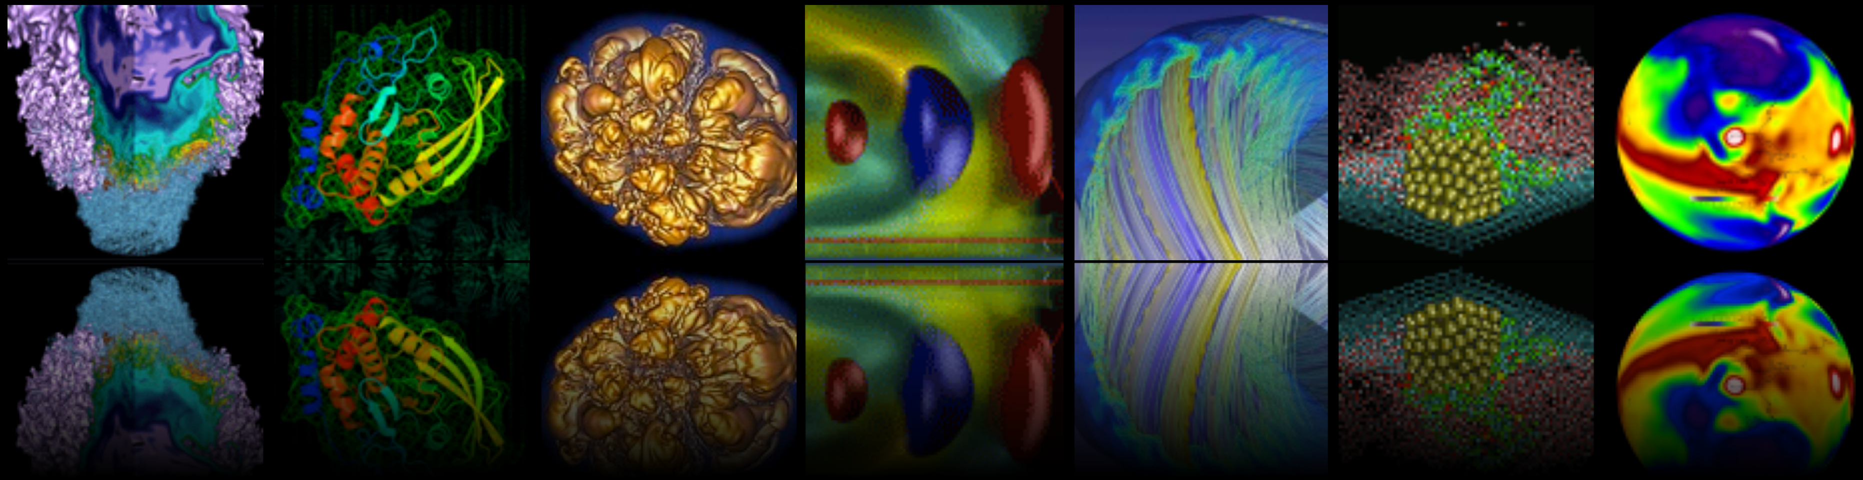

More Data, More Science, and … Moore's Law?

Kathy Yelick

Associate Laboratory Director for Computing Sciences
Lawrence Berkeley National Laboratory
Professor of Electrical Engineering and Computer Sciences
University of California at Berkeley

“Big Data” Changes Everything…What about Science?

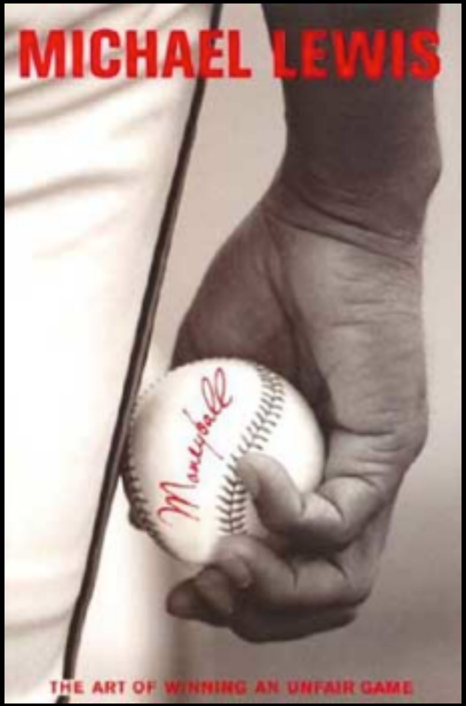

Combine simulation and observation for next Cosmology breakthrough

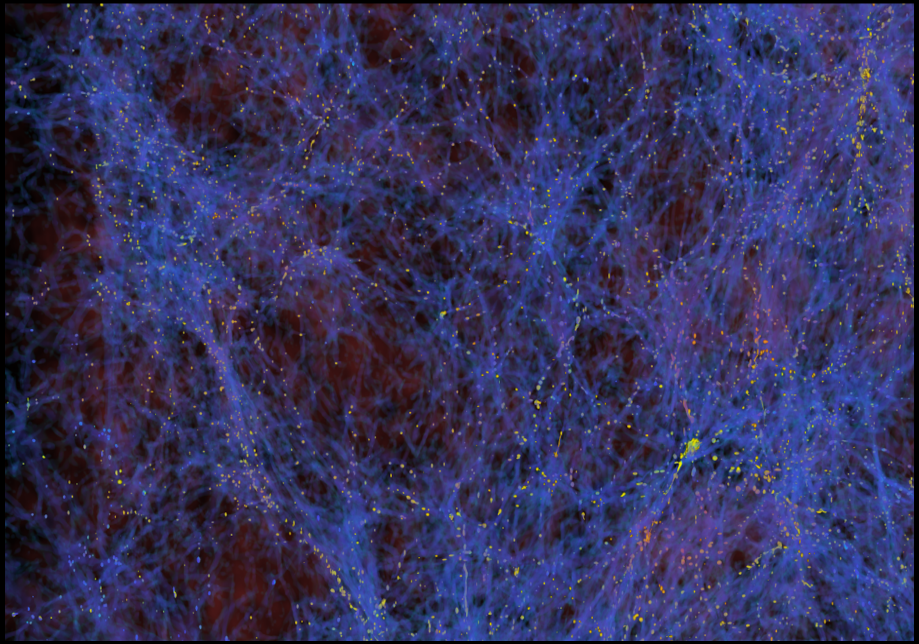

Nyx simulation of Lyman alpha forest using AMR

Kitt Peak National Observatory's Mayall 4-meter telescope, planned site of the DESI experiment

Reduce systematic bias in observation through simulation of ~1 Gigaparsec Baryon Acoustic Oscillations in the Lyman Alpha Forest and ~100 Gigaparsec simulation of galaxy clusters, both requiring adaptive mesh refinement (AMR).

Climate models and microbial analysis together to predict the future of the environment

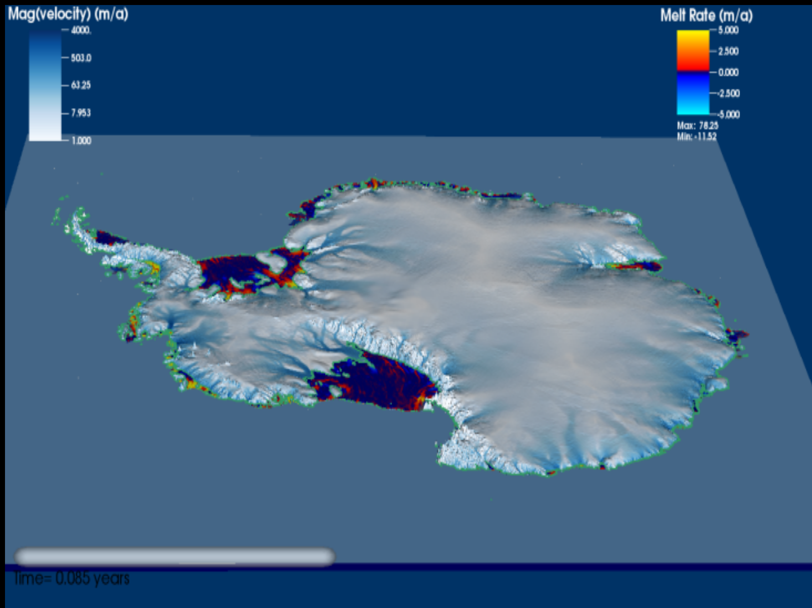

New climate modeling methods, including AMR “Dycore” produce new understanding of ice

Genomes to watersheds Scientific Focus Area

Understand interactions between environmental microbiomes and climate change with *kilometer resolution models* that track dynamic 3D features (with AMR) and *genome-enabled analysis* of environmental sensors.

Understand and control energy with advanced light sources and materials modeling

Materials Project

13,030 users hosted at NERSC with software co
developed by CRD

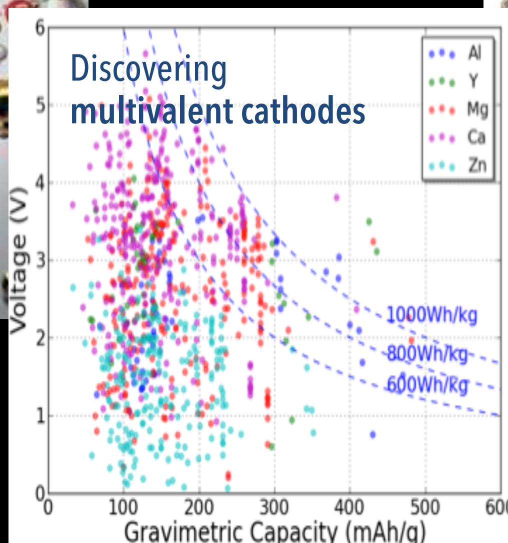

ALS-U Upgrade

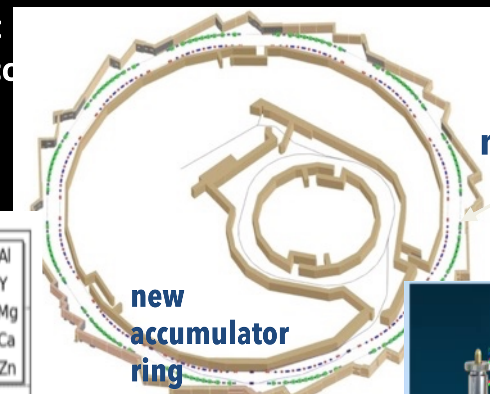

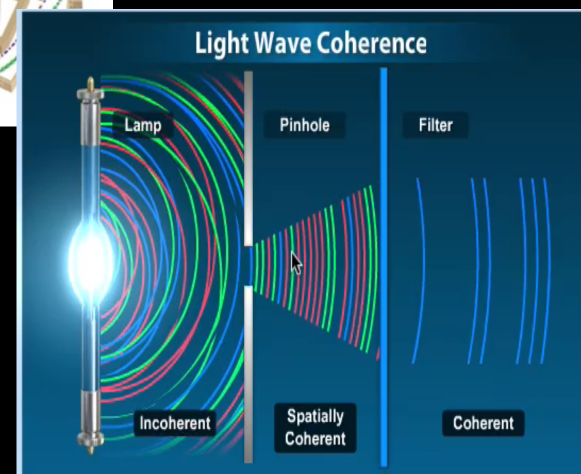

Understand and control the direction and flow of energy with minimal losses using *advanced instruments*, *high fidelity models*, and high throughput simulation and analysis for applications in energy, environment and computing,

Old School Scientific Workflow

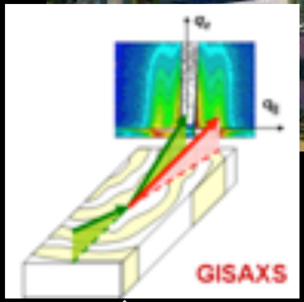

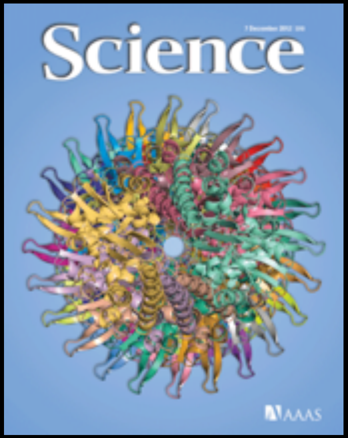

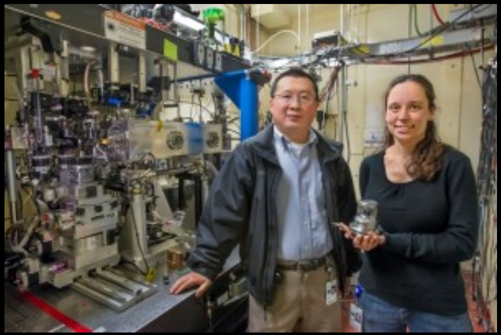

Computing, experiments, networking and expertise in a "Superfacility" for Science

Liu et al, "Fast printing and in situ morphology …". Adv Mater. 2015

Old School HPC: only for Simulation

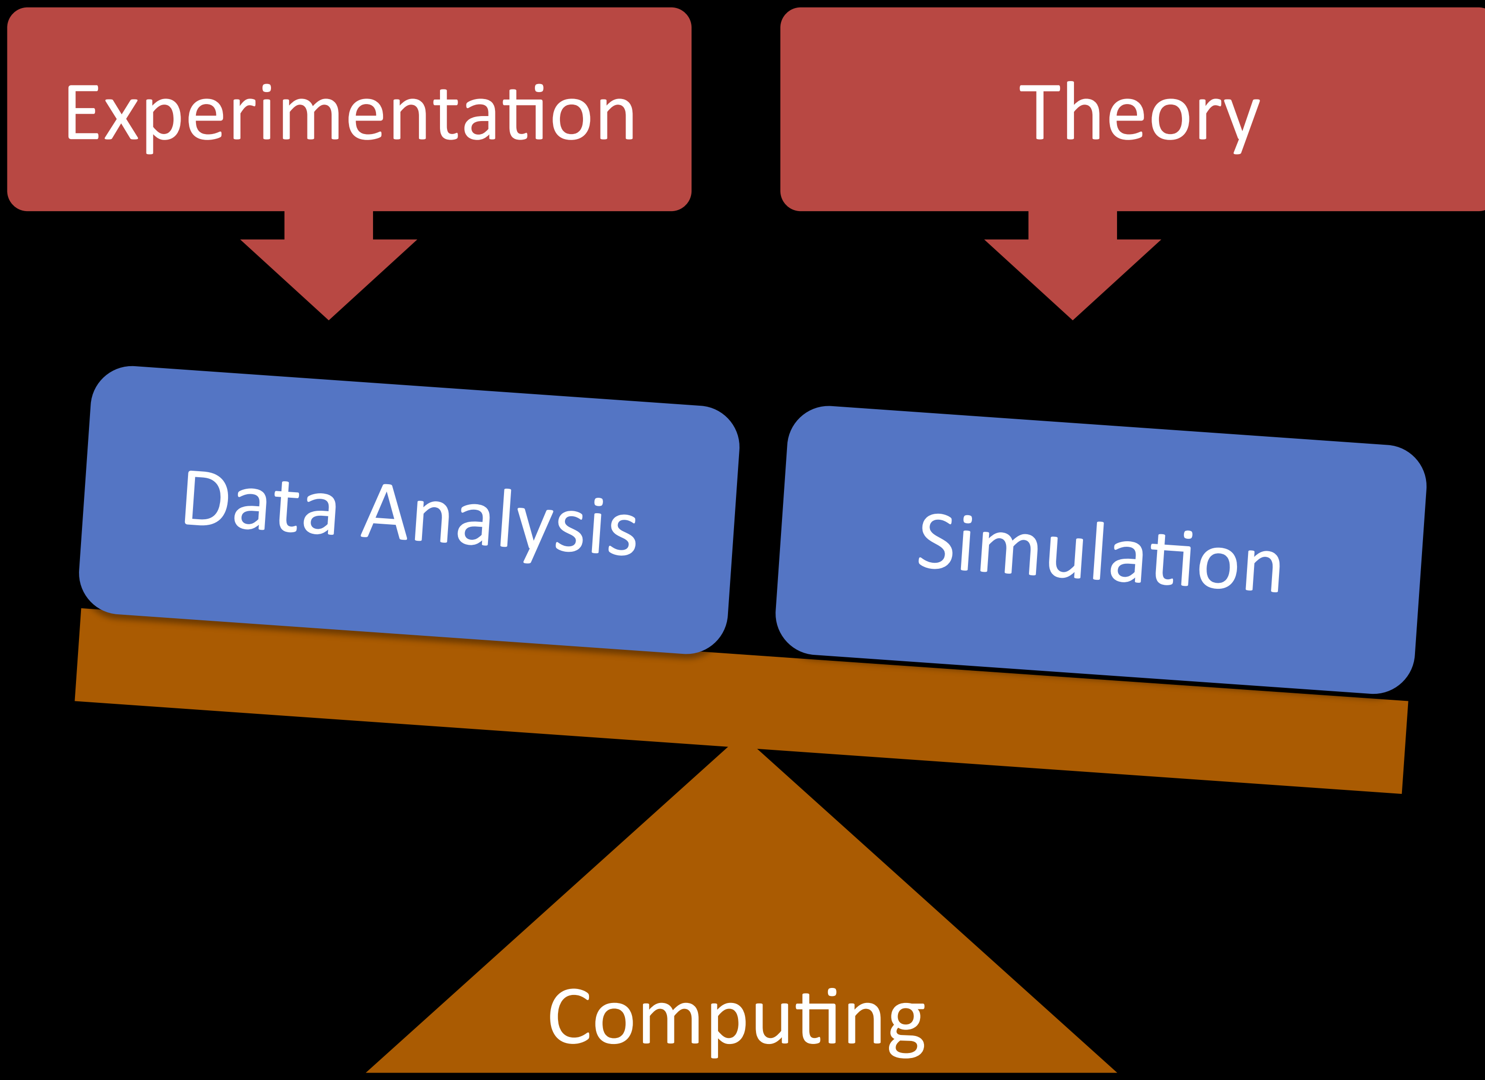

HPC is equally important in experimentation

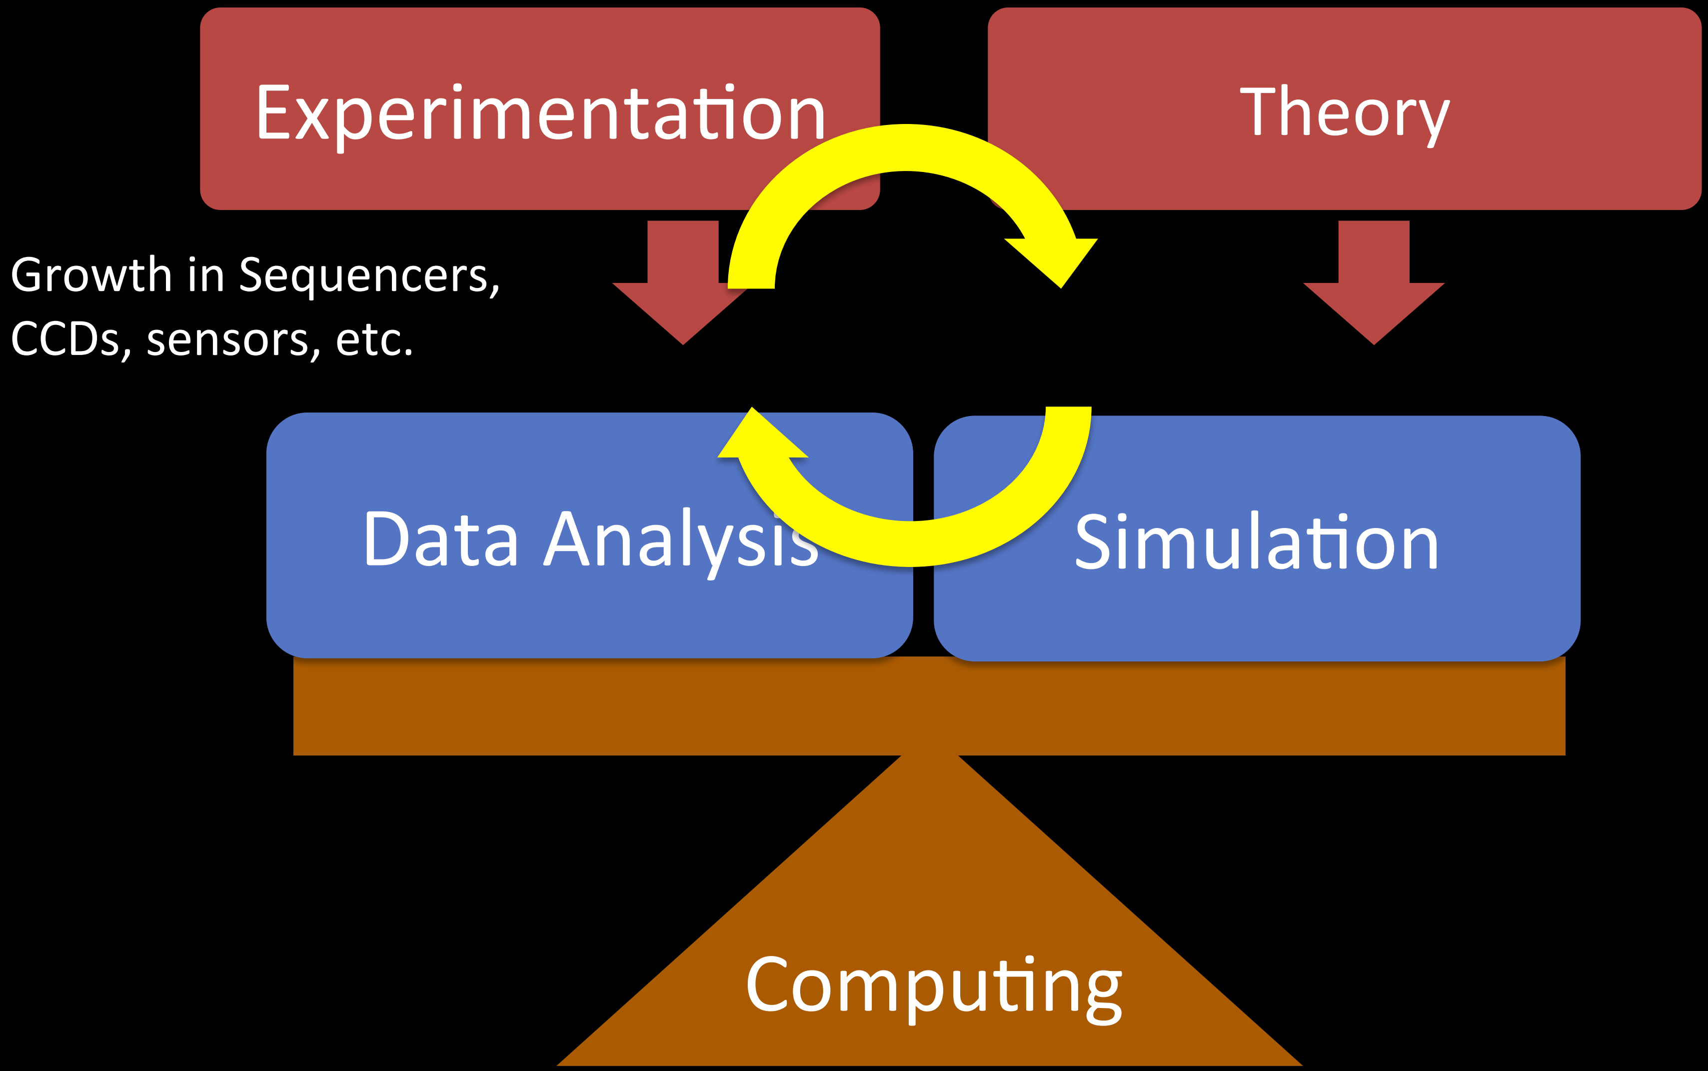

Questions?

1. **Are there MSU examples of "science at the boundary" of simulation and observation?**
   – How should you take advantage of these opportunities?

Part 2

The Data Tsunami

Science Data Growth is Outpacing Computing

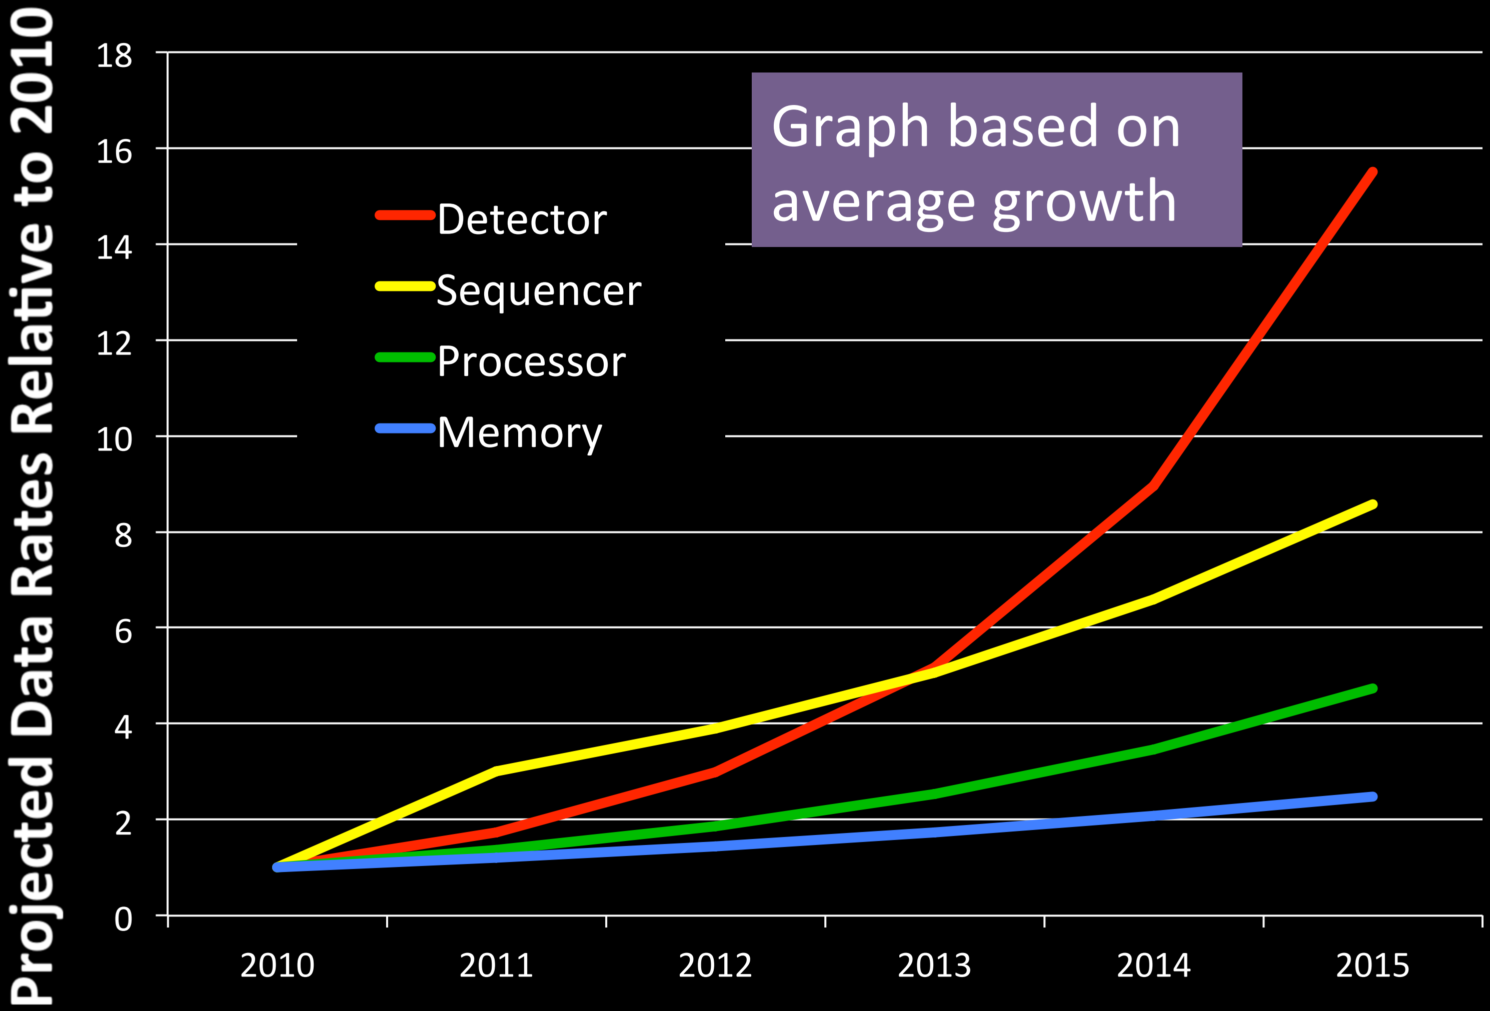

Old School Scientific Data Search

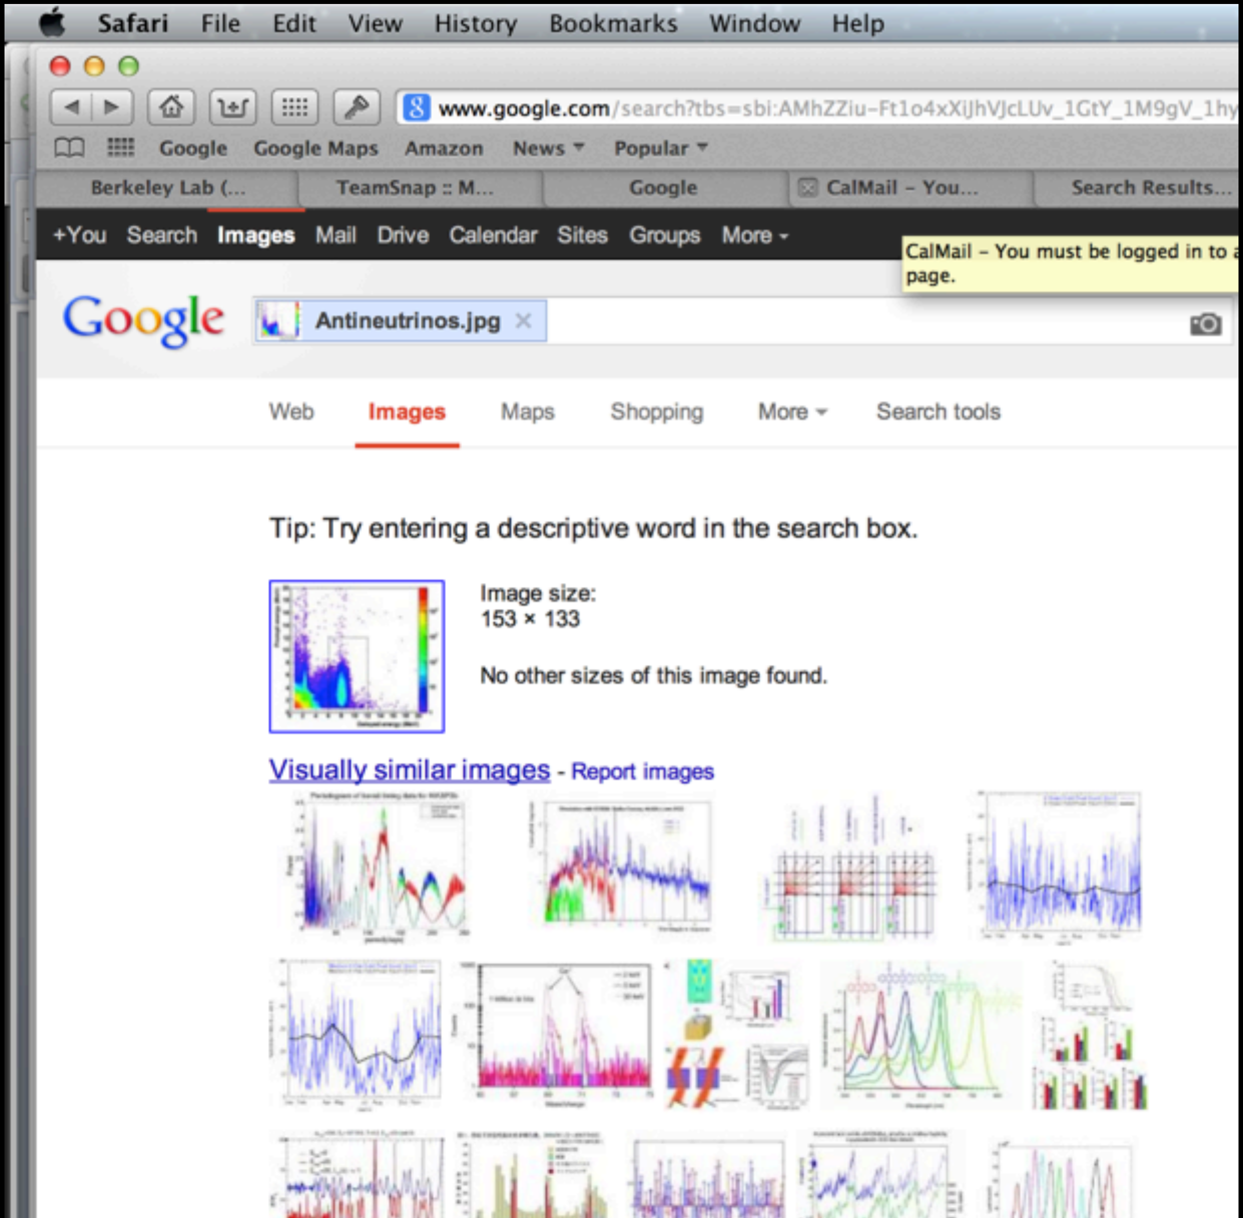

Automated Search, Meta-Data Analysis, and On-Demand Simulation

Automated metadata extraction using machine learning

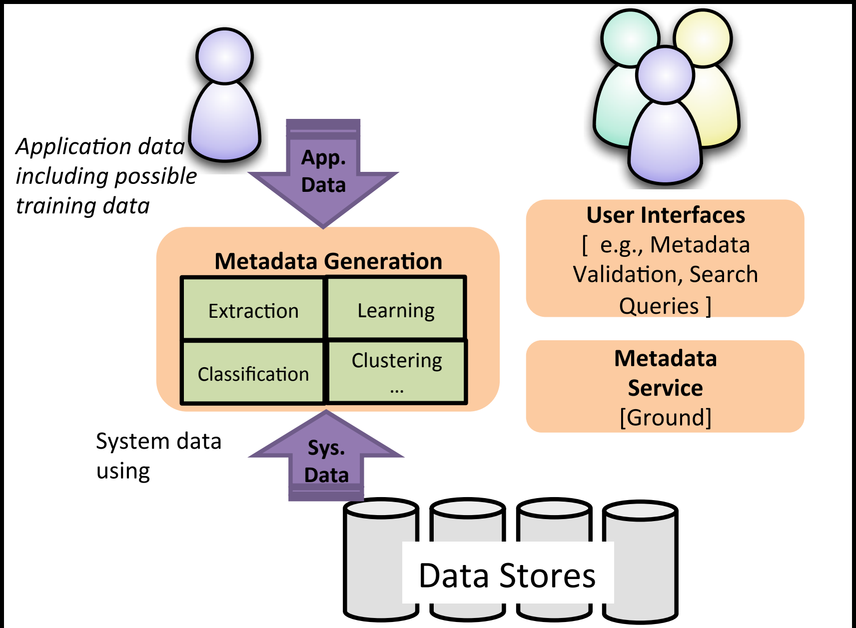

Jobs submitted by “bots” based on queries; algorithms extract informatics for design

Questions?

2. **What are the largest and most complex sources of research data at MSU?**
   – What types of data/CS/math/stat challenges are there?

Part 3

Networking and Computing Facilities Need to Adapt

ESnet: Exponential growth in networking

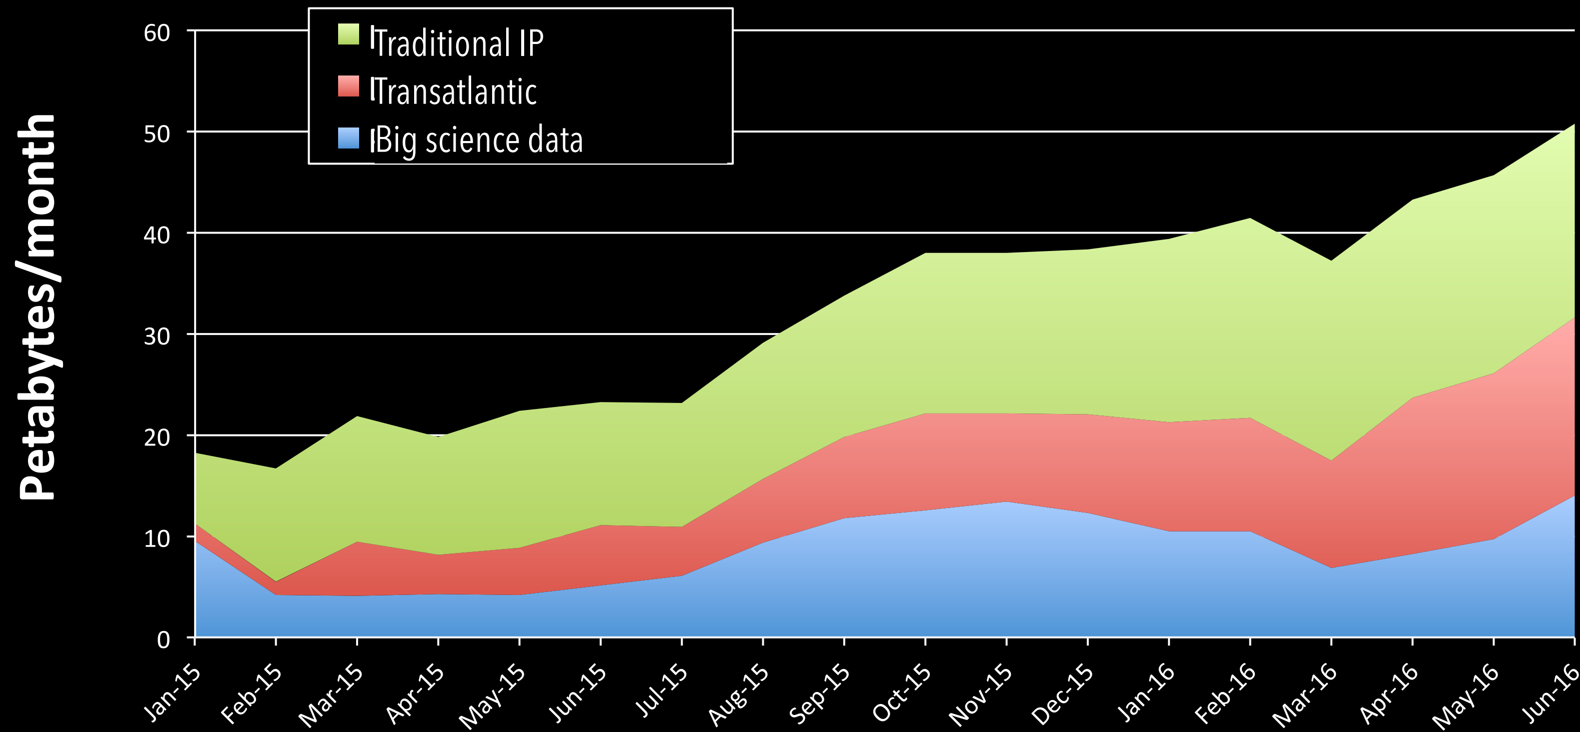

100 Exabytes/year by 2024!

Science DMZ to deliver bandwidth to the end users

OSCARS for bandwidth reservation

ESnet: Discovery Unconstrained by Geography

SLAC at Stanford

SLAC at Stanford

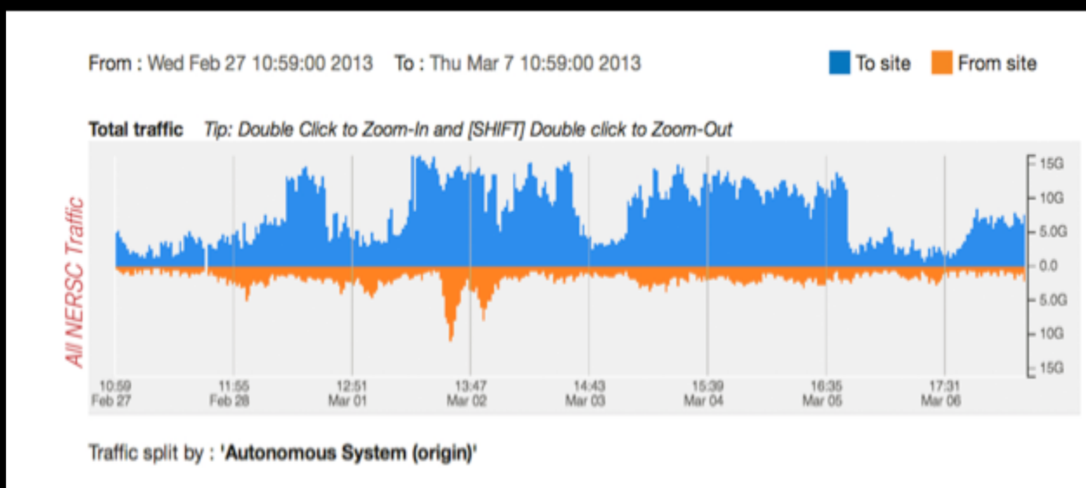

LCLS/NERSC/Esnet Superfacility demo for Photosystem II

→ 3x network traffic

Network performance enables efficiency of centralized computing

Systems configured for data-intensive science

NERSC Cori has data partition (Phase 1, Haswell) pre-exascale (Phase 2, KNL preproduction)
WAN-to-Cori optimized for streaming data: 100x faster from LCLS to Cori and Globus to CERN

Real-time queue prototyped at NERSC

- In 1998 dedicated hardware; now prototype queue on Cori
- <1% of NERSC allocation
- Cryo-Em, Mass spec, Telescopes, Accelerator, Light sources

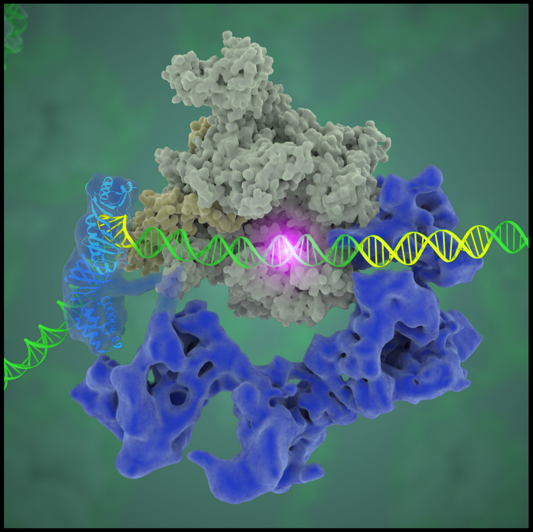

Cryo-EM: Image classification
Nogales Lab

PTF: Image subtraction pipeline

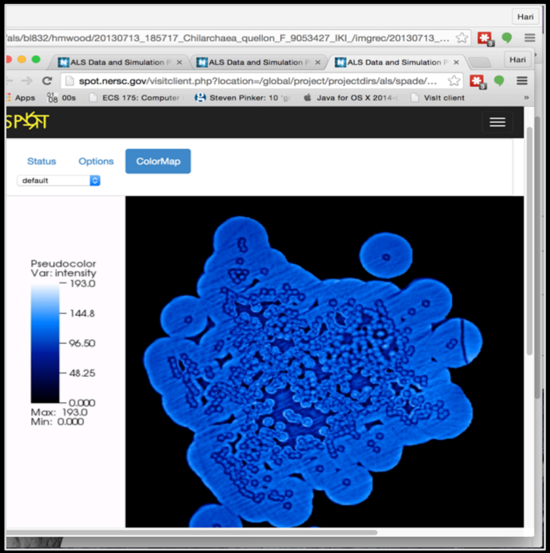

ALS: 3D Reconstruction,
rendered on SPOT web portal

Containers for HPC Systems

- Data analysis pipelines are often large, complex software stacks
- NERSC Shifter (with Cray), supports containers for HPC systems
- Used in HEP and NP projects
  (ATLAS, ALICE, STAR, LSST, DESI)

SHIFTER docker

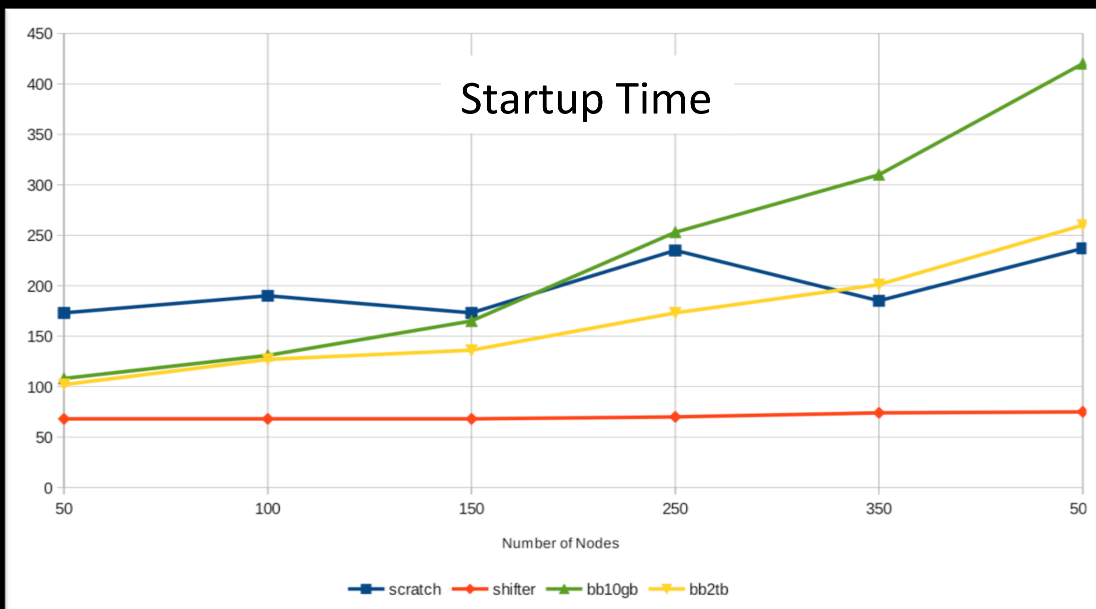


The Register
Biting the hand that feeds IT
DATA CENTER SOFTWARE NETWORKS SECURITY INFRASTRUCTURE DEVOPS BUSINESS HARDWA

Data Center ▸ HPC

Cray hoists Docker containers into supercomputers

Productivity gains without performance hits

18 Nov 2015 at 00:01, Drew Cullen 25 35

Questions?

3. **How should undergrad/grad programs be adapted to address data challenges in future careers?**
   – New courses, (joint) majors, research institutes?

Part 3

Computing, Mathematics and Statistics Research Challenges

CAMERA: Math for the Facilities

James Sethian, PI

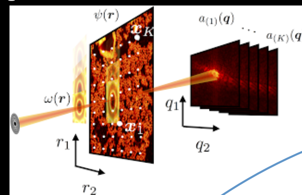

Designing mathematical algorithms to allow real-time analysis next to the equipment

Real-time streaming ptychography—ALS, delivered to NSLS2, LANL, BESSY,

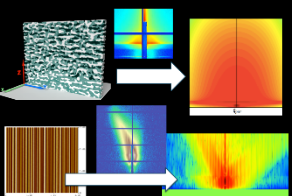

SFM/TEM + GISAXS

Multi-modal: Building the math that fuses information from multiple experiments

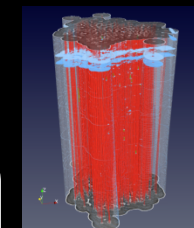

New algorithms to transform manual into automatic analysis

Automatic image processing for the ALS/GE

CAMERA workshop on Tomography: Joint with APS, ESRF, DIAMOND, LNLS, LLNL, SSRL,….,

Compare and integrate multiple analysis tools

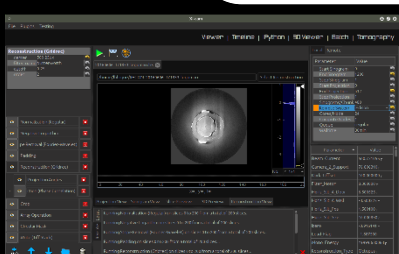

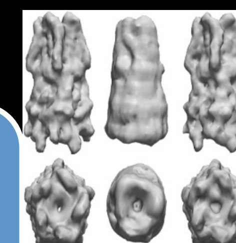

Inventing new math and models to match new acquisition technologies

Fluctuation scattering and single particle imaging for the LCLS

Cultural and Sociological Challenges

Robust and reliable codes and data flow: workflow environments

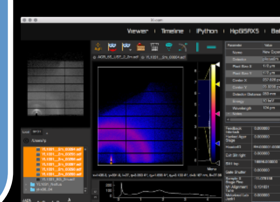

Workflow and access to remote supercomputers: XiCAM for ALS, SSRL, APS, NSLS2

Analytics vs. Simulation Kernels:

7 Giants of Data	7 Dwarfs of Simulation
Basic statistics	Monte Carlo methods
Generalized N-Body	Particle methods
Graph-theory	Unstructured meshes
Linear algebra	Dense Linear Algebra
Optimizations	Sparse Linear Algebra
Integrations	Spectral methods
Alignment	Structured Meshes

Machine Learning Mapping to Linear Algebra

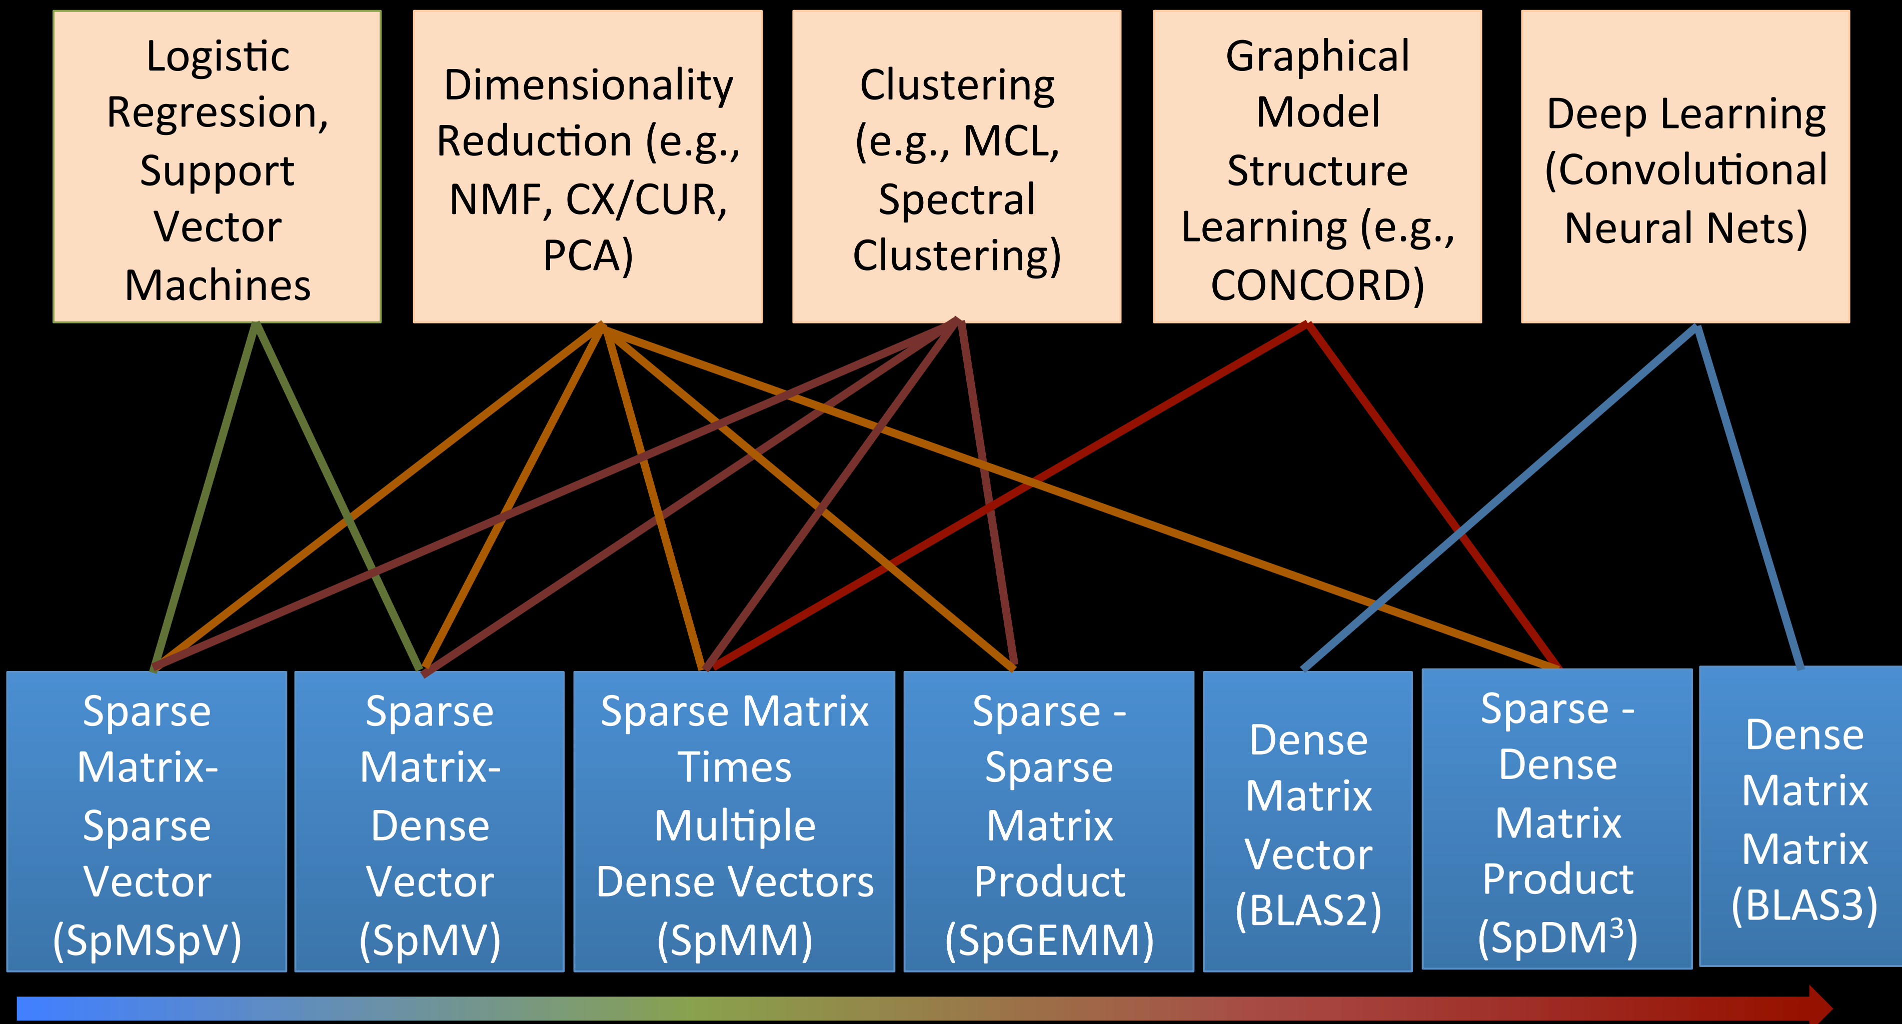

Aydin Buluc, Sang Oh, John Gilbert, Kathy Yelick

Challenge: Communication is expensive

Communication is expensive in time and energy

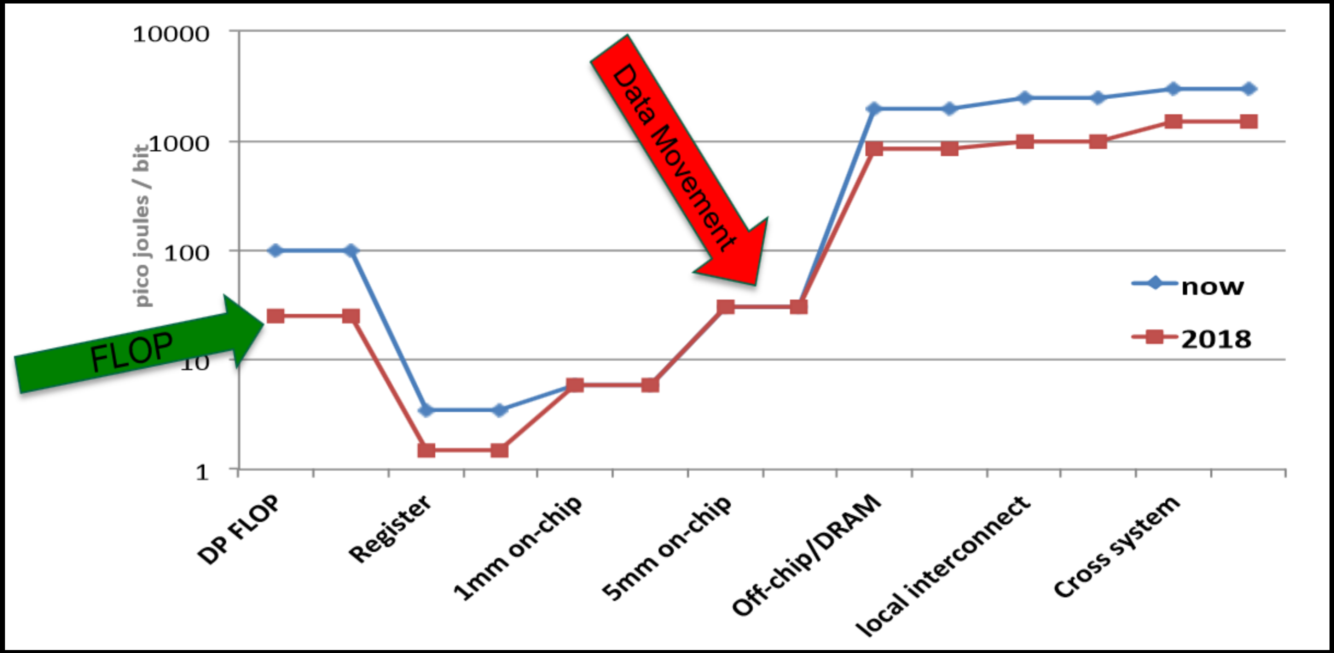

Hard to change: Latency is physics; bandwidth is money!

Communication-Avoiding Algorithms

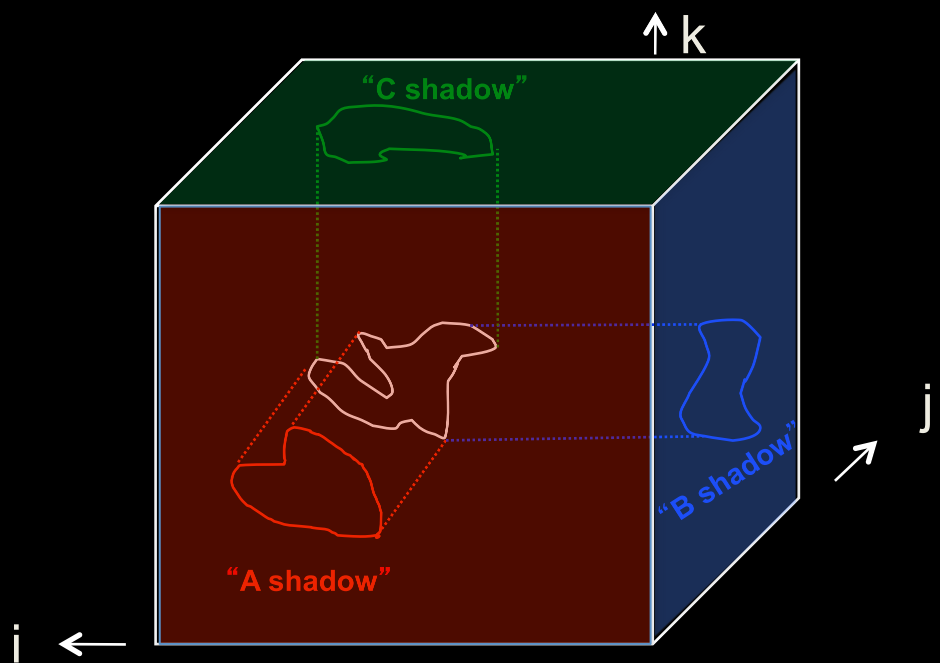

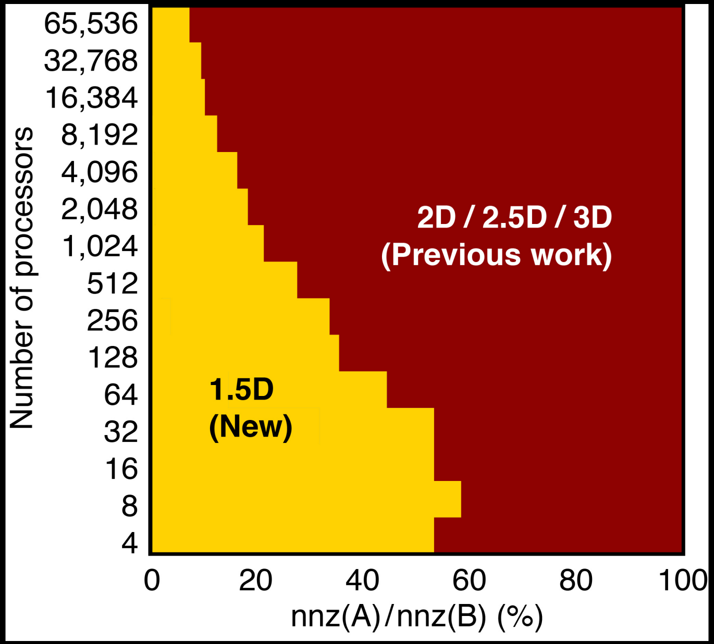

Matrix Multiplication code has a 3D iteration space; each point is a */+

Model for choosing communication-optimal algorithms for sparse matrices

for i
for j
for k C[i,j] … A[i,k] … B[k,j] …

Demmel et al on LA; Christ et al generalization

Koanantakool & Yelick

Interactive Analytics using Jupyter

```
In [10]: # overlaying the small H&E and MS images

registered_ms_image = ird.transform_img_dict(my_images[2], result)
big_registered_ms_image = imresize(registered_ms_image, optical_image.shape, interp='bicubic')

# cut out low intensity region of MS image for easy viewing of underlying H&E
masked_big_ms_image = np.ma.masked_where(big_registered_ms_image < 100, big_registered_ms_image)

# plot the two images overlayed
f = plt.figure(1, figsize=(20, 20))
plt.imshow(optical_monochrome, alpha=0.7, cmap=cm.Greys_r)
plt.imshow(masked_big_ms_image, alpha=0.3, cmap=cm.jet)
plt.axes().set_axis_off()
```

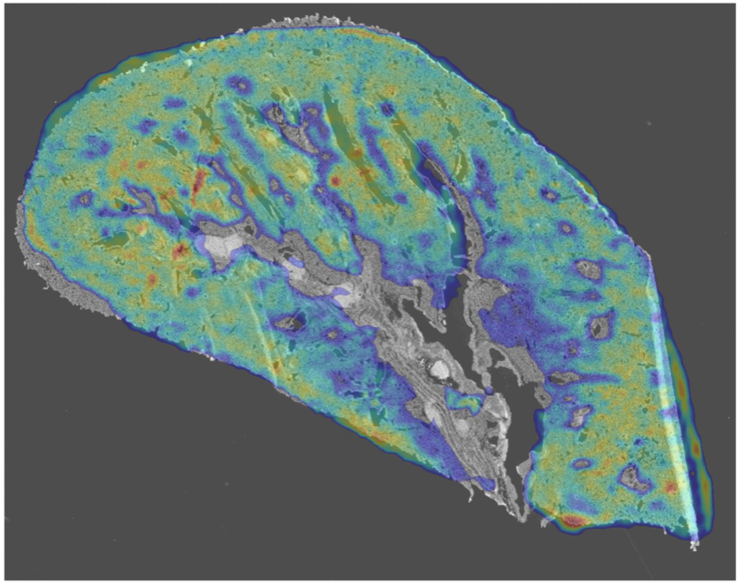

Science notebooks through Jupyter (iPython)

- Widely used in science
- Interactive HPC LDRD

Deployed at NERSC:

- >100 users pre-production

Fernando Perez et al

Random Access Analytics

- Genome assembly “needs shared memory”

Global Address Space

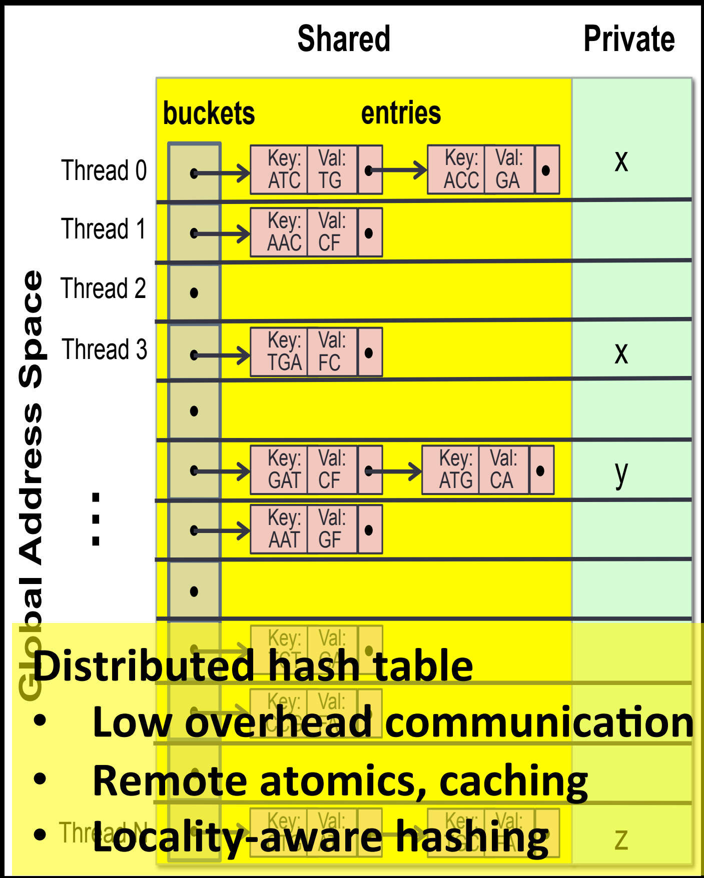

Distributed hash table
- Low overhead communication
- Remote atomics, caching
- Locality-aware hashing

Scales to 15K+ cores

4 minutes for human

First ever solution

E. Georganas, A. Buluc, J. Chapman, S. Hofmeyr, C. Aluru, R. Egan, L. Oliker, D. Rokhsar, K. Yelick

Data Fusion for Observation with Simulation

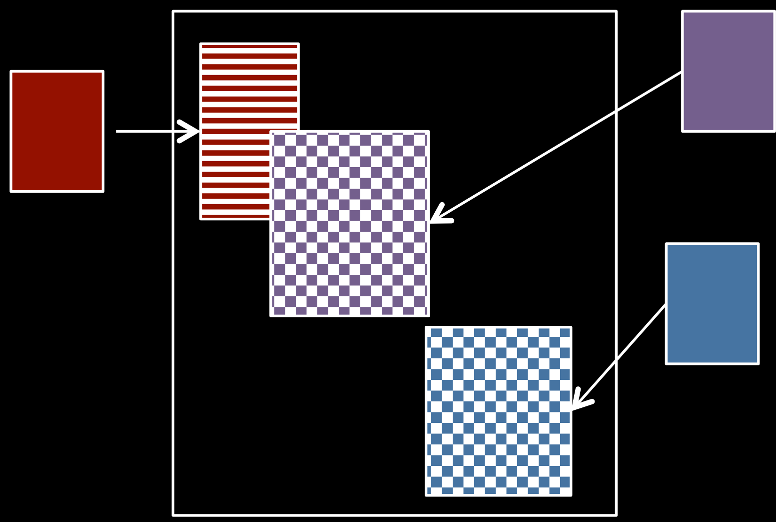

- **Unaligned data from observation**
- **One-sided strided updates**

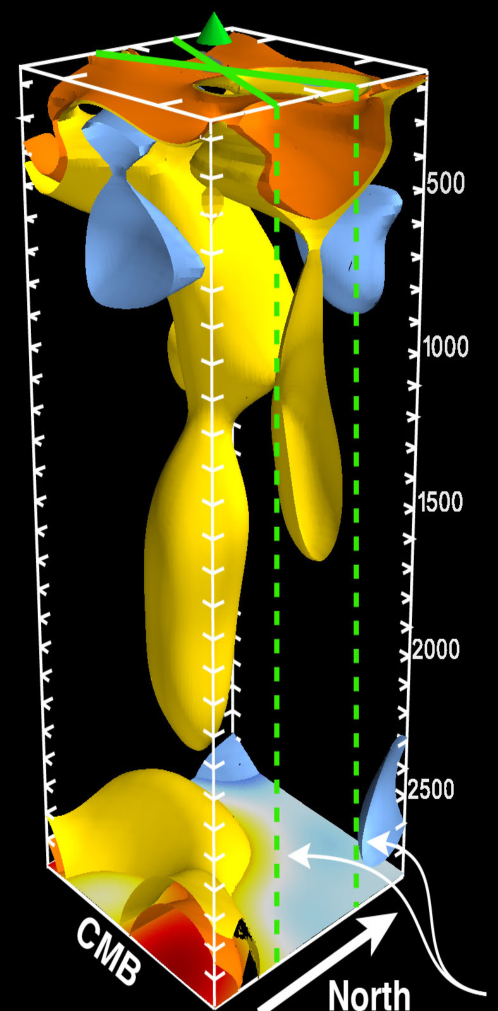

Hawaii hotspot geology

Scott French, Y. Zheng, B. Romanowicz, K. Yelick

Productive Programming

Spark™

Speed

Run programs up to 100x faster than Hadoop MapReduce in memory, or 10x faster on disk.

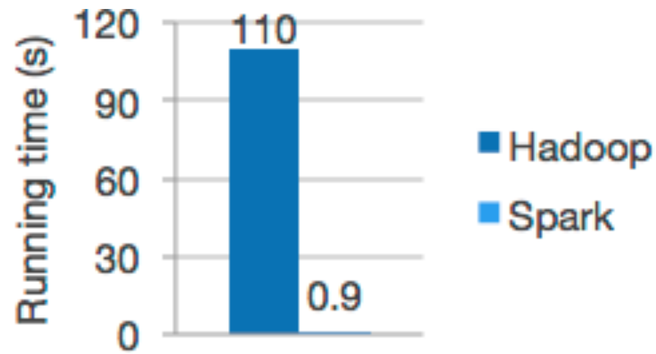

- **High failure rate**
- **Slow network**
- **Fast (local) disk**

And Spark is still 10x+ slower than MPI

Architectures for Data vs. Simulation

Massive Independent Jobs for Analysis and Simulation

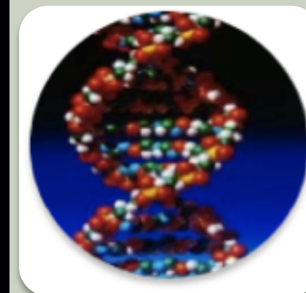

Random access, large data Analysis

Different architectures for simulation? Can simulation use data architectures?

Data processing with special purpose hardware

- General trend towards specialization for continued performance growth
- Data processing (on raw data) will be first in DOE

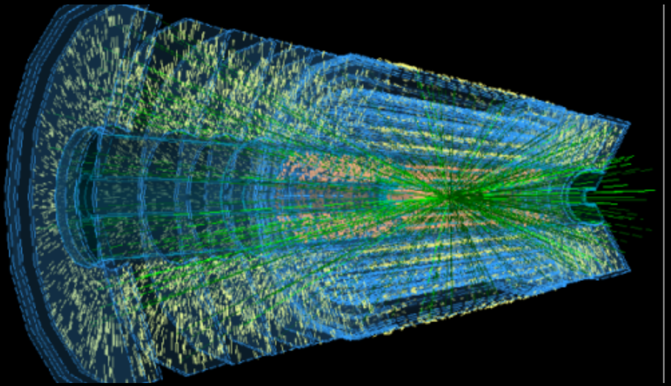

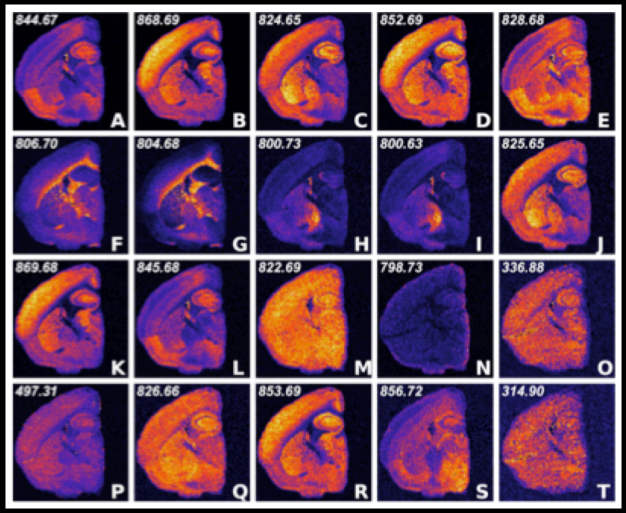

Particle Tracking with Neuromorphic chips

Deep learning processors for image analysis

Computing in Detectors

FPGAS for genome analysis

Questions?

4. Are there open problems or expertise gaps in computing/math/stat/data be addressed?

Questions?

1. **Are there MSU examples of "science at the boundary" of simulation and observation?**
   - How should you take advantage of these opportunities?
2. **What are the largest and most complex sources of research data at MSU?**
   - What types of data/CS/math/stat challenges are there?
3. **How should undergrad/grad programs be adapted to address data challenges in future careers?**
   - New courses, (joint) majors, research institutes?
4. **Are there open problems or expertise gaps in computing/math/stat/data be addressed?**

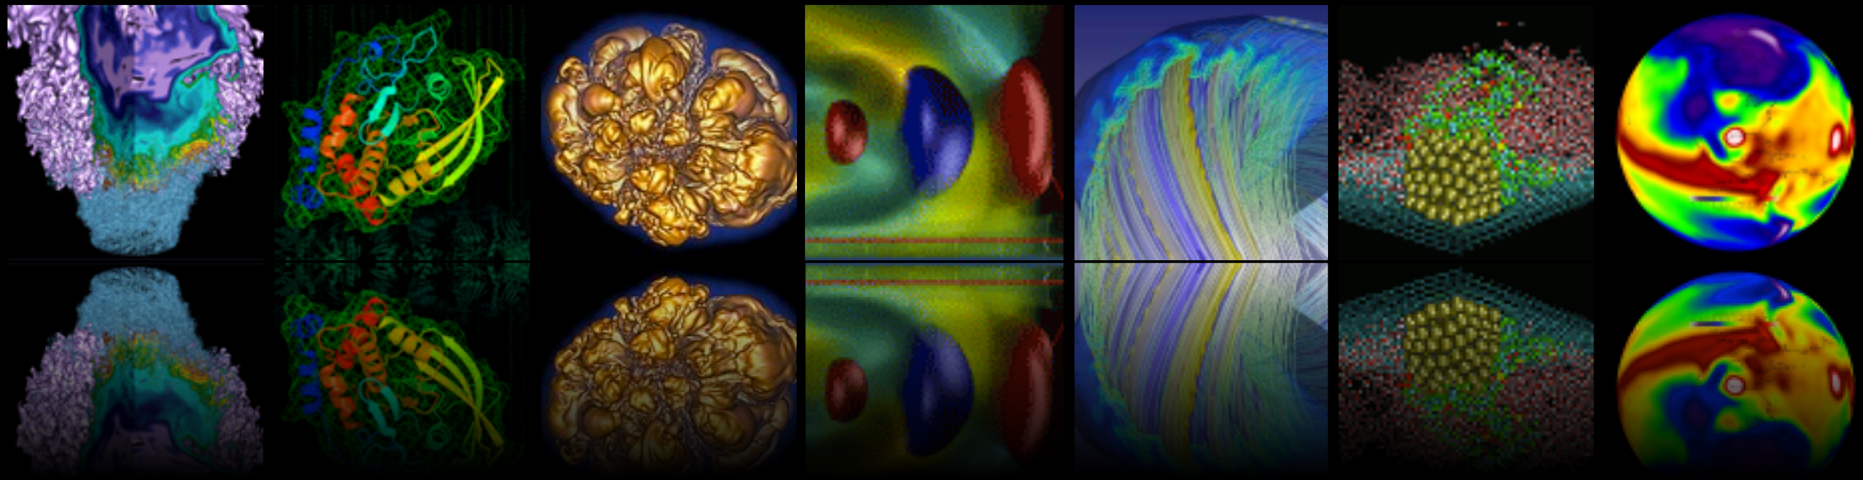

Extreme Data Science

The scientific process is poised to undergo a radical transformation based on the ability to access, analyze, simulate and combine large and complex data sets.

Slides: http://www.cs.berkeley.edu/~yelick/talks