U.S. DEPARTMENT OF ENERGY

Science in the Era of Extreme Heterogeneity

SSDBM

Chicago, IL

June 29, 2017

Lucy Nowell, PhD

Computer Science Program Manager

Advanced Scientific Computing Research

DOE Office of Science

DOE ORG Chart

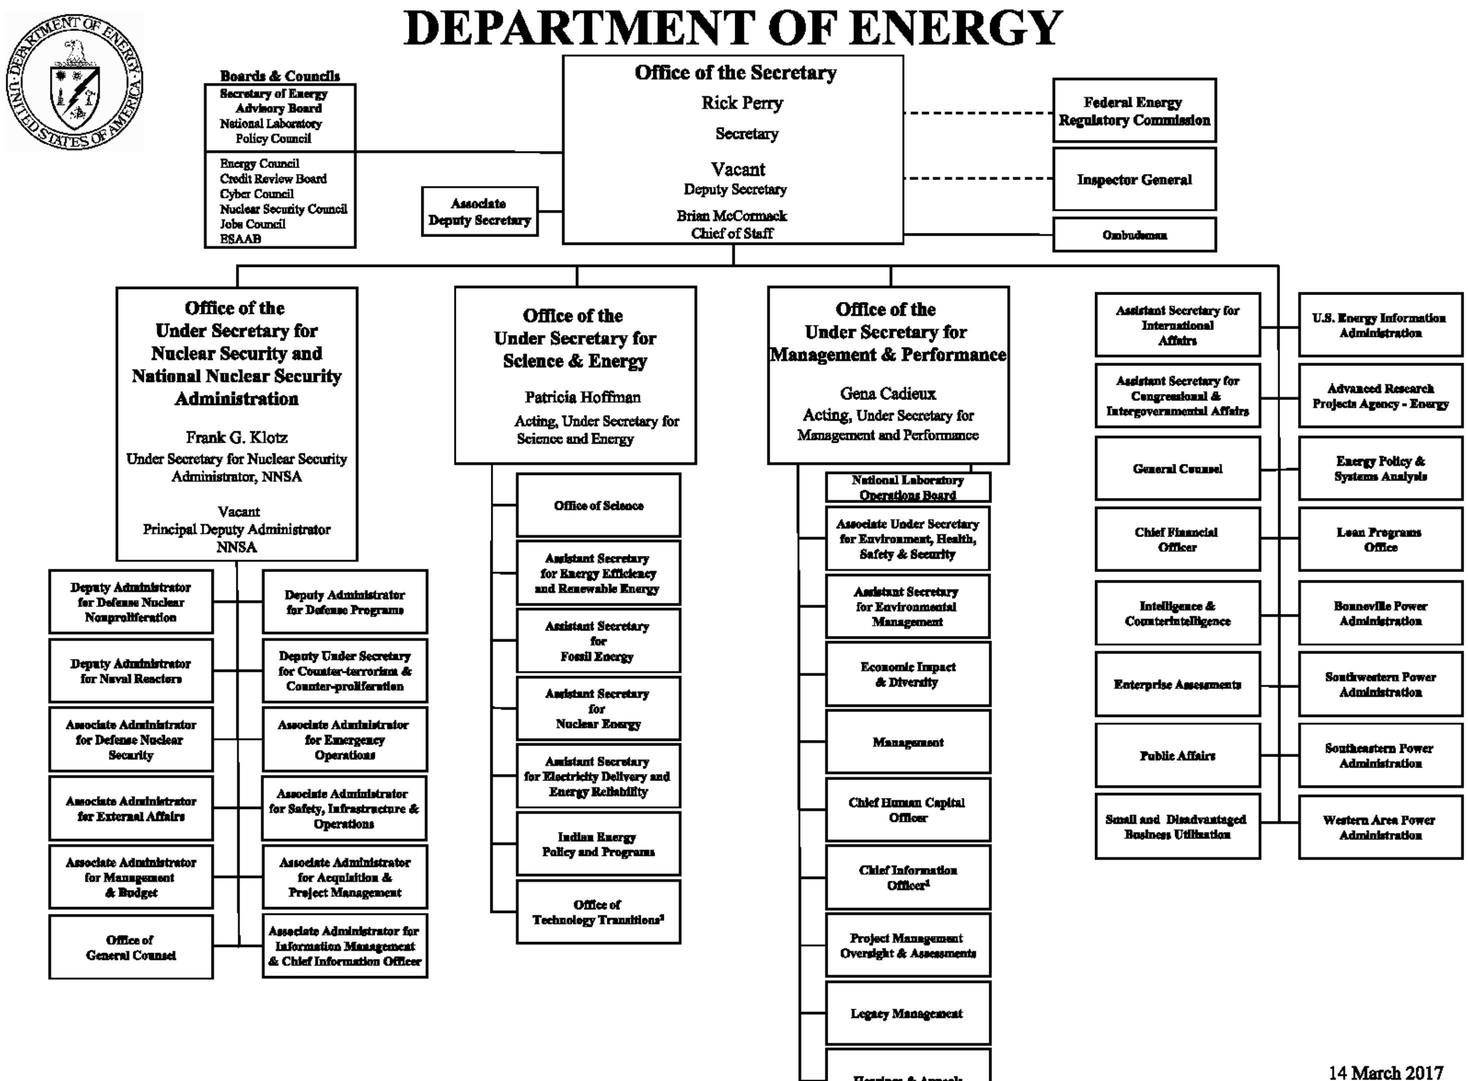

[1] The CIO reports directly to the Secretary for the purposes of carrying out responsibilities under Subchapter 44 U.S.C.§ 3506(a)(2)(A).
[2] The director of the Office of Technology Transitions also serves as DOE's Technology Transfer Coordinator who reports to the Secretary of Energy

U.S. DEPARTMENT OF ENERGY

Office of Science
By the numbers

Shown is a portion of SLAC's two-mile-long linear accelerator (or linac), which provides the electron beam for the new Linac Coherent Light Source (~~LCLS~~) – the world's first hard x-ray, free-electron laser. For nearly 50 years, SLAC's linac had produced high-energy electrons for physics experiments. Now researchers use the very intense X-ray pulses (more than a billion times brighter than the most powerful existing sources) much like a high-speed camera to take stop-motion pictures of atoms and molecules in motion, examining fundamental processes on femtosecond timescales.

SC delivers scientific discoveries and tools to transform our understanding of nature and advance the energy, economic, and national security of the U.S.

Research

- Provides about half of the U.S. Federal support for basic research in the physical sciences;
- Supports about 19,000 Ph.D. scientists, graduate students, engineers, and support staff at over 300 institutions and 10 DOE national laboratories;
- Maintains U.S. and world leadership in high-performance computing and computational sciences;
- Continues to be the major U.S. supporter of physics, chemistry, materials sciences, and biology for discovery and for energy sciences.

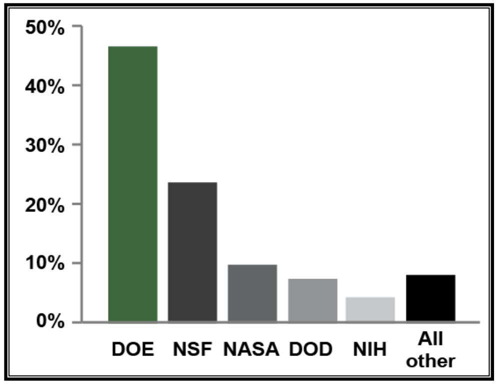

Support for basic research in the physical sciences by agency.

Source: *NSF Science and Engineering Indicators 2012*

Scientific User Facilities

- SC maintains the world's largest collection of scientific user facilities (aka research infrastructure) operated by a single organization in the world, used by more than 31,000 researchers each year.

Office of Science at a Glance

FY 2018 Request: $4.472B, (-16%)

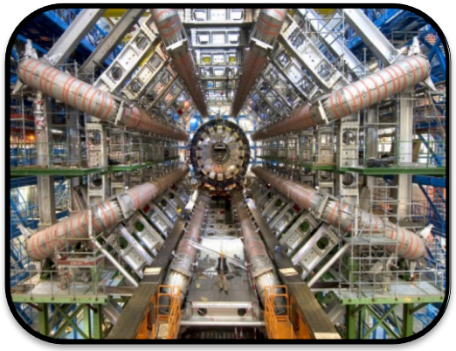

Largest Supporter of Physical Sciences in the U.S.

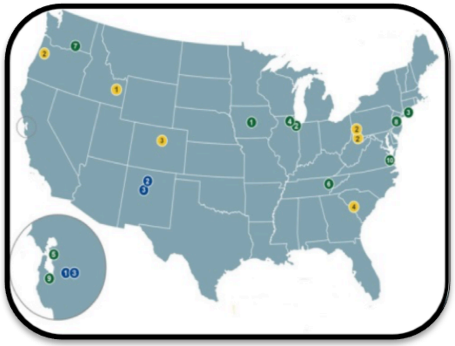

Funding at >300 Institutions including all 17 DOE Labs

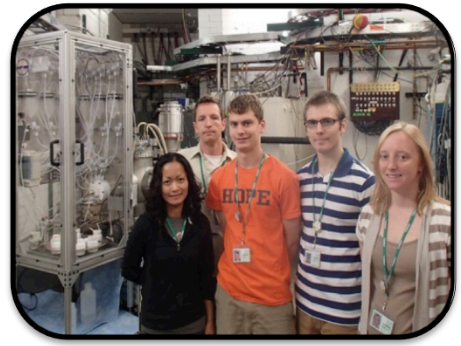

> 20,000 Scientists Supported

>31,000 Users of 25 SC Scientific Facilities

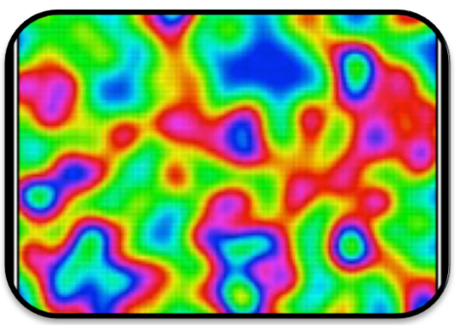

Research: 42%

~40% of Research to Universities

Facility Operations: 36%

Seventeen DOE National Laboratories

Office of Science Laboratories

1. **Ames Laboratory**
   Ames, Iowa
2. **Argonne National Laboratory**
   Argonne, Illinois
3. **Brookhaven National Laboratory**
   Upton, New York
4. **Fermi National Accelerator Laboratory**
   Batavia, Illinois
5. **Lawrence Berkeley National Laboratory**
   Berkeley, California
6. **Oak Ridge National Laboratory**
   Oak Ridge, Tennessee
7. **Pacific Northwest National Laboratory**
   Richland, Washington
8. **Princeton Plasma Physics Laboratory**
   Princeton, New Jersey
9. **SLAC National Accelerator Laboratory**
   Menlo Park, California
10. **Thomas Jefferson National Accelerator Facility**
    Newport News, Virginia

Other DOE Laboratories

1. **Idaho National Laboratory**
   Idaho Falls, Idaho
2. **National Energy Technology Laboratory**
   Morgantown, West Virginia
   Pittsburgh, Pennsylvania
   Albany, Oregon
3. **National Renewable Energy Laboratory**
   Golden, Colorado
4. **Savannah River National Laboratory**
   Aiken, South Carolina

NNSA Laboratories

1. **Lawrence Livermore National Laboratory**
   Livermore, California
2. **Los Alamos National Laboratory**
   Los Alamos, New Mexico
3. **Sandia National Laboratory**
   Albuquerque, New Mexico
   Livermore, California

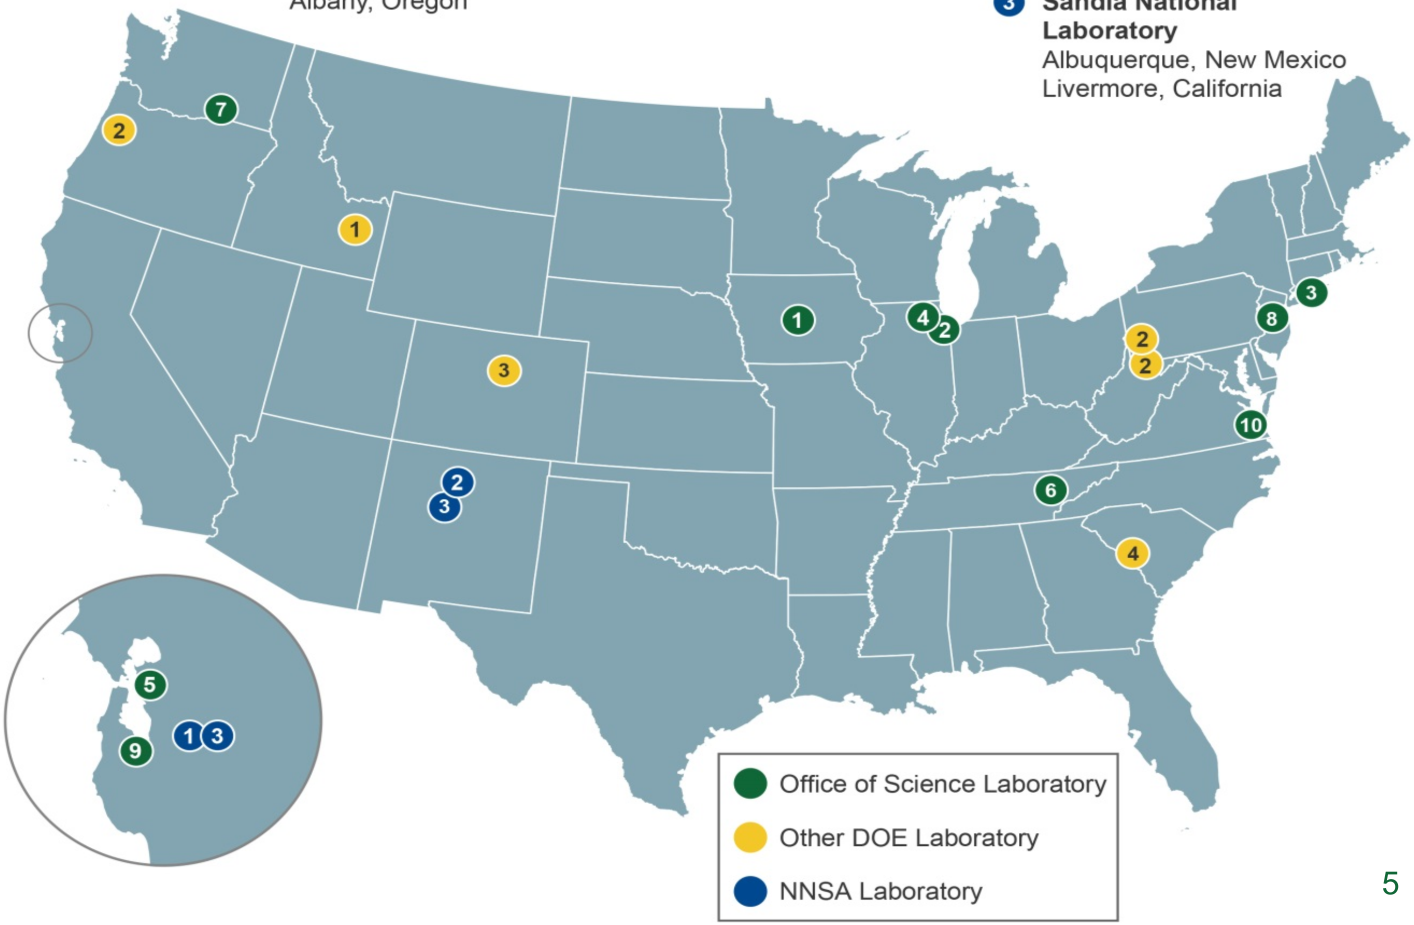

FY 2016
28 user facilities

U.S. DEPARTMENT OF
ENERGY
Office of Science

SC Research Programs

- **Advanced Scientific Computing Research (FY 2017 budget request $722 million) –** Advances applied mathematics, computer science, and computational research to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the U.S. Builds and operates some of the fastest computers in the world for open science. Leads the U.S. effort to develop the next generation of computing tools (exascale).
- **Basic Energy Sciences (FY 2017 budget request $1,554 million) –** Advances fundamental research to understand, predict, and ultimately control matter and energy at the electronic, atomic, and molecular levels to provide foundations for new energy technologies. Supports a large portfolio of core research in chemical sciences, geosciences, biosciences, and materials sciences and engineering, and targeted areas to advance DOE energy priorities. Constructs and supports scientific user facilities that enable atomic-level visualization and characterization of materials of all kinds, including biological materials.
- **Biological and Environmental Research (FY 2017 budget request $349 million) –** Advances fundamental research to achieve a predictive understanding of complex biological, climatic, and environmental systems for a secure and sustainable energy future. Supports core research in genomic sciences of plants and microbes, research to understand climate-relevant atmospheric and ecosystem processes and to understand the dynamic physical, biogeochemical, microbial, and plant processes interactions.
- **Fusion Energy Sciences (FY 2017 budget request $310 million) –** Advances the theoretical and experimental understanding of matter at high temperatures and density, including plasmas, plasma confinement, and fusion science.
- **High Energy Physics (FY 2017 budget request $673 million) –** Advances understanding of the basic constituents of matter, deeper symmetries in the laws of nature at high energies, and mysterious phenomena that are commonplace in the universe, such as dark energy and dark matter.
- **Nuclear Physics (FY 2017 budget request $503 million) –** Advances experimental and theoretical research to discover, explore, and understand all forms of nuclear matter. Supports DOE's isotopes production and applications program for production of stable and radioactive research isotopes.

SC Transition Leadership

U.S. DEPARTMENT OF ENERGY | Office of Science

THE OFFICE OF

ADVANCED SCIENTIFIC COMPUTING RESEARCH

Functional Organization Chart

The Office Of The Associate Director of Advanced Scientific Computing Research
- Barbara Helland, Associate Director
- Julie Stambaugh - Financial Analyst
- Lori Jernigan - Program Support Specialist
- Tameka Morgan (Contractor) - Administrative Assistant

ASCAC
- Christine Chalk - DFO
- Sally McPherson - Admin Support

Facilities Division
- Christine Chalk – Director (Acting)
- Sally McPherson - Administrative Support
- Facilities
  - Oak Ridge Leadership Computing Facility — Christine Chalk
  - Argonne Leadership Computing Facility — Betsy Riley
  - NERSC — Dave Goodwin
  - ESnet — Ben Brown
  - ALCC — Carolyn Lauzon
  - REP — Claire Cramer, Robinson Pino
  - CSGF — Christine Chalk

Computational Science Research and Partnerships (SciDAC) Division
- Steven Lee Director (Acting)
- Teresa Beachley - Grants and Contracts
- Angie Thevenot - Program Support Specialist
- Computer Science
  - Computer Science — Lucy Nowell
  - Data & Visualization
  - Extreme Scale — Sonia Sachs
  - File Systems — Vacant
  - HPC Sci & Apps — Vacant
- Applied Math
  - Base, Math: Algorithms, Models, Data — Steven Lee
  - Multiscale Mathematics — Abani Patra
  - Extreme Scale Science — Vacant
- Computational Partnerships
  - Biology SAPs — Vacant
  - Non-SC Program SAPs — Randall Laviolette
  - SC Program SAPs — Ceren Susut
  - Non-SC Program SAPs — Vacant
  - Math – SciDAC Centers & Institutes — Steven Lee-Acting
  - CS – SciDAC Centers & Institutes — Vacant
- Distributed Network Environment
  - Network Research — Thomas Ndousse-Fetter
  - Collaboratories/ Middleware — Rich Carlson

Advanced Computing Technologies Division
- Director - Vacant

February 2017

ASCR Investment Priorities

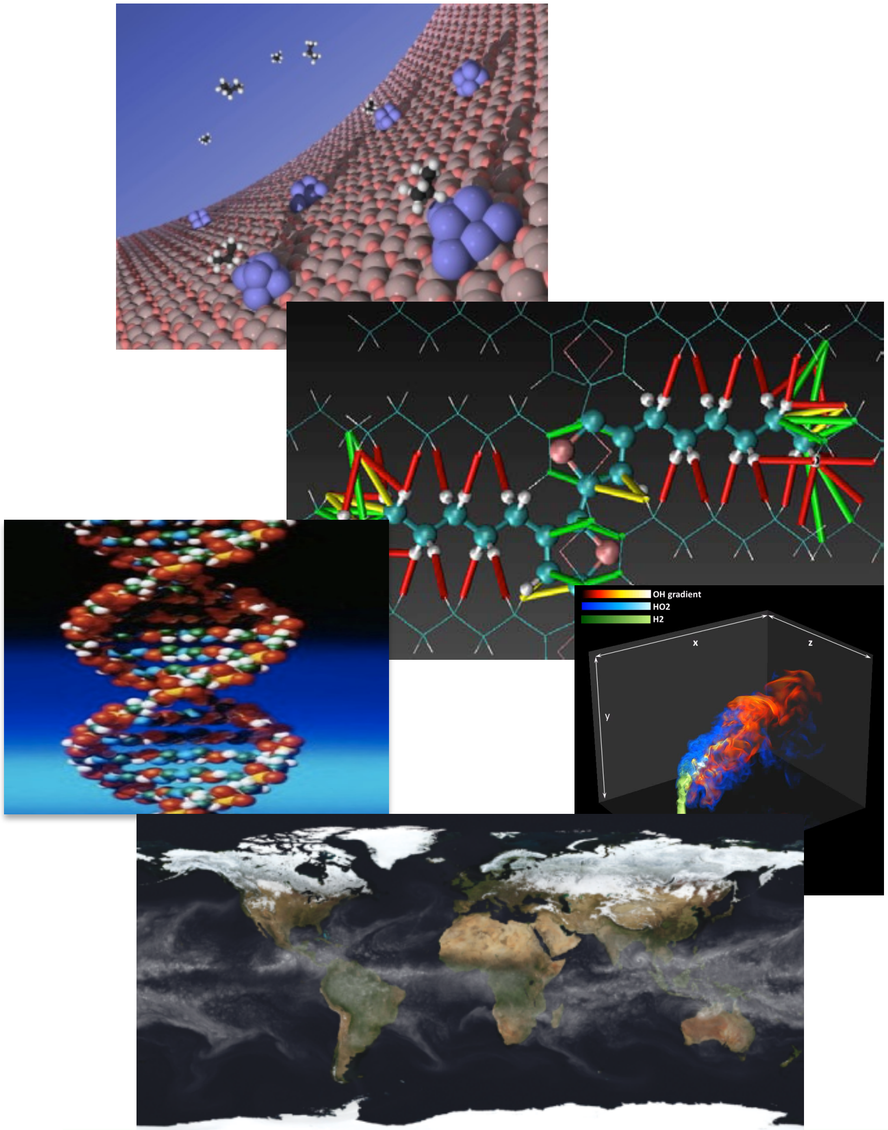

- **Exascale –** conduct research and development, and design efforts in hardware software, and mathematical technologies that will produce exascale systems for science applications
- **Facilities –** acquire and operate more capable computing systems, from multi-petaflop through exascale computing systems that incorporate technologies emerging from research investments
- **Large Scientific Data –** prepare today's scientific and data-intensive computing applications to migrate to and take full advantage of emerging technologies from research, development and design efforts
- **Begin R&D for post-Moore Era**

Top 10 Supercomputers in the World – June 2017

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
2	National University of Defense Technology	NUDT	Tianhe-2 NUDT TH-IVB-FEP, Xeon 12C 2.2GHz, IntelXeon Phi	China	3,120,000	33.9	17.8
3	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	361,760	19.6	2.27
4	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Opteron 16C 2.2GHz, Gemini, NVIDIA K20x	USA	560,640	17.6	8.21
5	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	17.2	7.89
6	Lawrence Berkeley National Laboratory	Cray	Cori Cray XC40, Intel Xeons Phi 7250 68C 1.4 GHz, Aries	USA	622,336	14.0	3.94
7	JCAHPC Joint Center for Advanced HPC	Fujitsu	Oakforest-PACS PRIMERGY CX1640 M1, Intel Xeons Phi 7250 68C 1.4 GHz, OmniPath	Japan	556,104	13.6	2.72
8	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	795,024	10.5	12.7
9	Argonne National Laboratory	IBM	Mira BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	786,432	8.59	3.95
10	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Xeon E5 16C 2.3GHz, Aries	USA	301,0564	8.10	4.23

ASCR Computing Upgrades At a Glance

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Name Planned Installation	**Edison**	**TITAN**	**MIRA**	**Cori** **2016**	**Summit** **2017-2018**	**Theta** **2016**	**Aurora** **2018-2019**
System peak (PF)	2.6	27	10	> 30	200	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	13.3	1.7	13
Total system memory	357 TB	710TB	768TB	~1 PB DDR4 + HBM+1.5PB persistent memory	> 2.4 PB DDR4 + HBM + 3.7 PB persistent memory	676 TB DDR4 + HBM	> 7 PB HBM, Local Memory and Persistent Memory
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~4,600 nodes	>3,200 nodes	>50,000 nodes
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

- **Demand exceeds capability by 2–6× across centers today**
- **Typical systems run at 80–90%+ utilization**

Innovative and Novel Computational Impact on Theory and Experiment

INCITE promotes transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage at the Argonne and Oak Ridge Leadership Computing Facilities (LCFs) for computationally intensive, large-scale research projects.

The 2018 INCITE Call for Proposals opened April 17, 2017 and will close June 23, 2017.

For more information visit http://www.doeleadershipcomputing.org/

2017 INCITE award statistics

- Request for Information helped attract new projects
- Call closed June 24th, 2016
- Total requests of more than 13 billion core-hours
- Awards of 5.8 billion core-hours for CY 2017
- **55 projects awarded of which 17 are renewals**

PIs by Affiliation (Awards)

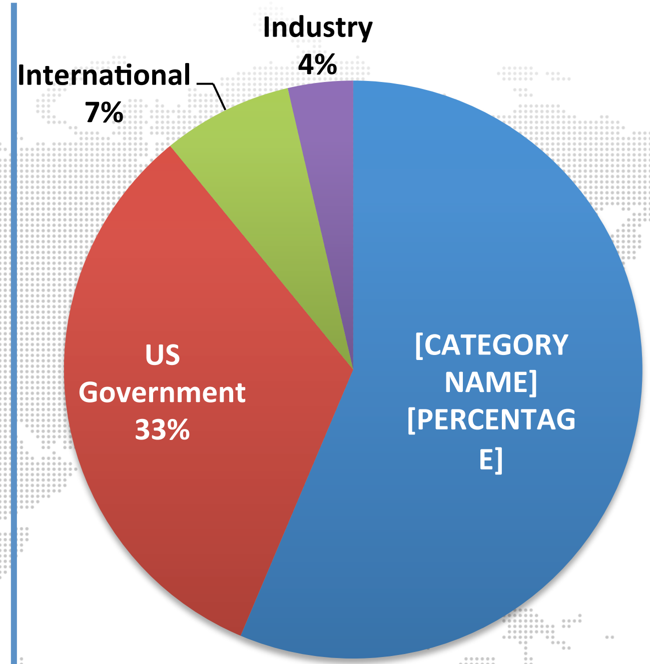

Acceptance rates

45% of nonrenewal submittals and 85% of renewals

Contact information
Judith C. Hill, INCITE Manager
hilljc@DOEleadershipcomputing.org

Summit on Extreme Heterogeneity

- **Held on June 8-9 at the DOE Germantown facility**
- **Organized by Jeffrey Vetter of Oak Ridge National Laboratory, in collaboration with**
  - John Shalf, Lawrence Berkeley National Laboratory
  - Pat McCormick, Los Alamos National Laboratory
  - Ron Brightwell, Sandia National Laboratory

- **Other participants:**
  - Maya Gokhale, Lawrence Livermore National Laboratory
  - David Donofrio, Lawrence Berkeley National Laboratory
  - Rob Ross, Argonne National Labortory
  - Catherine Schuman, Oak Ridge National Laboratory
  - Travis Humble, Oak Ridge National Laboratory (phone)
  - Katie Antypas, Lawrence Berkeley National Laboratory (phone)
  - Shinjae Yoo, Brookhaven National Laboratory (phone)

Summit Goals

- **Reach a shared understanding of what we mean by “extreme heterogeneity”**
- **Determine whether a workshop on this topic is needed**
- **If so, identify potential agenda items for the workshop**

- **Timeframe of interest**
  - 2025-2040 and beyond

Extreme Heterogeneity

- **Increased diversity in**
  - Processors and memory/storage technology
  - Workflows, including data intensive science, interactive supercomputing, and near-real-time requirements
  - Increasing complexity, volume and velocity of data, as Science User Facilities operated by other other Office of Science programs increasingly want to use ASCR supercomputers for data analytics, with requirements for near-real-time analysis to support steering experiments
  - Users and user expertise in HPC, as more scientists need the power of HPC

There have been multiple recent assaults on "Moore's Law" *(this isn't the first time)*

Moore's Law is an economic theory. ***There are ways to still scale digital technology after the end of classical lithographic scaling***

(e.g. end of Dennard Scaling in ~2004

No more exponential clock frequency scaling

Move to exponentially increasing parallelism)

Numerous Opportunities to Continue Moore's Law Technology!

(but winning solution is unclear)

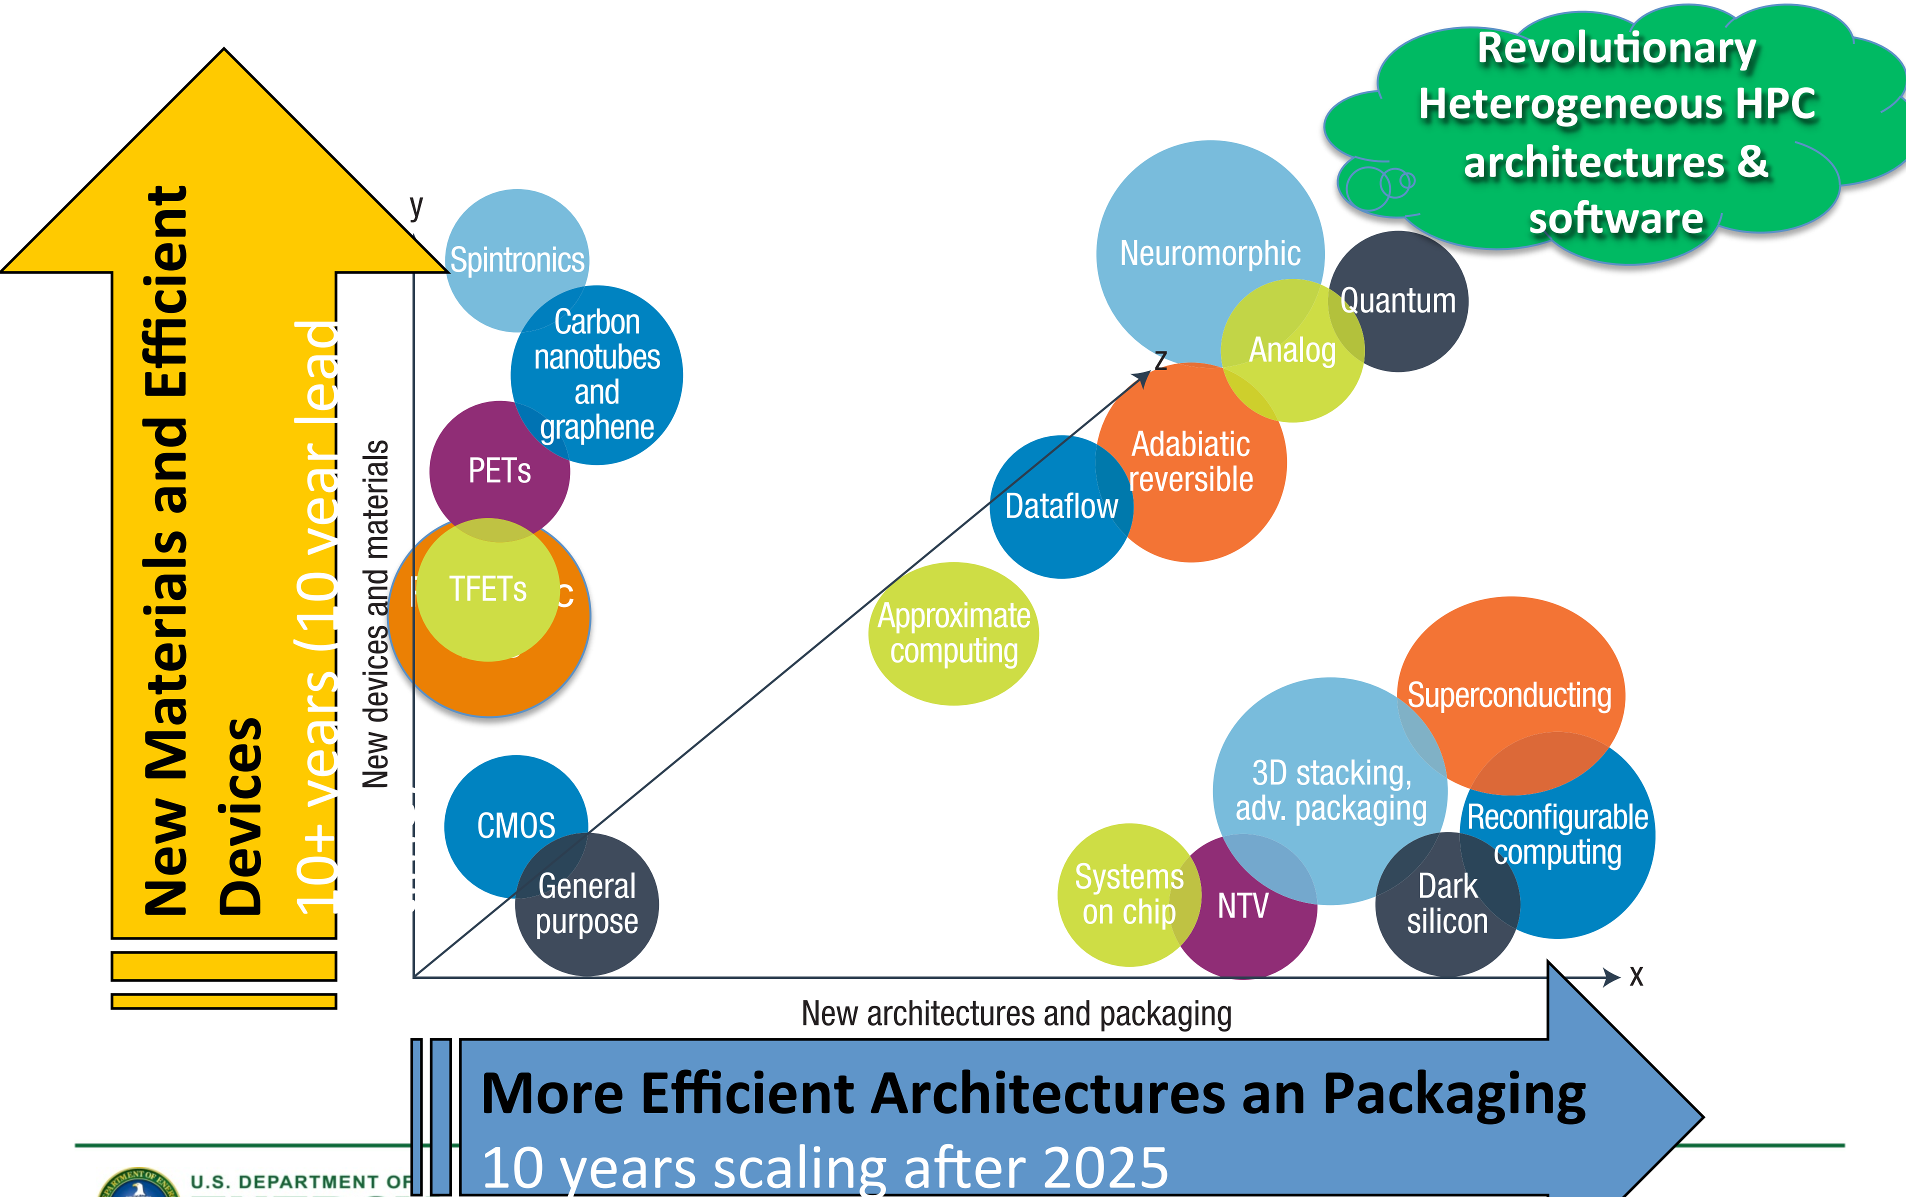

Slide courtesy of John Shalf

U.S. DEPARTMENT OF ENERGY

Computing Beyond Moore's Law

TABLE 1. Summary of techology options for extending digital electronics.

Improvement Class	Technology	Timescale	Complexity	Risk	Opportunity
Architecture and software advances	Advanced energy management	Near-Term	Medium	Low	Low
	Advanced circuit design	Near-Term	High	Low	Medium
	System-on-chip specialization	Near-Term	Low	Low	Medium
	Logic specialization/dark silicon	Mid-Term	High	High	High
	Near threshold voltage (NTV) operation	Near-Term	Medium	High	High
3D integration and packaging	Chip stacking in 3D using thru-silicon vias (TSVs)	Near-Term	Medium	Low	Medium
	Metal layers	Mid-Term	Medium	Medium	Medium
	Active layers (epitaxial or other)	Mid-Term	High	Medium	High
Resistance reduction	Superconductors	Far-Term	High	Medium	High
	Crystaline metals	Far-Term	Unknown	Low	Medium
Millivolt switches (a better transistor)	Tunnel field-effect transistors (TFETs)	Mid-Term	Medium	Medium	High
	Heterogeneous semiconductors/strained silicon	Mid-Term	Medium	Medium	Medium
	Carbon nanotubes and graphene	Far-Term	High	High	High
	Piezo-electric transistors (PFETs)	Far-Term	High	High	High
Beyond transistors (new logic paradigms)	Spintronics	Far-Term	Medium	High	High
	Topological insulators	Far-Term	Medium	High	High
	Nanophotonics	Near/Far-Term	Medium	Medium	High
	Biological and chemical computing	Far-Term	High	High	High

Topics Discussed

- Architectural trends
- Programming Systems
- Operating and Runtime Systems
- Libraries
- I/O, storage, vis, analytics
- Performance Portability
- Tools: performance, debugging, correctness

Lessons Learned

- **DOE supercomputers now in planning stages have increasingly heterogeneous processors and memory.**
- **Every ASCR supercomputing facility maintains a staff of 10+ experts whose role is to help science applications run (more efficiently) on our machines. At least in part, this is a reflection of the difficulty of programming the machines.**
- **Lack of code portability is already a significant issue and it is becoming worse.**
- **After years of effort, less than half of codes running on DOE supercomputers use GPUs and still fewer use them effectively.**

How Extreme Heterogeneity May Impact Science

- **Exponential increase in the difficulty of programming such extremely complex supercomputers**
  - More knowledge required to program amid extremely heterogeneous hardware
  - More devices to choose among means more risk of making a bad choice
  - Lost developer productivity
  - Lost performance

Summit Findings

- Users of supercomputing systems are being confronted with an exponential increase in architectural and programming complexity.
- Portable software stacks, tools, and programming systems and related software systems are critical to help users and facilities manage or at least mitigate this extreme heterogeneity.
- Computing markets are creating a greater diversity of problems that are being attacked with HPC resources.
- Extreme Heterogeneity in computer architectures is relatively new, and not well understood by the HPC community.

Why Do I Talk About This Here?

- **Extreme heterogeneity will result from market forces that are beyond our control. The changing technology will impact you, too!**
  - Do you write code? Do you care about your own productivity?
  - Do you ever use a "big" computer?
    - We expect petaflop systems to be common by 2030, if not sooner.
  - Do you care about application performance?

Next Steps

- **Workshop on Extreme Heterogeneity in very early planning stages.**
- **Goals:**
  - Increase our understanding of the potential impacts of Extreme Heterogeneity
  - Define a research agenda for developing a software stack and programming environment
  - Begin building a community of practice in this area

- **Stay tuned!**

Thank you!

Lucy Nowell

Lucy.Nowell@science.doe.gov