Big Ideas

WITH MICHAEL E. PAPKA

DATA, DEVICES AND INTERACTION LABORATORY

March 29, 2022

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)
- Continuously learning in other areas
  - University of Chicago - Business School (SLLP)
  - Harvard University - Business School (GMP)
  - Stanford University - Hasso Plattner Institute of Design (Design Thinking)

BIT ABOUT ME (CAREER)

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

BIT ABOUT ME (RESEARCH)

- Advanced Display Environments and Collaboration Technology
- High Performance Computing (Systems and Environments)
- Internet of Things (Computing Continuum) [Edge Computing]
- Information and Scientific Visualization
- Data Analysis
- Augmented/Virtual Reality

BIT ABOUT THE ddiLAB

- Joint with **School of Art and Design** and co-directed with Professor Joseph Insley (Time Arts)
- Focused on **visualization** and **data analysis** coupled to **high-performance computing** in the support of **science**, with side efforts involving the *Internet of Things* (edge computing) and interdisciplinary activities connected to computing
- Students
- - 2 PhD (Information Visualization, machine learning/edge computing)
- - 1 MS (HPC log analysis)
- - 8 Undergraduates (IoT, VR, and HPC)

BIT ABOUT ME (RESEARCH)

- Advanced Display Environments and Collaboration Technology
- **High Performance Computing (Systems and Environments)**
- **Internet of Things (Computing Continuum) [Edge Computing]**
- Information and Scientific Visualization
- Data Analysis
- **Augmented/Virtual Reality**

SUPERCOMPUTERS

Cray X-MP/4 supercomputer, 1985
(www.computerhistory.org/revolution/supercomputers/10/25/23)

SUPERCOMPUTERS

Argonne's Aurora Supercomputer, 2022/23

800megaFLOPS to ~2exaFLOPS

800,000,000 to ~2,000,000,000,000,000,000 FLOPS

2,500,000,000x

SO WHAT DOES ALL THAT COMPUTE POWER ENABLE?

- Large-Scale Computing on the Connectomes of the Brain
- Global Radiation MHD Simulations of Massive Star Envelopes
- Molecular Design of Dye-Sensitized Solar Cells
- Real-time Computing in support of DIII-D National Fusion Facility

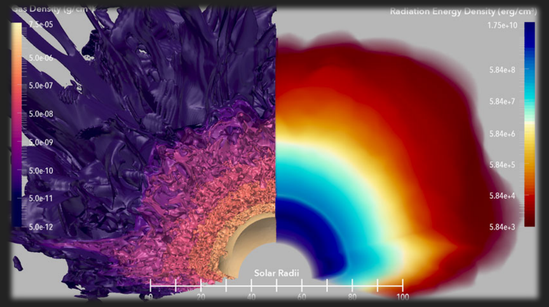

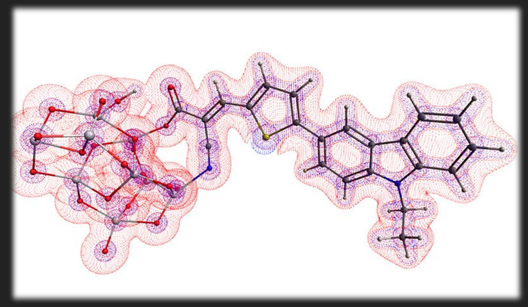

Doga Gursoy, Argonne National Laboratory | Lars Bildsten, University of California, Santa Barbara | Jaqueline M. Cole, University of Cambridge | David Schissel, General Atomics

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you translate theory into simulation?

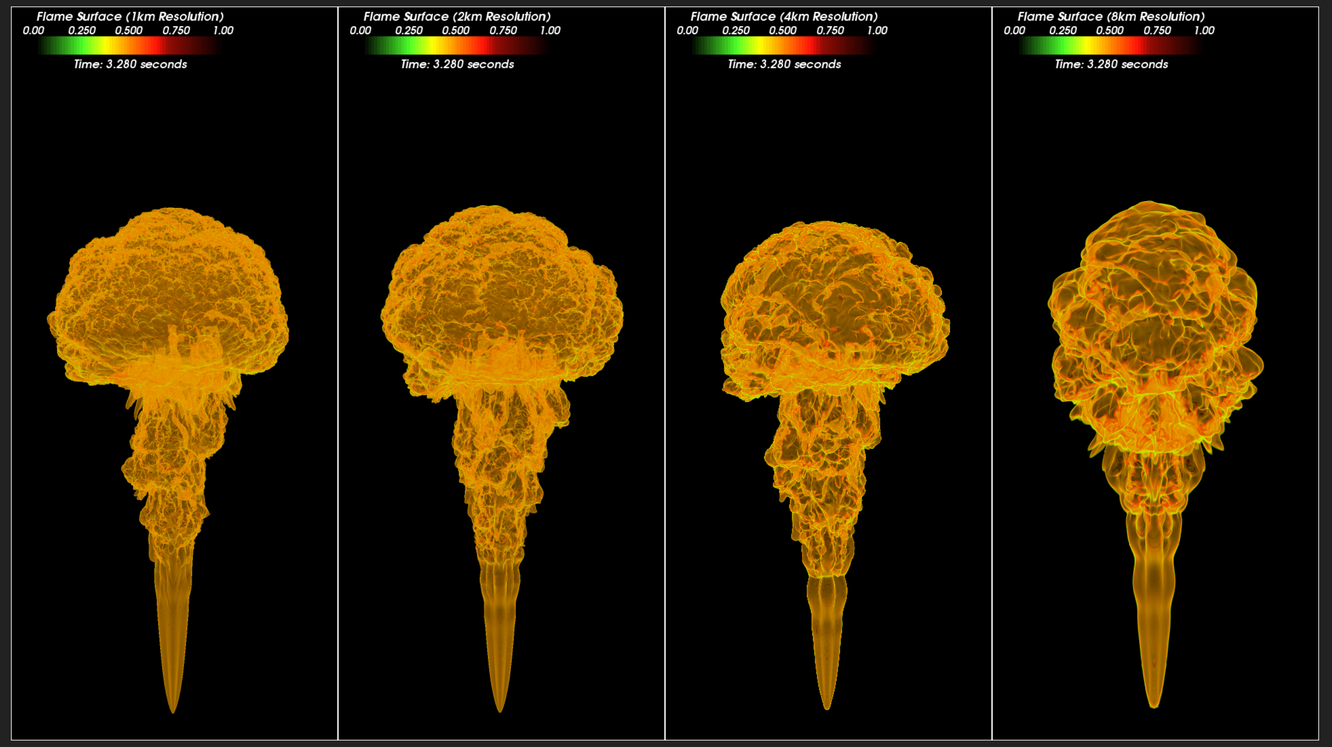

R. Fisher, L. Kadanoff, D. Lamb, A. Dubey, T. Plewa, A. Calder, F. Cattaneo, P. Constantin, I. Foster, M. E. Papka, S. I. Abarzhi, S. M. Asida, P. M. Rich, C. C. Glendenin, K. Antypas, D. J. Sheeler, L. B. Reid B. Gallagher, and S. G. Needham, *Terascale Turbulence Computation Using the FLASH3 Application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

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you more efficiently operate a facility?

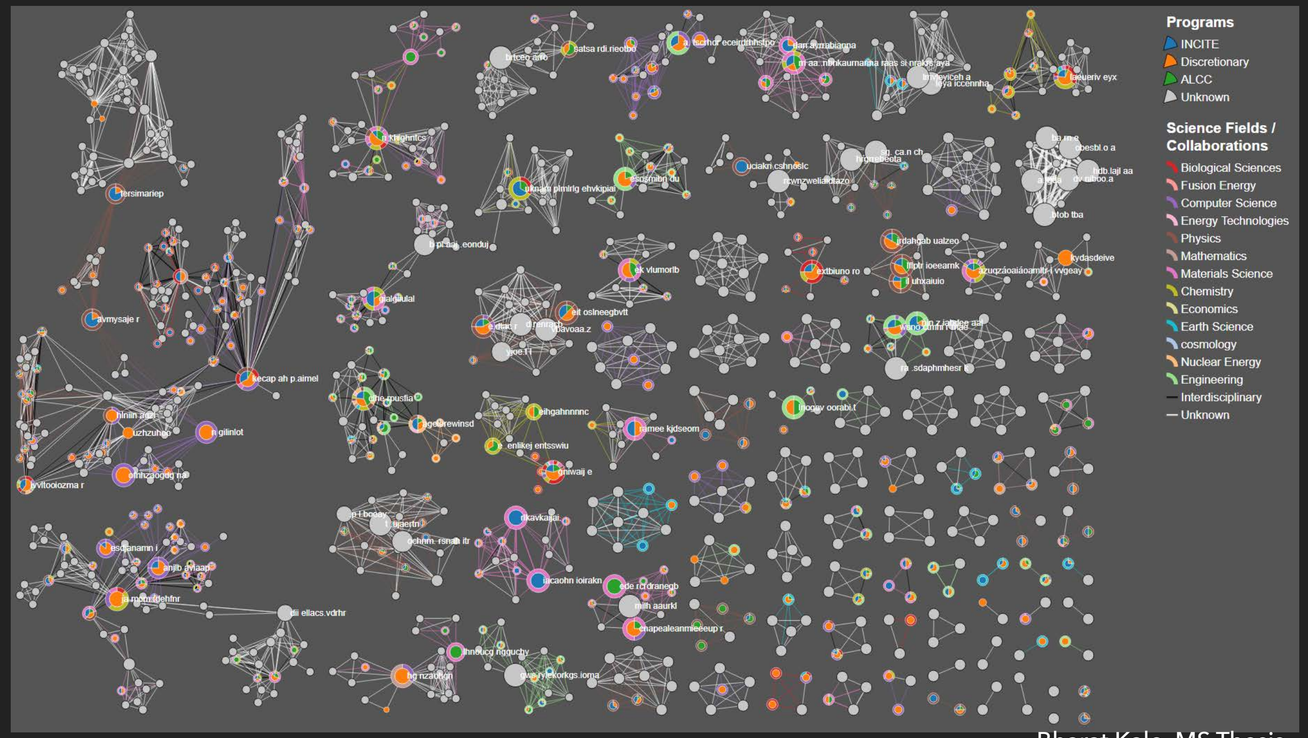

Bharat Kale, MS Thesis

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 RESEARCH OPPORTUNITIES

- How do you effectively schedule and operate a resource?

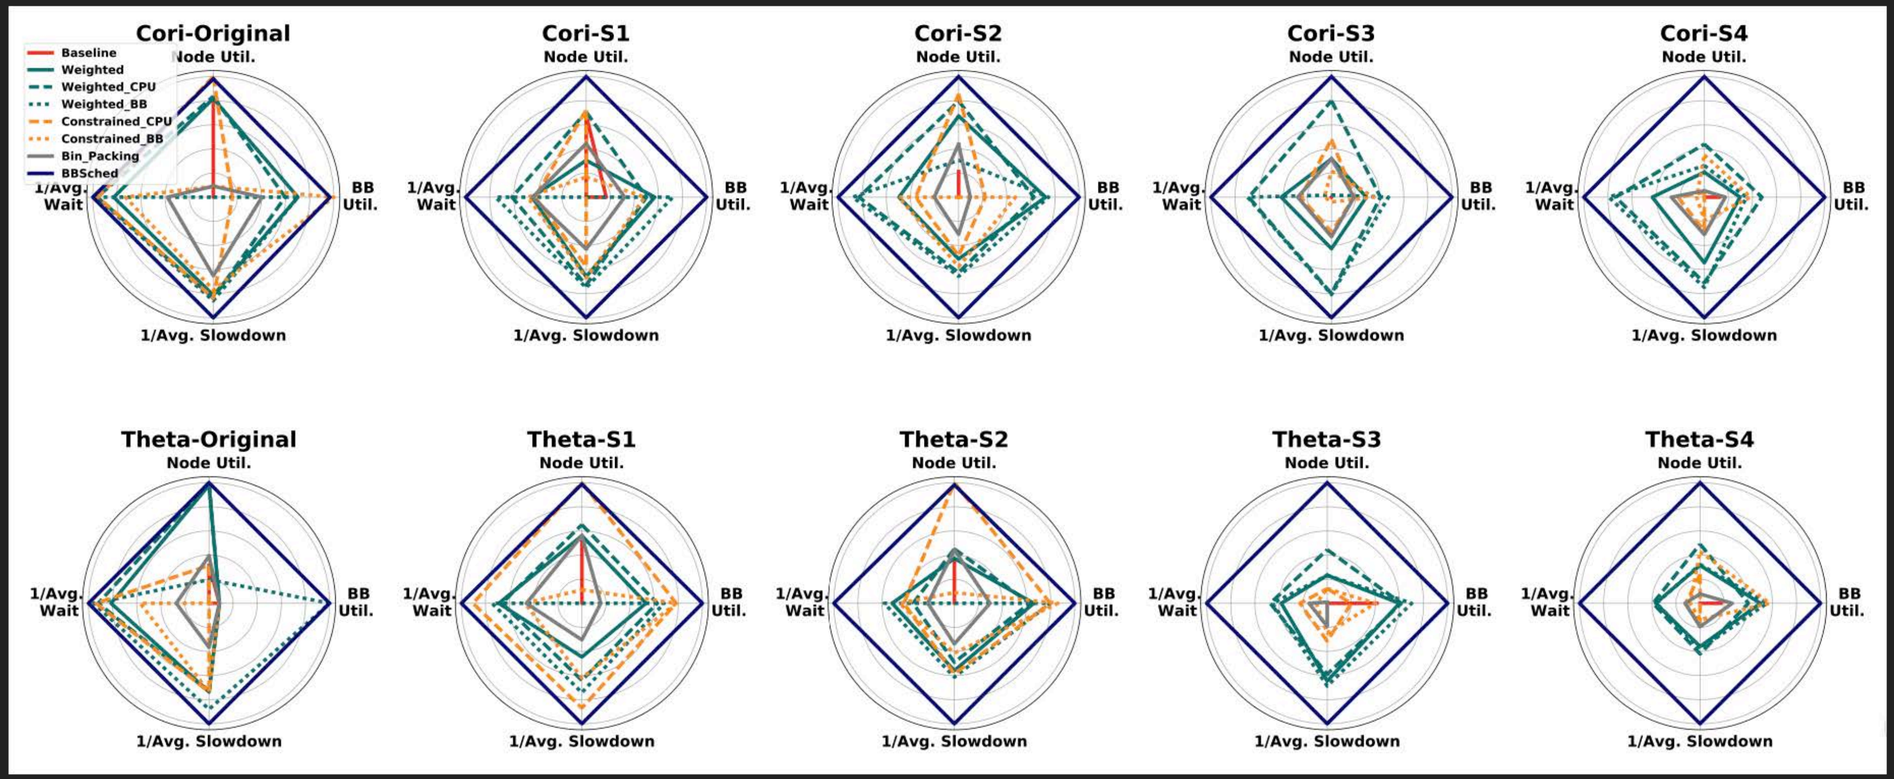

Y. Fan, Z. Lan, P. Rich, W. E. Allcock, M. E. Papka, B. Austin, D. Paul, *Scheduling Beyond CPUs for HPC*, **Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing**, pp. 97-108, June 2019.

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you evolve traditional HPC environment?

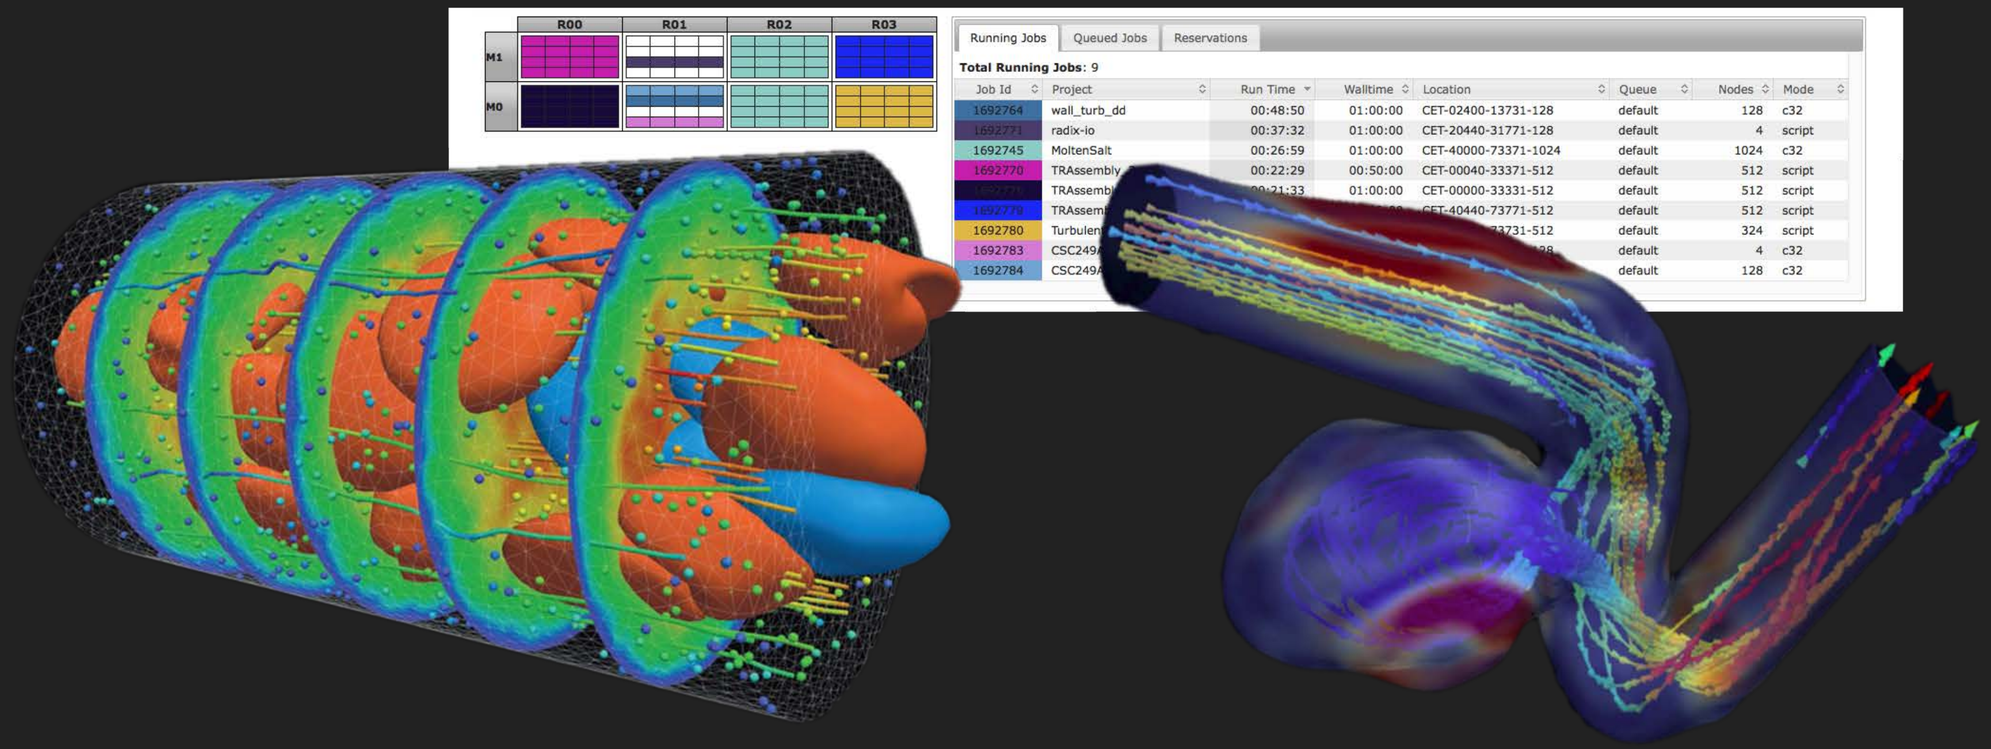

L. Grinberg, J. A. Insley, D. Fedosov, V. A. Morozov, M. E. Papka, G. E. Karniadakis, *Tightly Coupled Atomistic-Continuum Simulations of Brain Blood Flow on Petaflop Supercomputers*, **Computing in Science and Engineering**, 14(6):58-67, 2012.

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you evolve traditional HPC environment to address real-time needs?

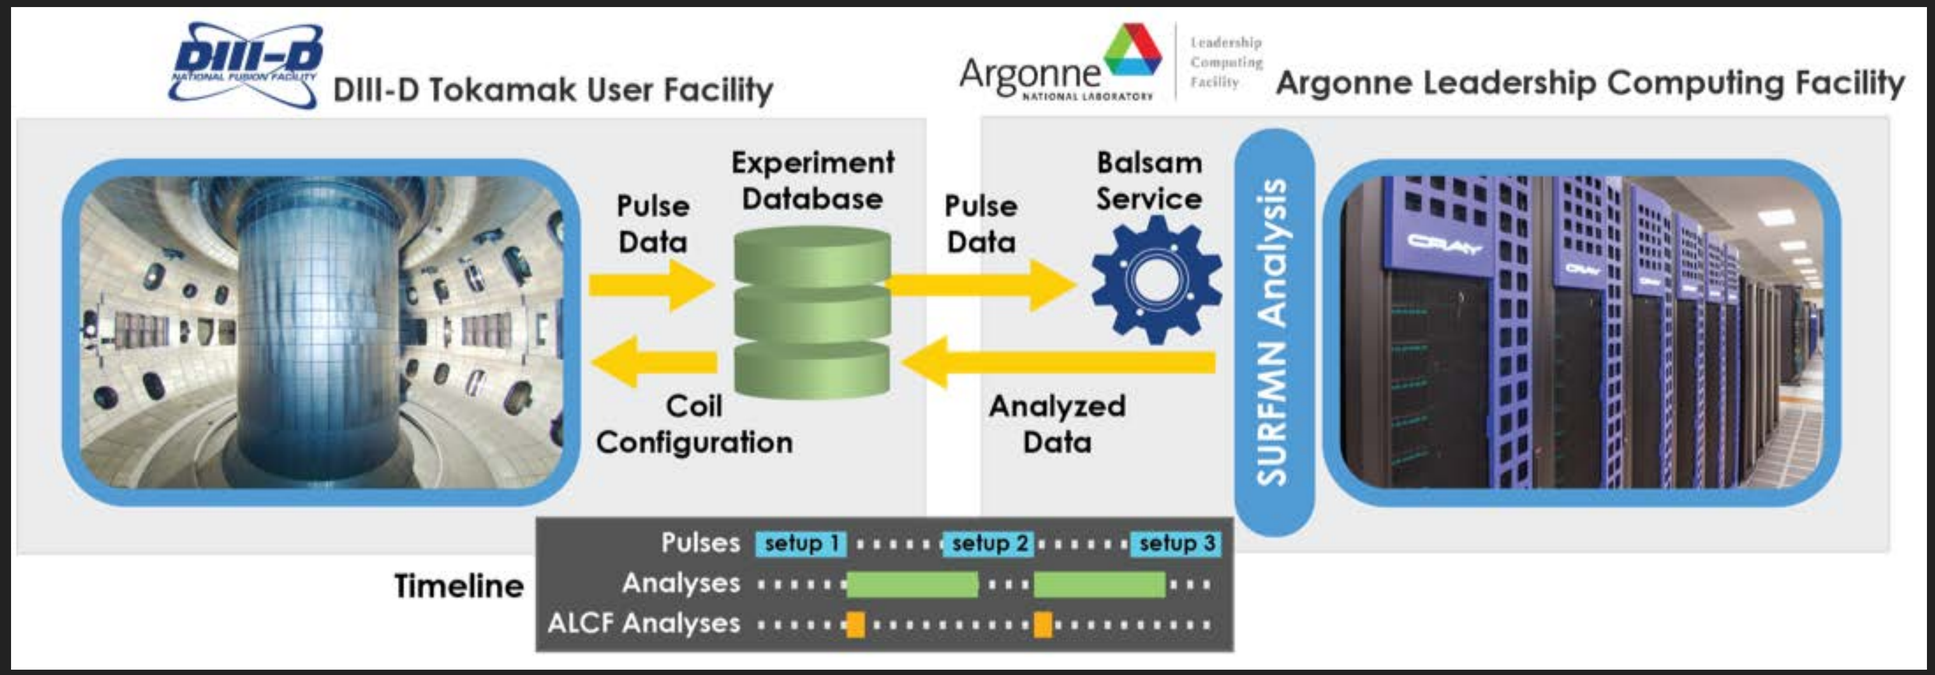

M. Kostuk, T. D. Uram, T. Evans, D. M. Orlov, M. E. Papka, D. Schissel, *Automatic Between-Pulse Analysis of DIII-D Experimental Data Performed Remotely on a Supercomputer at Argonne Leadership Computing Facility*, **Fusion Science and Technology**, September 2017.

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you evolve traditional HPC environment handle complex workloads?

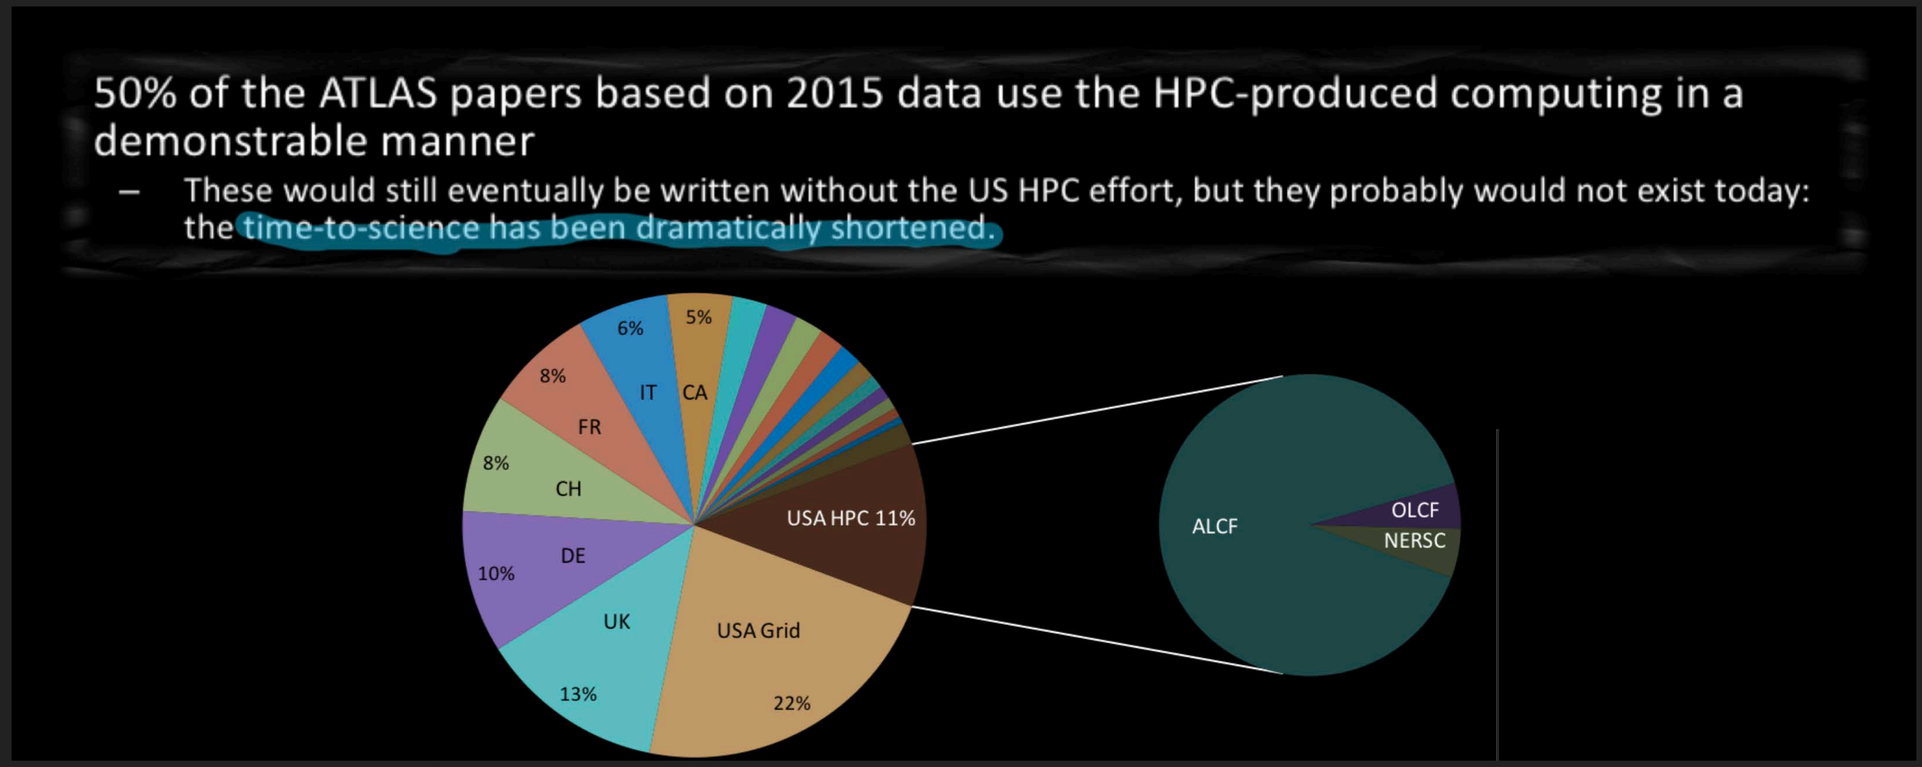

T. LeCompte(HEP){06/02/2016} and J. T. Childers, T. D. Uram, D. Benjamin, T. J. LeCompte, M. E. Papka, *An Edge Service for Managing HPC Workflows*, **Proceedings of the Fourth International Workshop on HPC User Support Tools (HUST'17)**, Denver, CO, November 12, 2017

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do you evolve traditional HPC environment to increase engagement?

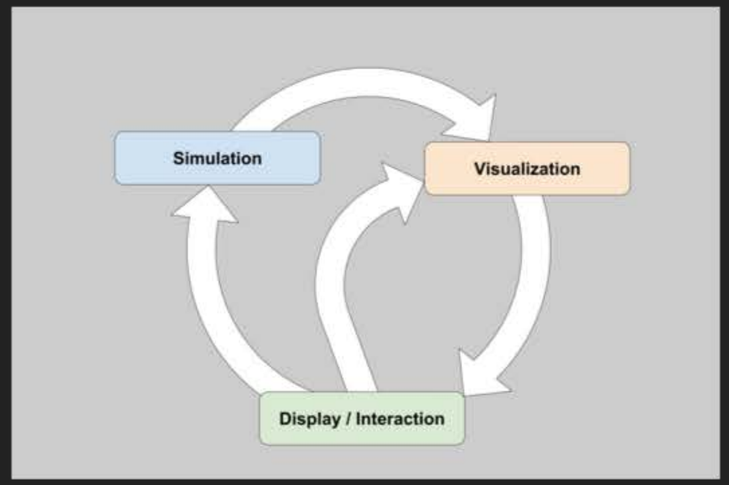

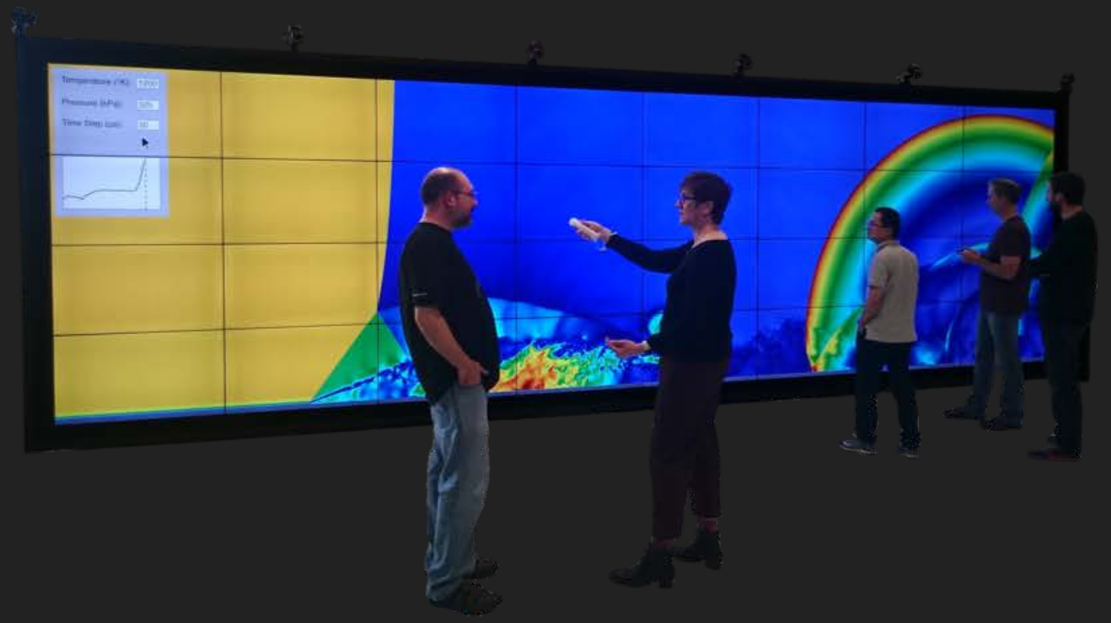

T. Marrinan, A. Nishimoto, J. A. Insley, S. Rizzi, A. Johnson, M. E. Papka, *Interactive Multi-Modal Display Spaces for Visual Analysis*, **Proceedings of the 2016 ACM on Interactive Surfaces and Spaces**, pp. 421-426, Niagara Falls, Canada, November 6, 2016.

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do we **enable** scientists to be the most **productive** from **start to finish**?
  - How do we improve **usability**?
  - How do **simplify** supercomputing?

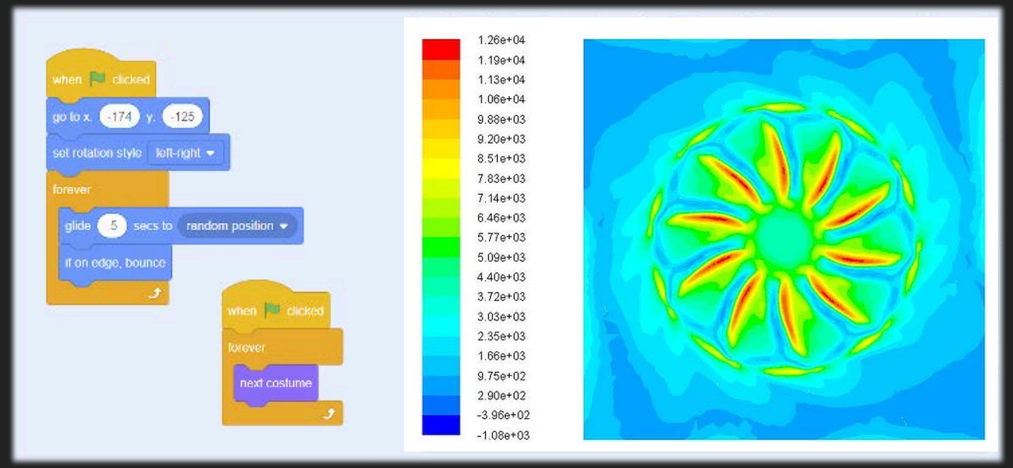

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do we enable scientists to be the most productive from start to finish?
- Problems with science management:
  - Tracking simulations and output: **burdensome**
  - Finding and reproducing old simulations: **difficult**
  - Monitoring live simulations: **inconvenient, idiosyncratic**
  - Post-processing, analysis and archival of results: **haphazard**
  - Assessing simulation behavior/performance: **challenging**

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do we enable scientists to be the most productive from start to finish?
  - **S**imulation **m**anagement and **a**nalysis system for Fl**ash** (Smaash)
    - 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

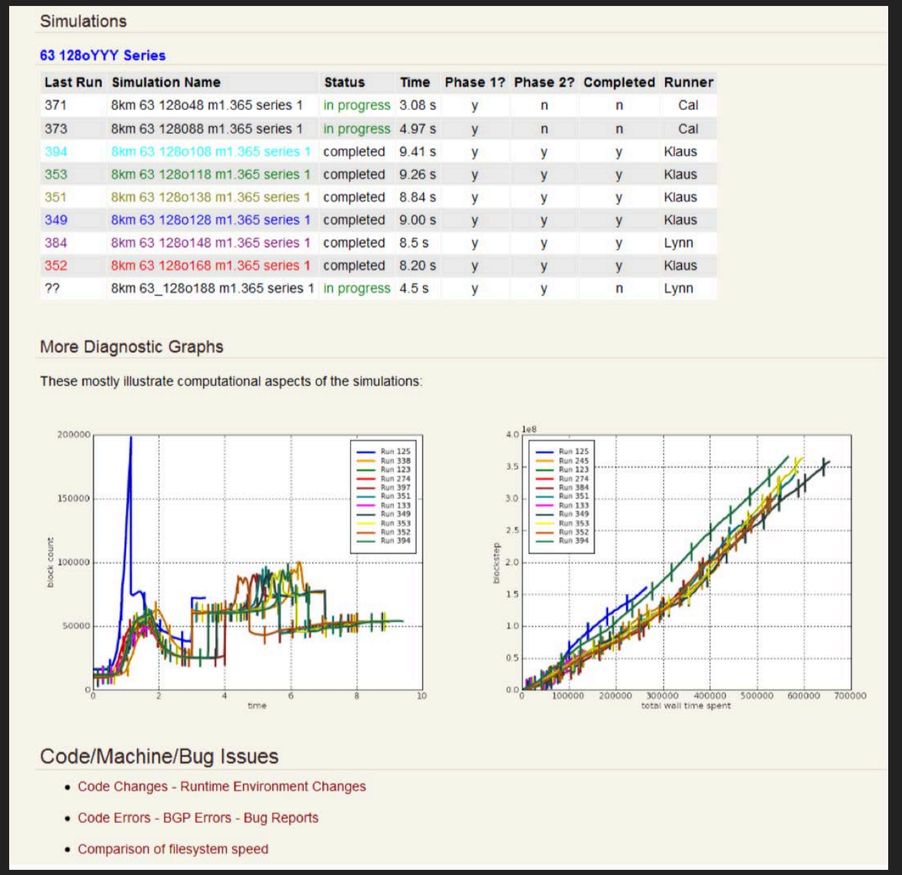

R. Hudson, J. Norris, L. B. Reid, K. Weide, G. C. Jordan, and M. E. Papka, *Experiences Using Smaash to Manage Data-Intensive Simulations*, Proceedings of the 20th International Symposium on High-Performance Parallel and Distributed Computing, pp. 205-15, San Jose, CA, June 2011.

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do we enable scientists to be the most productive from start to finish?

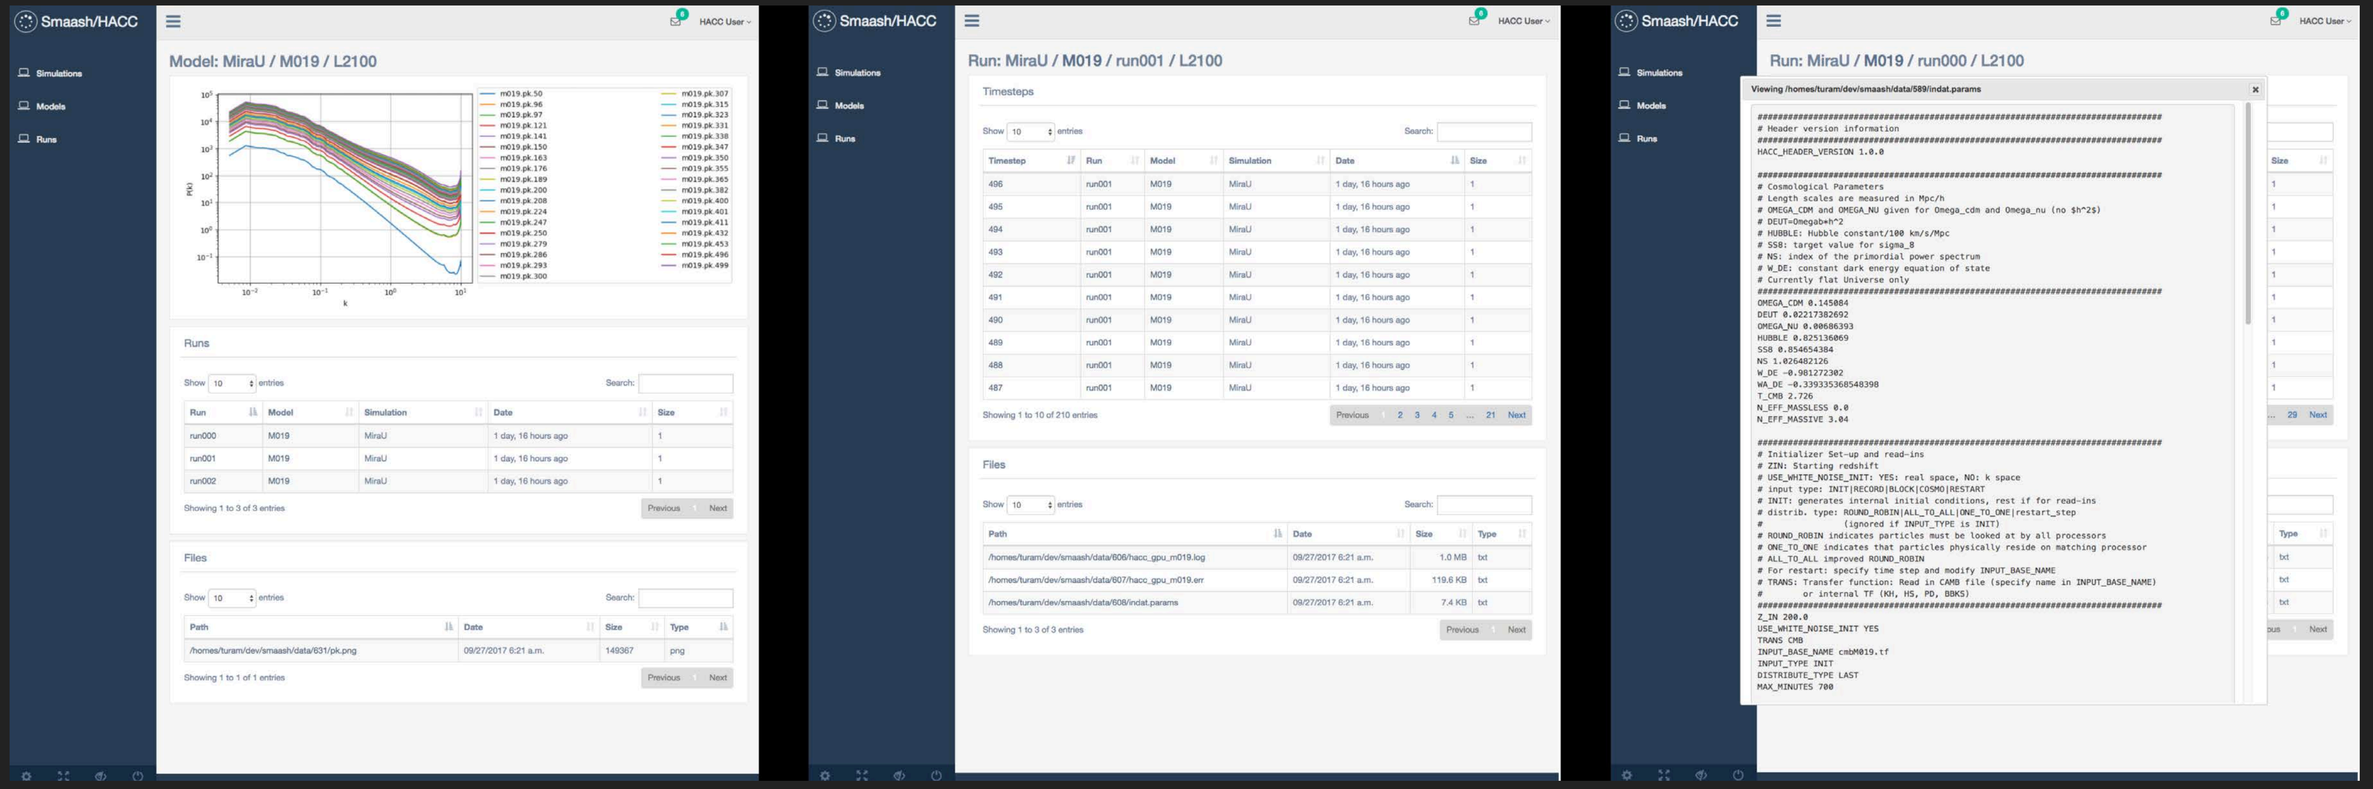

HIGH PERFORMANCE COMPUTING RESEARCH OPPORTUNITIES

- How do we enable scientists to be the most productive from start to finish?

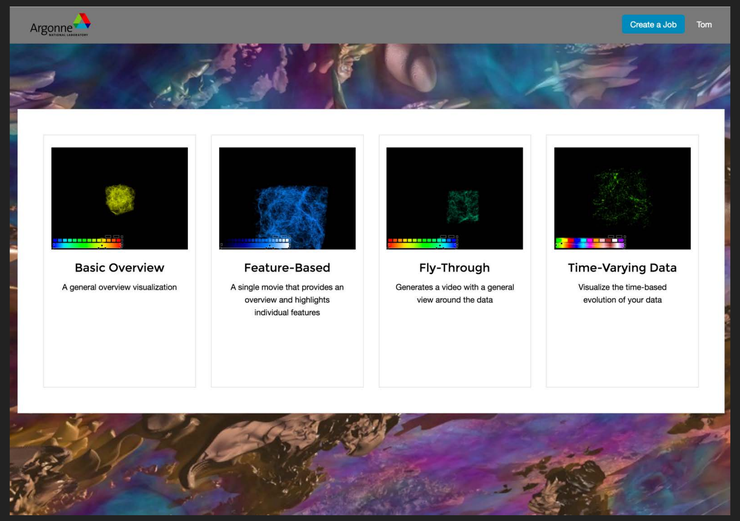

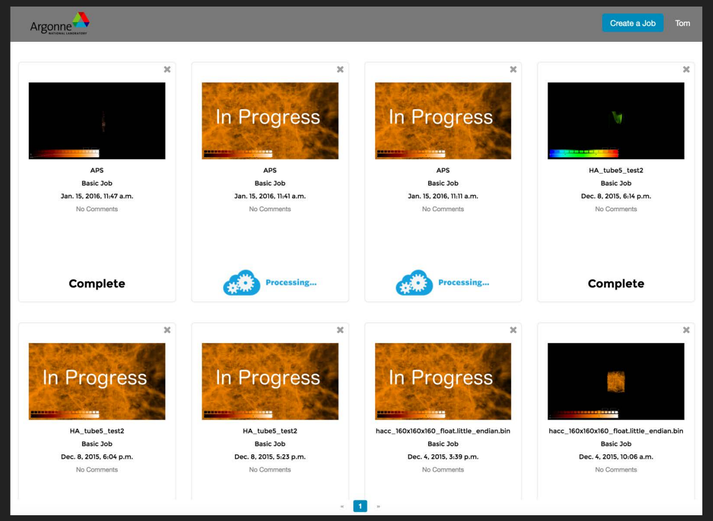

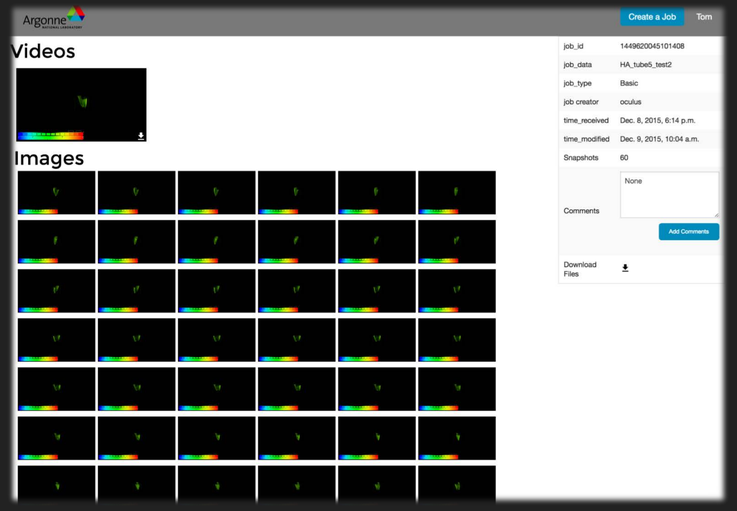

INTERNET OF THINGS (COMPUTING CONTINUUM) [EDGE COMPUTING]

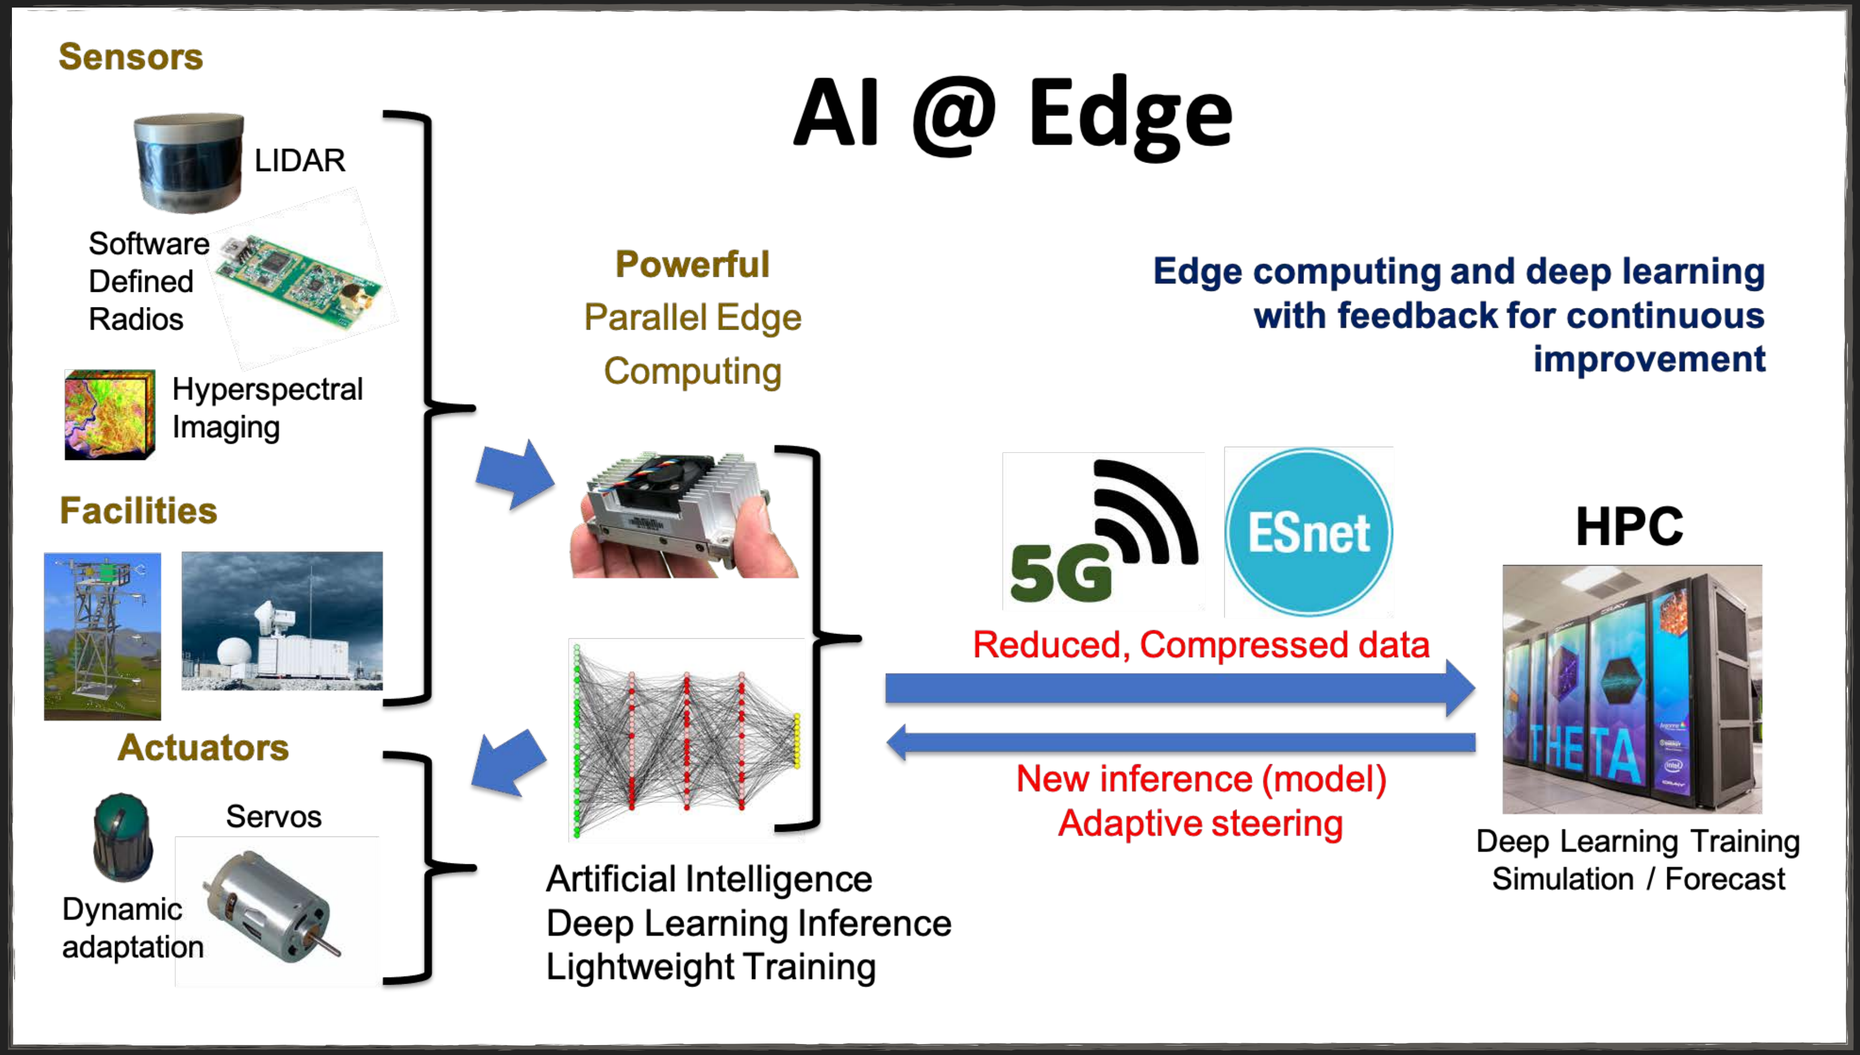

Slide: Pete Beckman (ANL/NWU)

INTERNET OF THINGS (COMPUTING CONTINUUM) [EDGE COMPUTING]

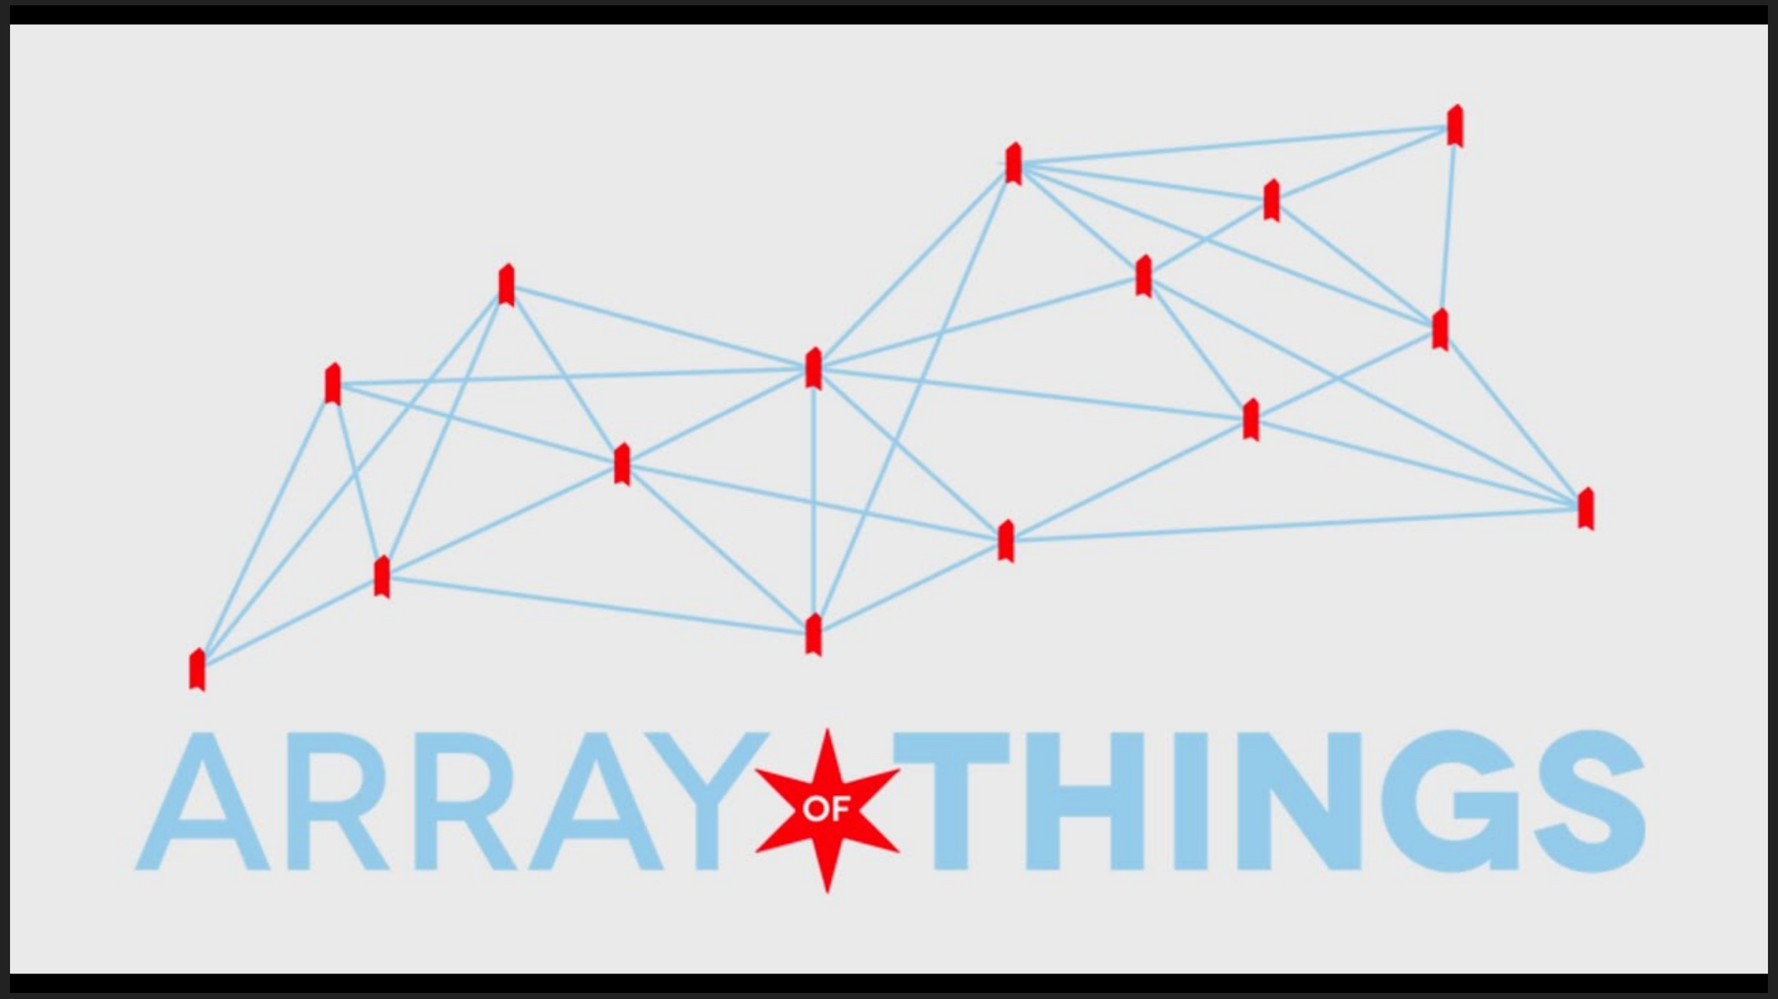

INTERNET OF THINGS (COMPUTING CONTINUUM) [EDGE COMPUTING]

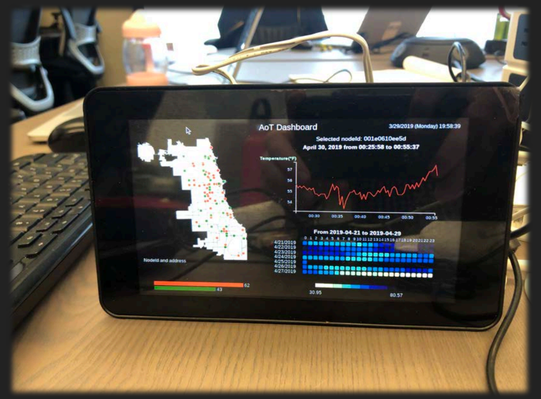

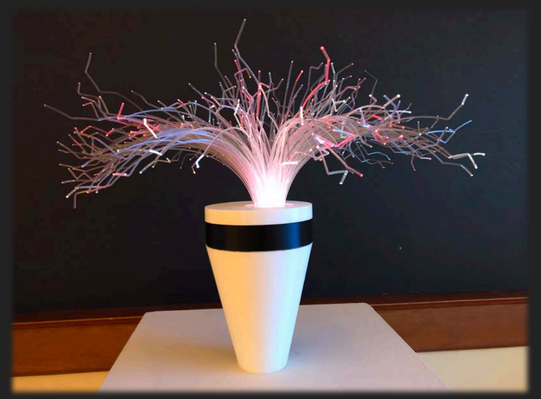

C. Catlett, P. Beckman N. Ferrier, H. Nusbaum, M. E. Papka, M. G. Berman, R. Sankaran, *Measuring Cities with Software-Defined Sensors, Journal of Social Computing, 1(1), pp. 14 - 27, September 2020.*

VIRTUAL REALITY

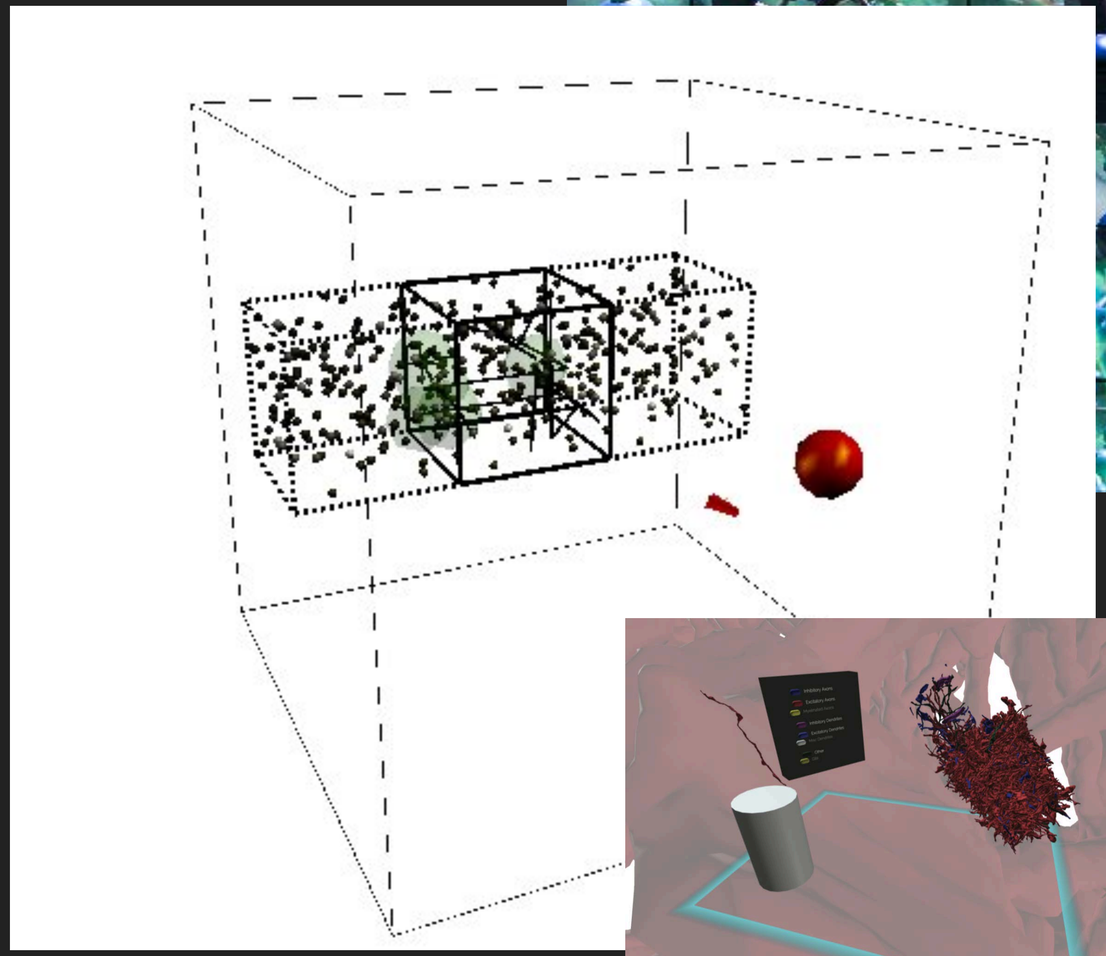

Use of **Virtual Reality** in Science

- Usability, virtual interactions
- Rendering, realistic data in VR time

VIRTUAL REALITY

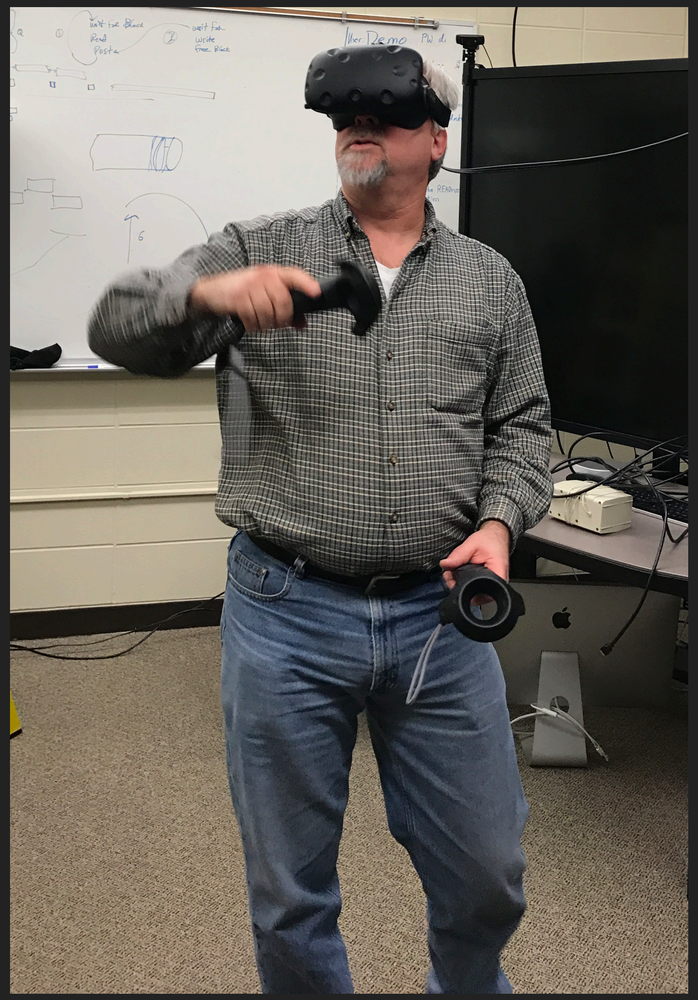

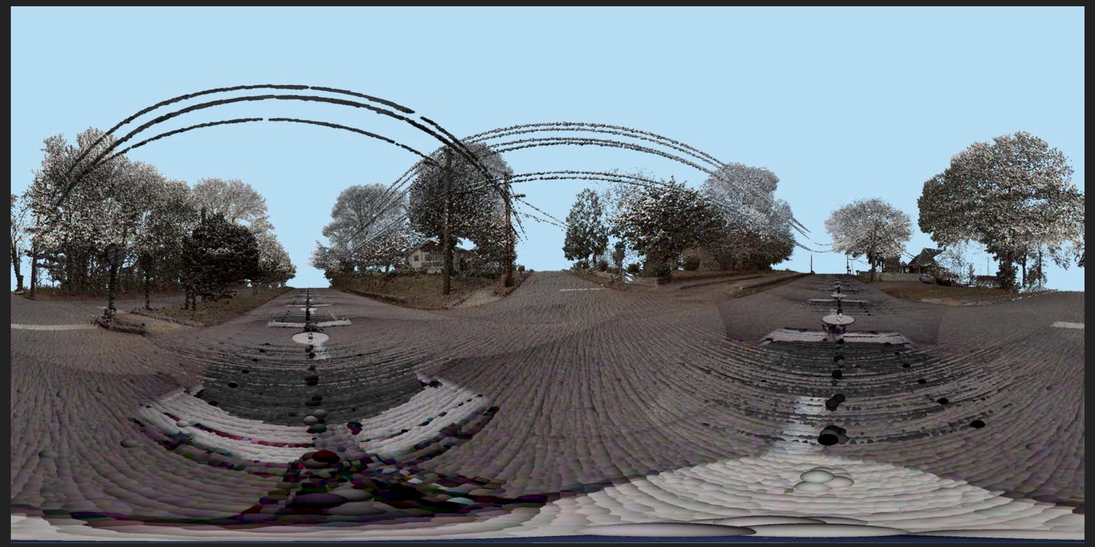

T. Marrinan, M. E. Papka, *Real-Time Omnidirectional Stereo Rendering: Generating 360° Surround-View Panoramic Images for Comfortable Immersive Viewing,* to appear **Proceedings of IEEE VR 2021**, Lisbon, Portugal, 03/27 - 04/01/2021.

SUMMER RESEARCH OPPORTUNITIES (IN PERSON NORMALLY, CURRENTLY VIRTUAL)

www.anl.gov/hr/argonne-career-and-internship-opportunities

ACKNOWLEDGMENTS

- 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 grant OAC-1935984.
- Thanks to all the students of the ddiLab and my colleagues at NIU and ANL.

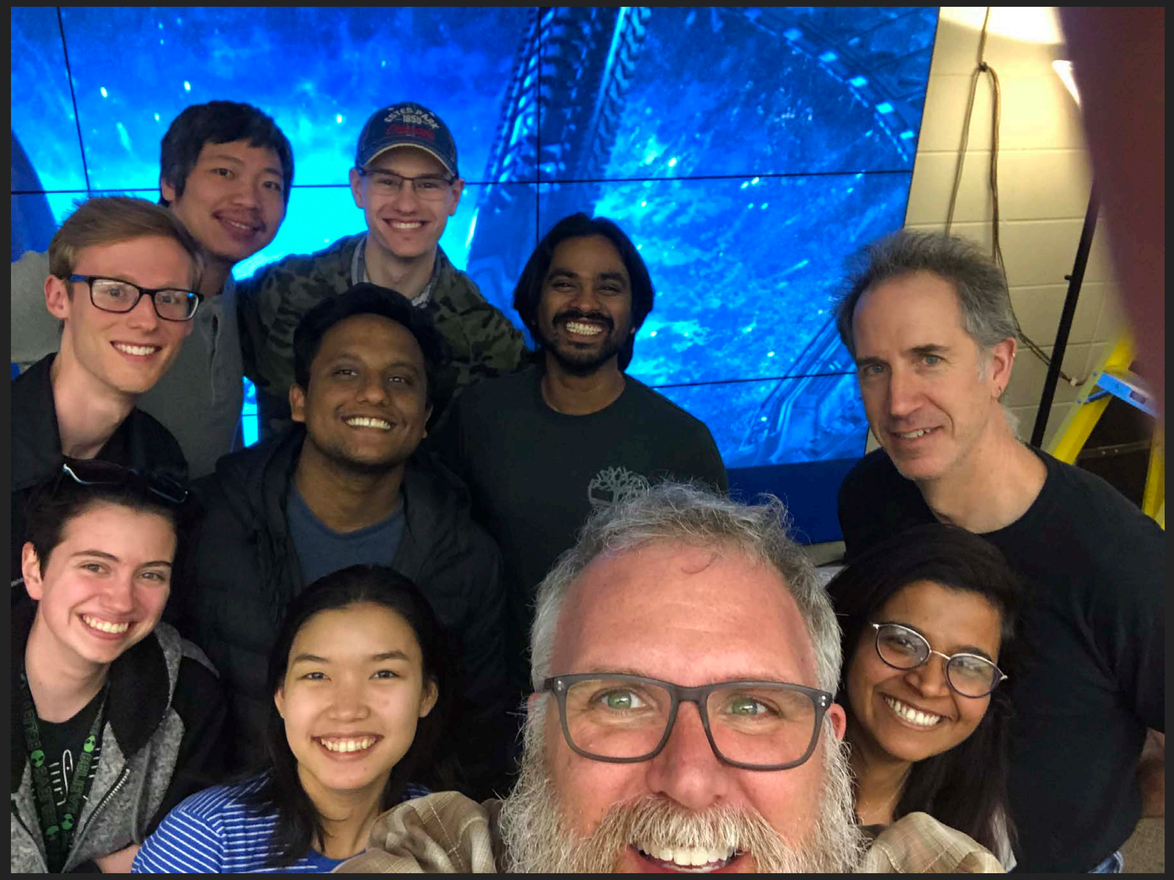

Spring 2019 ddiLab

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

Sir Isaac Newton