Big Ideas!

from the Data, Devices and Interaction Laboratory

Thursday, October 3 2019

Michael E. Papka

Northern Illinois University

What is a Big Ideas Class?

What is a Big Ideas Class?

Lectures and discussions of current research and technical developments in computer science for beginning graduate research students. Topics will emphasize open problems and recent scientific advances. Content may vary to reflect research advances in areas such as data analytics, scientific computing, graphics and visualization.

What is a Big Ideas Class?

Lectures and **discussions** of current **research** and technical developments in computer science for beginning graduate research students. Topics will emphasize **open problems** and recent scientific advances. Content may vary to reflect research advances in areas such as data analytics, scientific computing, graphics and visualization.

What is a Big Ideas Class?

Lectures and **discussions** of current **research** and technical developments in computer science for beginning graduate research students. Topics will emphasize **open problems** and recent scientific advances. Content may vary to reflect research advances in areas such as data analytics, scientific computing, graphics and visualization.

- Who has active research projects?
- What are NIU CS faculty interests?
- Where do I get more information?
- How do I get involved?

Bit About Me (Education)

- Northern Illinois University - Physics (BS)
- University of Illinois @ Chicago - Computer Science (MS)
- University of Chicago - Computer Science (MS, PhD)
- University of Chicago - Business School (SLLP)
- Harvard University - Business School (GMP)
- Stanford University - Hasso Plattner Institute of Design (Design Thinking)

Bit About Me (Professional)

- Fermi National Accelerator Laboratory (Undergraduate/Graduate)
- Argonne National Laboratory
- Northern Illinois University

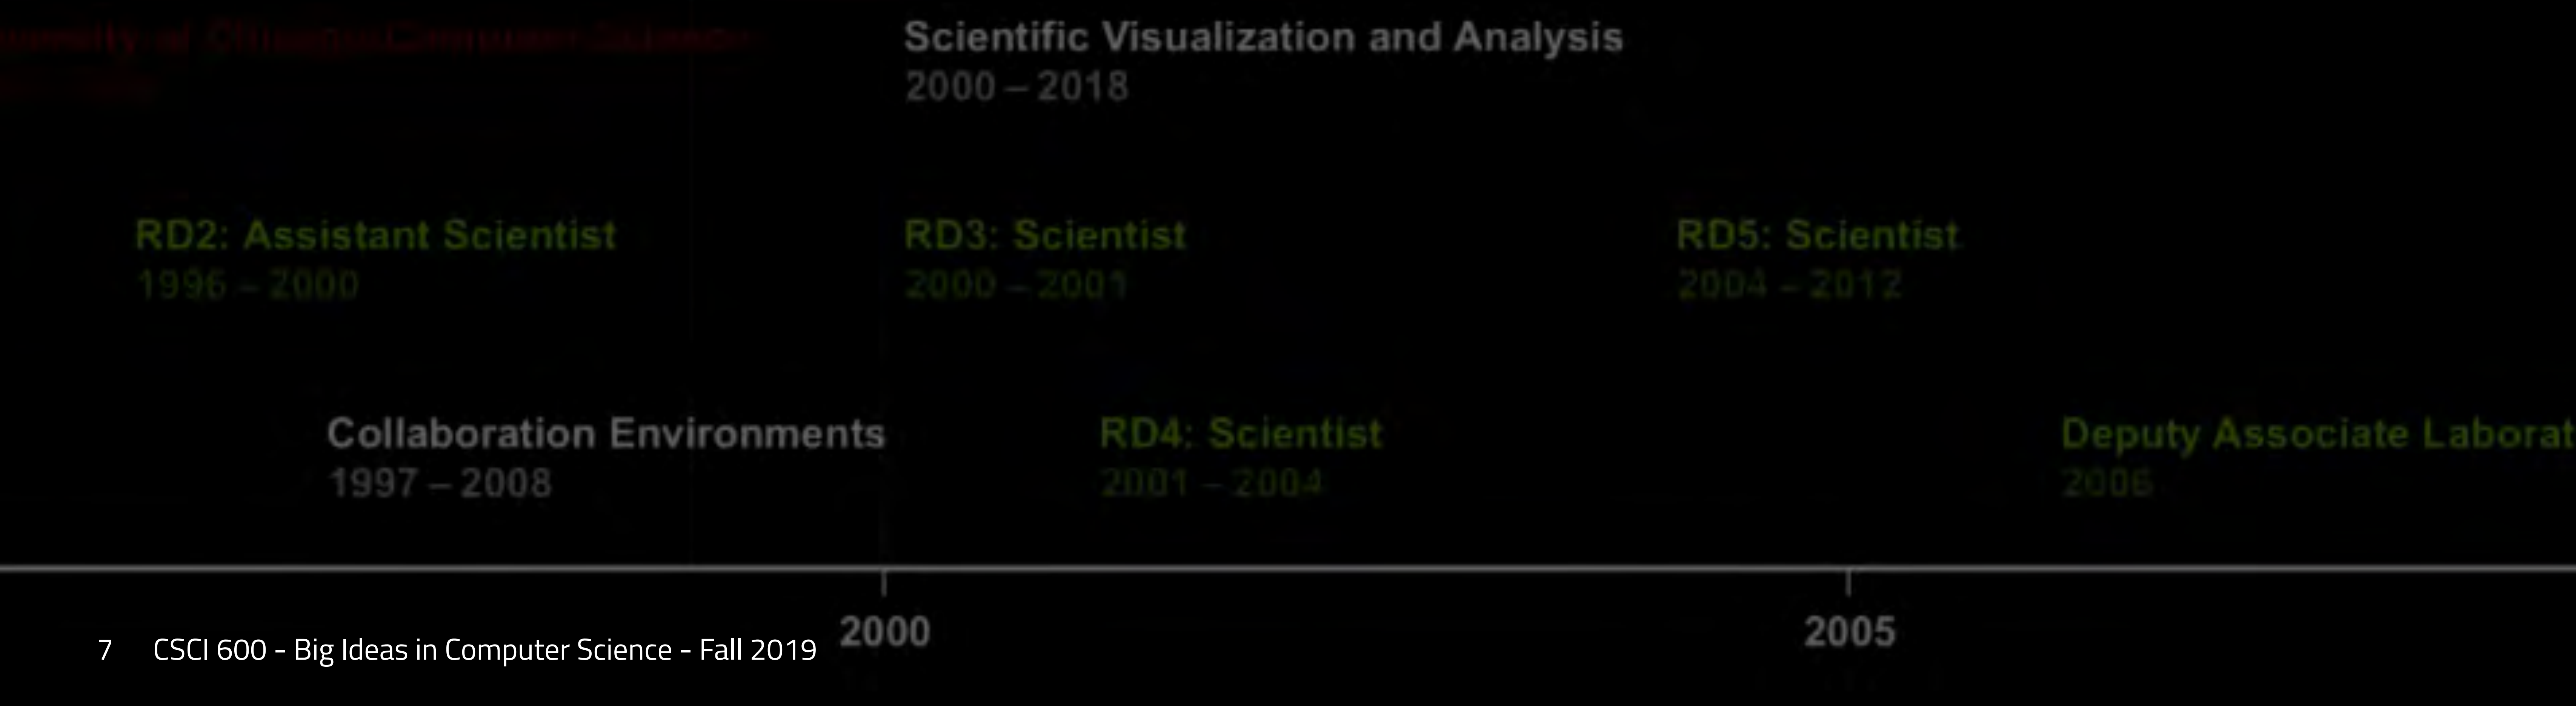

Bit About Me (Research - *Areas/Interests*)

- Advanced Display Environments
- Collaboration Technology
- High Performance Computing (Environments)
- Information Visualization
- Scientific Visualization and Analysis
- Augmented/Virtual Reality

Bit About the ddiLab

- Joint lab with School of Art and Design - Time Arts (Professor Joseph Insley)
- Emphasis on visualization and data analysis coupled to high-performance computing in the support of science
- Students
  - 1 PhD (Information Visualization)
  - 3 MS (HPC Log Analysis, Authentication Infrastructure and Machine Learning)
  - 3 Undergraduates (Virtual Reality and HPC Log Analysis)

Scientific Visualization and Analysis

- vl3: volume rendering library
  - Parallel volume rendering library that exploits GPU hardware
  - Uses native data formats
- Integration with virtual and augmented reality
- Usability and collaboration
- Domain specific visualizations

vl3: Volume Rendering Library

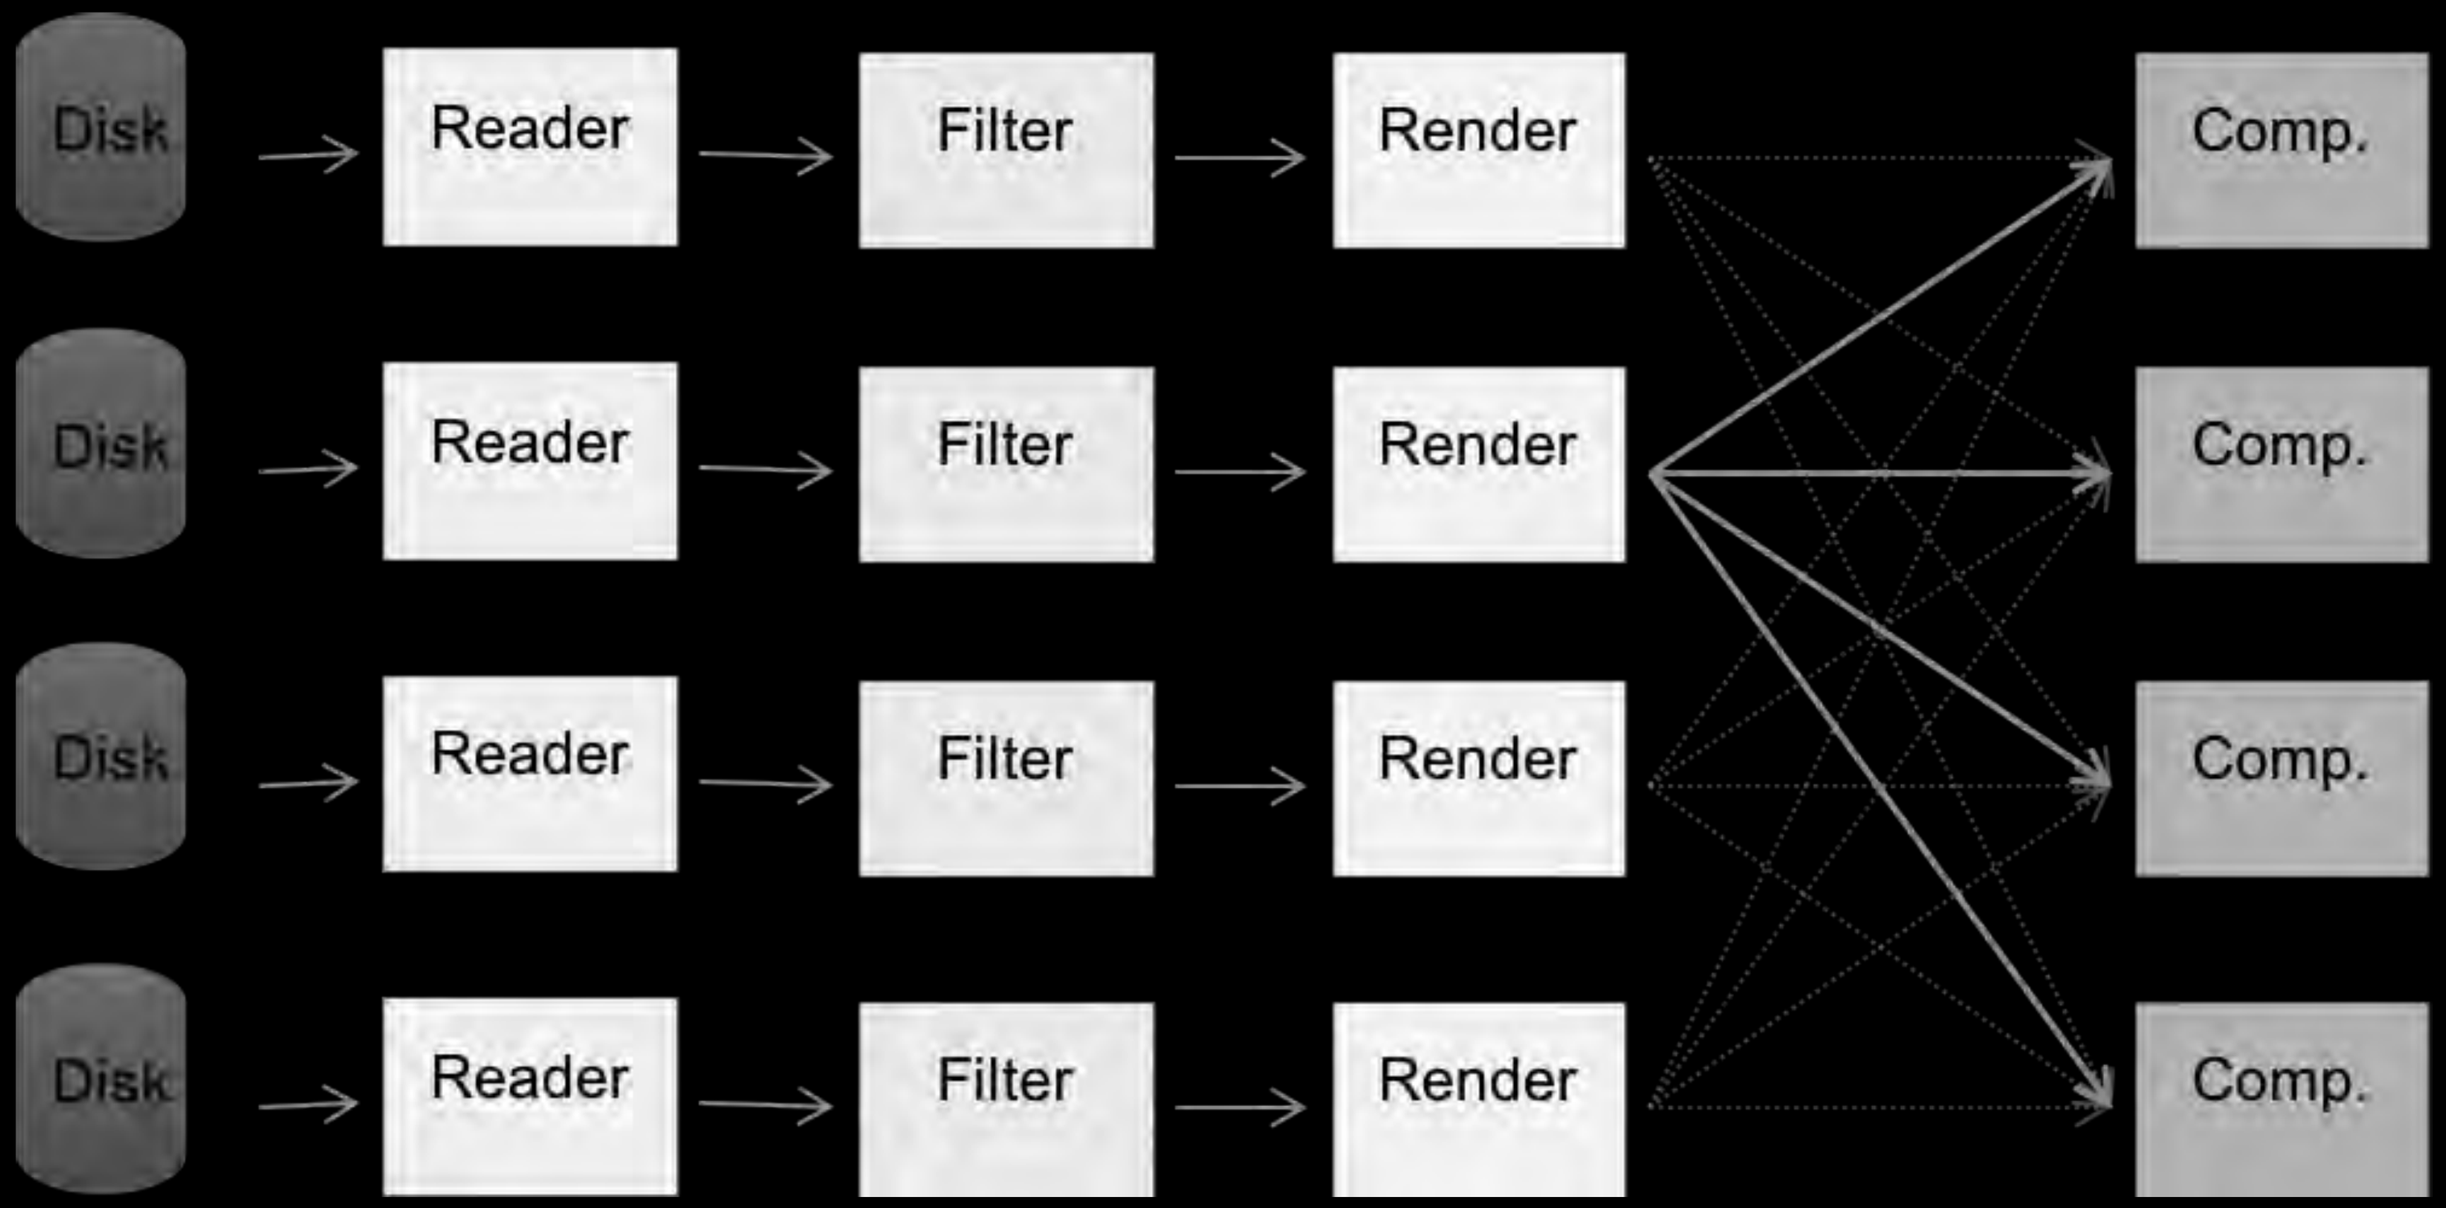

Thermonuclear Flame Plume Rising in a Column

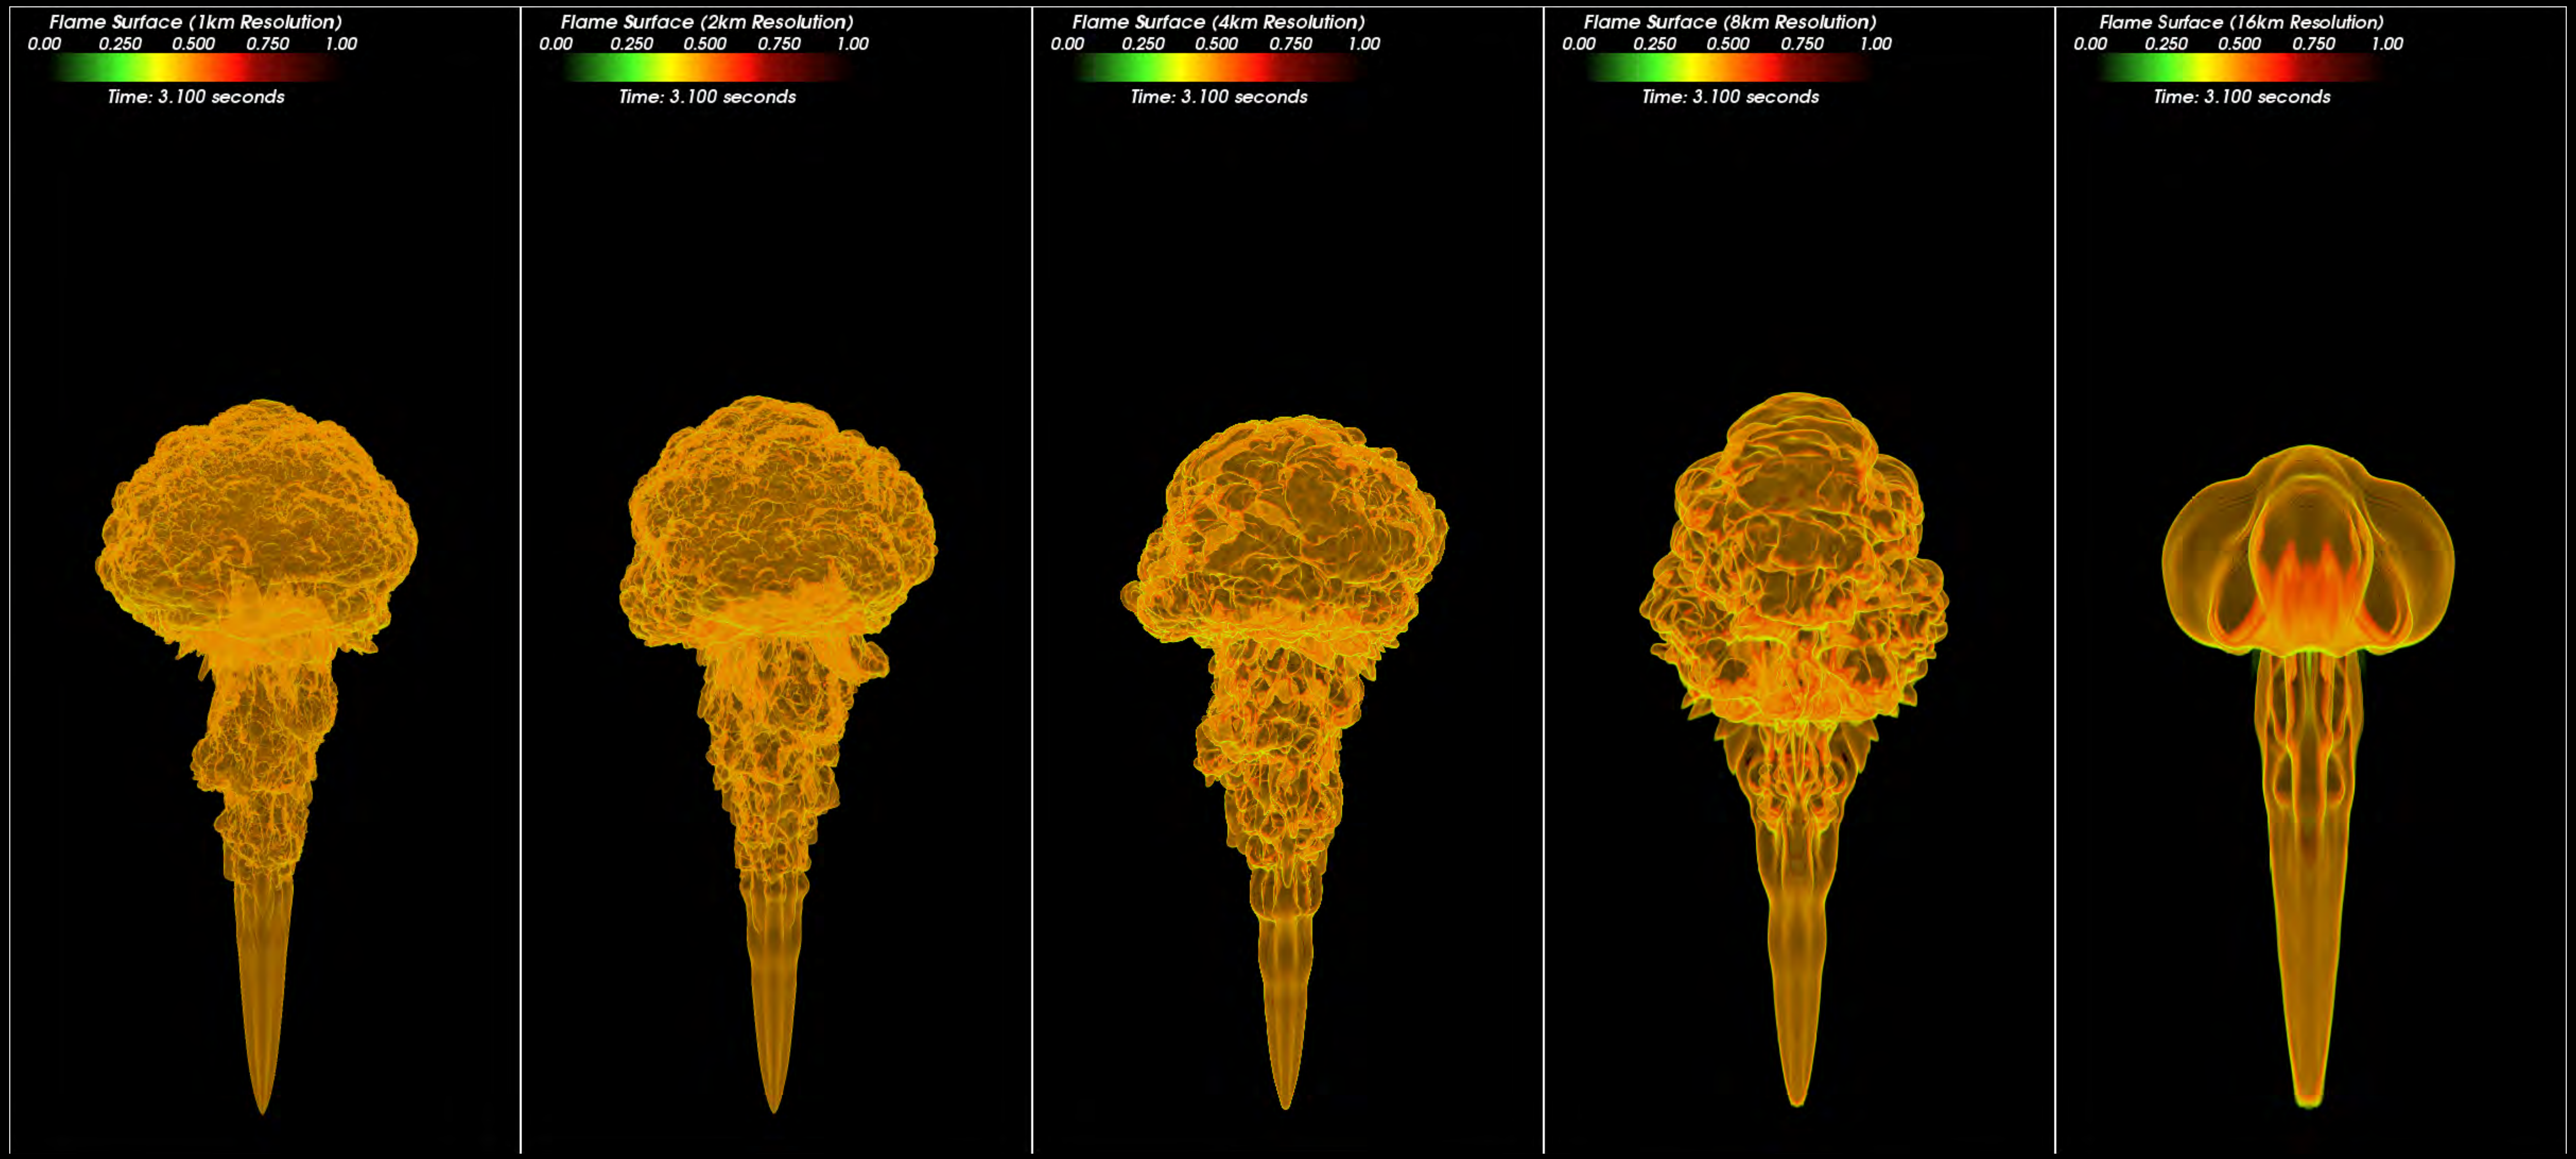

Improving the Quality of Visualizations[a]

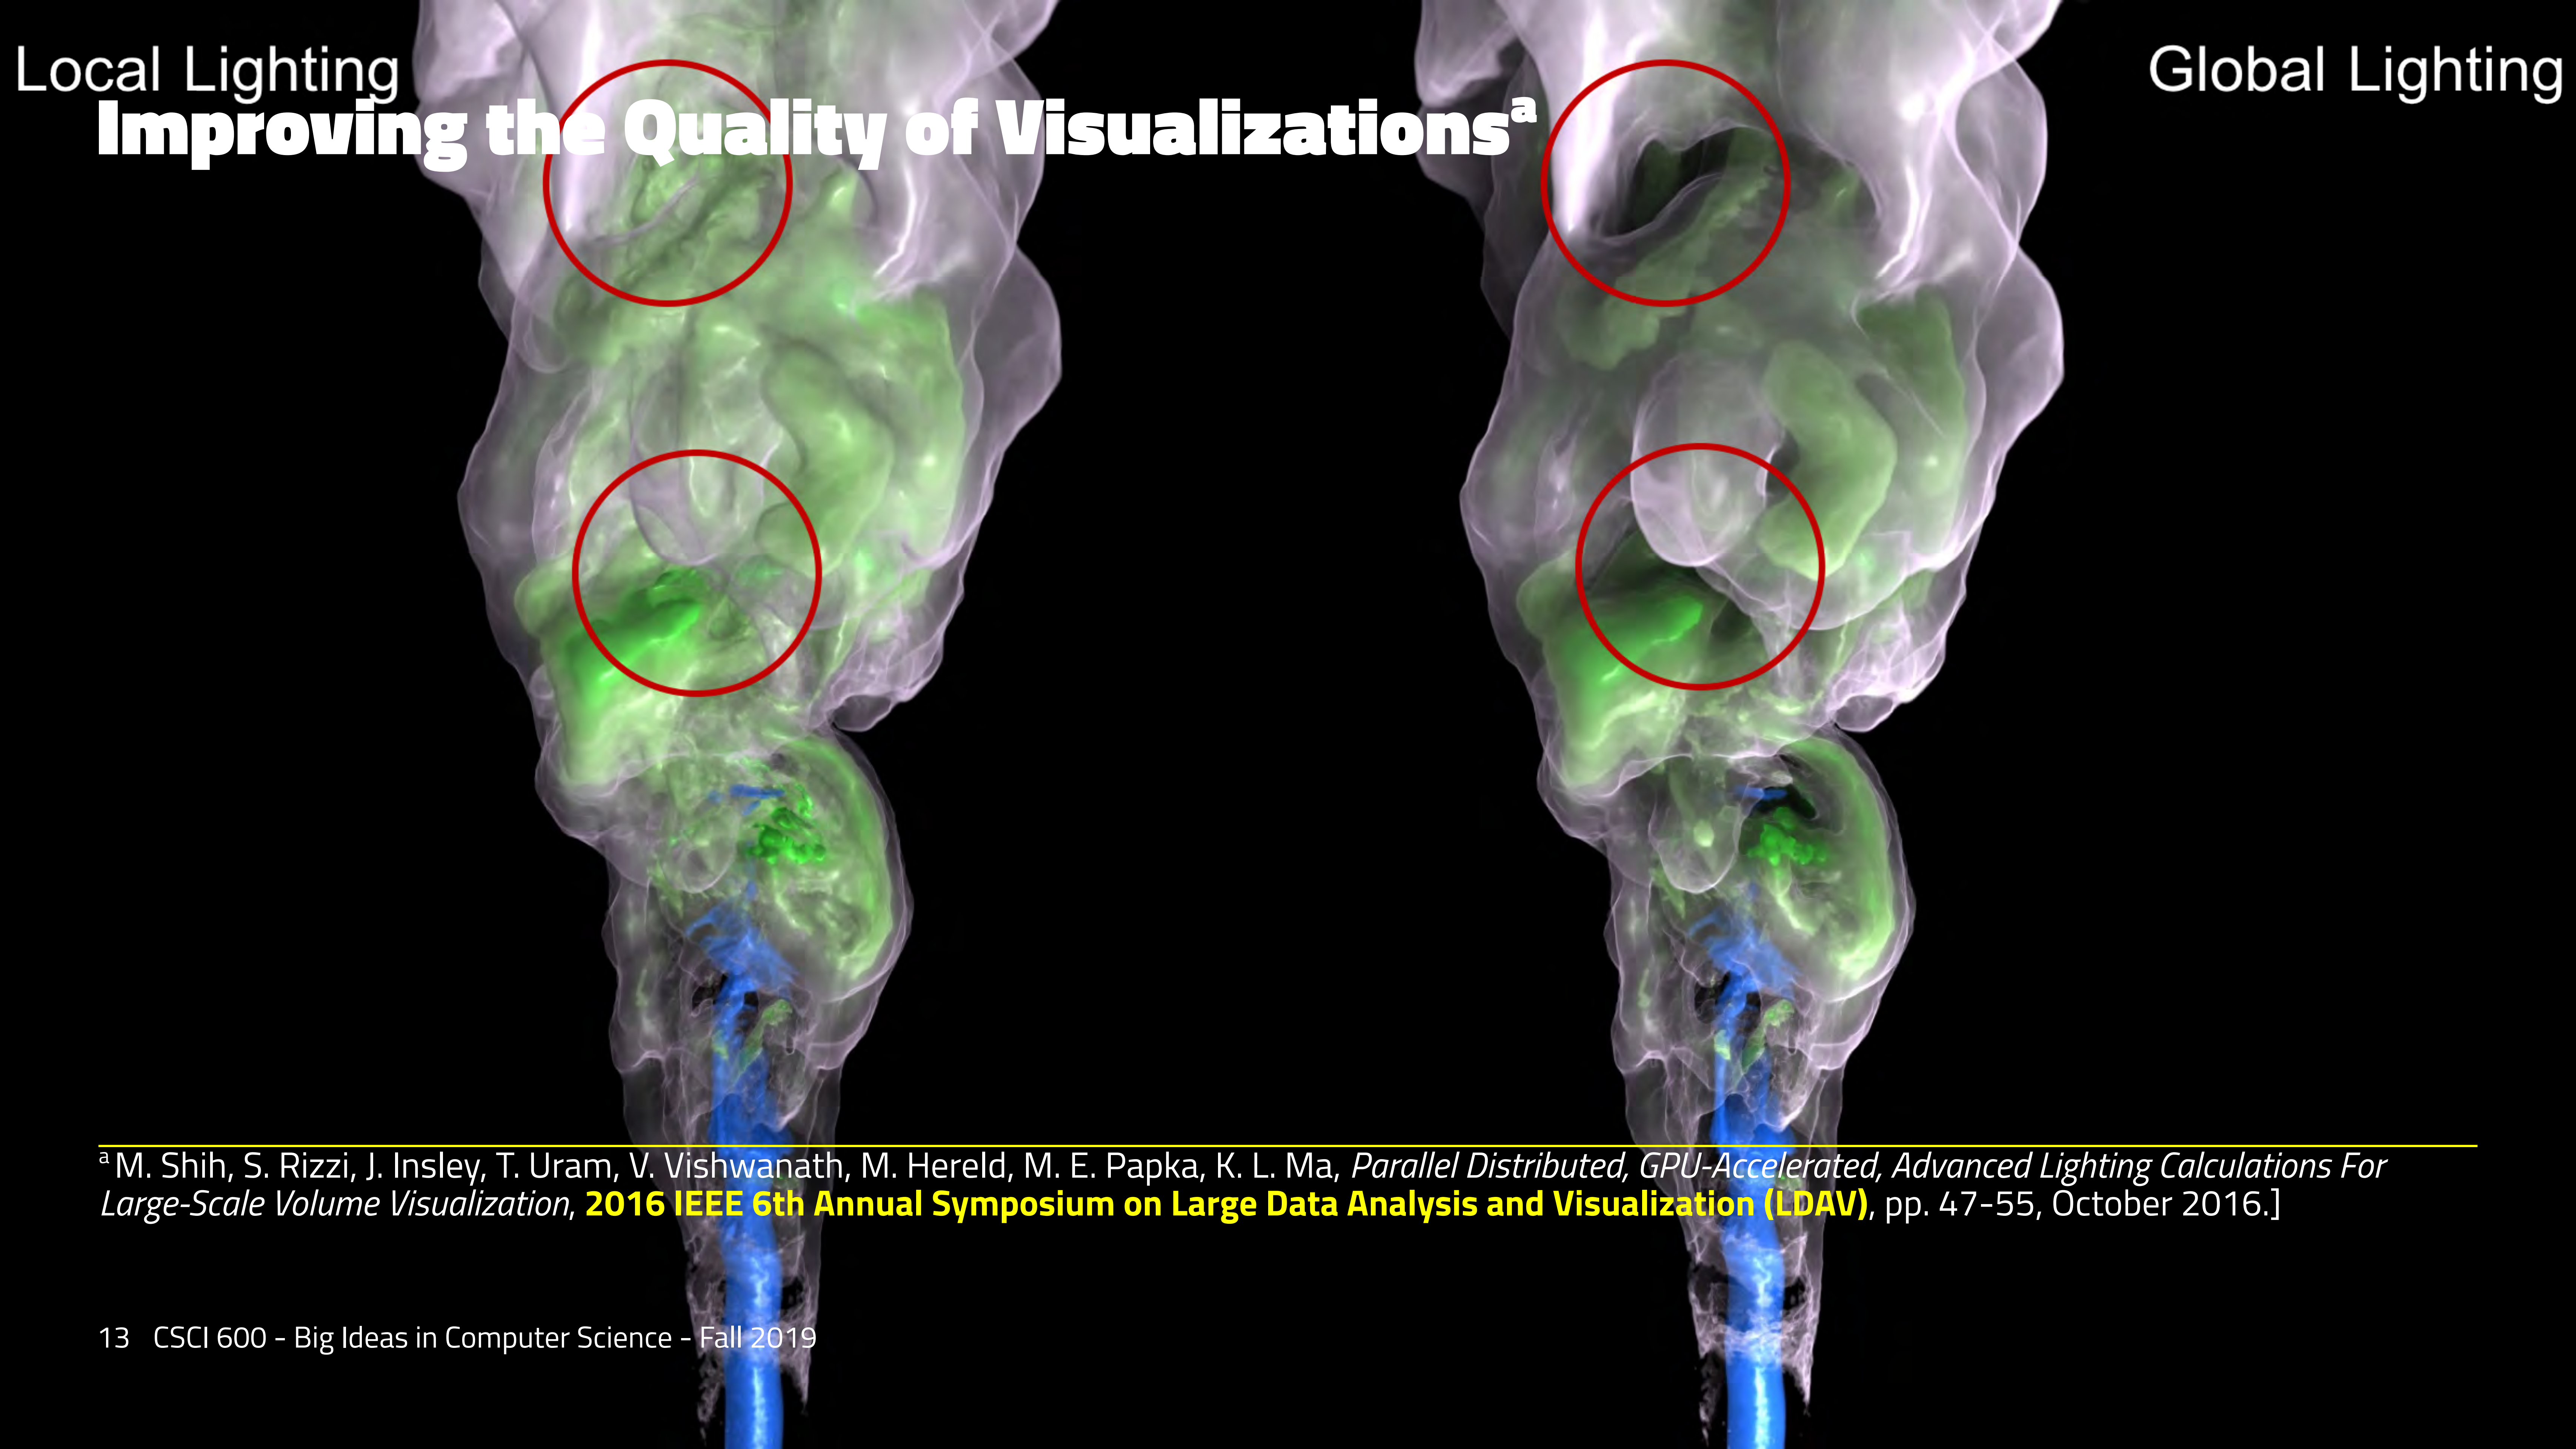

[a] M. Shih, S. Rizzi, J. Insley, T. Uram, V. Vishwanath, M. Hereld, M. E. Papka, K. L. Ma, *Parallel Distributed, GPU-Accelerated, Advanced Lighting Calculations For Large-Scale Volume Visualization*, **2016 IEEE 6th Annual Symposium on Large Data Analysis and Visualization (LDAV)**, pp. 47-55, October 2016.]

Virtual Reality[b,c]

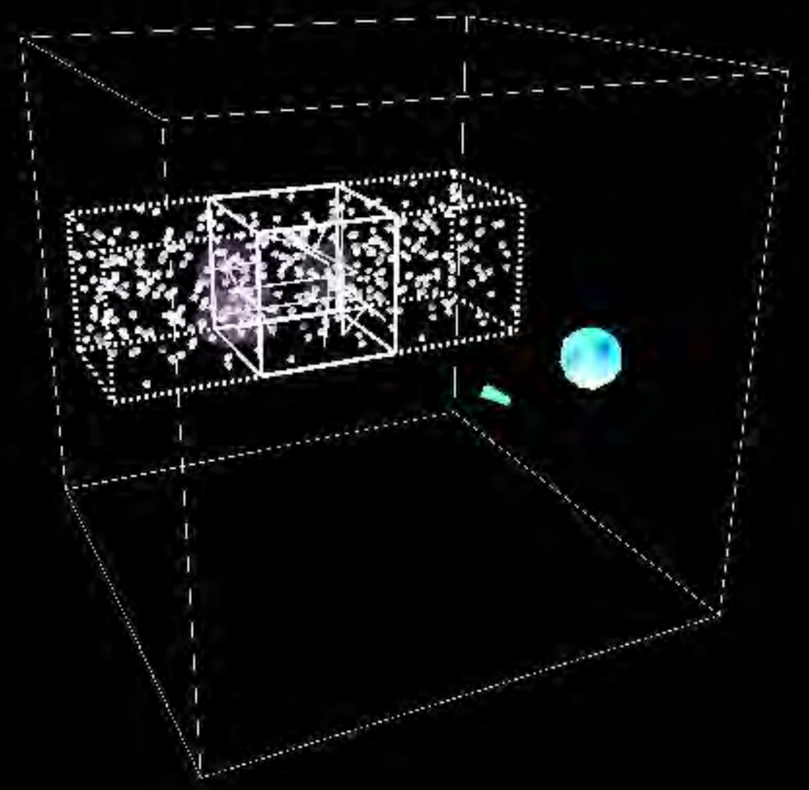

[b] T. Disz, M. E. Papka, R. Stevens, M. Pellegrino, V. Taylor, *Virtual Reality Visualization of Parallel Molecular Dynamics Simulation*, **1995 Simulation Multiconference Symposium**, pp. 483-87, Phoenix, AZ, April 1995.

[c] K. Reda, A. Knoll, K. Nomura, M. E. Papka, A. E. Johnson, J. Leigh, *Visualizing Large-Scale Atomistic Simulations in Ultra-resolution Immersive Environments*, **Proceedings of the 2013 IEEE Symposium on Large Data Analysis and Visualization (LDAV 2013)**, pp. 59-66, Atlanta, GA, October 13-14, 2013.

Virtual Reality[d]

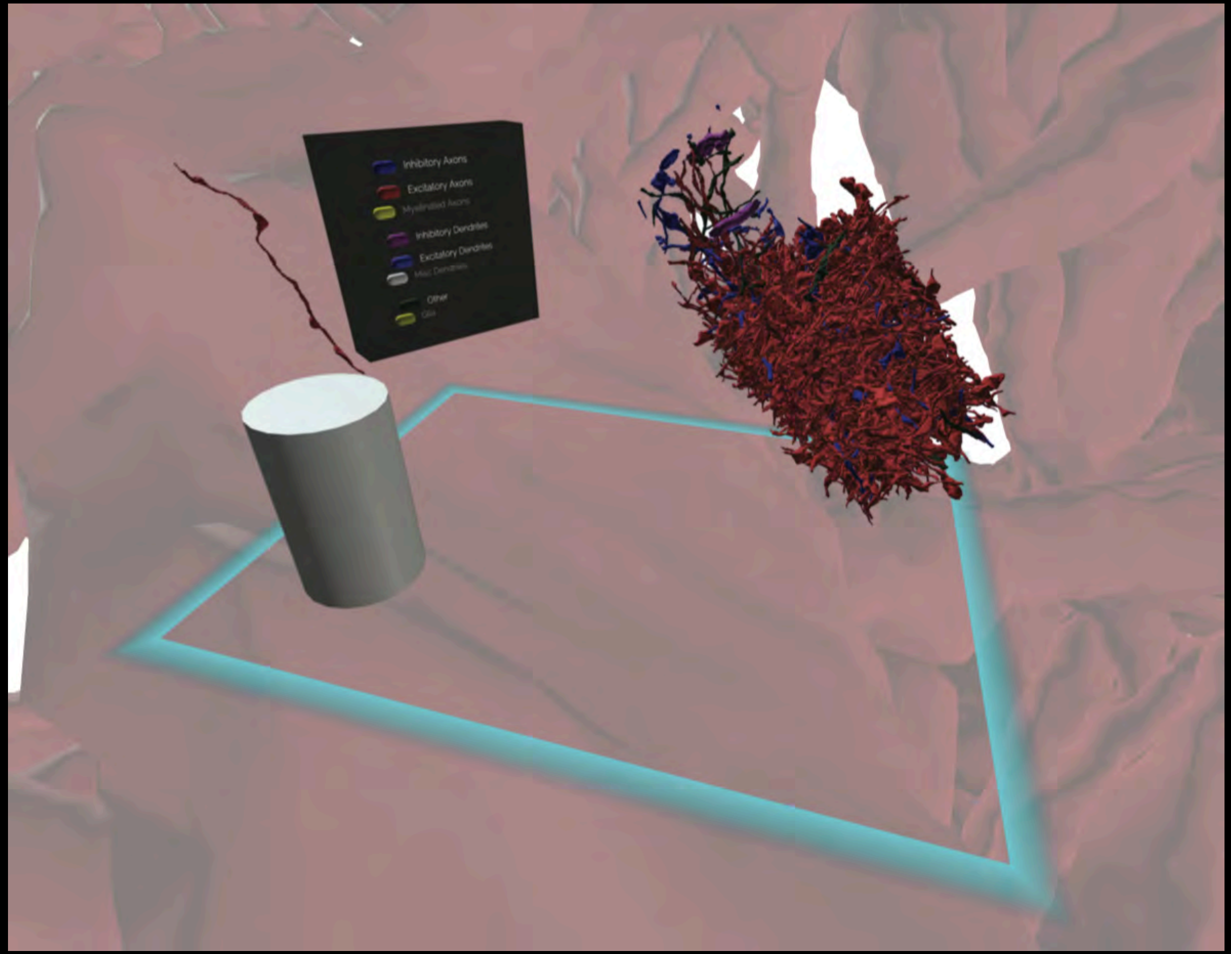

[d] E. B. Brooks, J. A. Insley, M. E. Papka, S. Rizzi, *Virtual reality tools for the correction of automated volume segmentation errors using dense surface reconstructions*, **2017 IEEE 7th Symposium on Large Data Analysis and Visualization (LDAV)**, pp. 92-93, October 2, 2017. [POSTER]

Usability and Collaboration[e]

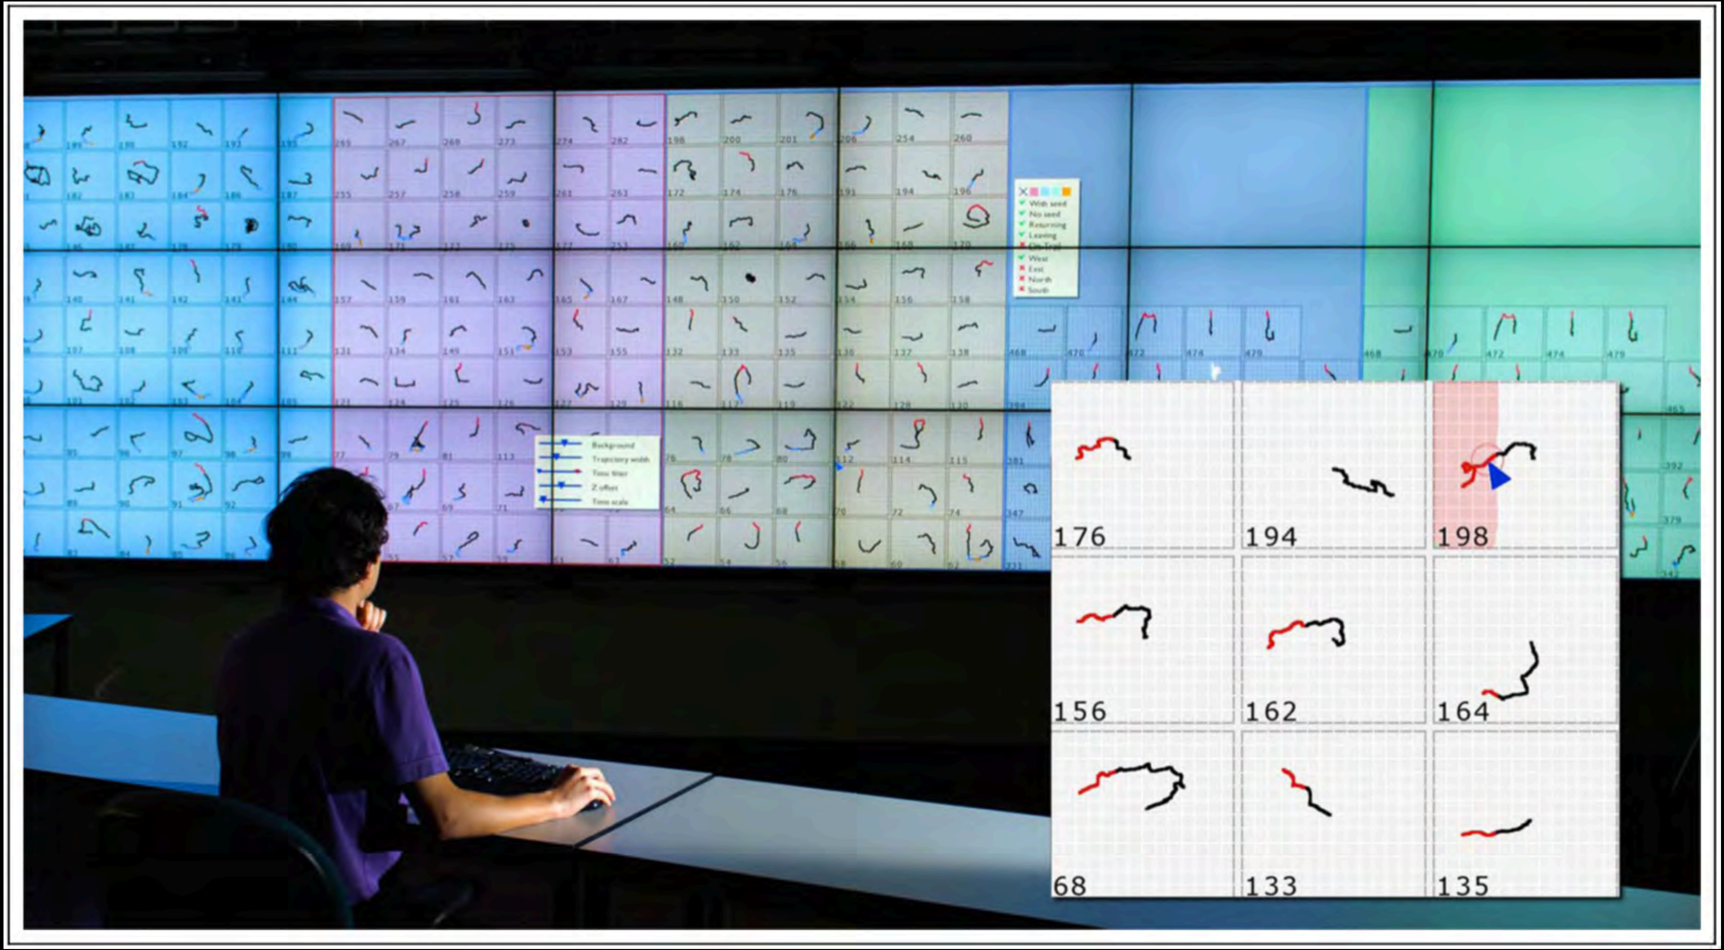

[e] K. Reda, A. E. Johnson, M. E. Papka, J. Leigh, *Modeling and Evaluating User Behavior in Exploratory Visual Analysis*, **Information Visualization** 15(4), pp. 325-339, October 2016.

Usability and Collaboration[f]

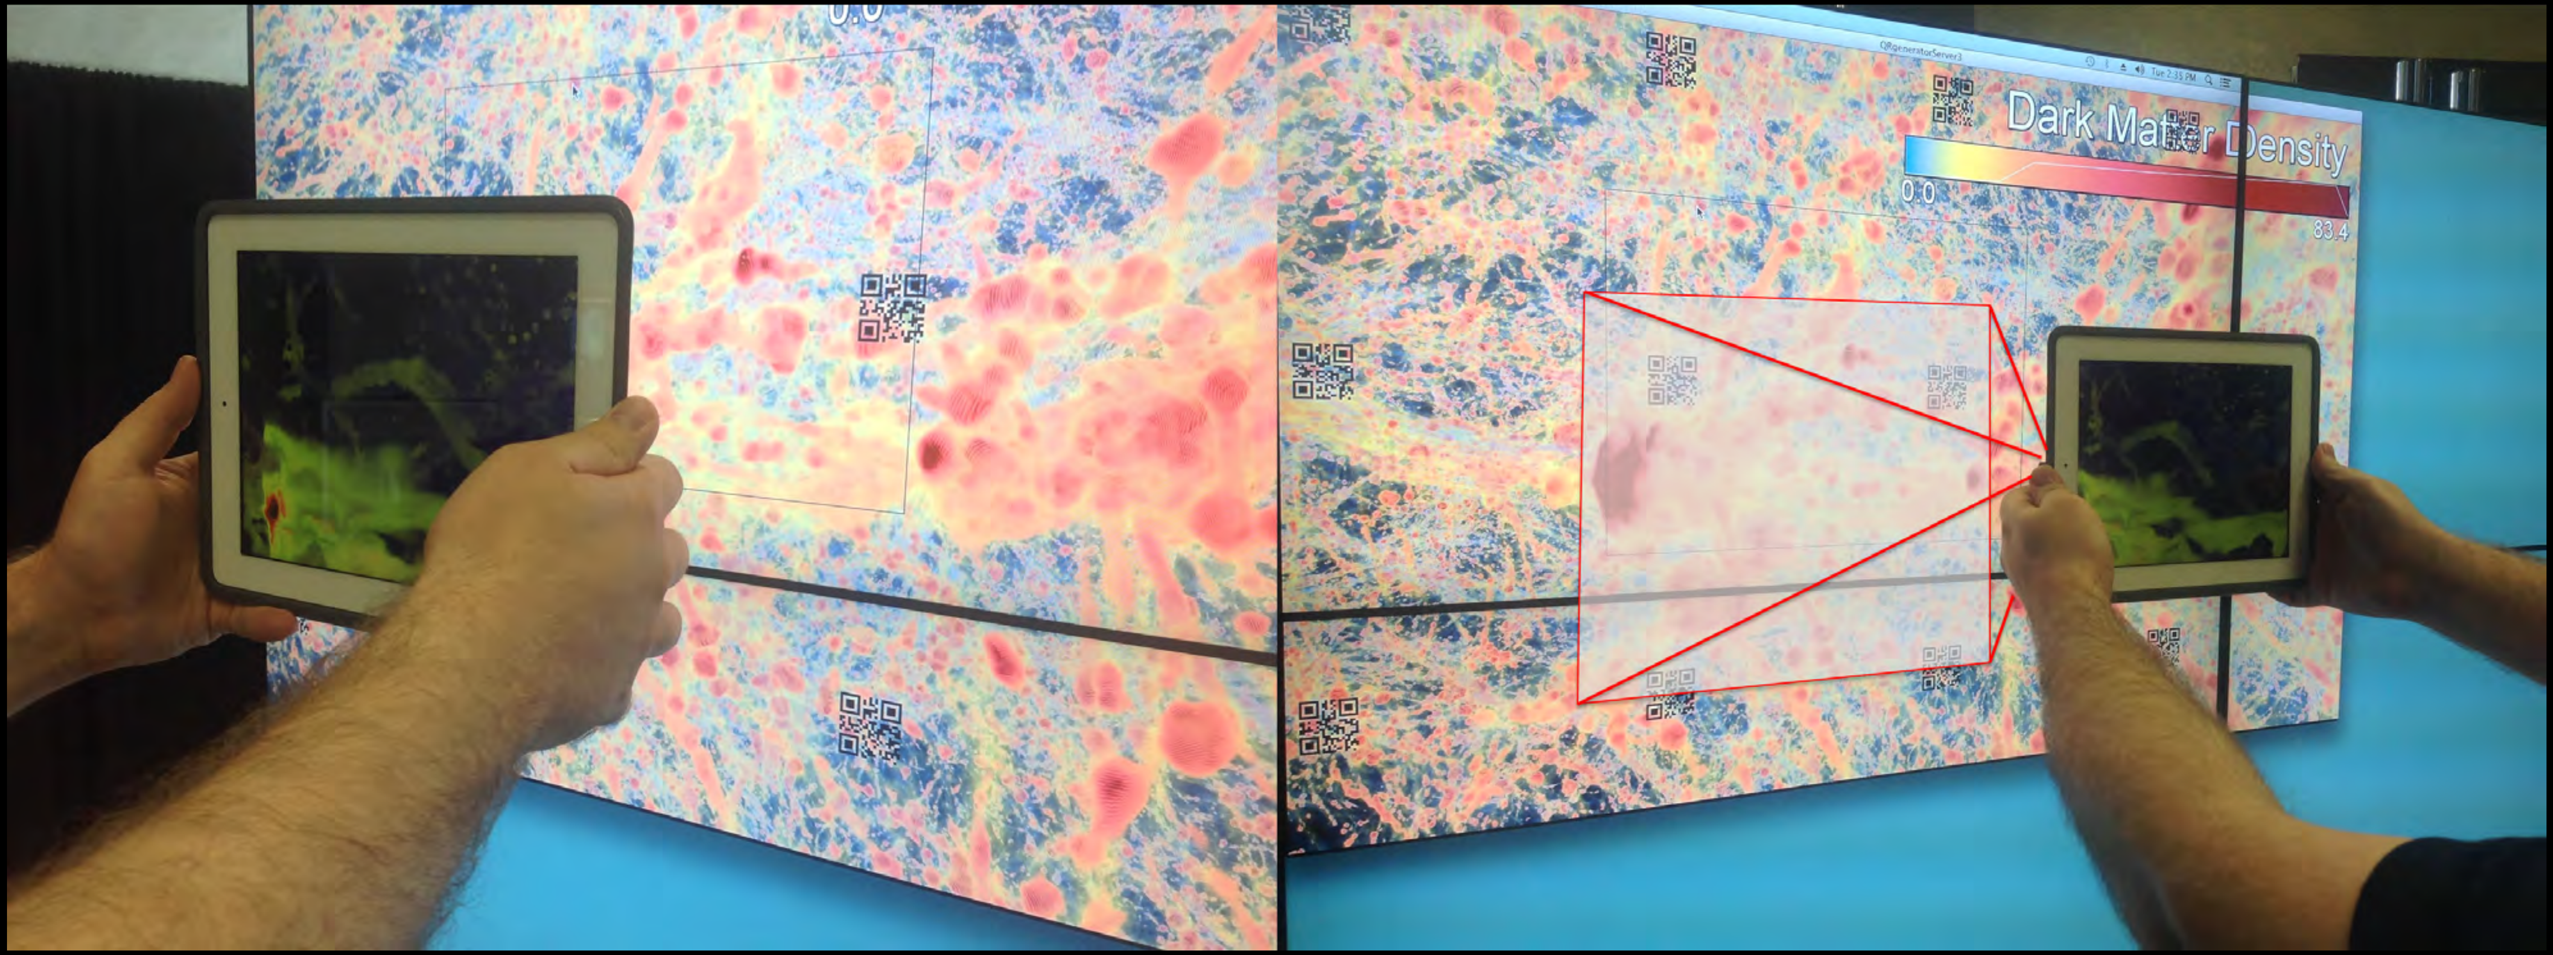

[f] **P. Lindner**, **A. Rodriguez**, T. Uram, M. E. Papka, *Augmenting Views on Large Format Displays with Tablets*, **Proceedings of the 2nd ACM Symposium on Spatial User Interaction (SUI 2014)**, Honolulu, HI, October 4-5, 2014. [Poster]

Domain Specific Visualizations

- Applied solutions to specific problems within domain
- Deep partnership with domain experts
- Current effort with NIU Chemistry
  - *Visualizing and Quantifying Structural Ordering Underlying Static Structure Factor Peaks from Molecular Dynamics Simulations* Travis Mackoy, Bharat Kale, Ralph Wheeler
  - *Comparison Visualizations of Electron Density Approximation Methods* Anouar Benali (ANL), Joe Insley, Ralph Wheeler

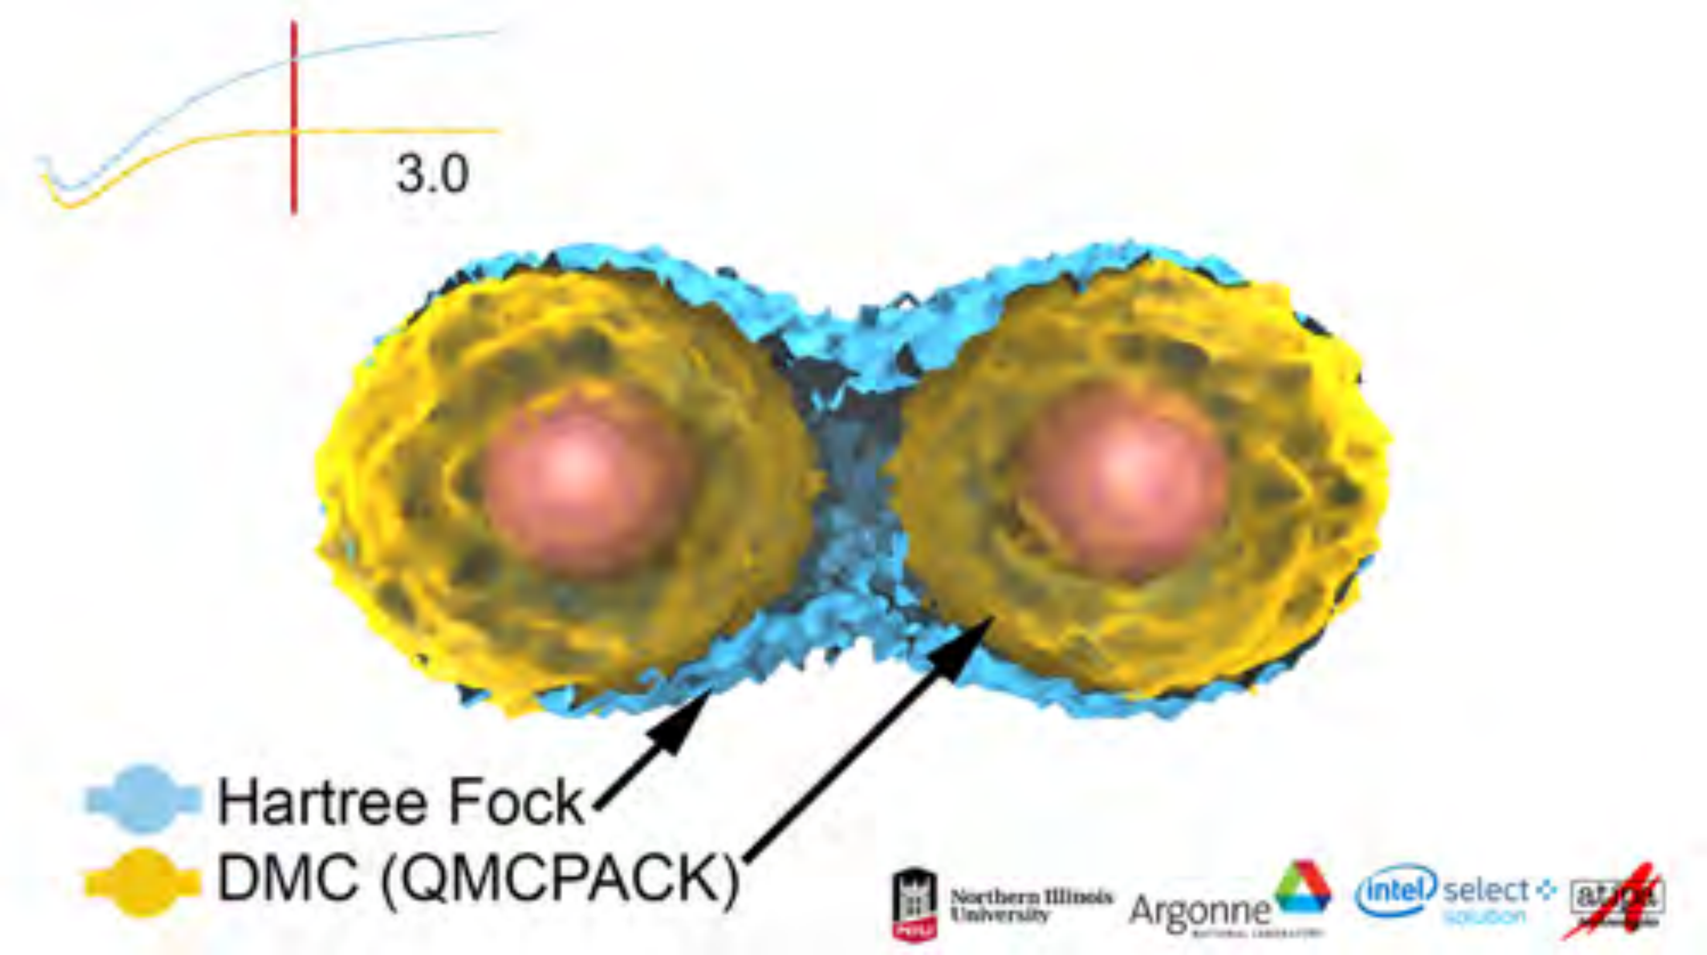

Domain Specific Visualizations[g]

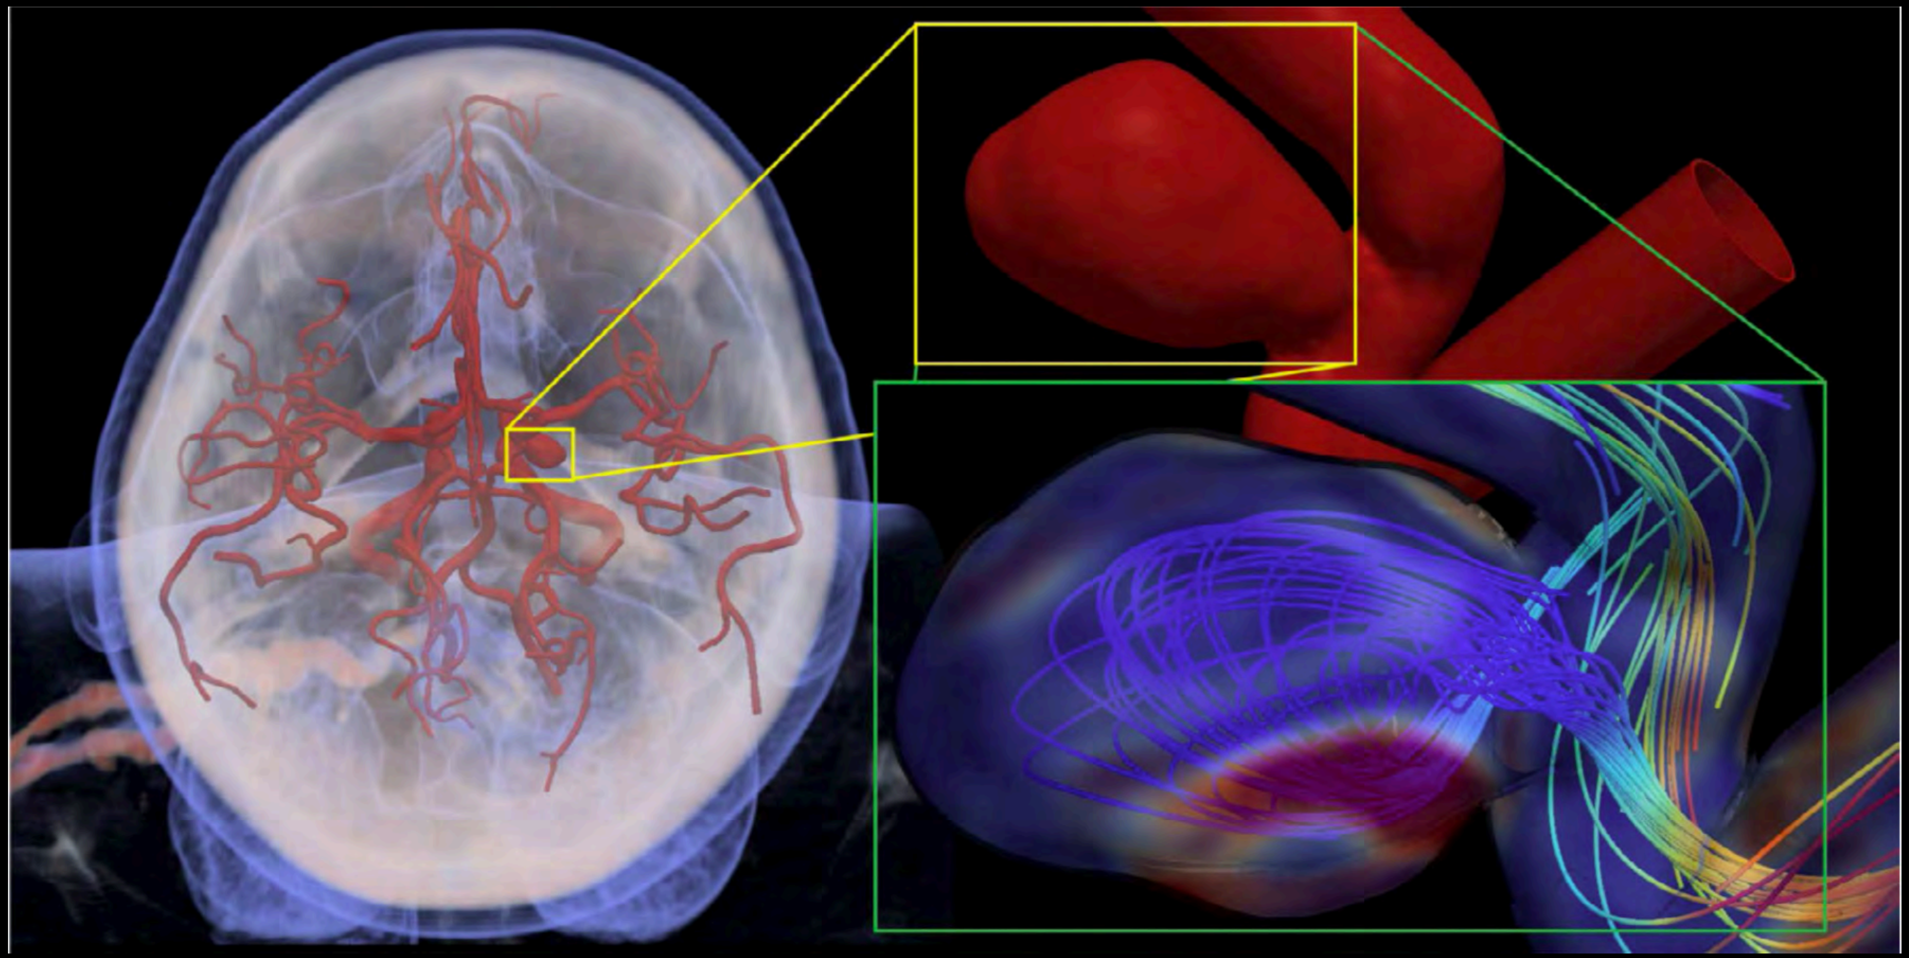

[g] P. Perdikaris, J.A. Insley, L. Grinberg, Y. Yu, M. E. Papka, G. E. Karniadakis, *Visualizing Multiphysics, Fluid-Structure Interaction Phenomena in Intracranial Aneurysms*, **Parallel Computing**, 55, pp. 9-16, July 2016.

Domain Specific Visualizations[h]

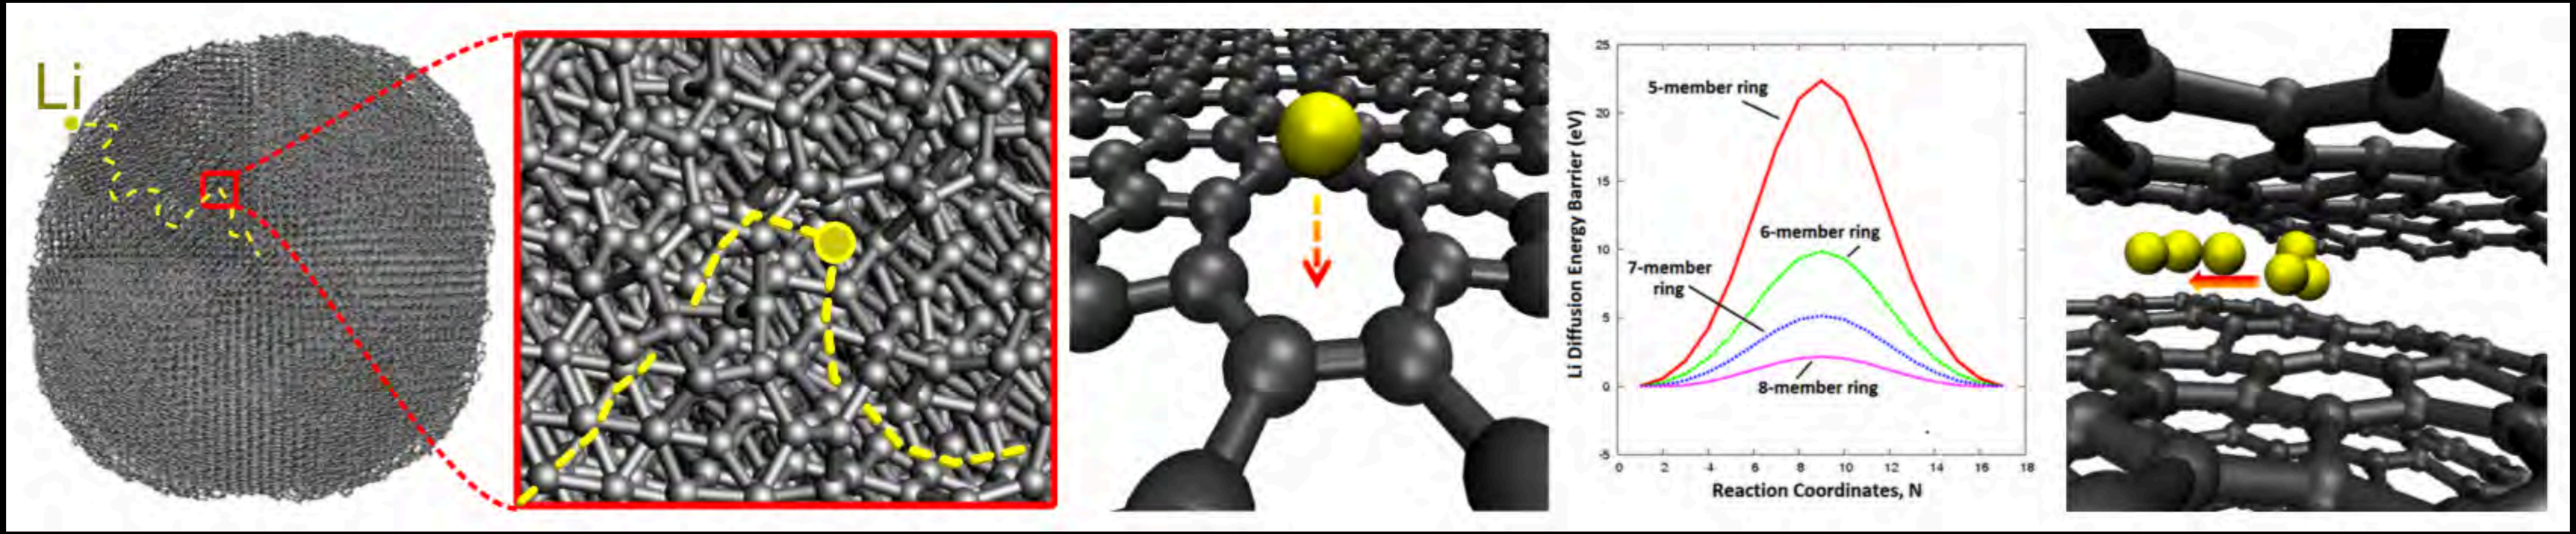

[h] A. Gyulassy, A. Knoll, K. C. Lau, B. Wang, P.-T. Bremer, M. E. Papka, L. Curtiss, V. Pascucci, *Interstitial and Interlayer Ion Diffusion Geometry Extraction in Graphitic Nanosphere Battery Materials*, **IEEE Transactions on Visualization and Computer Graphics**, 22(1):916-925, January 2016.

Domain Specific Visualizations

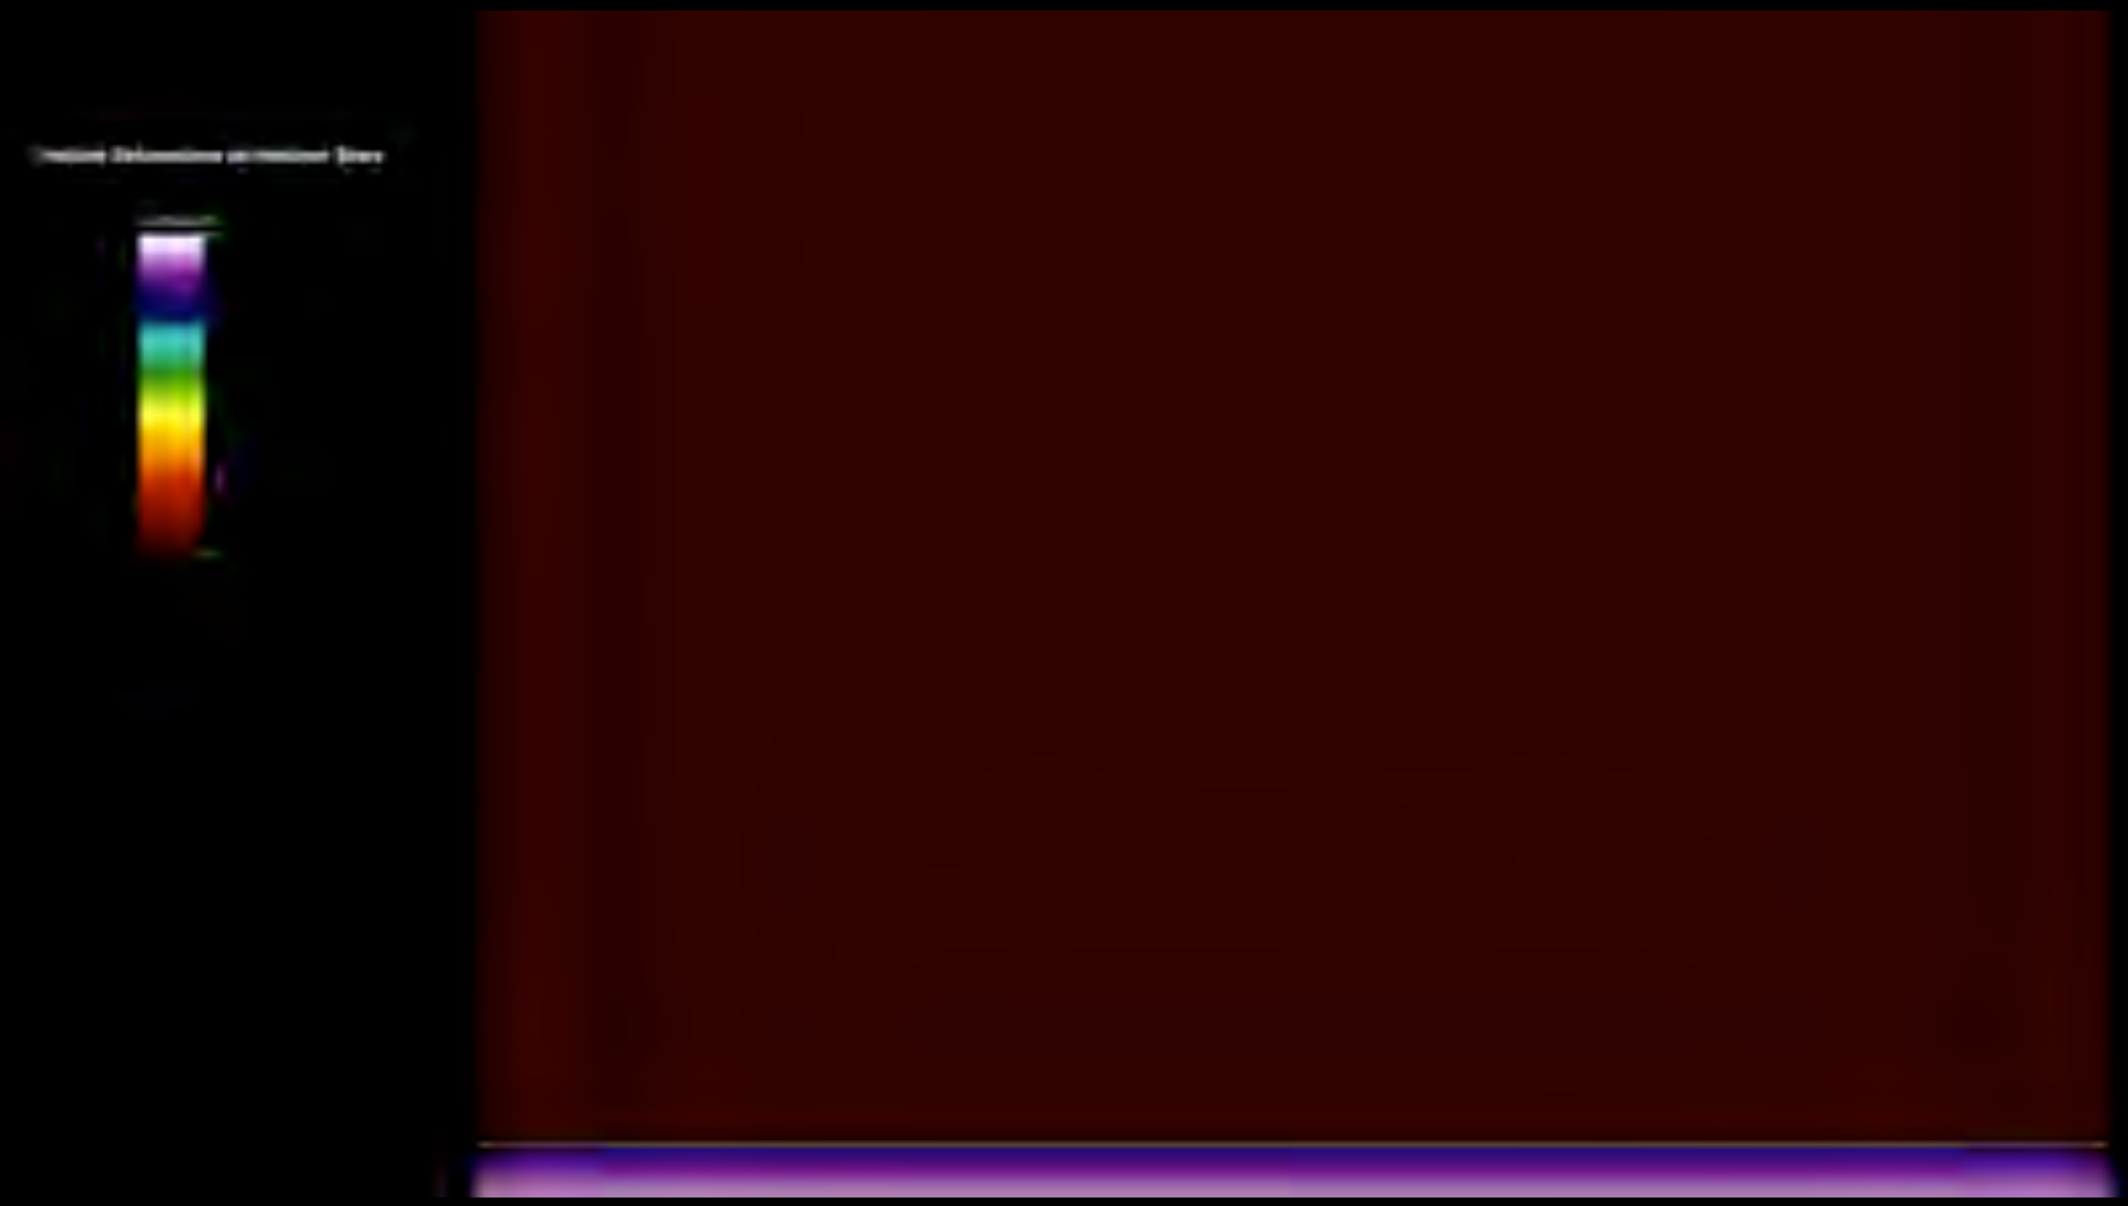

Domain Specific Visualizations

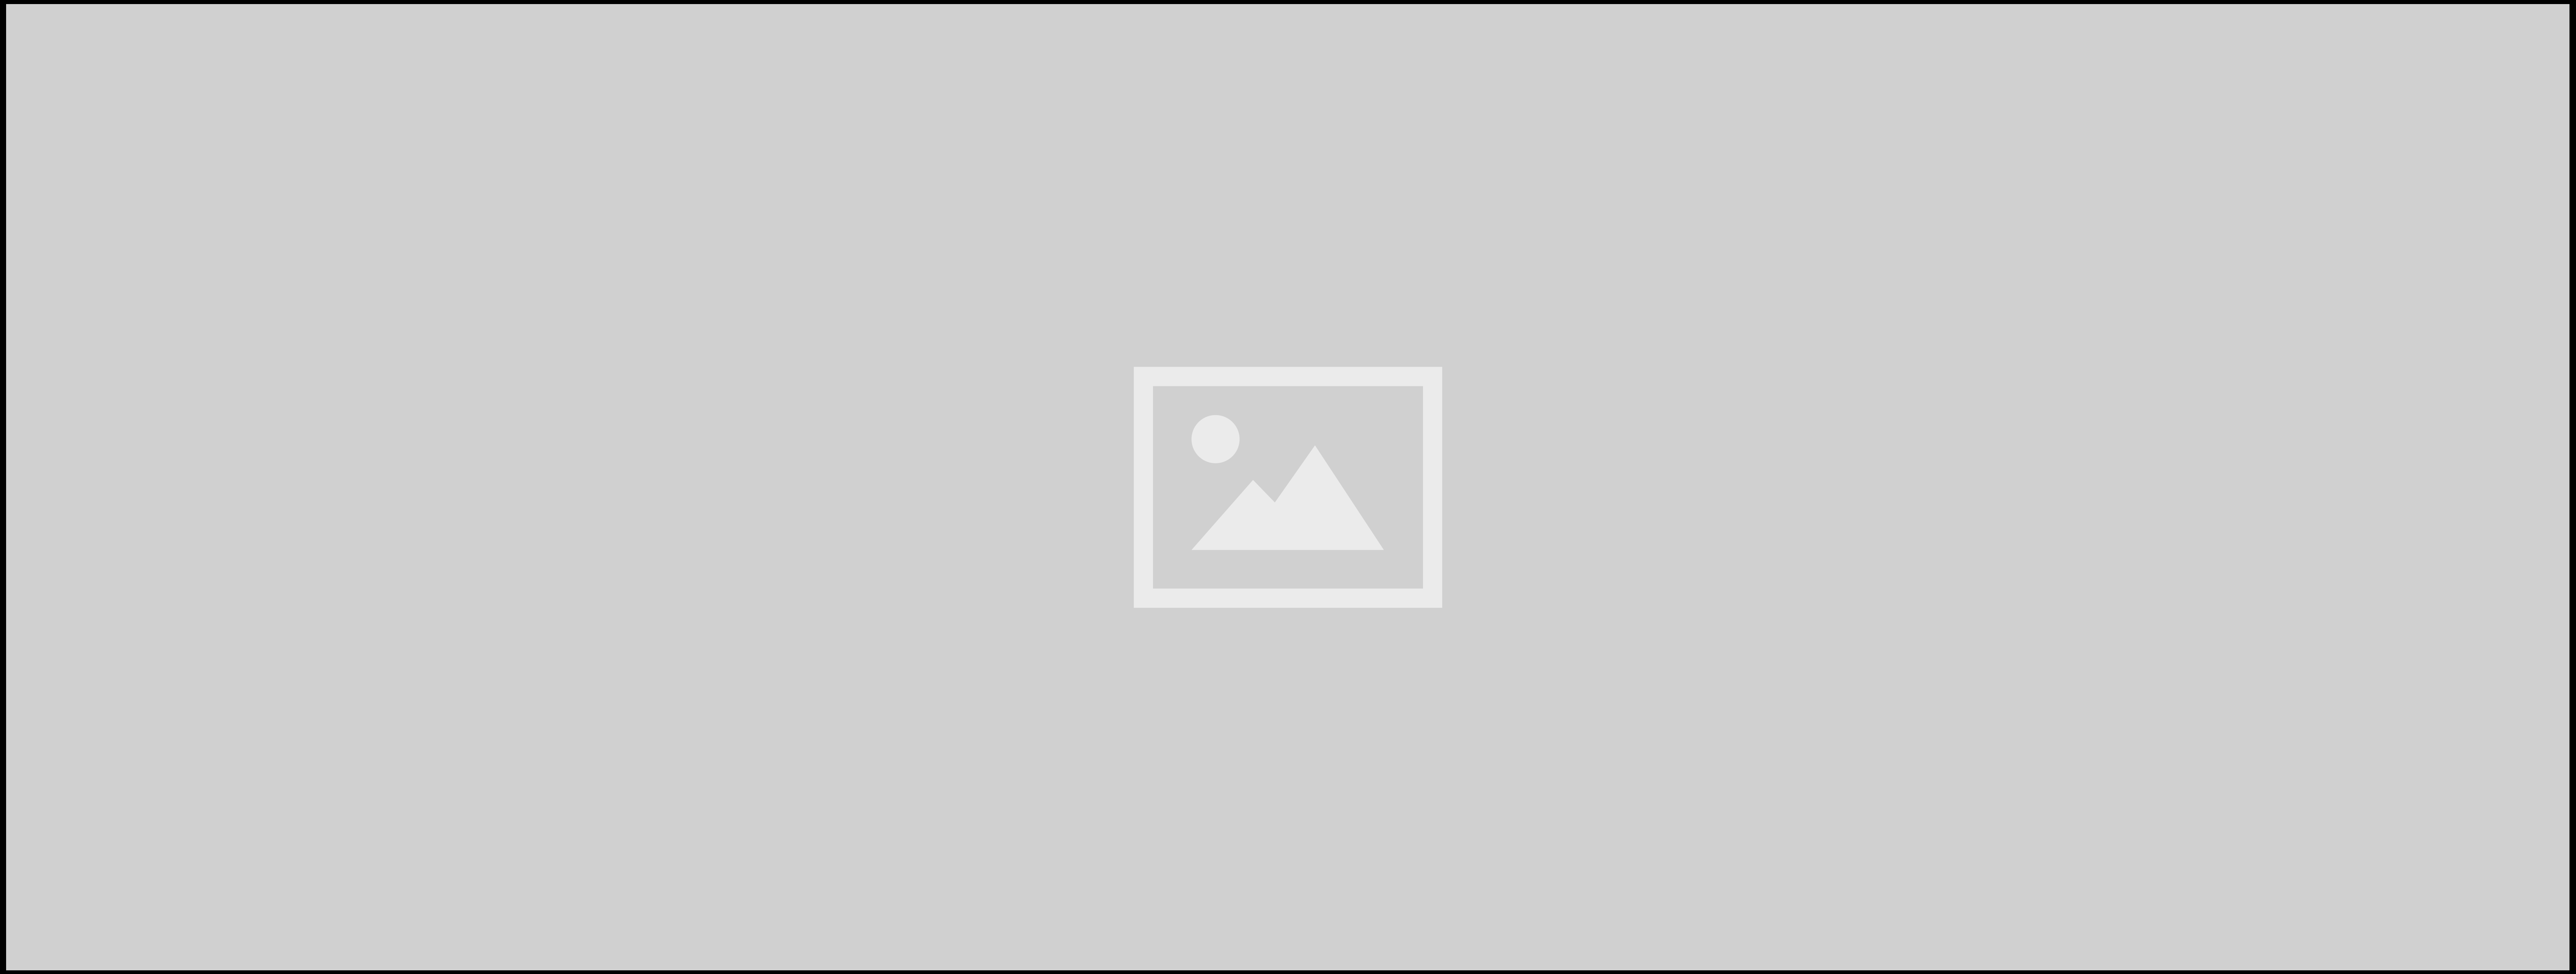

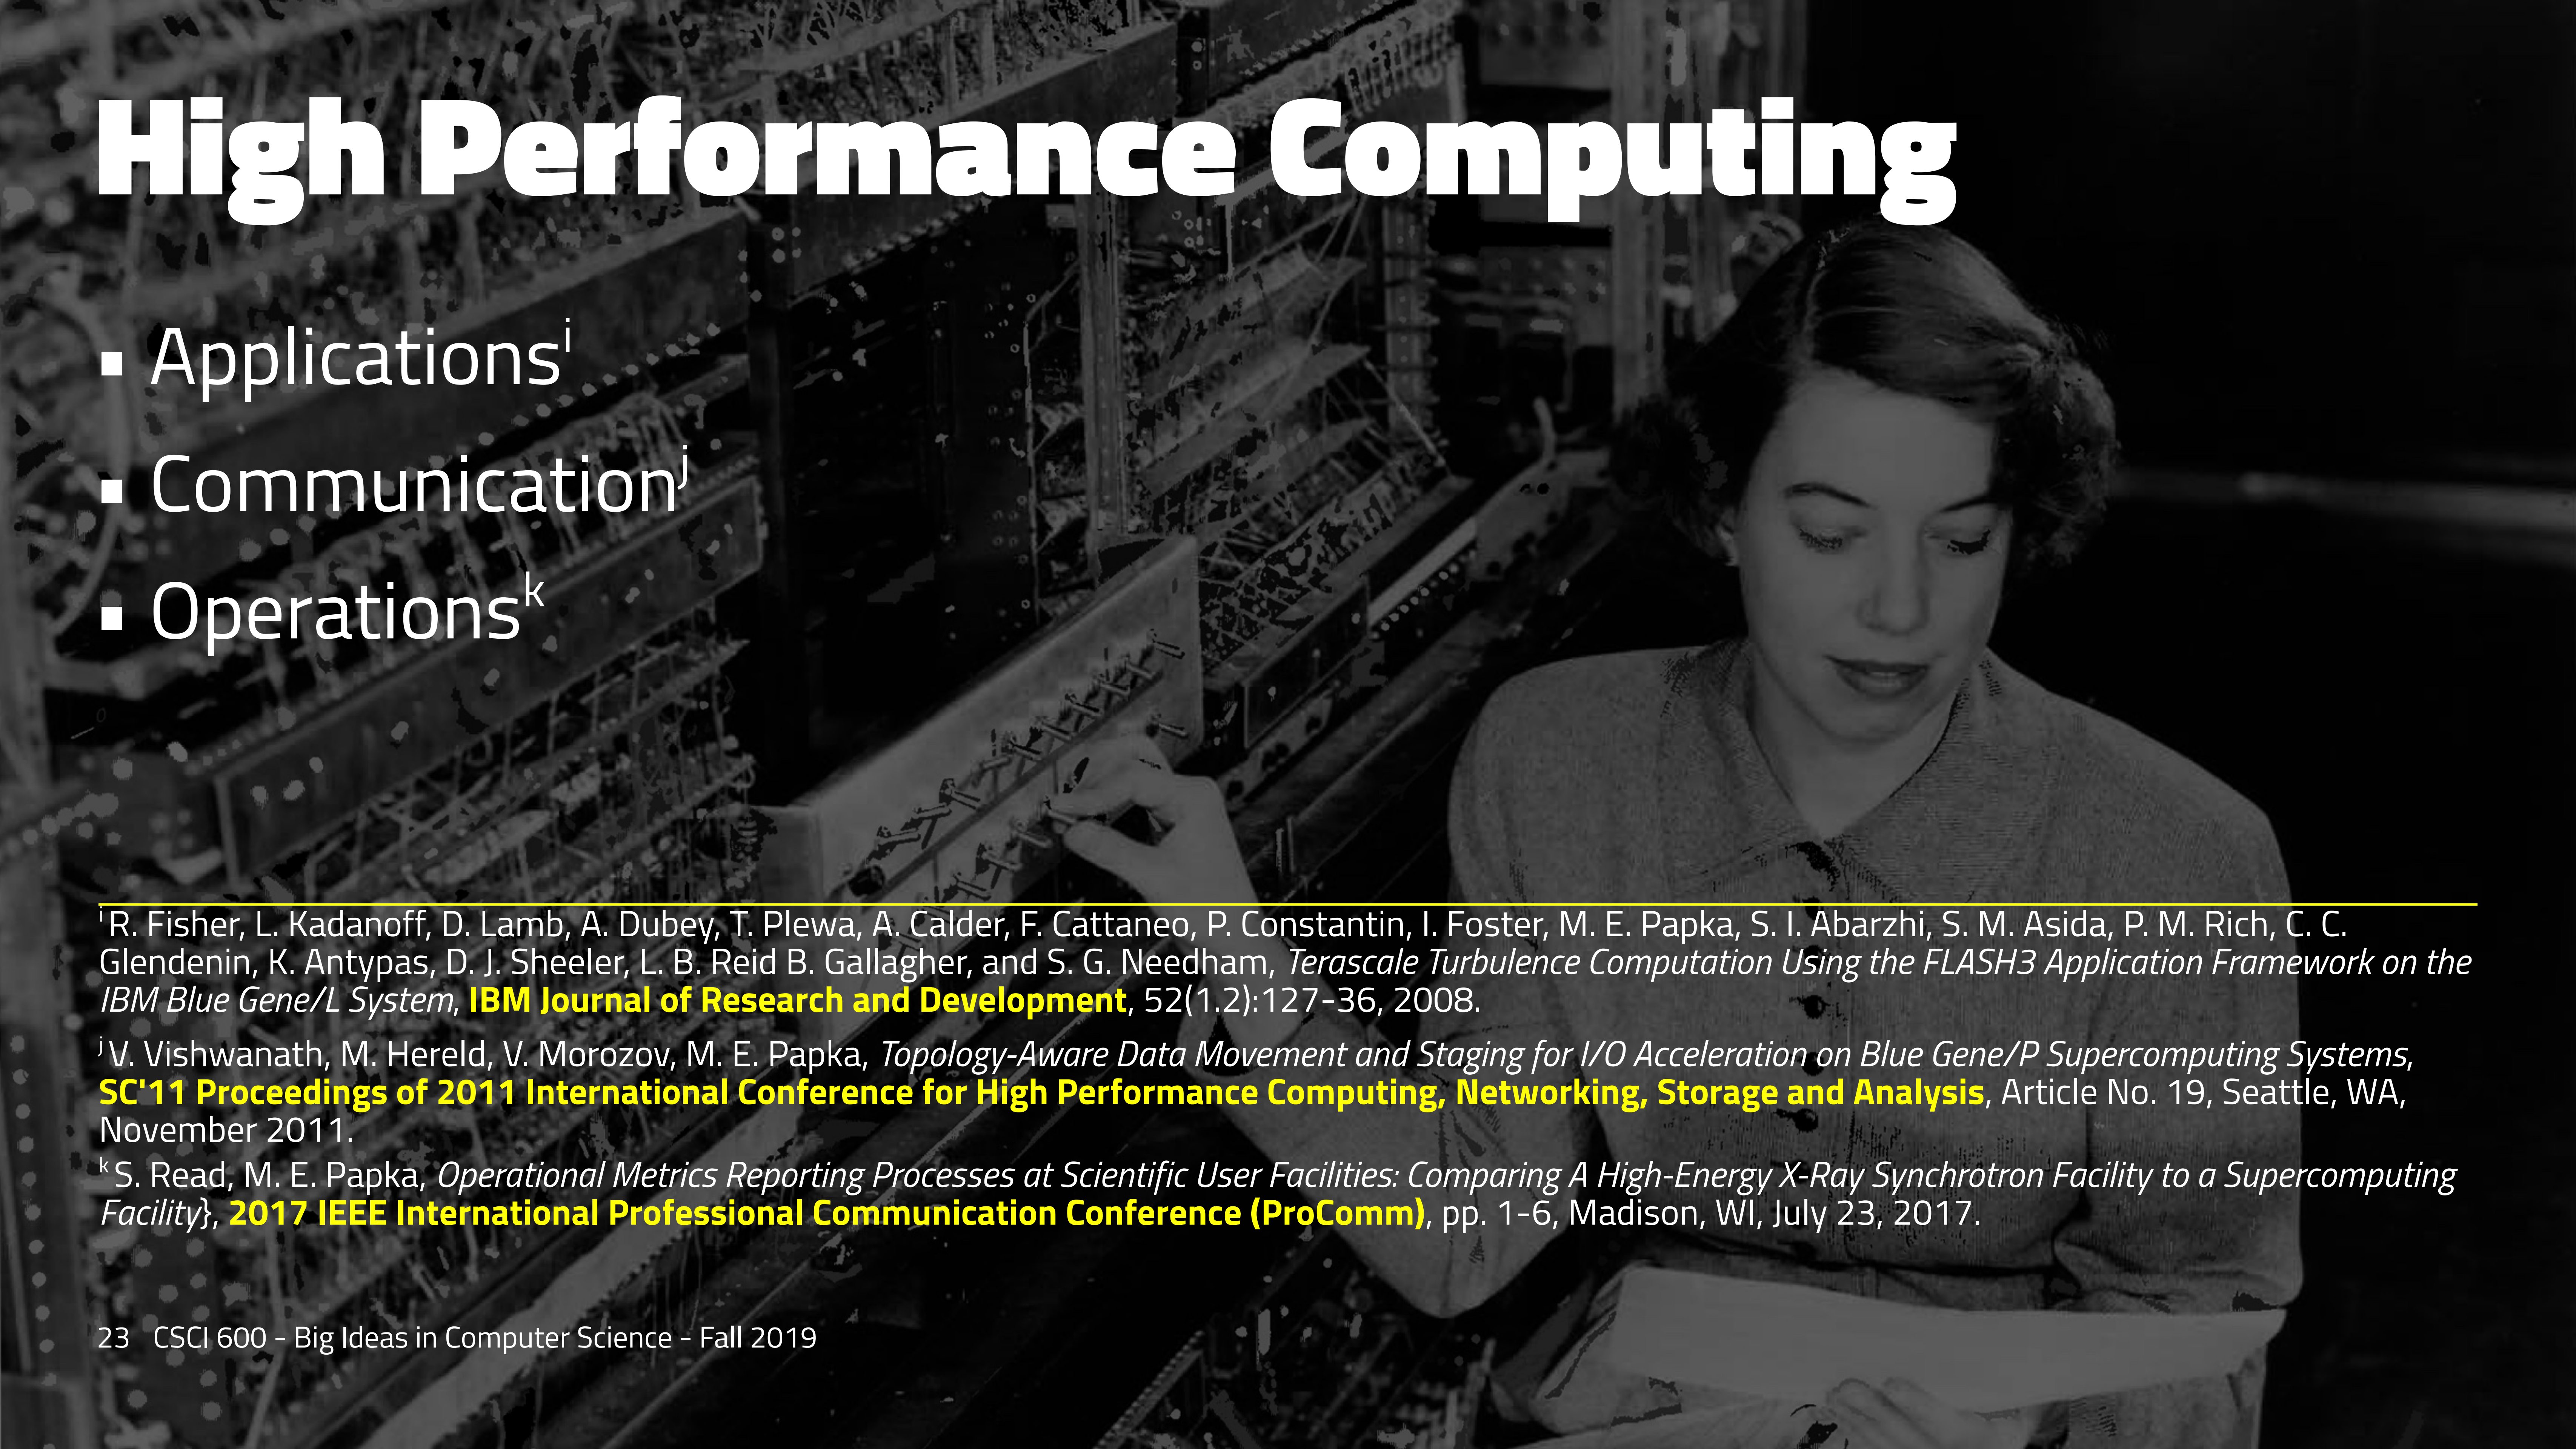

High Performance Computing

- Applications[i]
- Communication[j]
- Operations[k]

[i] R. Fisher, L. Kadanoff, D. Lamb, A. Dubey, T. Plewa, A. Calder, F. Cattaneo, P. Constantin, I. Foster, M. E. Papka, S. I. Abarzhi, S. M. Asida, P. M. Rich, C. C. Glendenin, K. Antypas, D. J. Sheeler, L. B. Reid B. Gallagher, and S. G. Needham, *Terascale Turbulence Computation Using the FLASH3 Application Framework on the IBM Blue Gene/L System*, **IBM Journal of Research and Development**, 52(1.2):127-36, 2008.

[j] V. Vishwanath, M. Hereld, V. Morozov, M. E. Papka, *Topology-Aware Data Movement and Staging for I/O Acceleration on Blue Gene/P Supercomputing Systems*, **SC'11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis**, Article No. 19, Seattle, WA, November 2011.

[k] S. Read, M. E. Papka, *Operational Metrics Reporting Processes at Scientific User Facilities: Comparing A High-Energy X-Ray Synchrotron Facility to a Supercomputing Facility*}, **2017 IEEE International Professional Communication Conference (ProComm)**, pp. 1-6, Madison, WI, July 23, 2017.

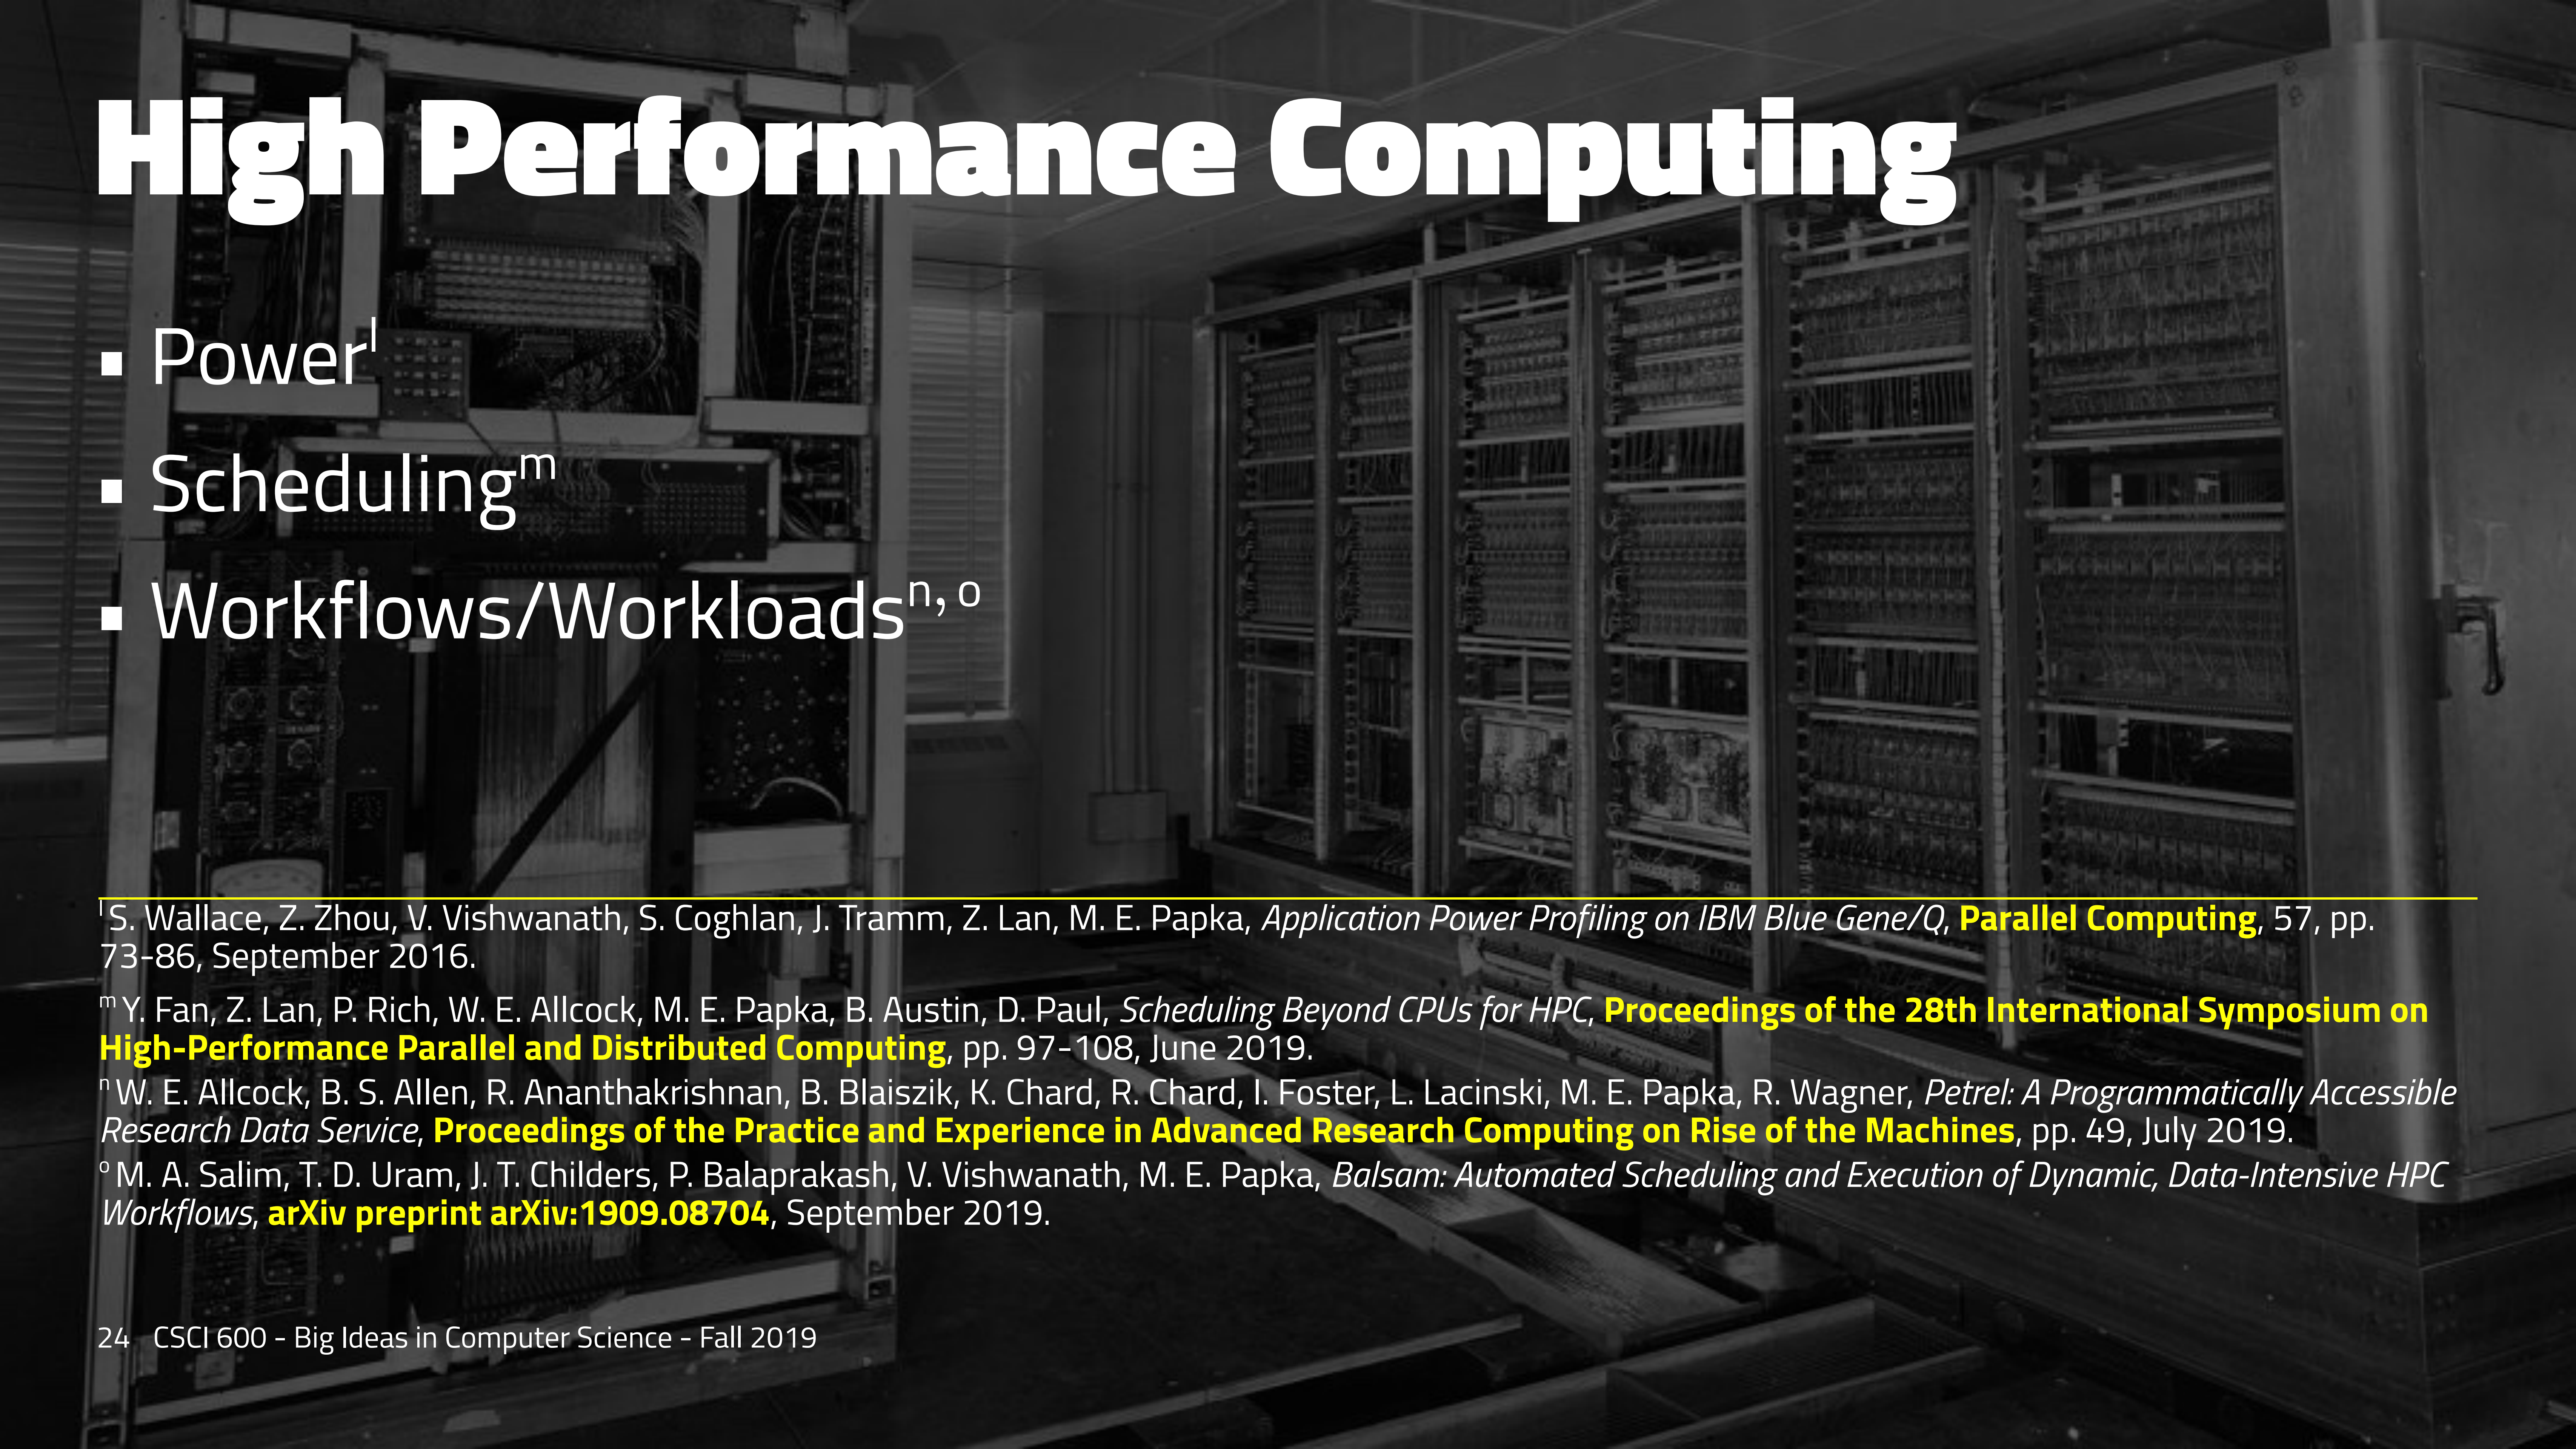

High Performance Computing

- Power[l]
- Scheduling[m]
- Workflows/Workloads[n, o]

[l] S. Wallace, Z. Zhou, V. Vishwanath, S. Coghlan, J. Tramm, Z. Lan, M. E. Papka, *Application Power Profiling on IBM Blue Gene/Q*, **Parallel Computing**, 57, pp. 73-86, September 2016.

[m] Y. Fan, Z. Lan, P. Rich, W. E. Allcock, M. E. Papka, B. Austin, D. Paul, *Scheduling Beyond CPUs for HPC*, **Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing**, pp. 97-108, June 2019.

[n] W. E. Allcock, B. S. Allen, R. Ananthakrishnan, B. Blaiszik, K. Chard, R. Chard, I. Foster, L. Lacinski, M. E. Papka, R. Wagner, *Petrel: A Programmatically Accessible Research Data Service*, **Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines**, pp. 49, July 2019.

[o] M. A. Salim, T. D. Uram, J. T. Childers, P. Balaprakash, V. Vishwanath, M. E. Papka, *Balsam: Automated Scheduling and Execution of Dynamic, Data-Intensive HPC Workflows*, **arXiv preprint arXiv:1909.08704**, September 2019.

Traditional[P]

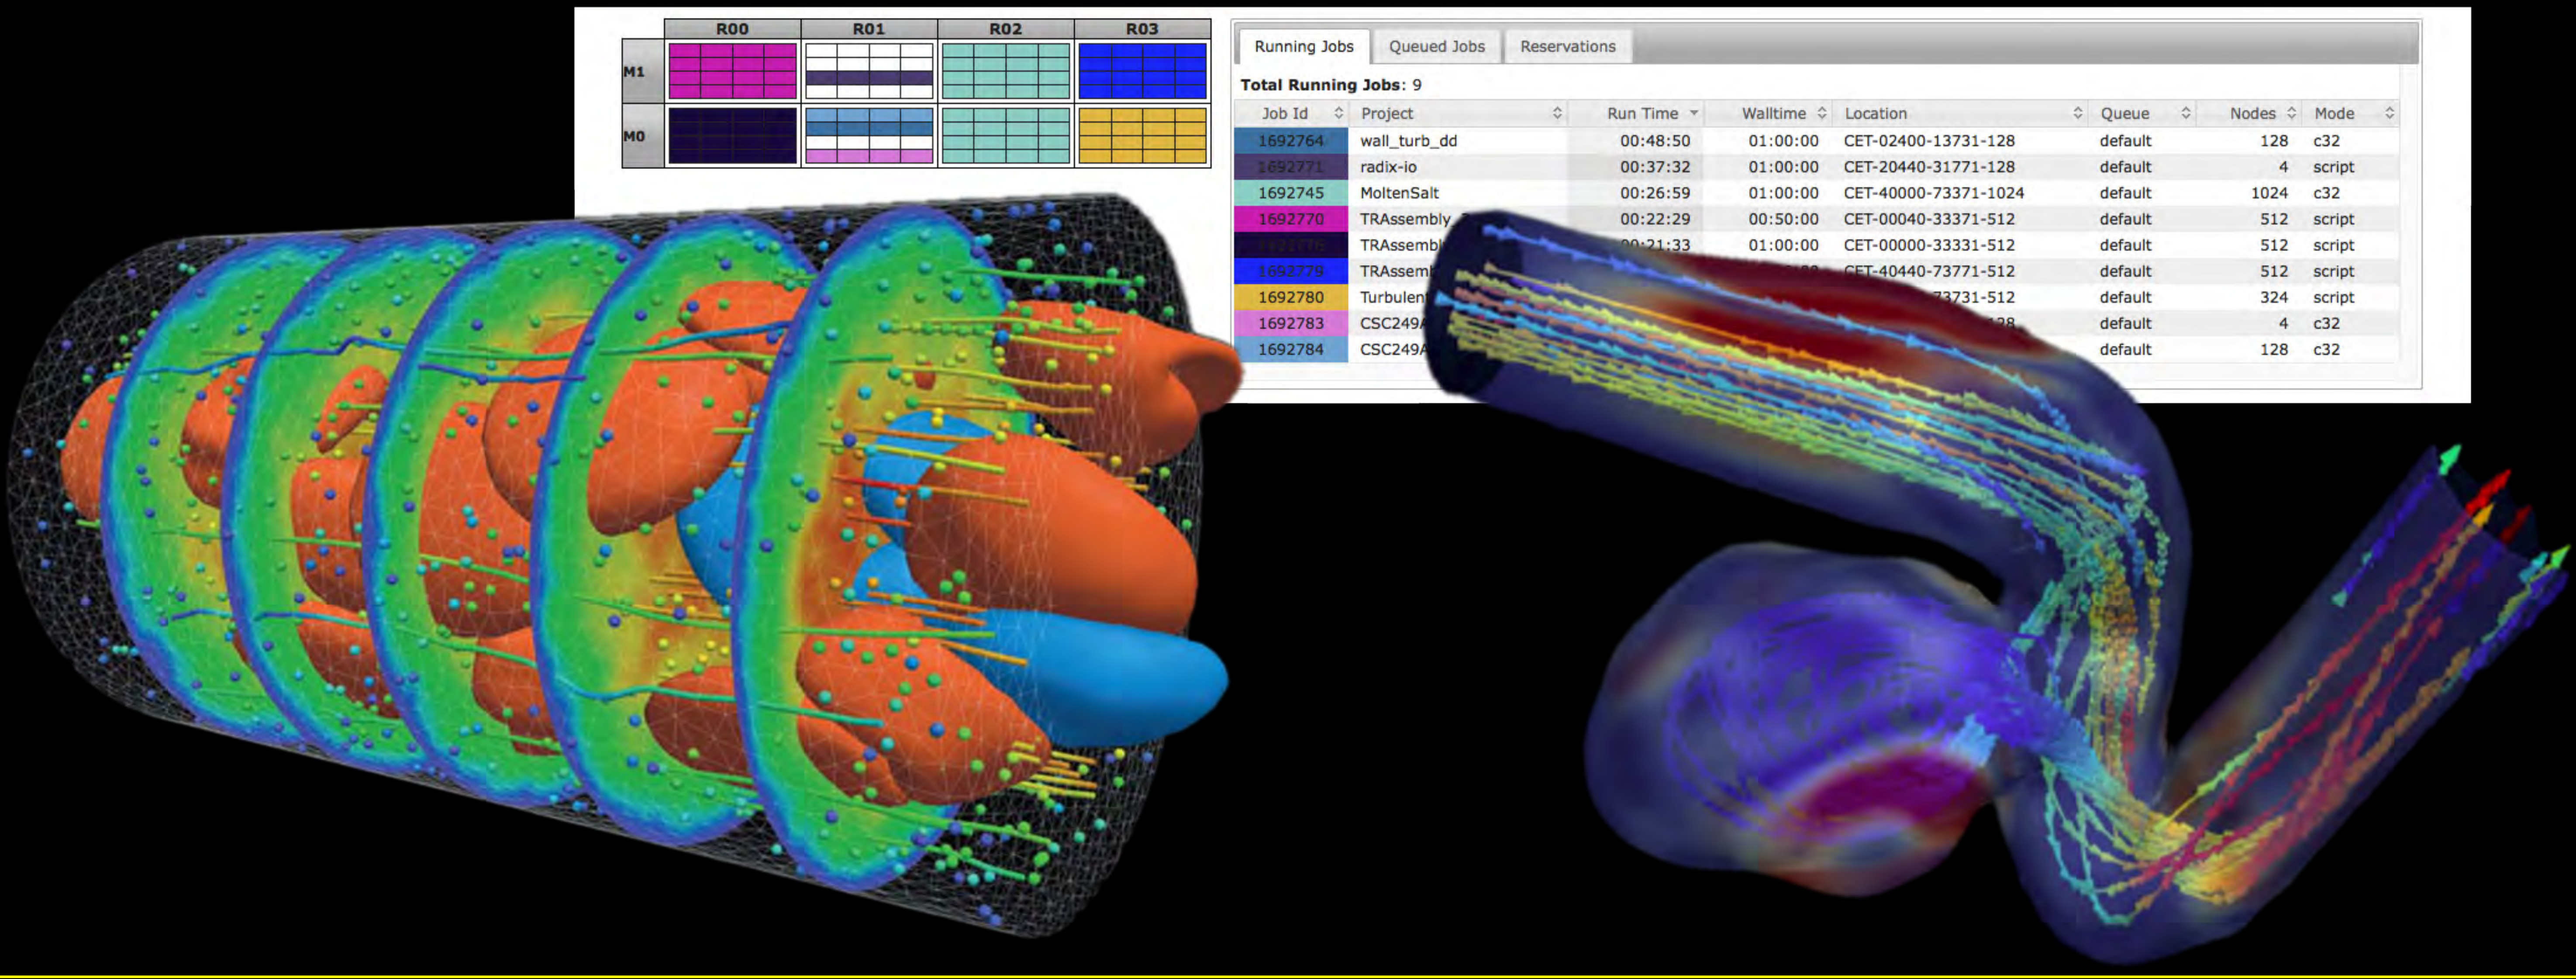

[P] L. Grinberg, J. A. Insley, D. Fedosov, V. A. Morozov, M. E. Papka, G. E. Karniadakis, *Tightly Coupled Atomistic-Continuum Simulations of Brain Blood Flow on Petaflop Supercomputers*, **Computing in Science and Engineering**, 14(6):58-67, 2012.]

Evolving (scheduling constraints)[q]

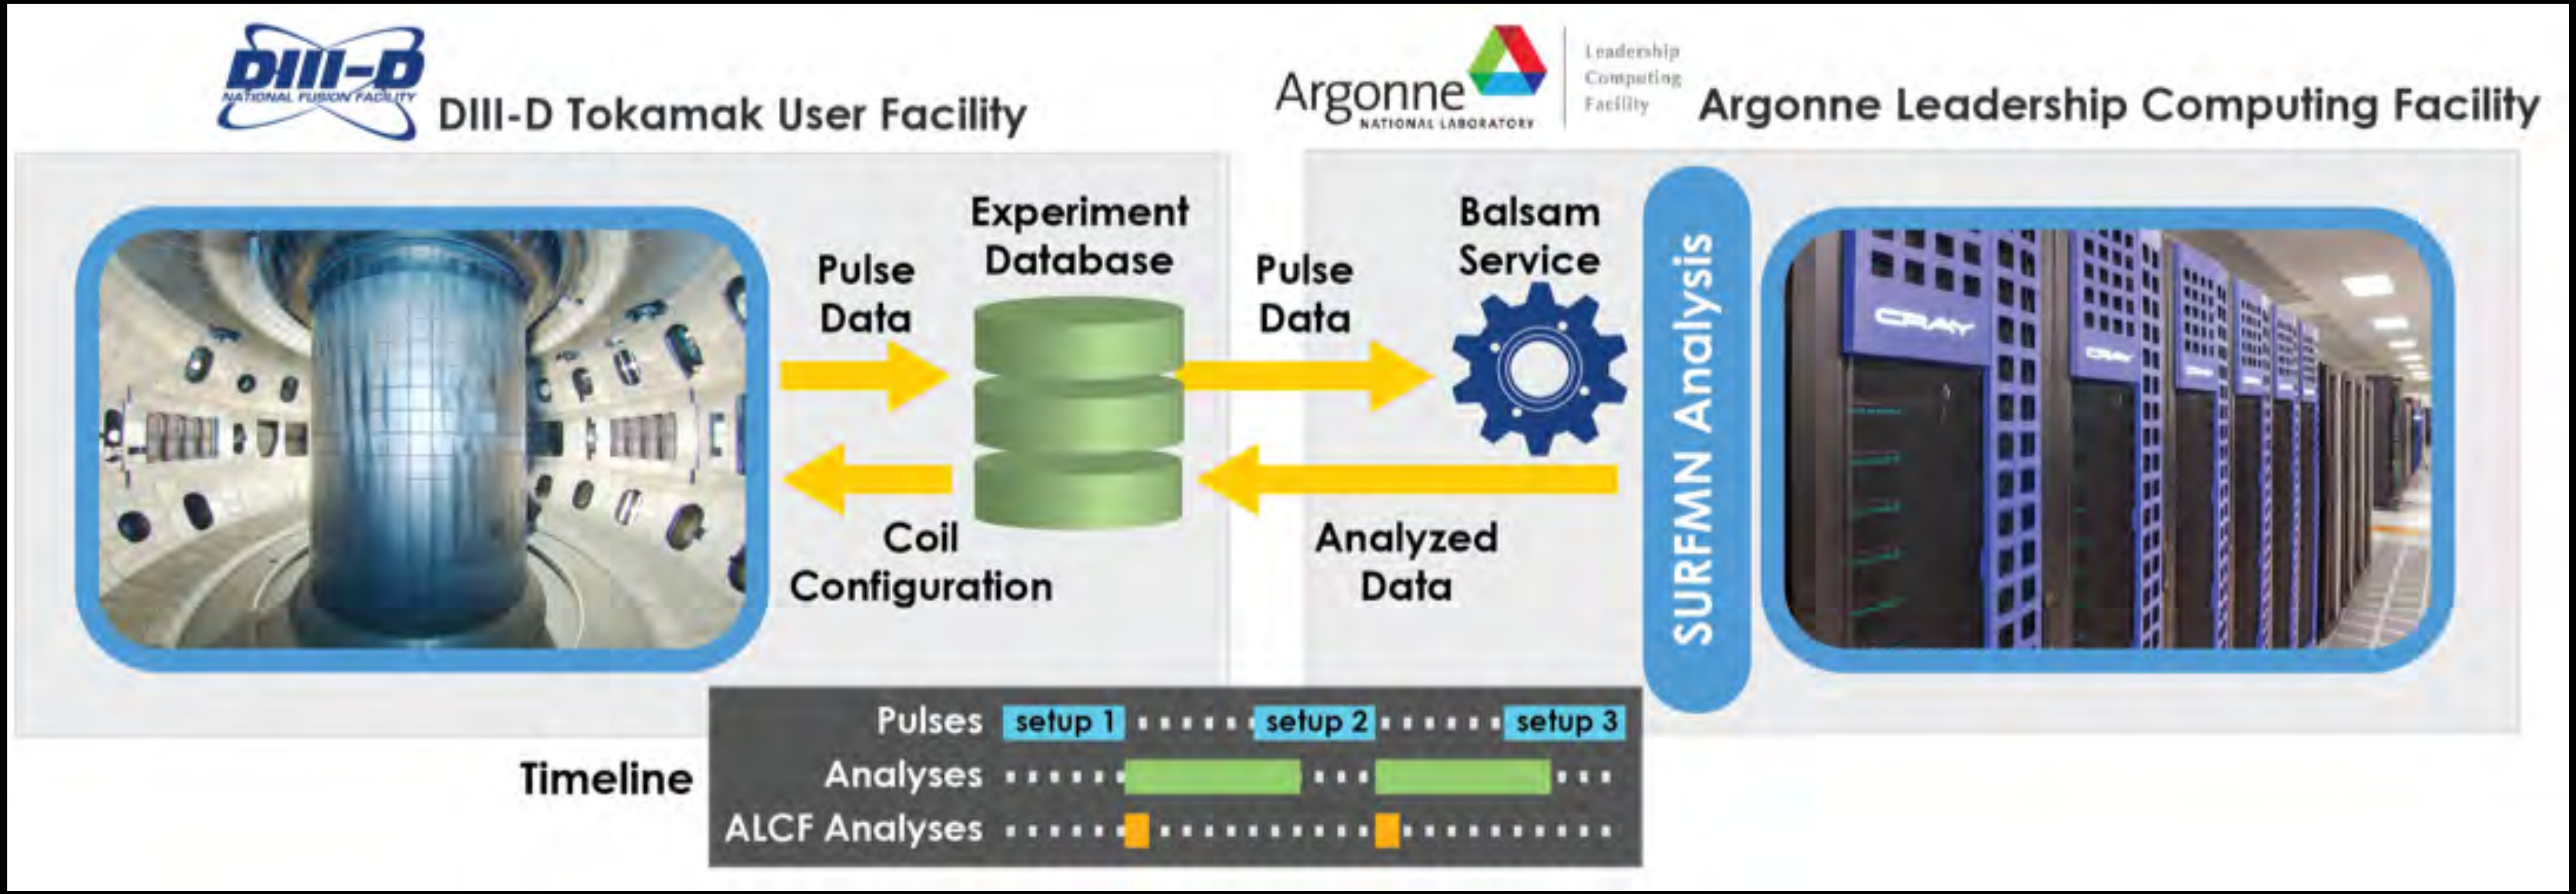

[q] M. Kostuk, T. D. Uram, T. Evans, D. M. Orlov, M. E. Papka, D. Schissel, *Automatic Between-Pulse Analysis of DIII-D Experimental Data Performed Remotely on a Supercomputer at Argonne Leadership Computing Facility*, **Fusion Science and Technology**, September 2017.]

Evolving (complex workflows)[r]

50% of the ATLAS papers based on 2015 data use the HPC-produced computing in a demonstrable manner

– These would still eventually be written without the US HPC effort, but they probably would not exist today: the time-to-science has been dramatically shortened.

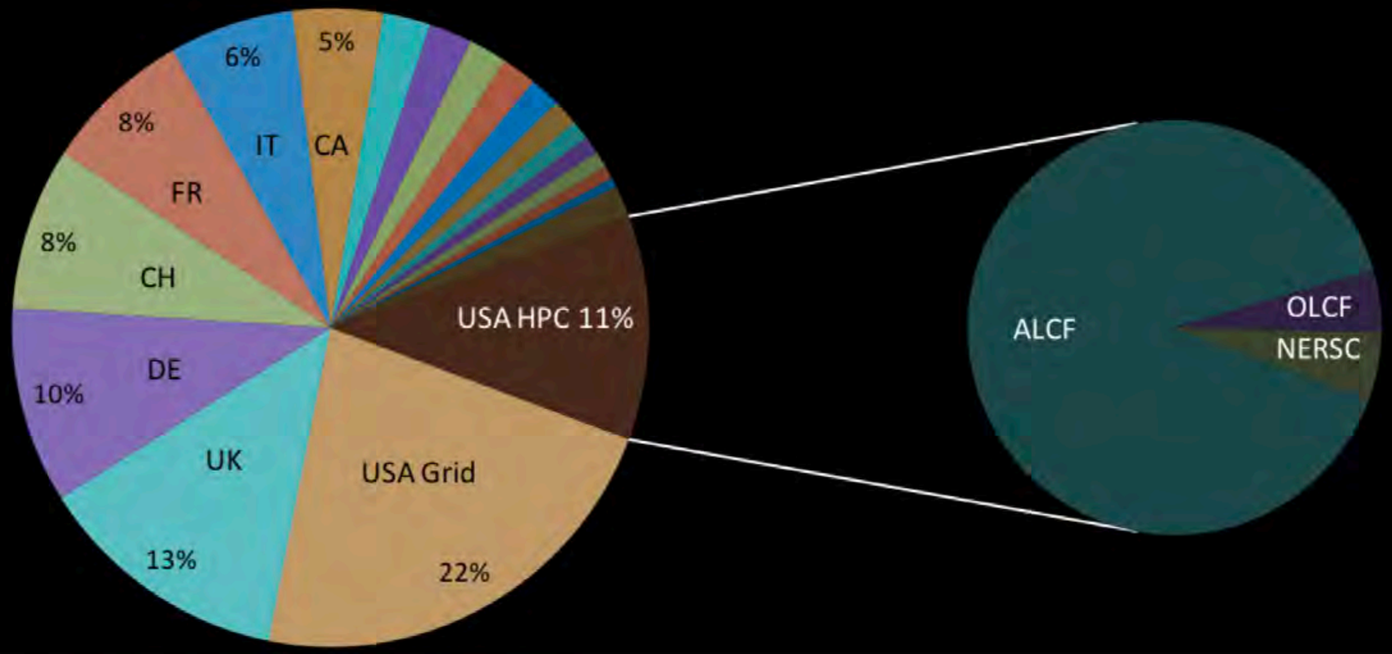

[r] T. LeCompte(HEP){06/02/2016} and J. T. Childers, T. D. Uram, D. Benjamin, T. J. LeCompte, M. E. Papka, *An Edge Service for Managing HPC Workflows*, **Proceedings of the Fourth International Workshop on HPC User Support Tools (HUST'17)**, Denver, CO, November 12, 2017.]

Evolving (increased engagement)[s]

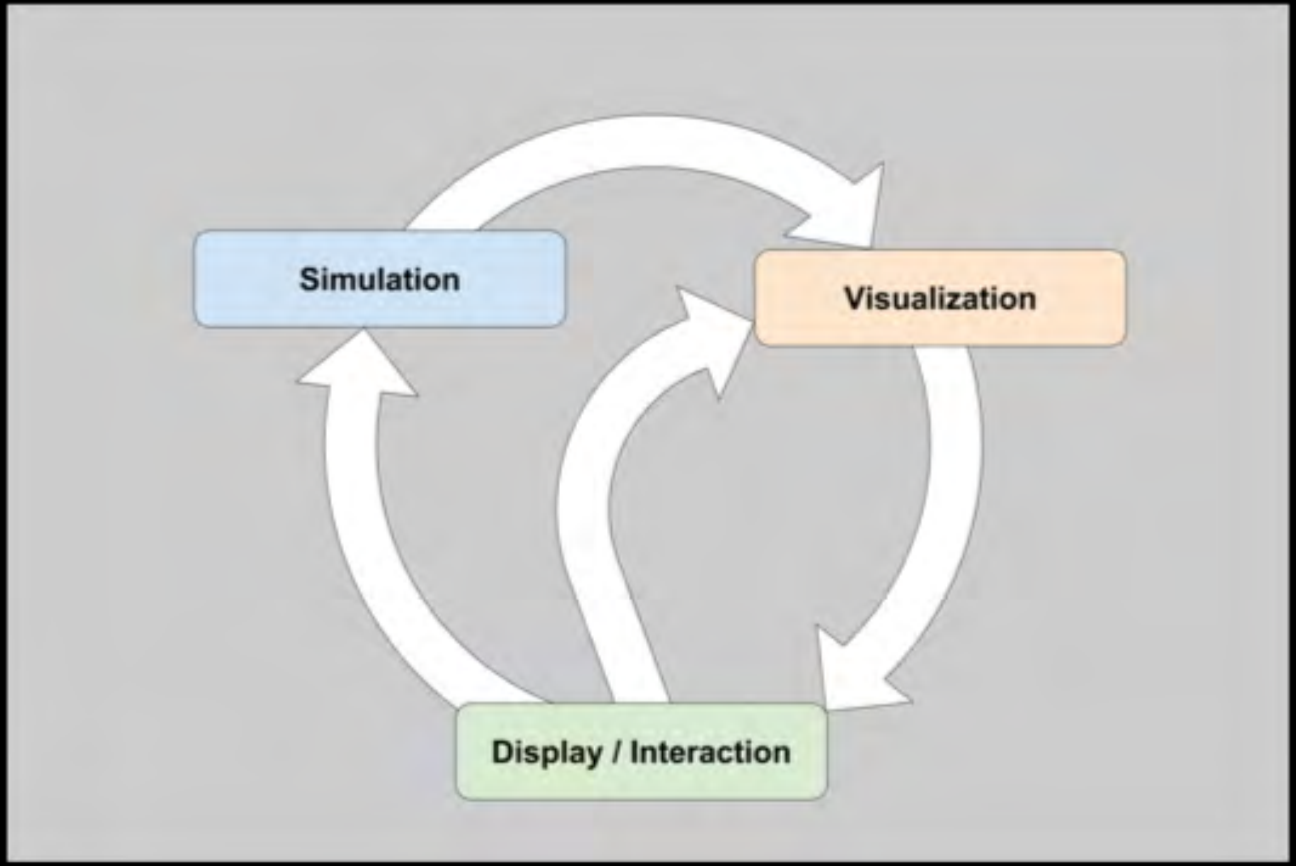

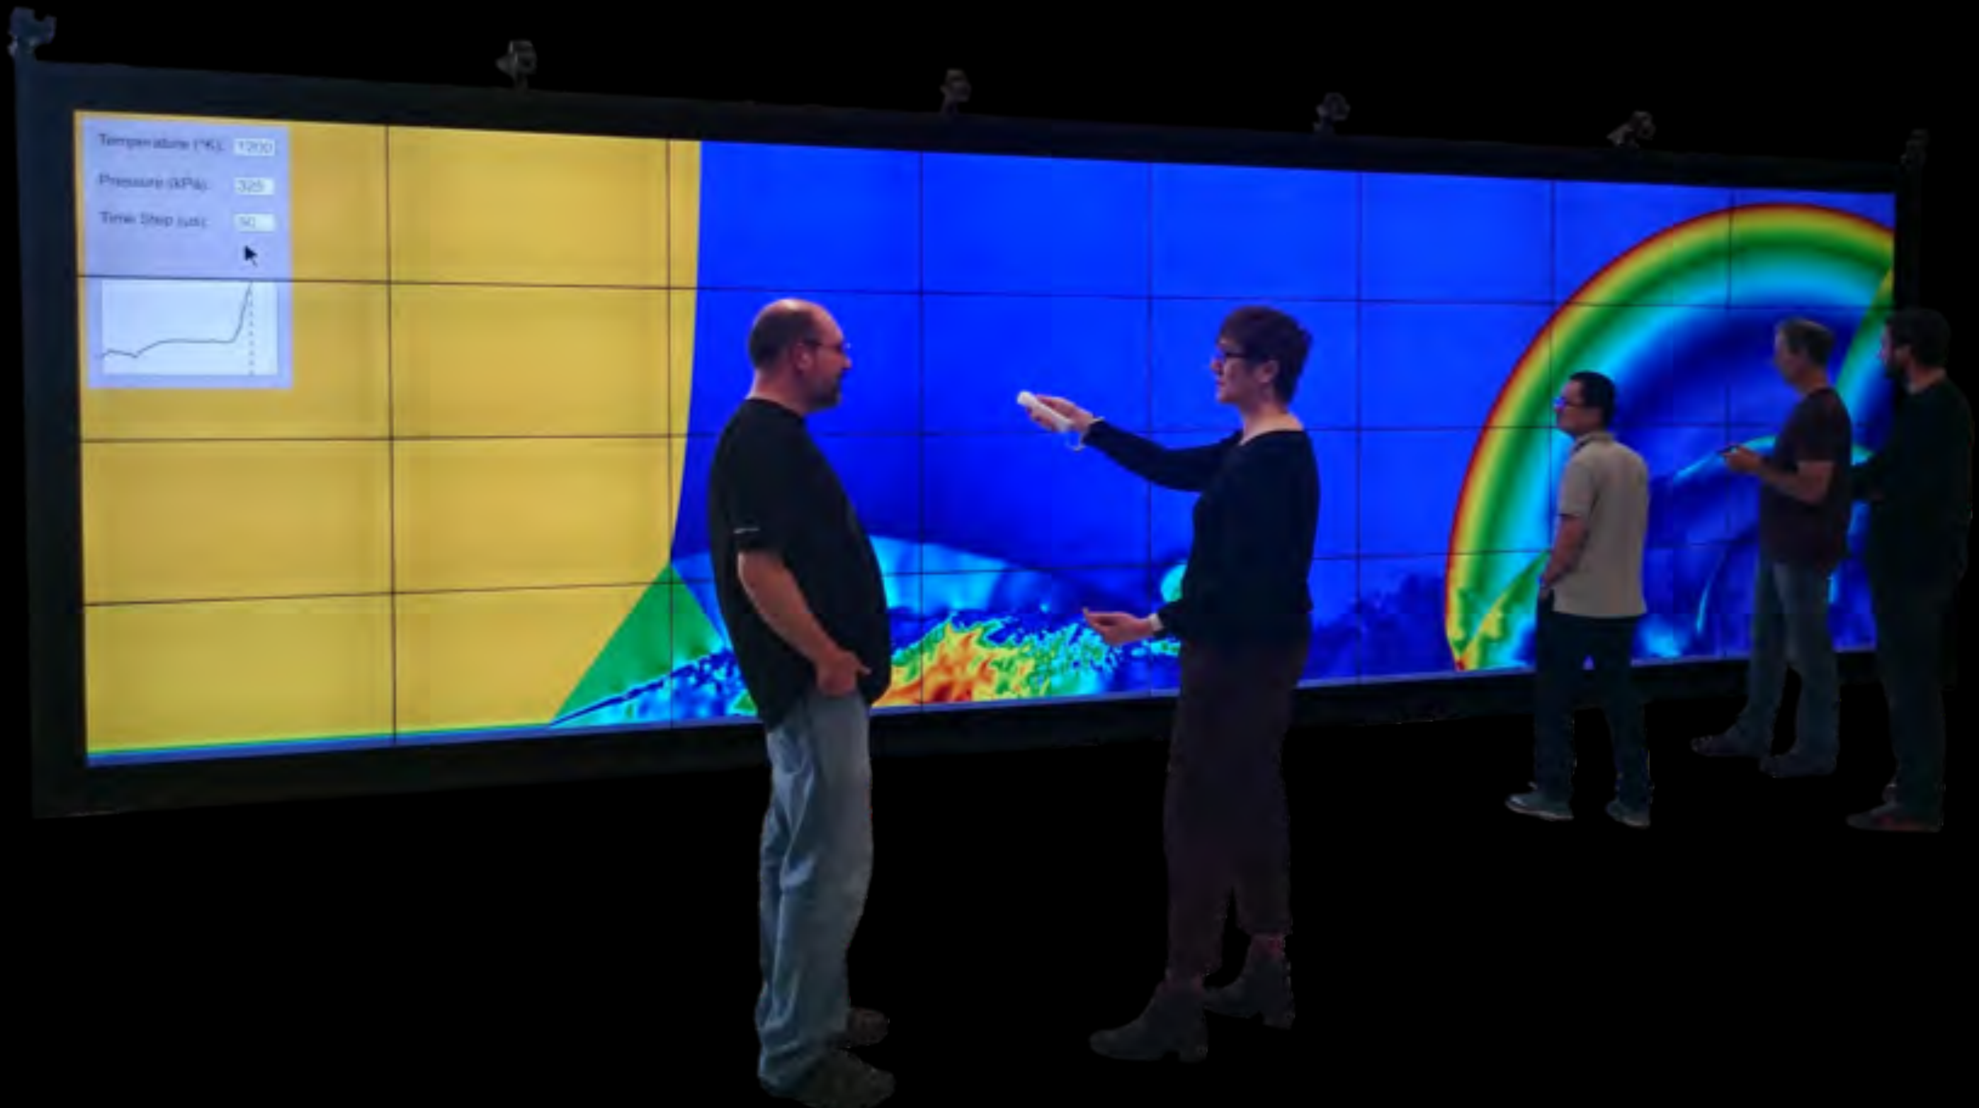

[s] T. Marrinan, A. Nishimoto, J. A. Insley, S. Rizzi, A. Johnson, M. E. Papka, *Interactive Multi-Modal Display Spaces for Visual Analysis*, **Proceedings of the 2016 ACM on Interactive Surfaces and Spaces**, pp. 421-426, Niagara Falls, Canada, November 6, 2016.]

HPC Environments

How do we **enable** scientists (users) to be the most **productive** from *start to finish*?

- How do we improve *usability*?

- How do we *enable* users of all levels?

- How do *simplify* supercomputing?

Workflow of Today

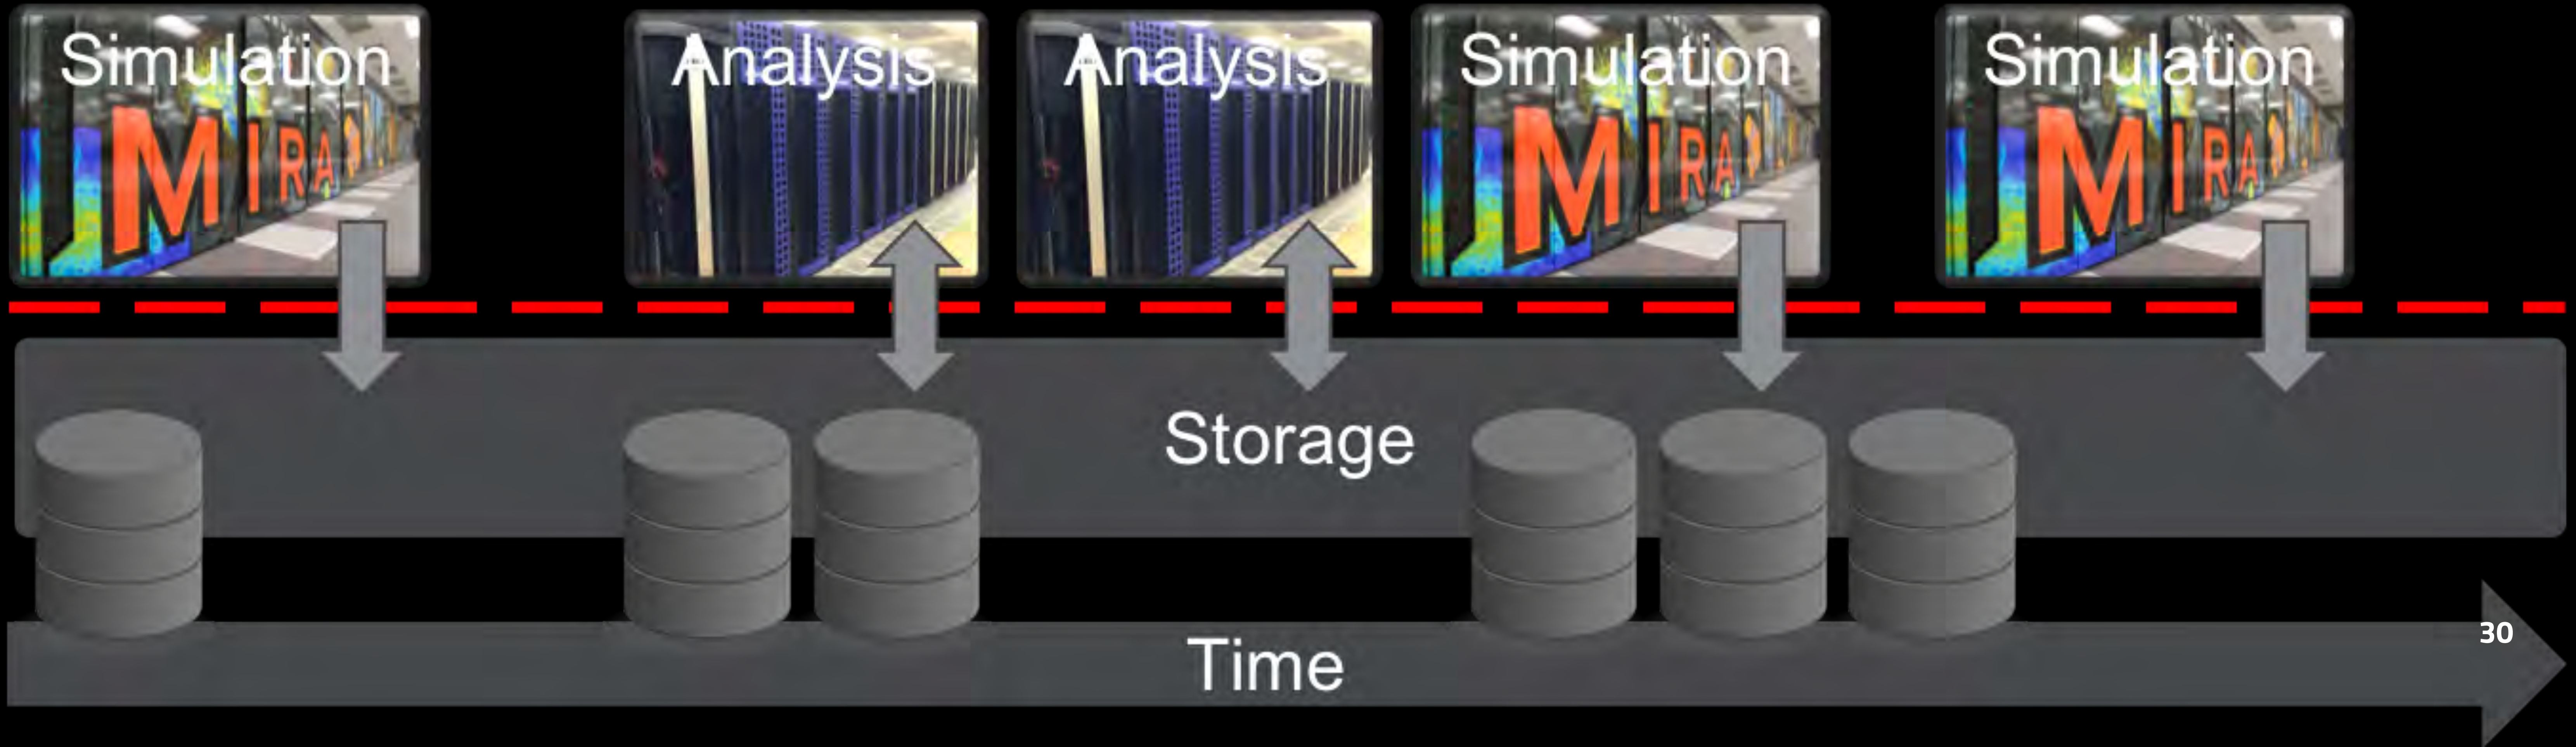

Workflow of Tomorrow (Today)

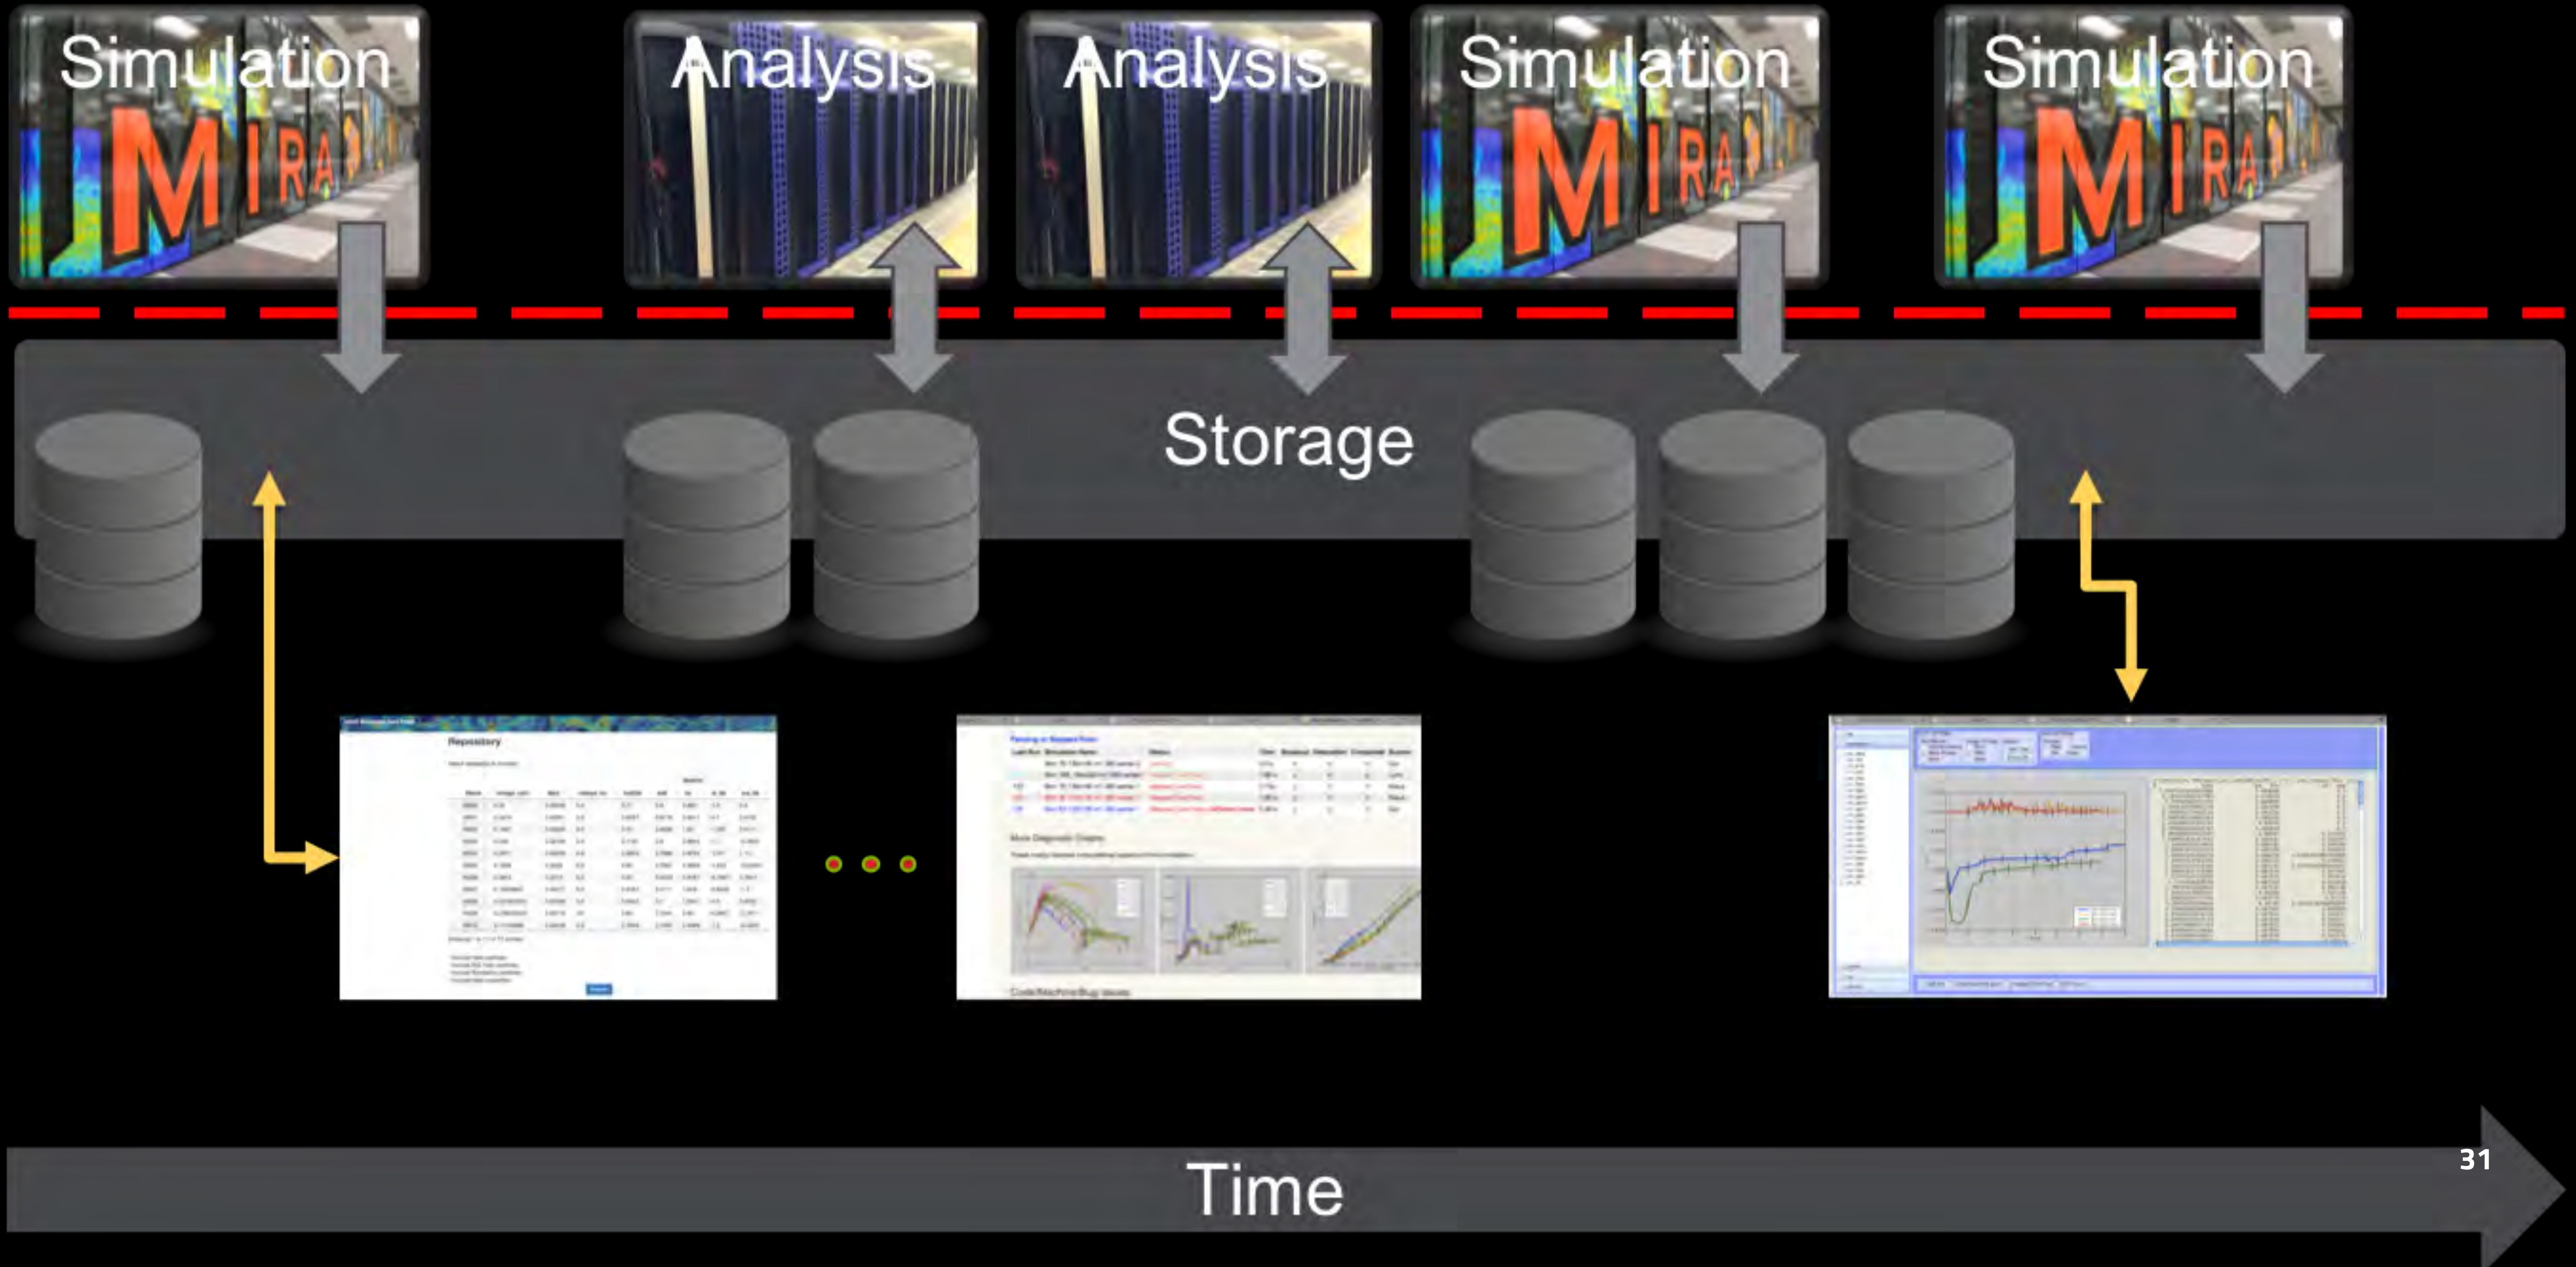

Facility of Tomorrow

Simulation

Analysis

User Community via Web and Experimental / Observational-Specific Clients

Jupyter Notebooks

Globus Online

PDACS

Machine Learning

Smaash

DESCQA

job submission/adaptation layer

ALCF-hosted
Experimental / Observational-Specific Resources
(servers like JupyterHub, databases, web, indexed search, visualization)

Facility Storage

Experimental / Observation Datasets

Cosmology

APS

SNS

LHC

Climate

Brain Imaging

Observations *(Science Management)*

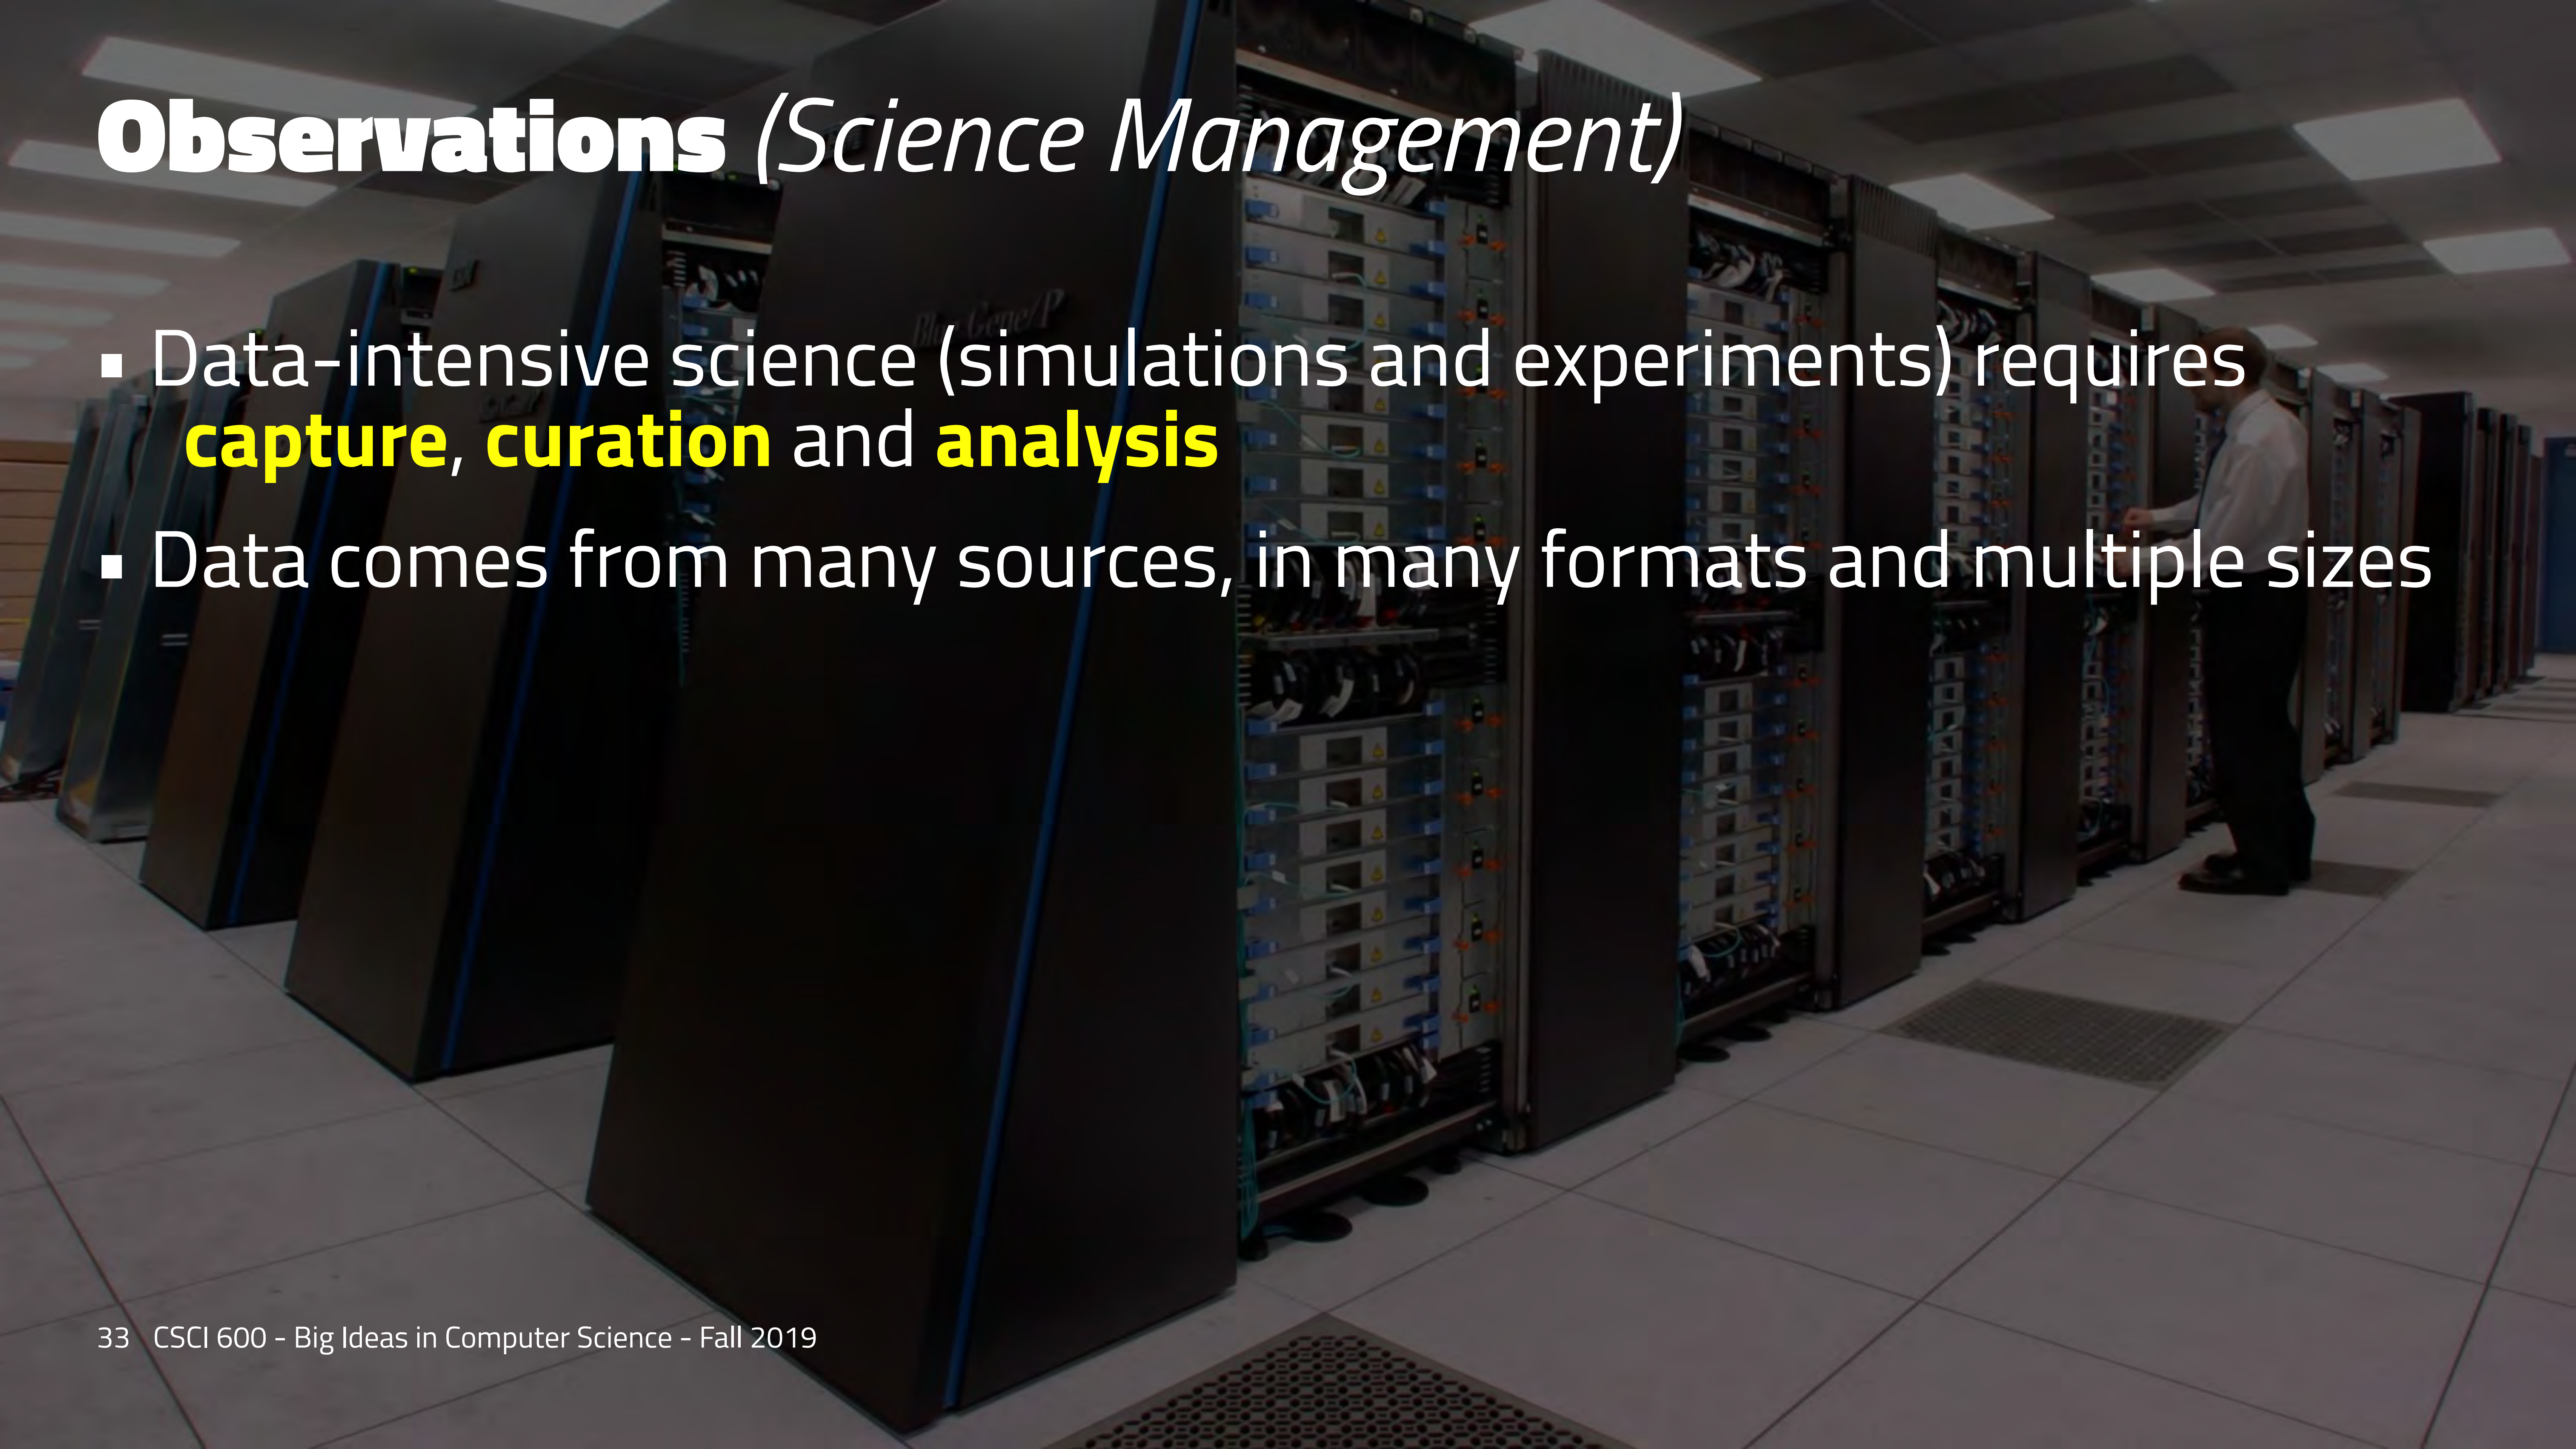

- Data-intensive science (simulations and experiments) requires **capture**, **curation** and **analysis**
- Data comes from many sources, in many formats and multiple sizes

Observations *(Science Management)*

- Problem with science management:
  - Tracking simulations and output **[difficult]**
  - Finding and reproducing old simulations: **[difficult]**
  - Monitoring live simulations: **[inconvenient, idiosyncratic]**
  - Post-processing, analysis and archival of results: **[haphazard]**
  - Assessing simulation behavior/performance: **[difficult]**

Increased Access to Scientific Communities

Support for Application Teams

Simulation **m**anagement and **a**nalysis system for Fl**ash** (Smaash)[t]

- Tracking and coordination of data (simulation and meta)
- Run-time monitoring of simulations and automated analysis of simulation output
- Method for managing / executing common workflows

[t] R. Hudson, J. Norris, L. B. Reid, K. Weide, G. C. Jordan, and M. E. Papka, *Experiences Using Smaash to Manage Data-Intensive Simulations*, **Proceedings of the 20th International Symposium on High-Performance Parallel and Distributed Computing**, pp. 205-15, San Jose, CA, June 2011.

Prototype Partner - Flash

- Multi-physics
- Adaptive-mesh

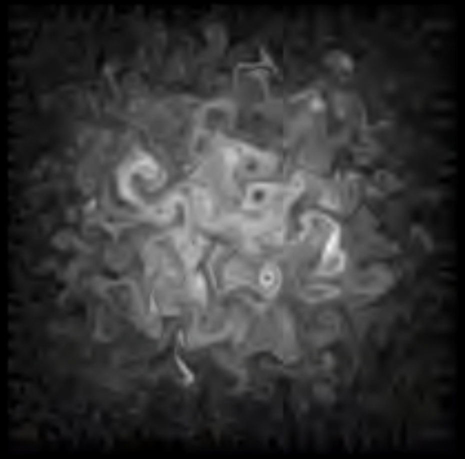

Compressible turbulence

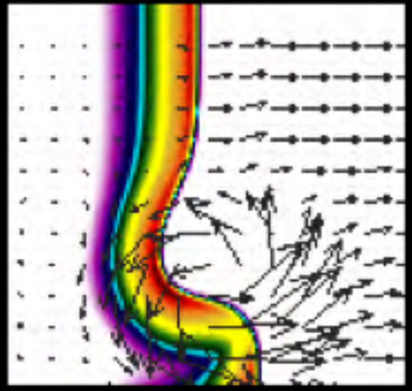

Flame-vortex interactions

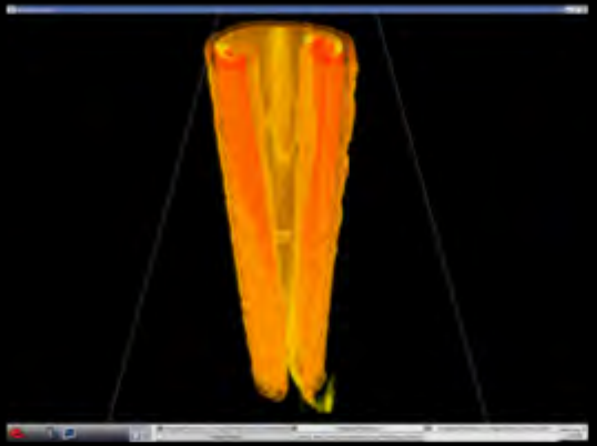

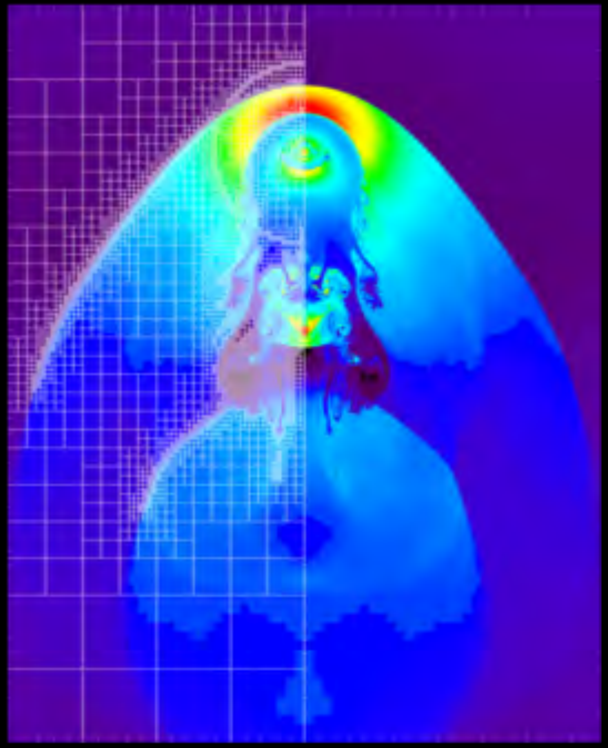

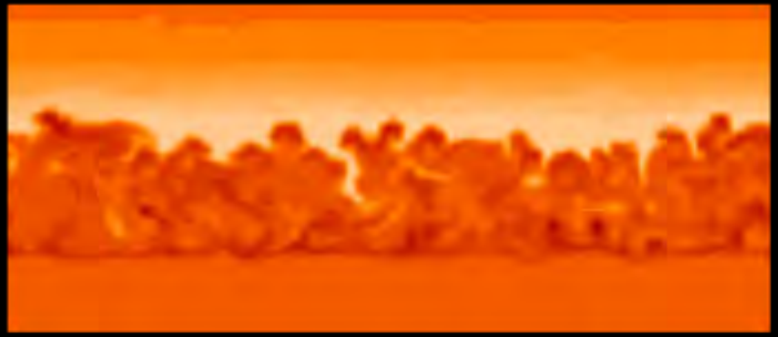

Nova outbursts on white dwarfs

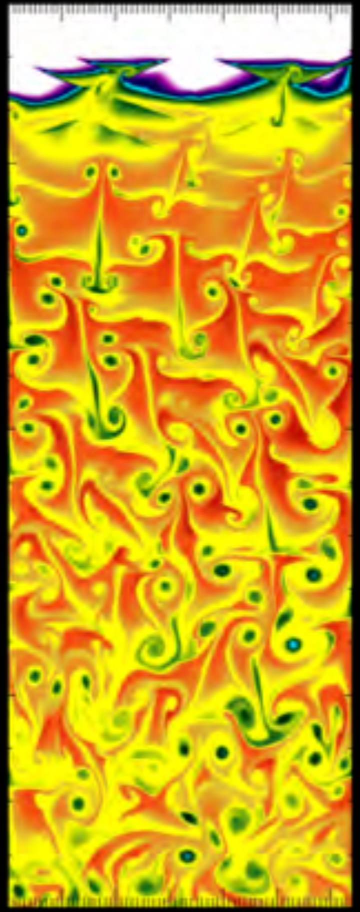

Cellular detonations

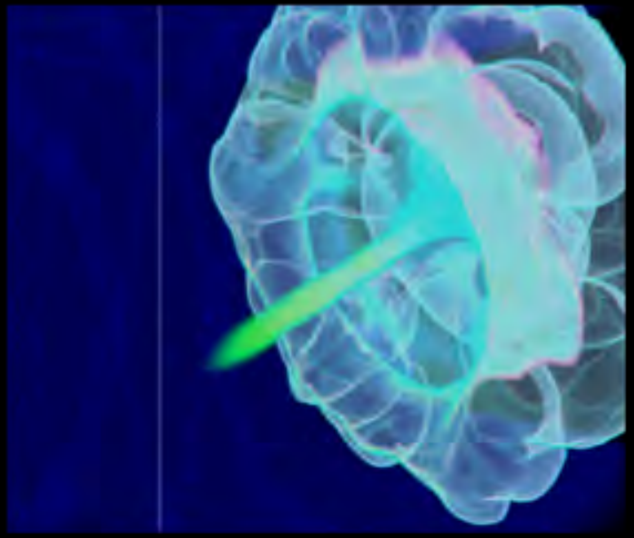

White

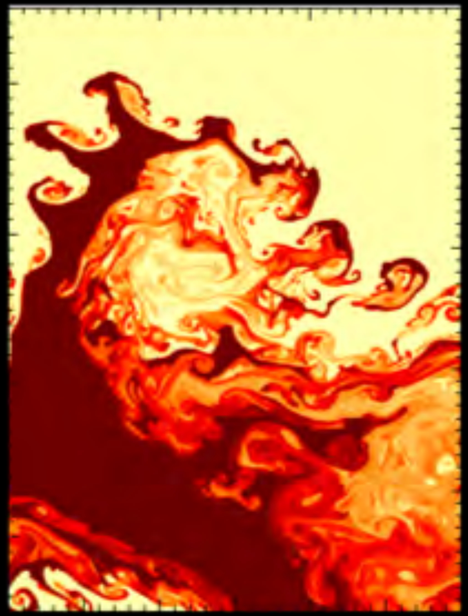

Rayleigh-Taylor instability

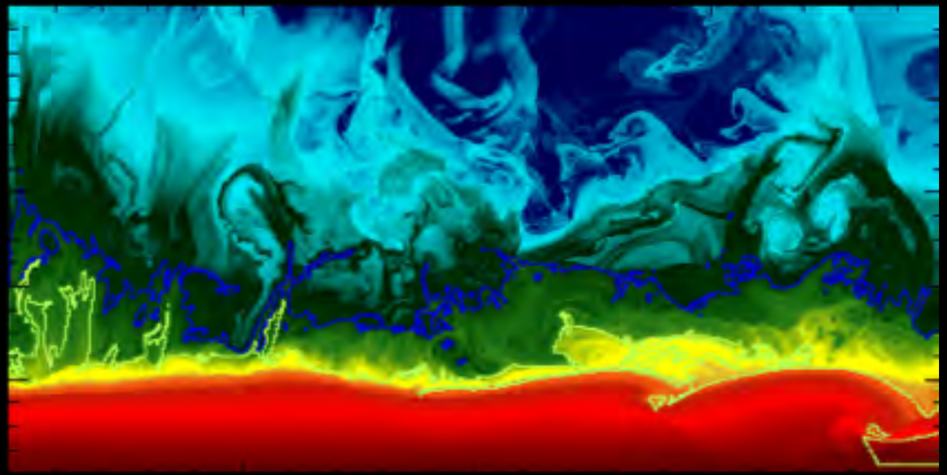

Helium burning on neutron stars

Prototype Partner - Flash

- Meta-data output
  - .log: simulation progress, warnings, errors, resource use
  - .dat: integrated grid quantities
- Scientific data output (HDF5)
  - Checkpoint: complete information needed to restart simulation
  - Plotfile: data values of interest for analysis
  - Particle files: tracer particles of interest during analysis

Smaash Components

- Database (manages meta-data)
- Back end services (co-located with compute resources and scientific data)
- Front end interfaces (user facing)

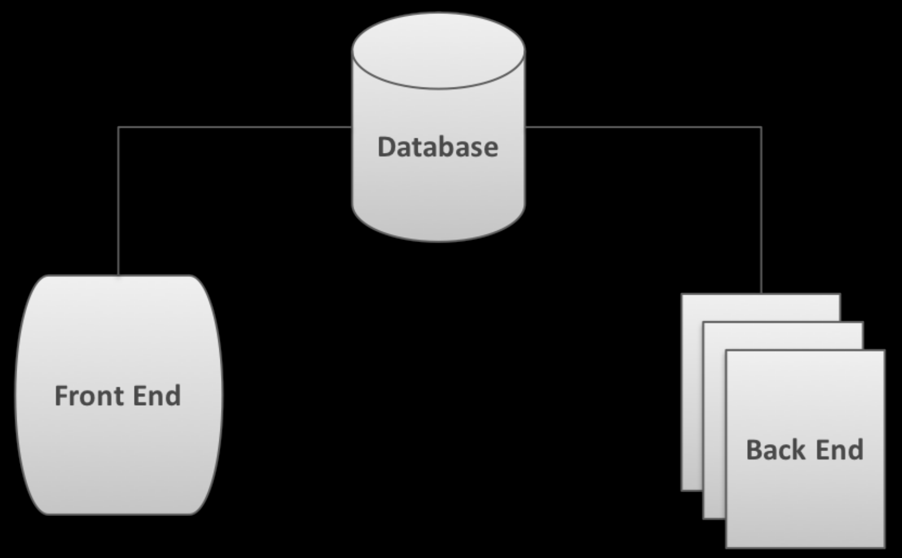

Smaash Implementation

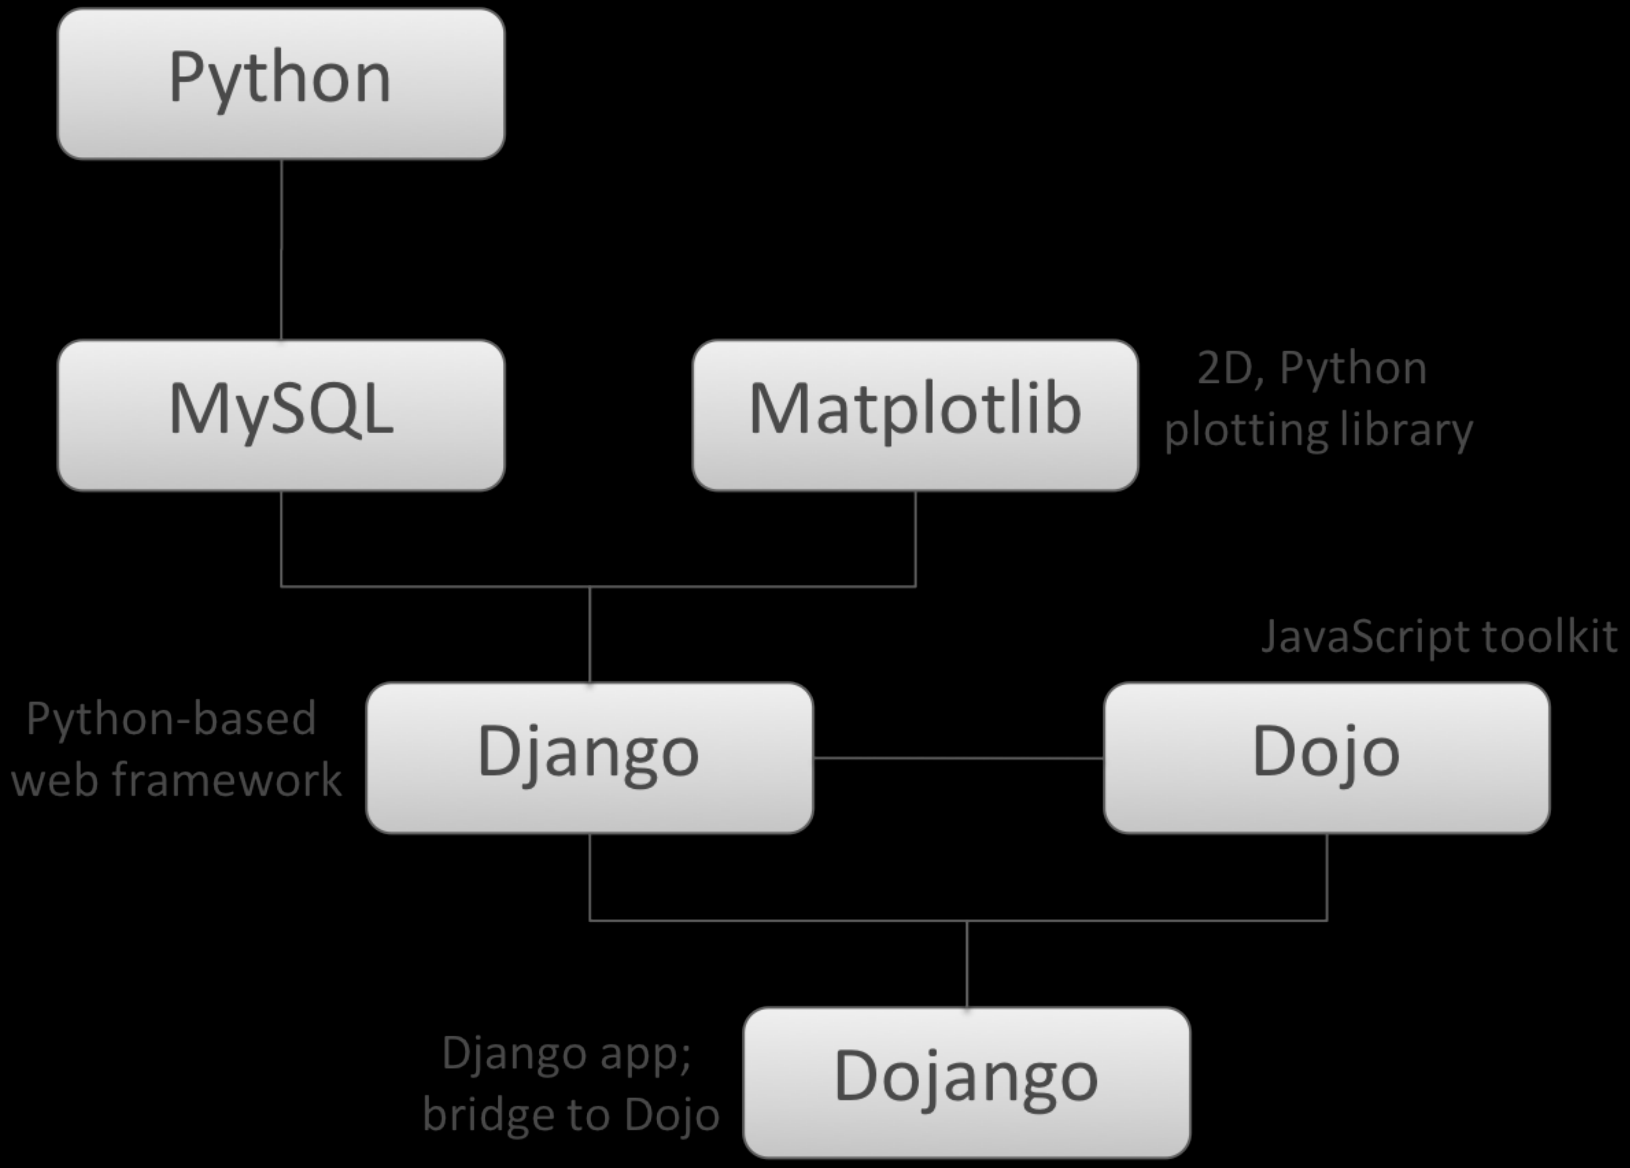

Smaash Back End Services

- Collector - captures and stores meta-data in database about simulation
- Archiver - automates the archiving of data
- Verifier - cross checks output and database entries
- Associator - connects a current simulation with campaign
- Observer - responsible for updates to user (email)
- Visualizer - automatic running of user specified visualization scripts

Smaash Front End Interfaces (Views)

- Tree - collection of campaigns, simulations and runs
- Graph - quick graphs of results
- Monitor - automated visualizations
- Summary - details and notes

Tree View

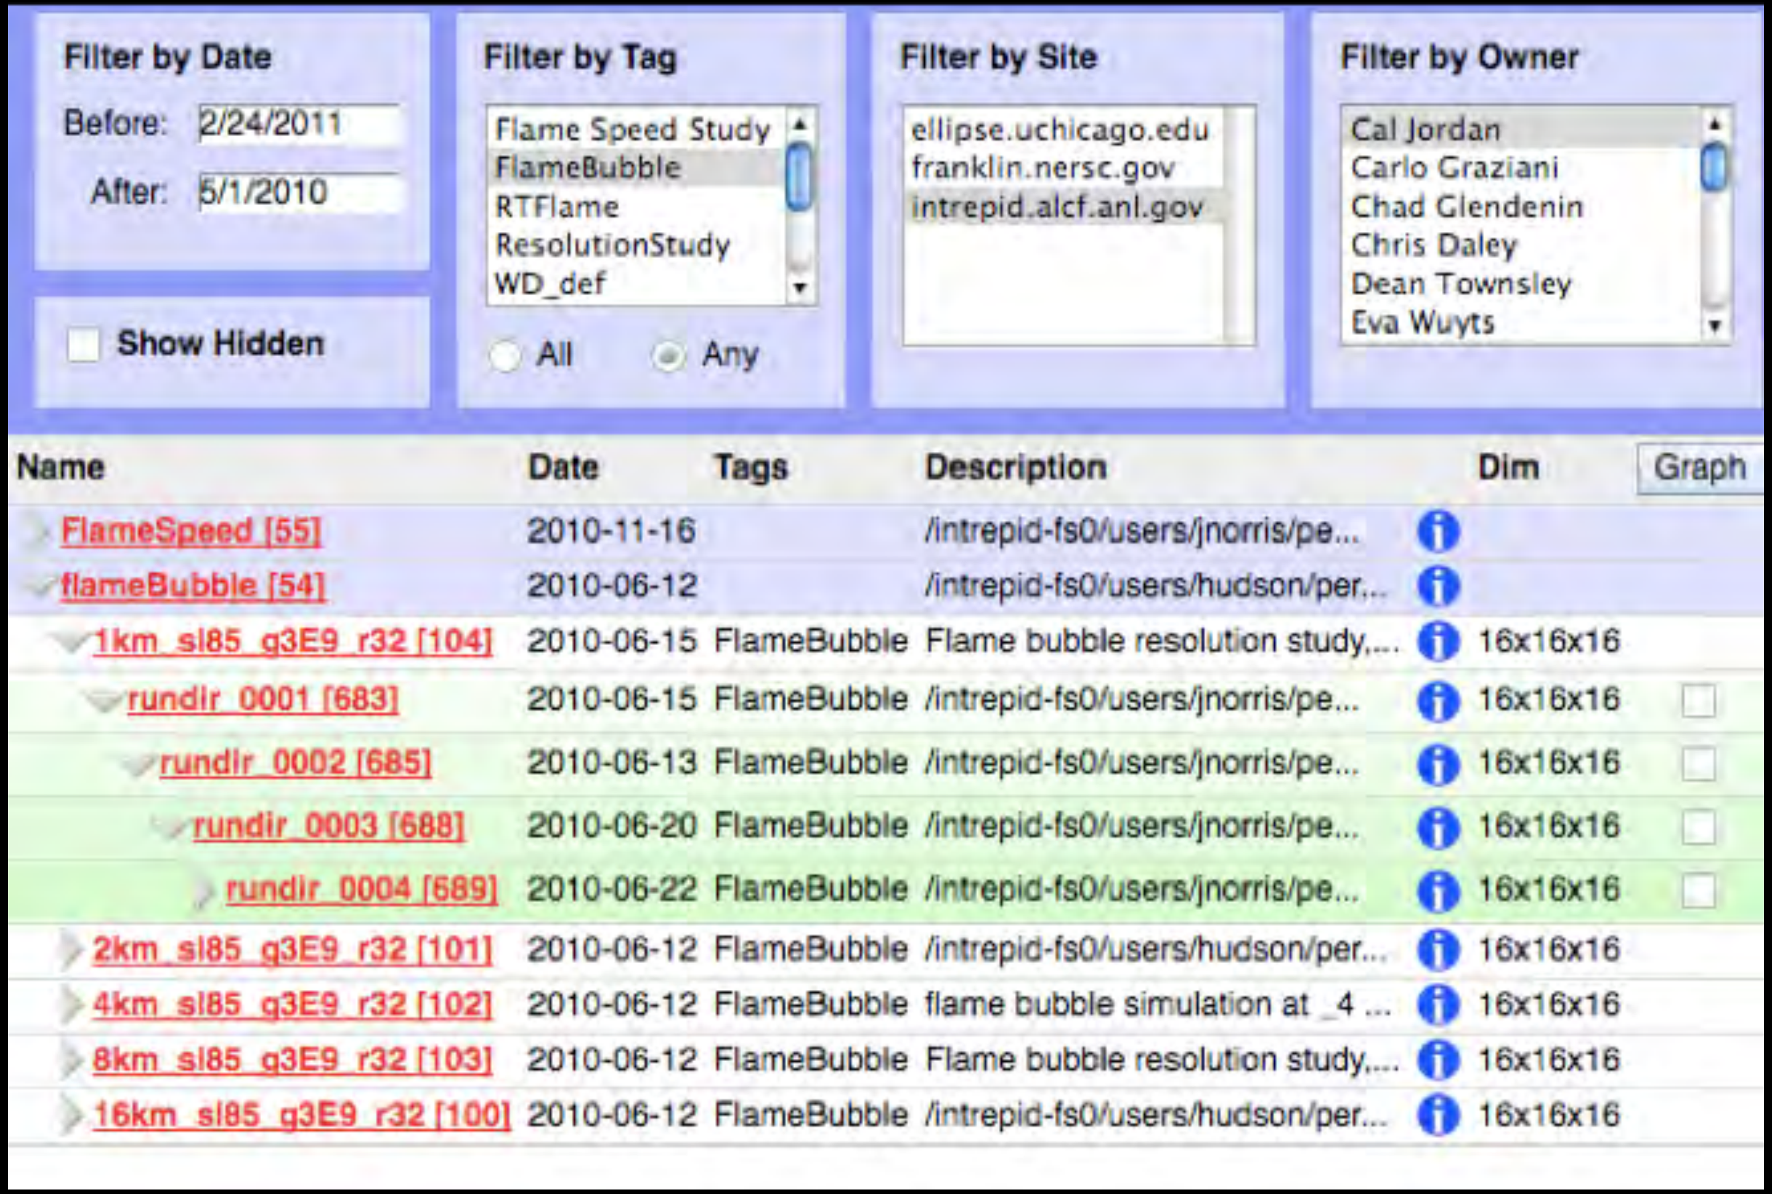
Filter by Date
Before: 2/24/2011
After: 5/1/2010
Show Hidden
Filter by Tag
Flame Speed Study
FlameBubble
RTFlame
ResolutionStudy
WD_def
All
Any
Filter by Site
ellipse.uchicago.edu
franklin.nersc.gov
intrepid.alcf.anl.gov
Filter by Owner
Cal Jordan
Carlo Graziani
Chad Glendenin
Chris Daley
Dean Townsley
Eva Wuyts
Name
Date
Tags
Description
Dim
Graph
FlameSpeed [55] 2010-11-16 /intrepid-fs0/users/jnorris/pe...
flameBubble [54] 2010-06-12 /intrepid-fs0/users/hudson/per...
1km_sl85_g3E9_r32 [104] 2010-06-15 FlameBubble Flame bubble resolution study,... 16x16x16
rundir_0001 [683] 2010-06-15 FlameBubble /intrepid-fs0/users/jnorris/pe... 16x16x16
rundir_0002 [685] 2010-06-13 FlameBubble /intrepid-fs0/users/jnorris/pe... 16x16x16
rundir_0003 [688] 2010-06-20 FlameBubble /intrepid-fs0/users/jnorris/pe... 16x16x16
rundir_0004 [689] 2010-06-22 FlameBubble /intrepid-fs0/users/jnorris/pe... 16x16x16
2km_sl85_g3E9_r32 [101] 2010-06-12 FlameBubble /intrepid-fs0/users/hudson/per... 16x16x16
4km_sl85_g3E9_r32 [102] 2010-06-12 FlameBubble flame bubble simulation at _4 ... 16x16x16
8km_sl85_g3E9_r32 [103] 2010-06-12 FlameBubble Flame bubble resolution study,... 16x16x16
16km_sl85_g3E9_r32 [100] 2010-06-12 FlameBubble /intrepid-fs0/users/hudson/per... 16x16x16

Graph View

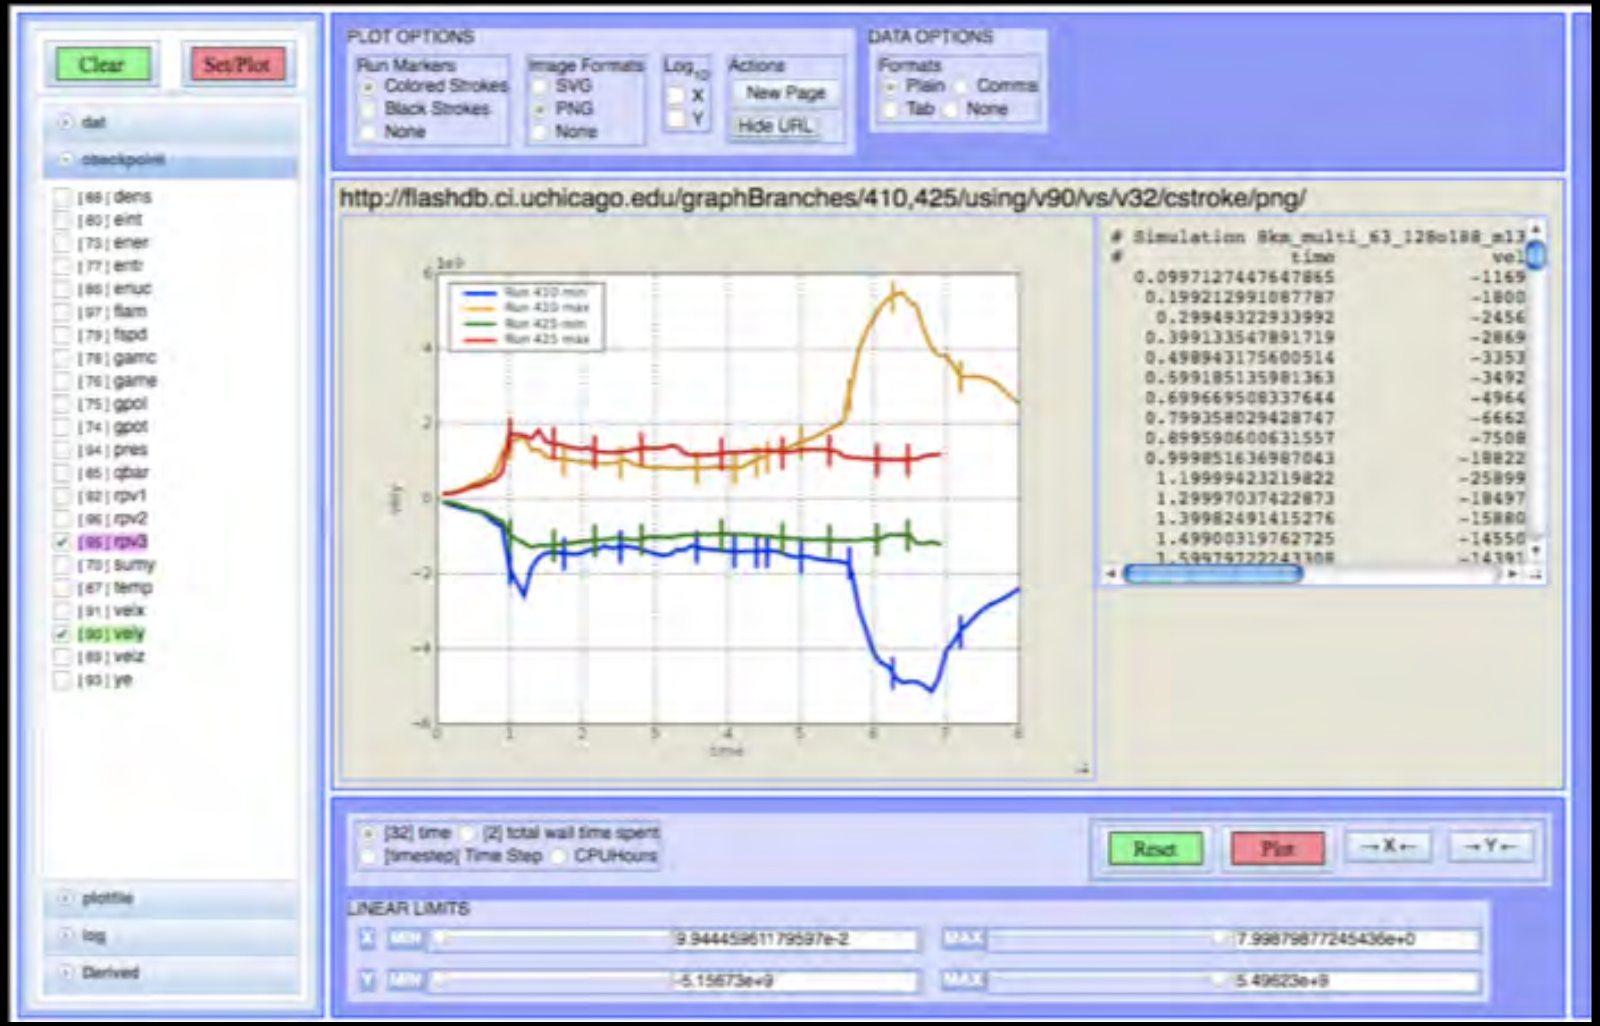

http://flashdb.ci.uchicago.edu/graphBranches/410,425/using/v90/vs/v32/cstroke/png

Monitor View

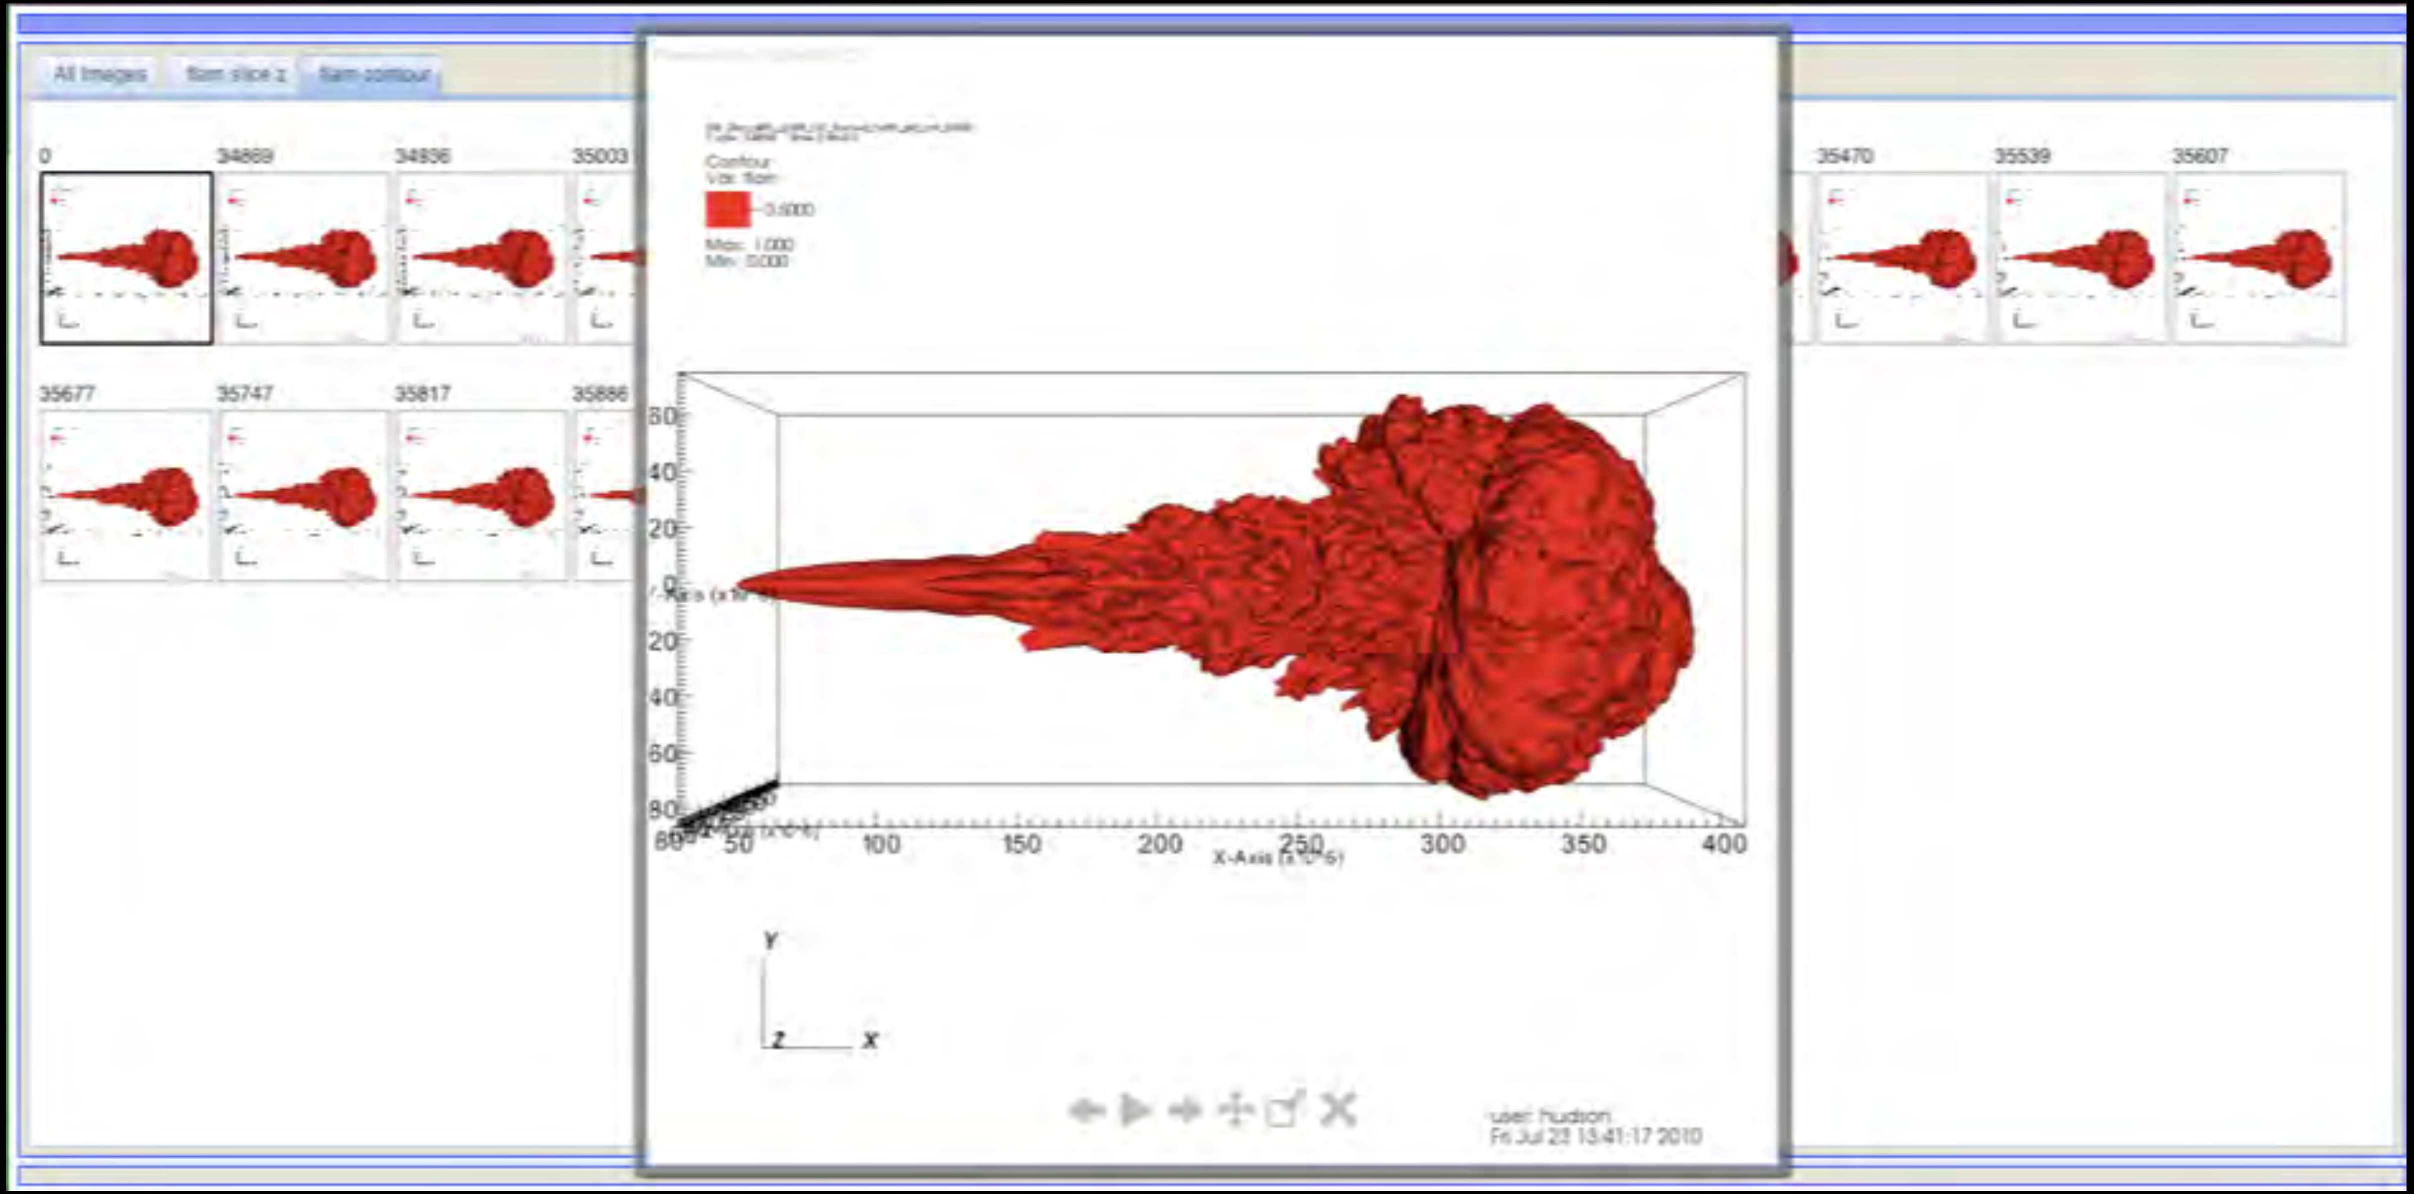

Summary View

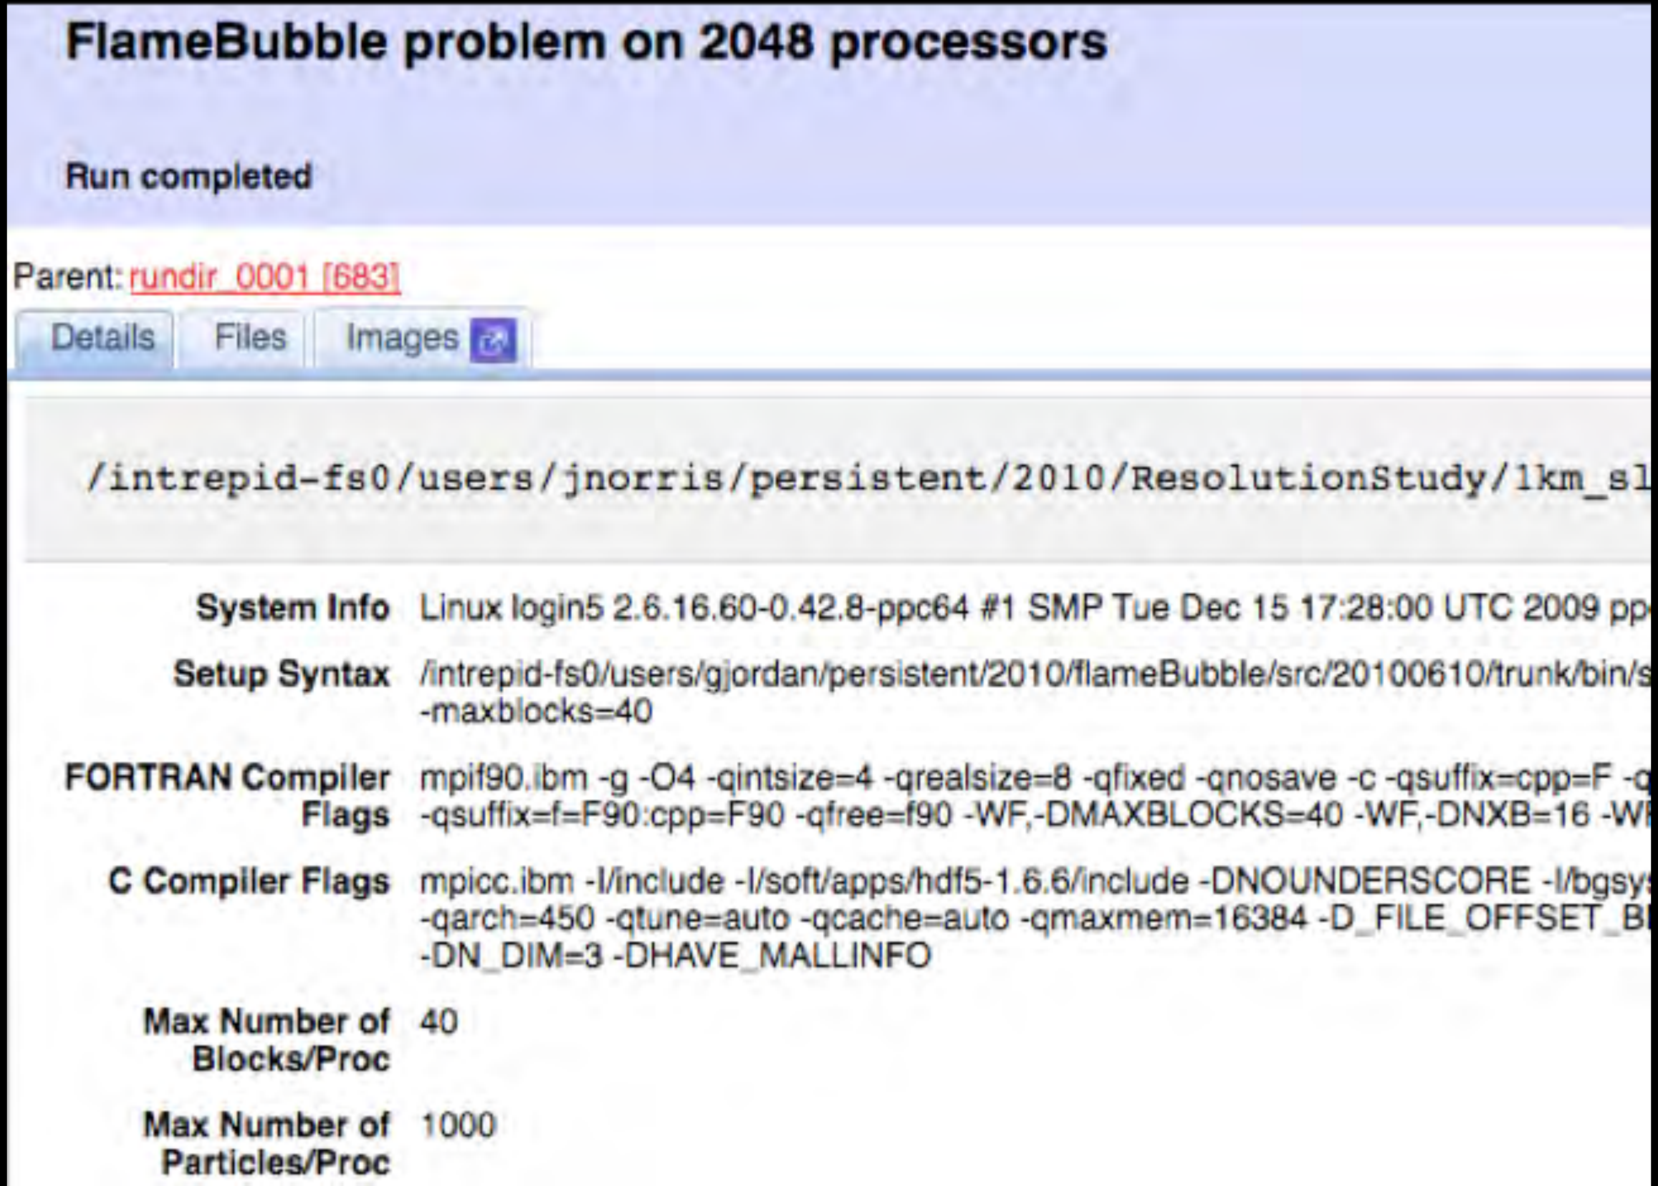

Smaash Outcomes (Simulation State)

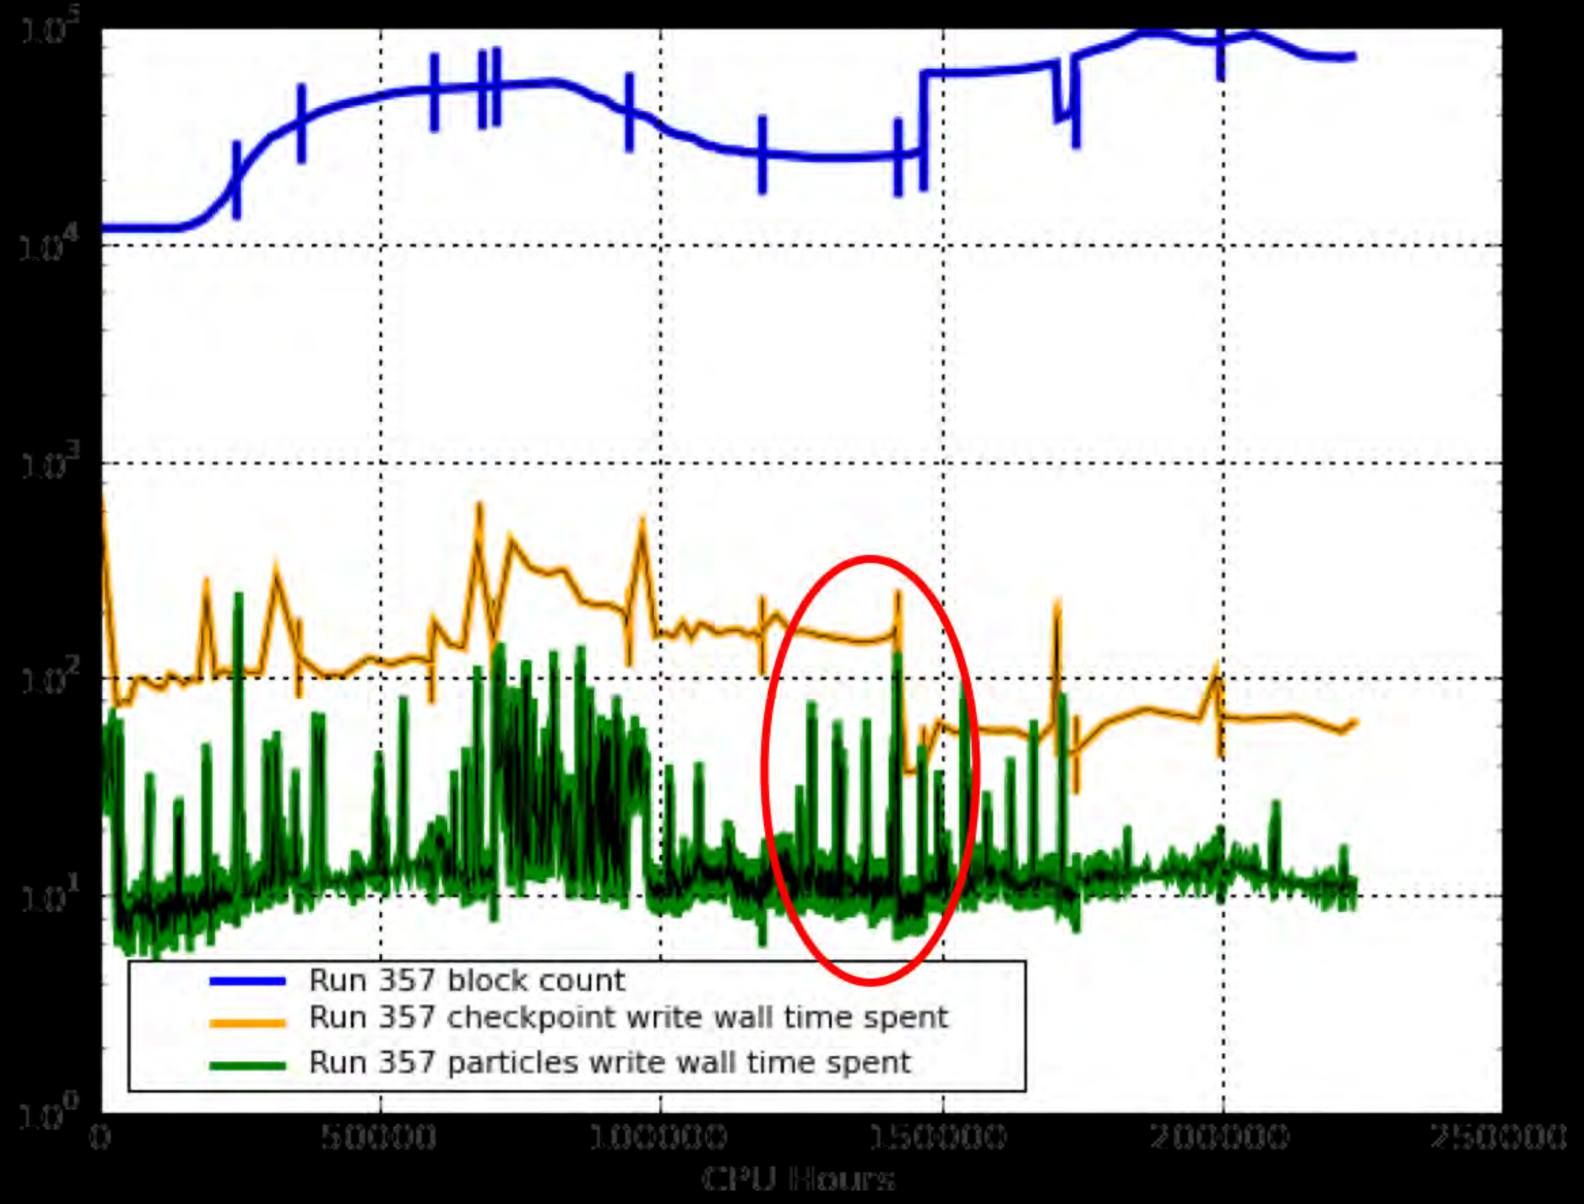

Smaash Outcomes (Analysis)

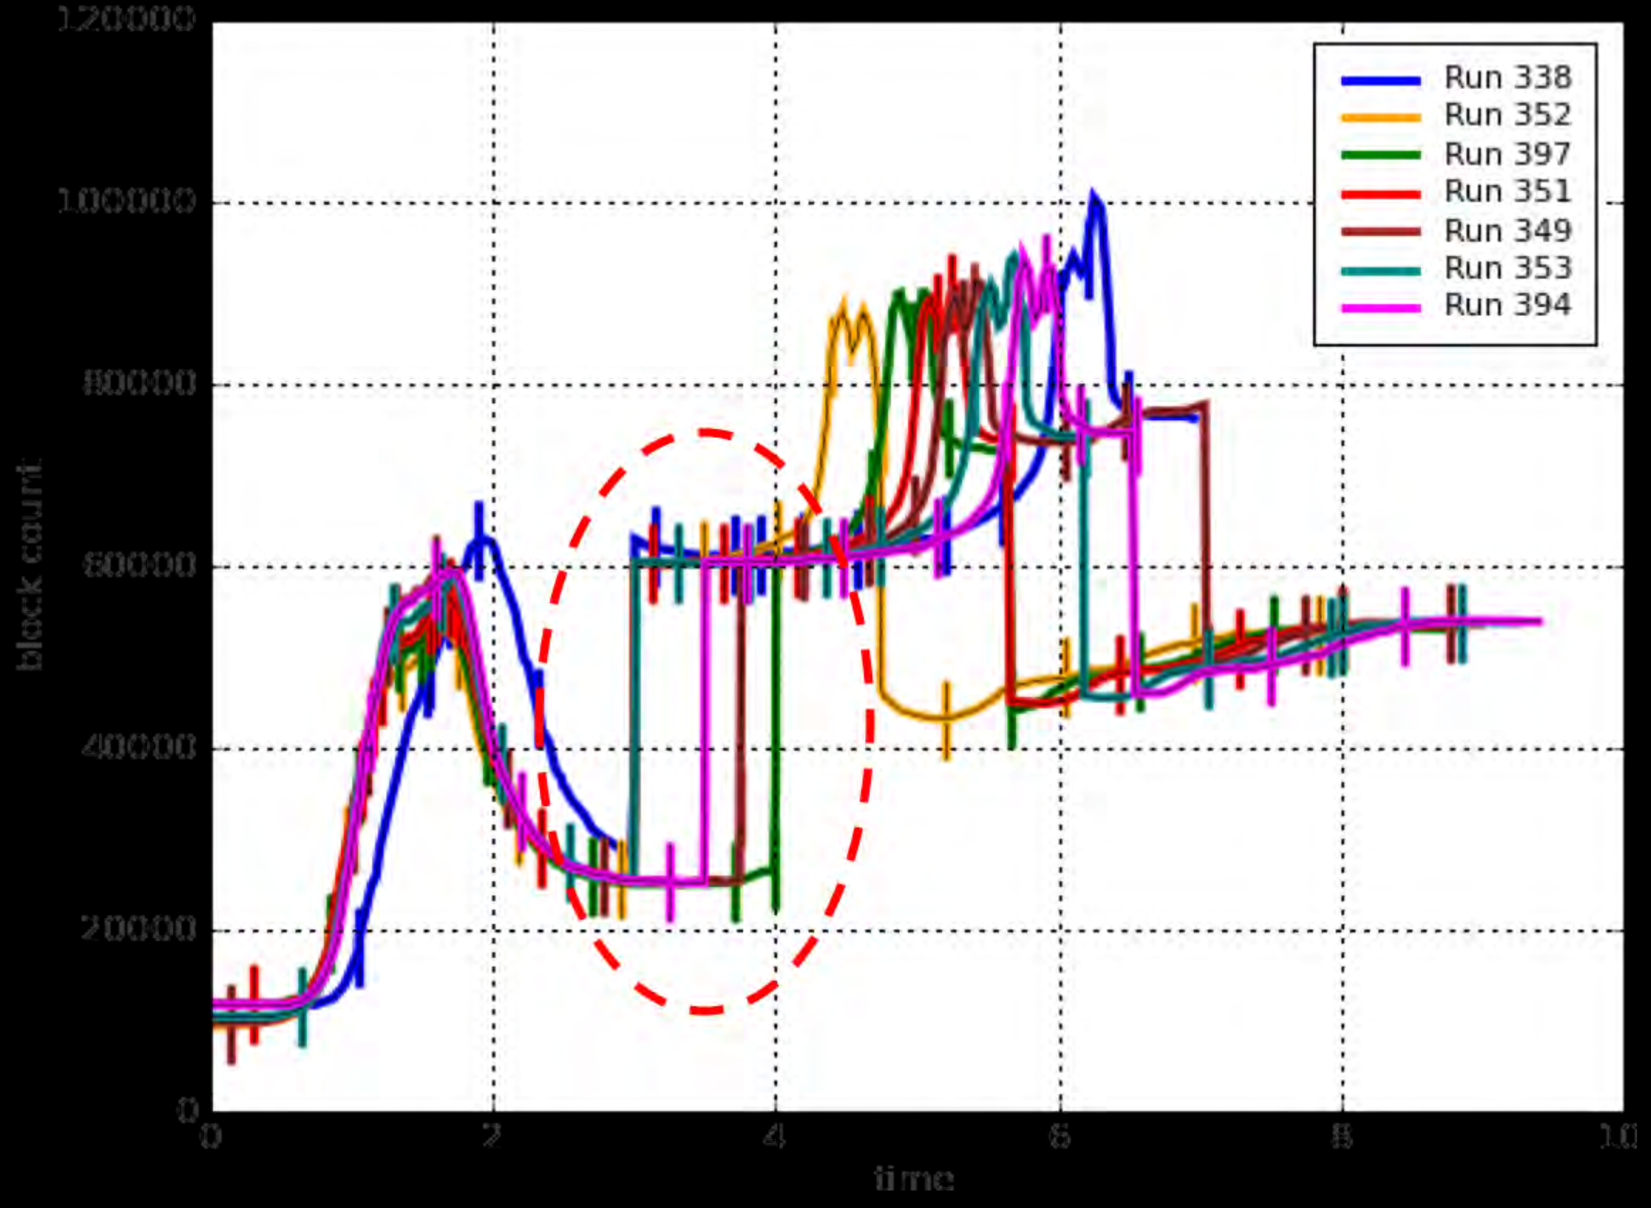

Smaash Outcome (Notebook)

ulation Tree | Graph | Flash Simulation Tree | Graph | MultipleBubbles < FlashDB < Cl...

Pending or Stopped Runs

Last Run	Simulation Name	Status	Time	Breakout	Detonation	Completed	Runner
-	8km 79 138o148 m1.365 series 2	pending	0.0 s	n	n	n	Cal
274	8km 188_184o220 m1.365 series 1	stopped, low Enuc	1.65 s	y	n	y	Lynn
133	8km 79 138o148 m1.365 series 1	stopped, low Enuc	2.15s	y	n	n	Klaus
123	8km 56 123o148 m1.365 series 1	stopped, low Enuc	1.80 s	y	n	n	Klaus
125	8km 63 128o148 m1.385 series 1	Stopped, Lynn rerun, **different mass**	5.40 s	y	y	n	Cal

More Diagnostic Graphs

These mostly illustrate computational aspects of the simulations:

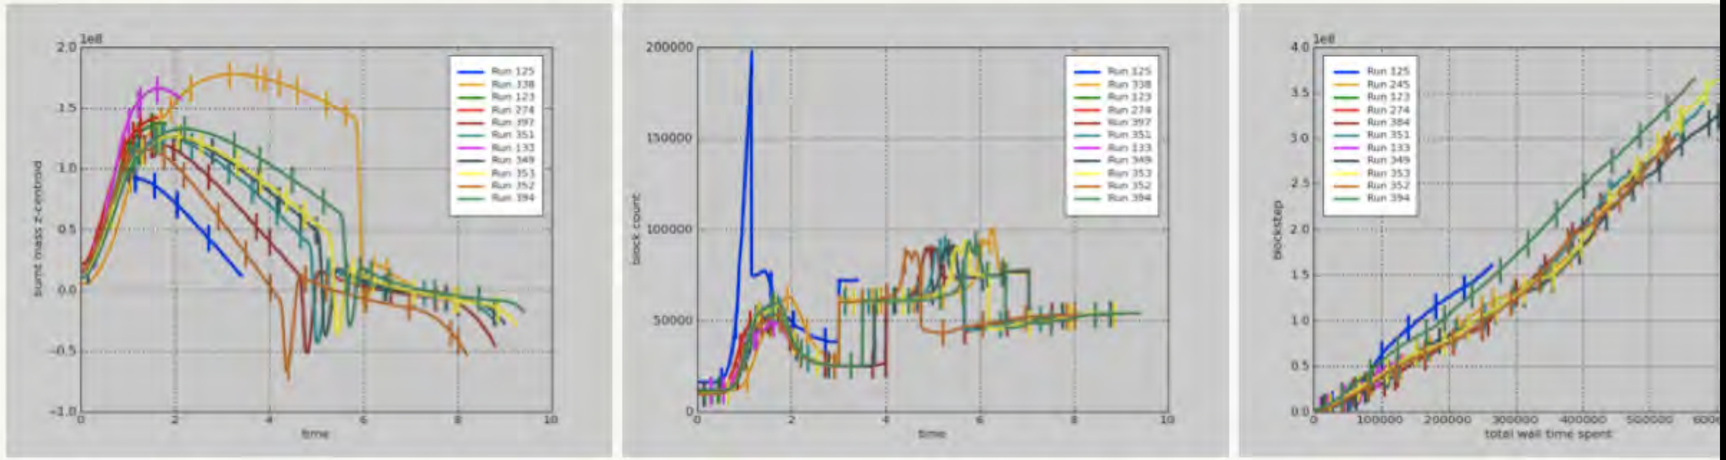

Code/Machine/Bug Issues

Smaash Today

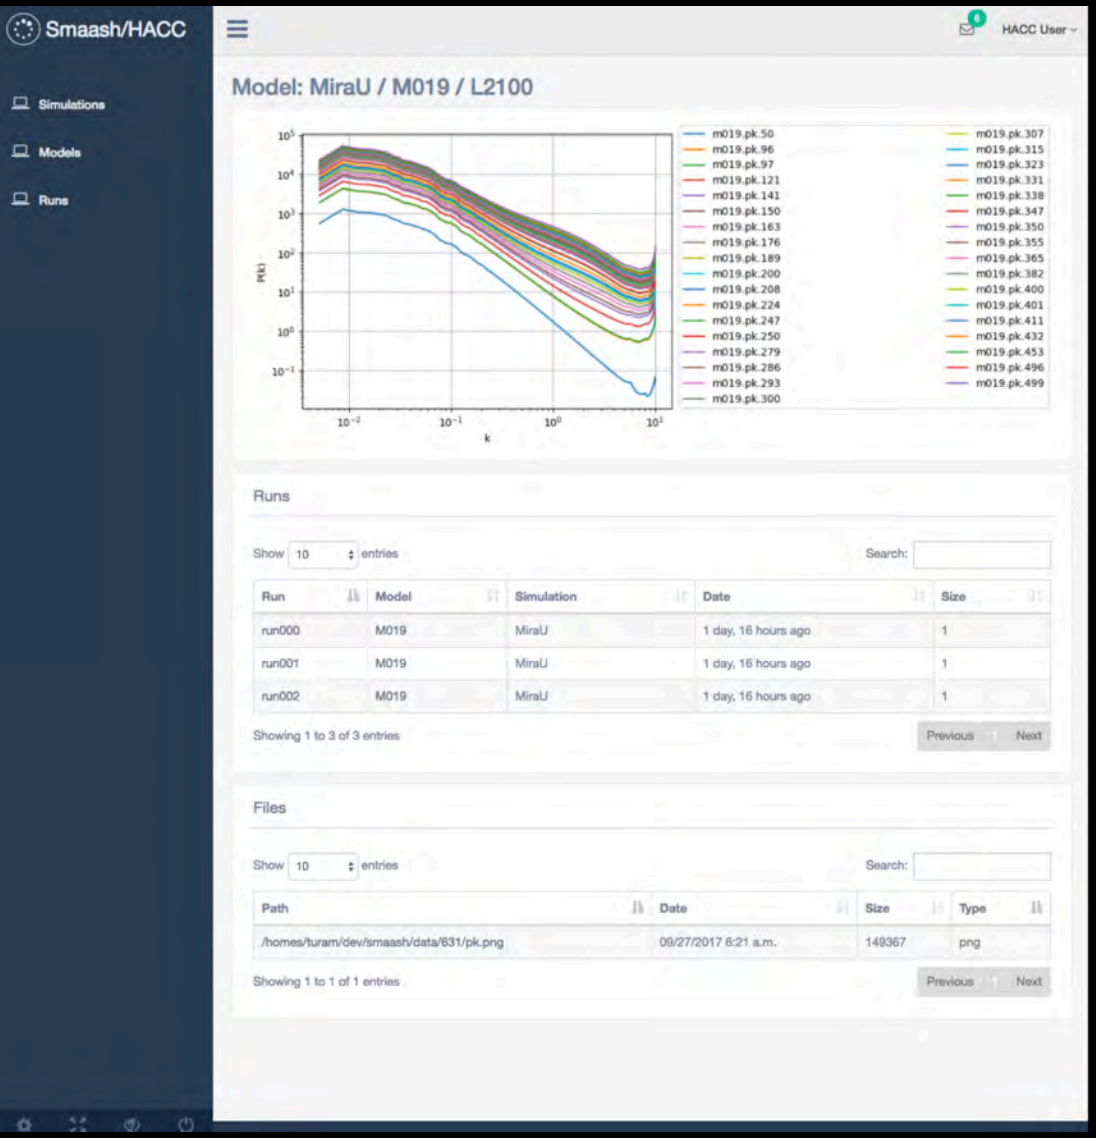

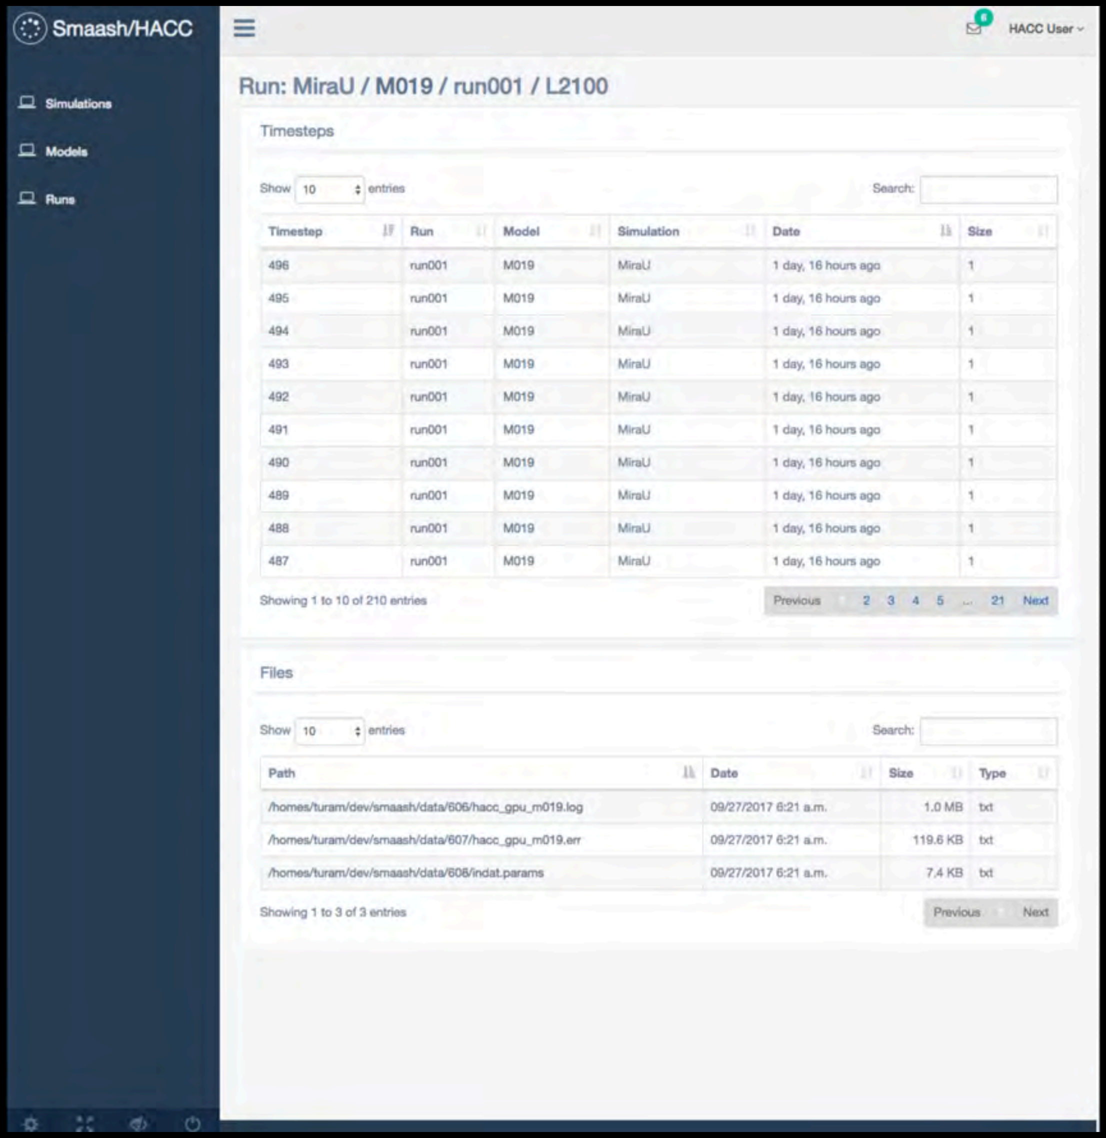

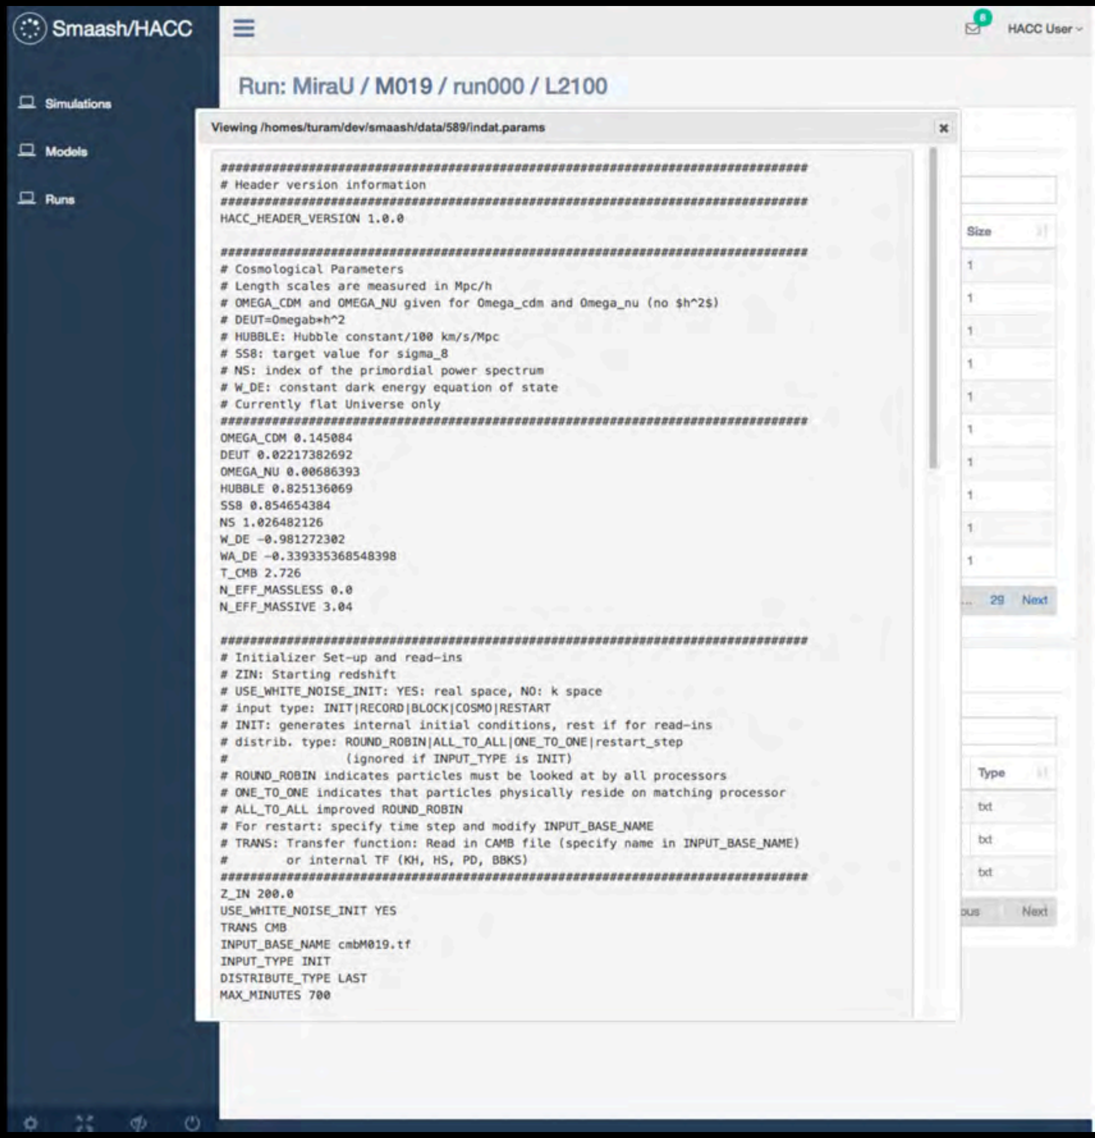

New Efforts in Science

Yesterday

APS

Beamline PC

DTN

Turnaround time depends on availability of APS Compute

APS Compute

4000 cores

Today

APS

ALCF Theta (11.7 PetaFLOPs)

Beamline PC

DTN

Turnaround time depends on availability of ALCF Compute

Balsam

Balsam establishes a deep job reservoir, and flows jobs to ALCF machines based on availability

- Individual APS application runs to individual ALCF jobs
- Many small APS jobs flow into a single large/long ALCF job
- Balsam provisions new jobs as needed to satisfy workload
- Balsam leverages backfill where possible; otherwise, optimizes jobs for scheduler

- Job monitoring
- Results view
- Workflow extension

New Efforts in Science[u]

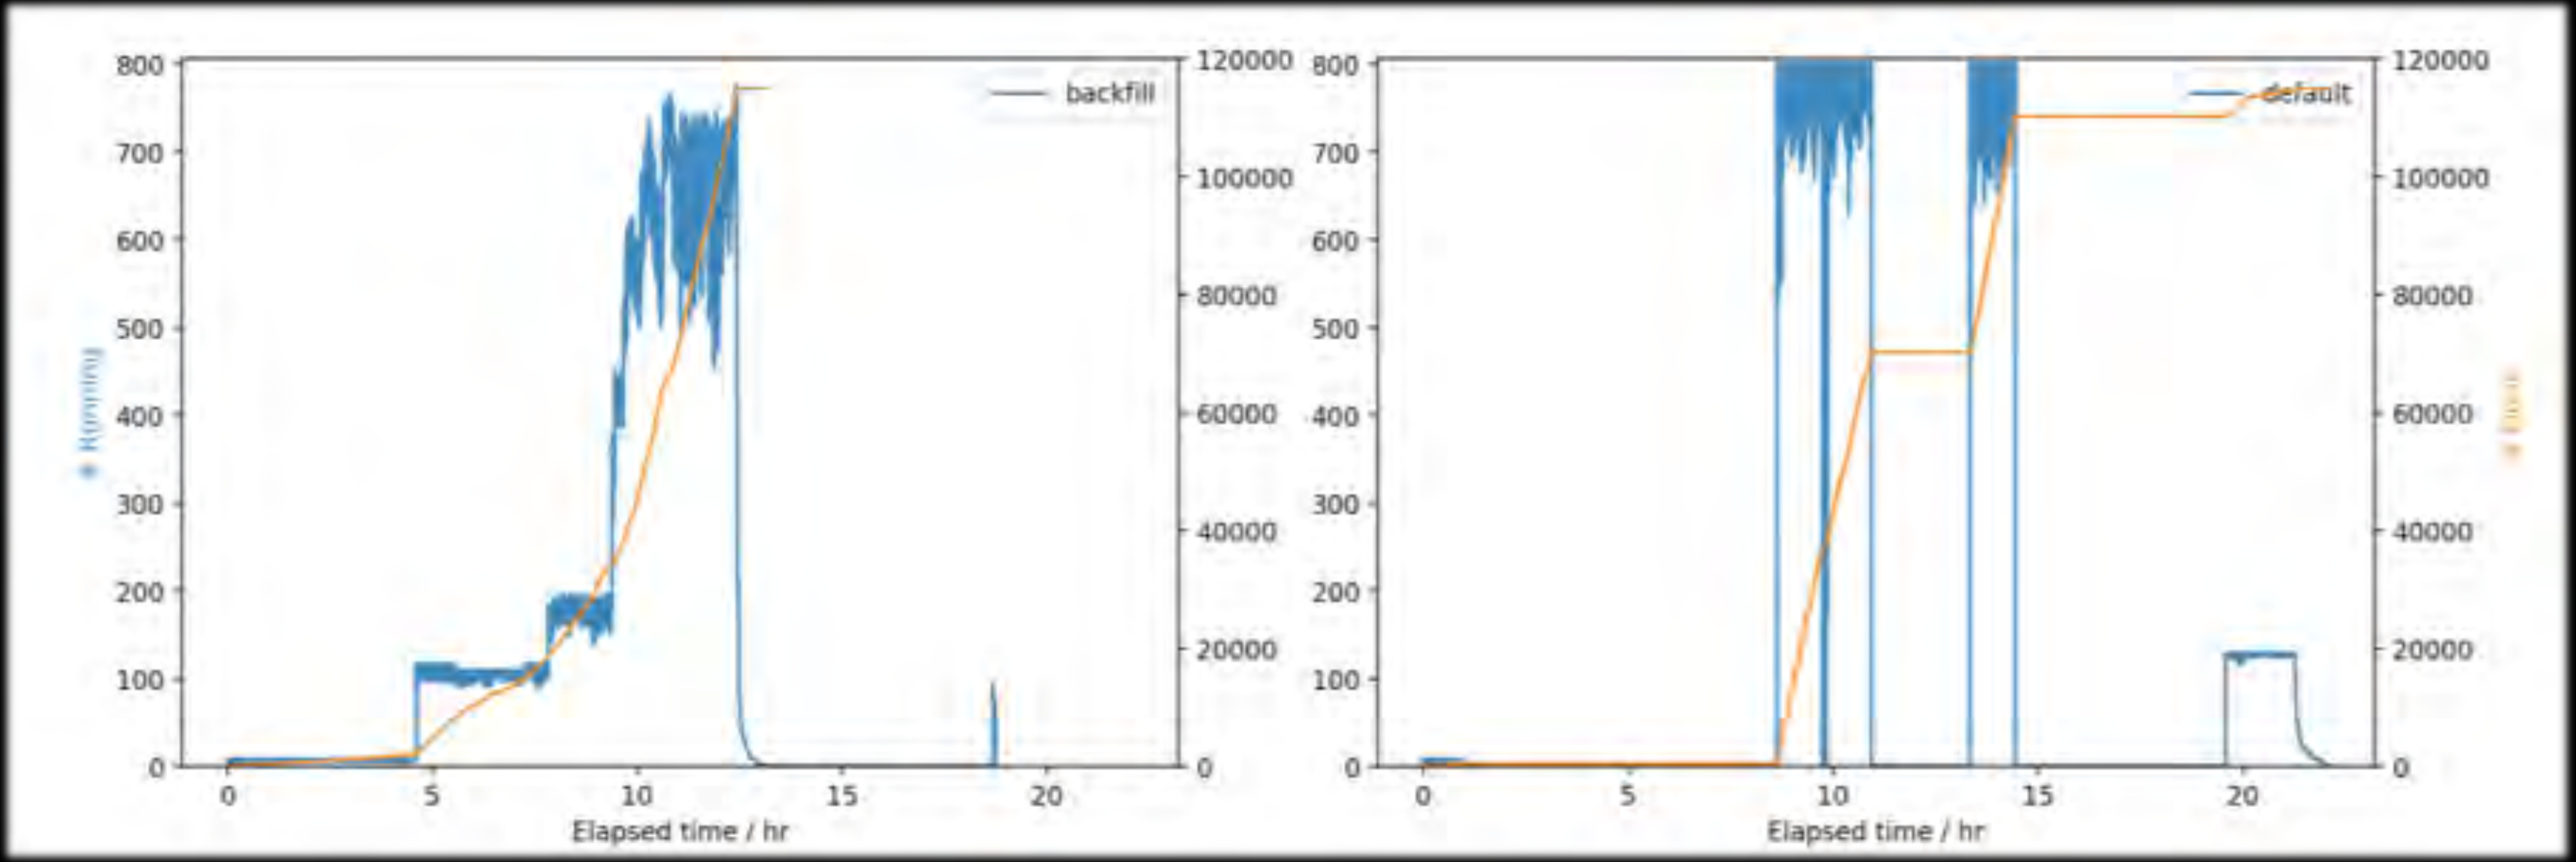

[u] M. A. Salim, T. D. Uram, J. T. Childers, P. Balaprakash, V. Vishwanath and M. E. Papka, *Balsam: Automated Scheduling and Execution of Dynamic, Data-Intensive HPC Workflows*, to appear **1st Annual Workshop on Large-scale Experiment-in-the-Loop Computing**, Denver, CO, November 2019.

Last Topic

Information Visualization

- Connection to **X** science

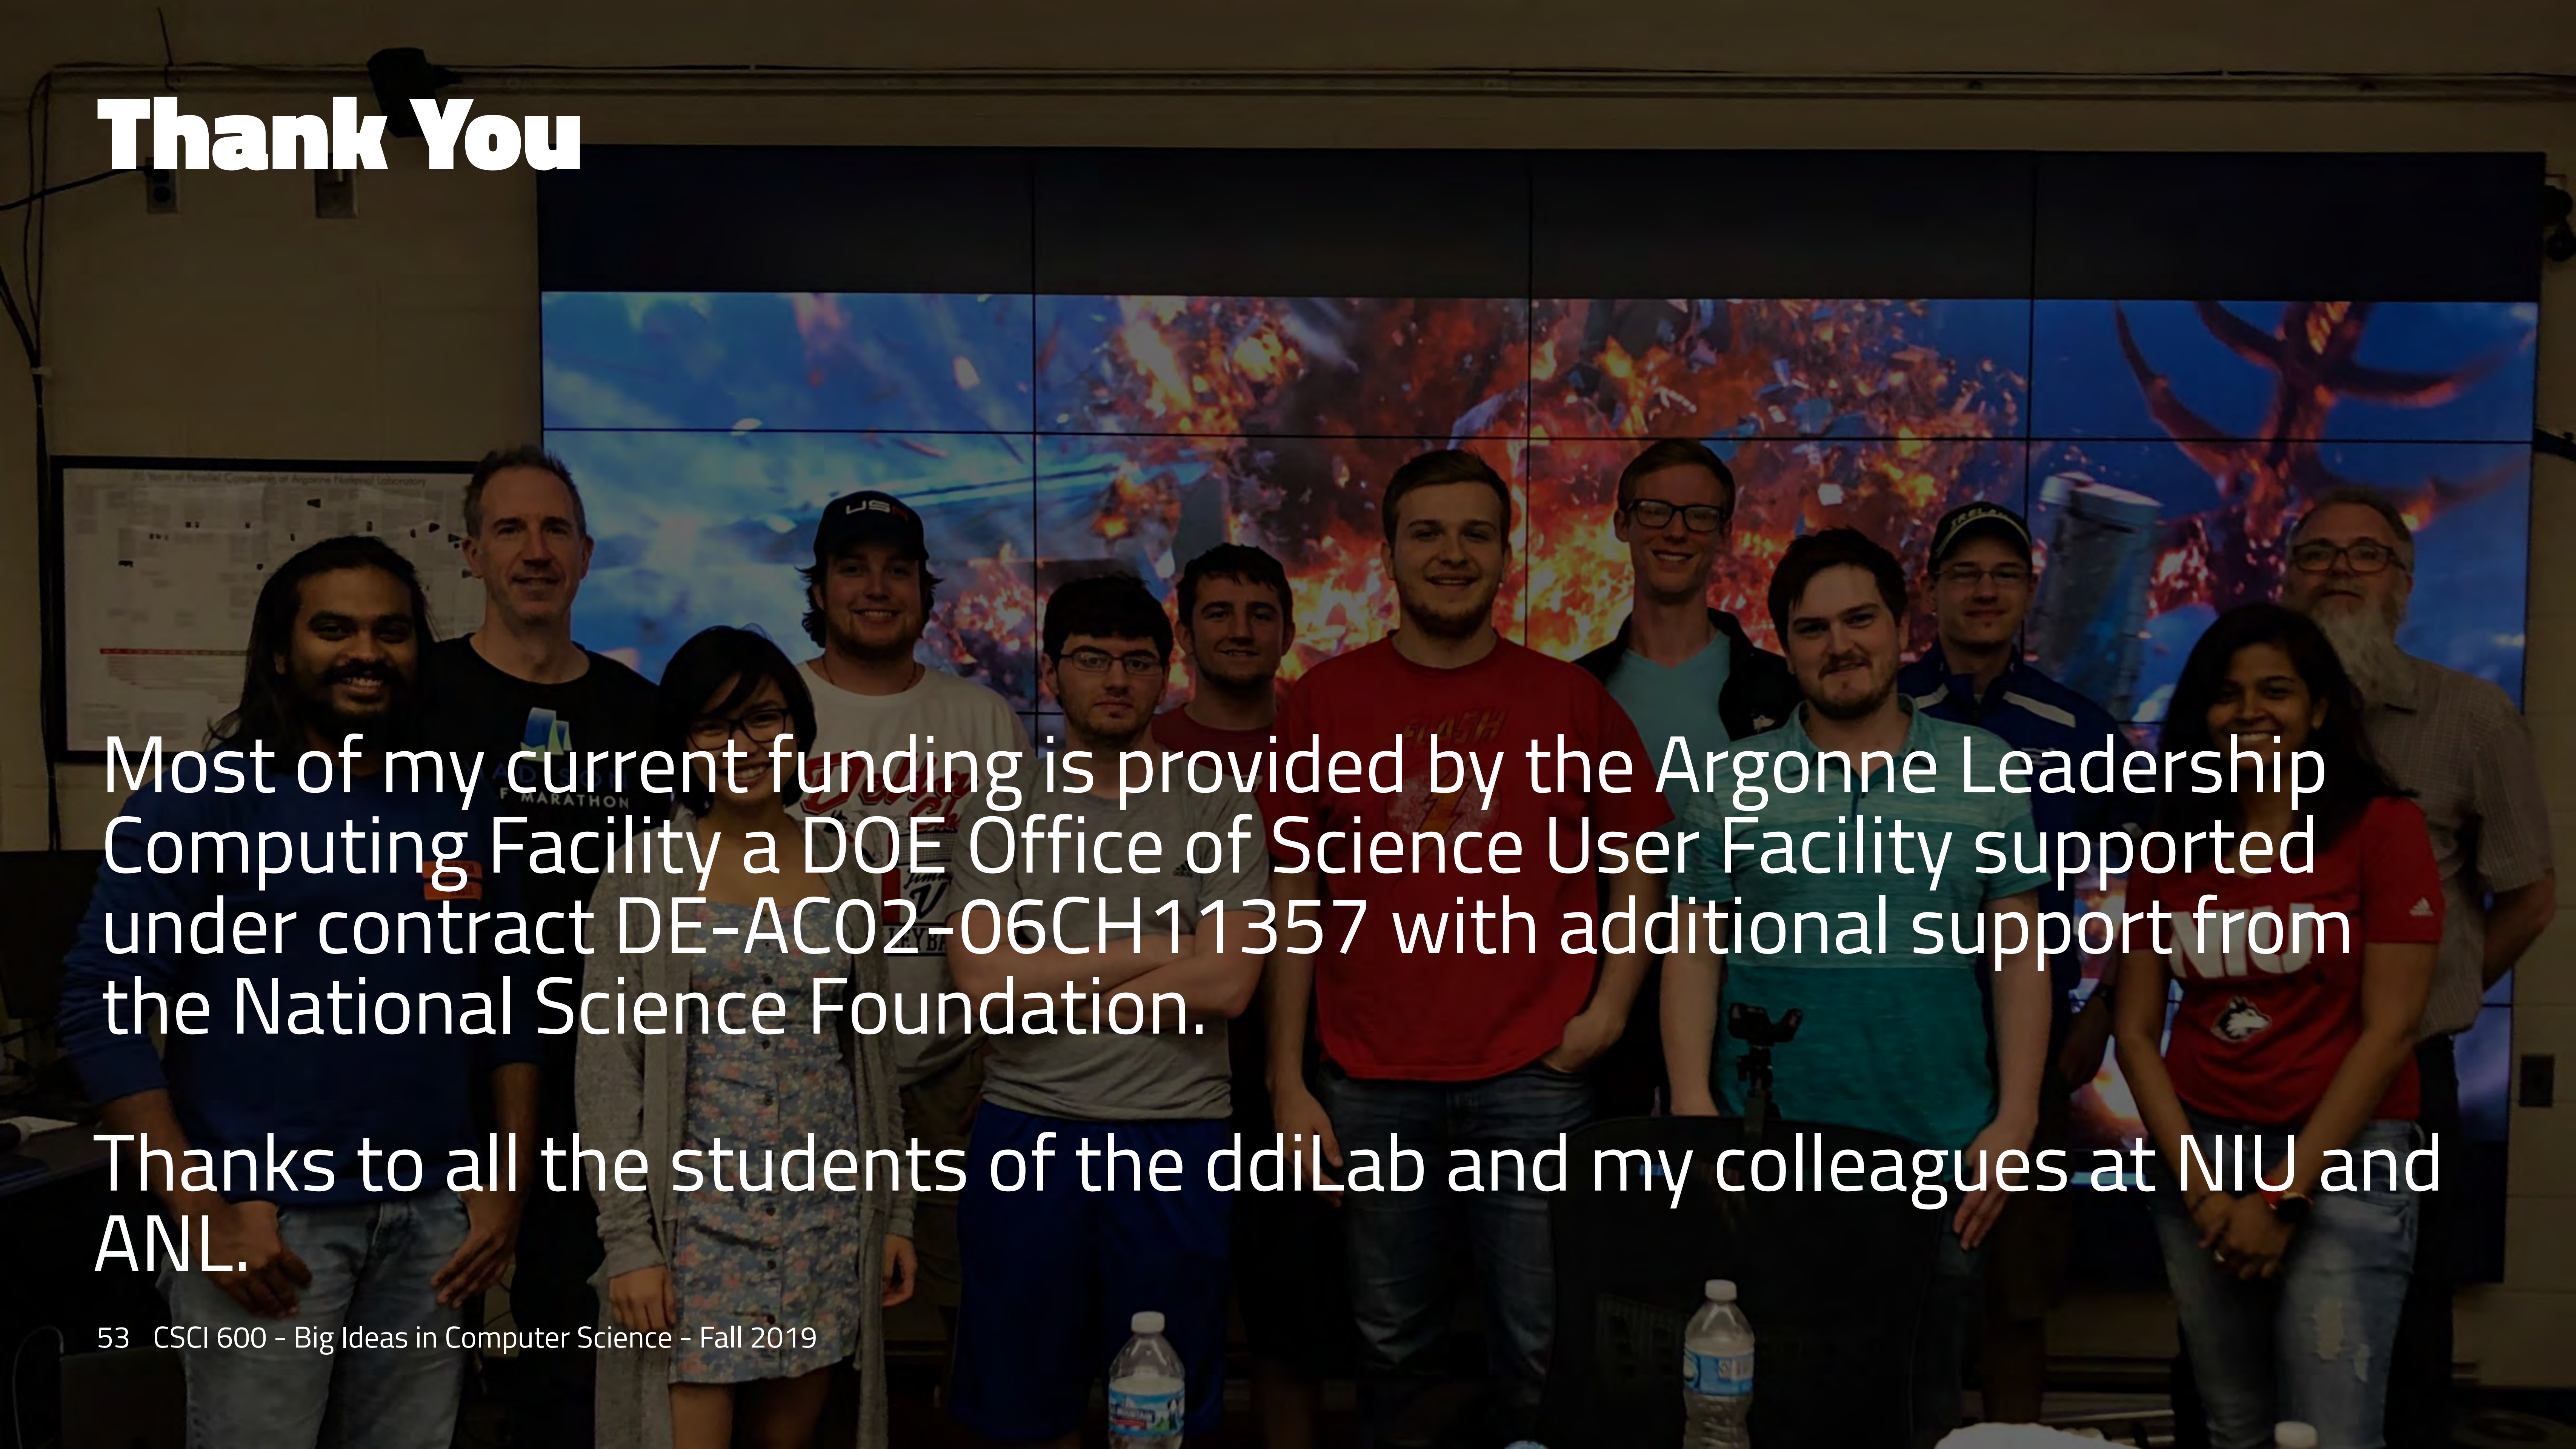

Thank You

Most of my current funding is provided by the Argonne Leadership Computing Facility a DOE Office of Science User Facility supported under contract DE-AC02-06CH11357 with additional support from the National Science Foundation.

Thanks to all the students of the ddiLab and my colleagues at NIU and ANL.

If I have seen further it is by standing on the shoulders of giants.

— Sir Isaac Newton

Extra Slides

HPC Landscape (Yesterday)

Simulation Applications
64bit floating point
memory bandwith
random access to memory
sparse matrices
distributed memory jobs
synchronous input/output multinode
scalability limited communication
low latency high bandwidth
large coherency domains (sometimes)
output typically greater than input
output rarely read
output is data

HPC Landscape (Today)

Simulation Applications	Big Data Applications	Deep Learning Applications
64bit floating point	64bit and integer important	lower precision <= 32bit
memory bandwith	data analysis pipelines	inferencing can be 8bit (TPU)
random access to memory	databases including NoSQL	scaled integer possible
sparse matrices	MapReduce/SPARK	training dominates development
distributed memory jobs	millions of jobs	inference dominates pro
synchronous input/output multinode	input/output bandwidth limited	reuse of training data
scalability limited communication	data management limited	data pipelines needed
low latency high bandwidth	many task parallelism	dense float point typical SGEMM small DFT, CNN
large coherency domains (sometimes)	large-data in and large-data out	ensembles and search
output typically greater than input	input and output both important	single models small
output rarely read	output is read and used	input more important than output
output is data	output is data	output is models