte Carlo: (Titan & Edison)
  - Experiment: 1 frame per second + 600 frames total/sample + 15 min sample change = total of 25 min
  - Analysis: 12-20 minutes = Time to fit single frame per node
  - **Real-time is possible!**
- GISAXS with HIPGISAXS: (Titan: 1 K20X GPU/node)
  - Complexity of the material description led to intractably large calculations
  - Printing demo experiments created 36,000 frames in 3 days => requires 1/2 TITAN-year to analyze
  - **Need much better analysis and selection algorithms**

Photon Science is poised for a new, transformative era of discovery.

- BES Facilities are facing a data tsunami that will either bury us, or yield unprecedented scientific output.
- **BES Facility Data Demos – "extraordinarily ambitious"**
  - Loosely-coupled, workflow-driven architecture allowed agile integration of new formats, algorithms, visualizations.
  - Ability to quickly contribute attracted many participants.
- **Super-facilities are the necessary next step in the evolution of BES light-source facilities and must be strategically supported.**
  - **Delivering scientific insight for all lightsource users.**
  - **Concept generalizable to other science facilities.**

BES Facilities Data Demos Participants:

- ANL: APS, Globus*
  - Rachana Ananthakrishnan*, **Ian Foster***
  - **Francesco De Carlo**, Arthur Glowacki, Doga Gursoy, Nicholas Schwarz
- BNL: NSLS*, CSC
  - Wei Xu*, Shun Yao, **Dantong Yu**
- LBNL: ALS✞, CRD✪, ESnet✣, NERSC*, PBD/CCI
  - Aaron Brewster, Shane Canon*, Eli Dart✣, Jack Deslippe*, Abdellilah Essiari✪, Enrique Gomez✞, Alex Hexemer✞, Dinesh Kumar✞, Jon Dugan✣, Sherry Li✪, Feng Liu✞, Harinarayan Krishnan✪, Dmitry Morozov✪, Dula Parkinson✞, Simon Patton✪, Prabhat*, Thomas Russell✞, Abhinav Sarje✪, Polite Stewart✞, Taghrid Salmak✪, Nicholas Sauter, Eric Schaible✞, **Craig E. Tull✪**, Singanallur Venkatakrishnan✞, Chenhui Zhu✞, Daniela Ushizima✪
- ORNL: CADES*, OLCF
  - **Galen M. Shipman***, Don Maxwell, James Rogers, Jack C. Wells
- SLAC: LCLS*, SCA
  - **Amber Boehnlein**, Antonio Ceseracciu, Igor Gaponenko*, Wilko Kroeger, Christopher O'Grady, Amedeo Perazzo*

http://spot.nersc.gov/docs.php for references, papers, presentations.