U.S. DEPARTMENT OF ENERGY | Office of Science

Advanced Scientific Computing Research

Presented to the

Advanced Scientific Computing Advisory Committee

by

Barbara Helland
Associate Director

September 26, 2017

Where we were in April

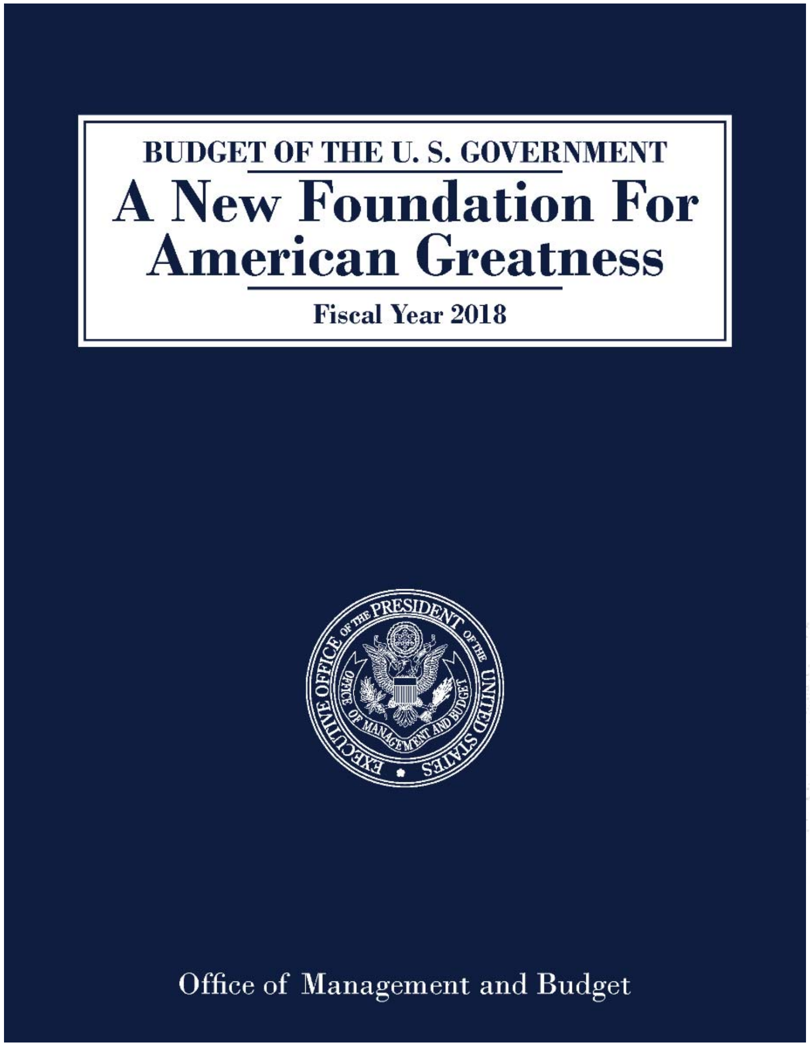

"This Budget, therefore, includes $639 billion for the Department of Defense—a $52 billion increase from the 2017 annualized continuing resolution level. This increase will be offset by targeted reductions elsewhere." page 2

Office of Science FY 2018 President's Request

(Dollars in thousands)

	FY 2016 Enacted	FY 2016 Current w/SBIR-STTR[a]	FY 2017 Enacted	FY 2018 President's Request	FY 2018 Request vs. FY 2017 Enacted	
Science						
Advanced Scientific Computing Research	621,000	621,000	647,000	722,010	+75,010	+11.6%
Basic Energy Sciences	1,849,000	1,849,000	1,871,500	1,554,500	-317,000	-16.9%
Biological and Environmental Research	609,000	609,000	612,000	348,950	-263,050	-43.0%
Fusion Energy Sciences	438,000	438,000	380,000	309,940	-70,060	-18.4%
High Energy Physics	795,000	795,000	825,000	672,700	-152,300	-18.5%
Nuclear Physics	617,100	617,100	622,000	502,700	-119,300	-19.2%
Workforce Development for Teachers and Scientists	19,500	19,500	19,500	14,000	-5,500	-28.2%
Science Laboratories Infrastructure	113,600	113,600	130,000	76,200	-53,800	-41.4%
Safeguards and Security	103,000	103,000	103,000	103,000	...	...
Program Direction	185,000	185,000	182,000	168,516	-13,484	-7.4%
Subtotal, Science	5,350,200	5,350,200	5,392,000	4,472,516	-919,484	-17.1%
Rescission of Prior Year Balances	-3,200	-3,200	-239	...	+239	-100.0%
Total, Science Appropriation	**5,347,000**	**5,347,000**	**5,391,761**	**4,472,516**	**-919,245**	**-17.0%**

[a]The FY 2016 Enacted column printed in the FY 2018 Congressional Budget Justification (President's Request) includes SBIR/STTR funding in the program lines

and reflects programmatic updates through the end of the fiscal year.

[b]This column provides the Annualized CR amount (CR through April 28, 2017; P.L. 114-254). It is calculated by reducing the FY 2016 Enacted by 0.1901%

ASCR FY 2018 President's Request
(in thousands)

	FY 2016	FY 2017	FY 2018		
	Approp	Approp.	President's Request	House Mark	Senate Mark
Mathematical, Computational, and Computer Sciences Research					
Applied Mathematics	49,229	29,229	30,104	34,104	34,104
Computer Science	56,848	29,296	29,296	32,608	32,608
Computational Partnerships (SciDAC)	47,918	32,596	41,268	45,268	46,395
Next Generation Networking for Science	19,000	16,000	...	...	...
SBIR/STTR	6,181	10,271	11,261	4,242	4,285
Total, Mathematical, Computational, and Computer Sciences Research	179,995	117,392	111,929	116,222	117,392
High Performance Computing and Network Facilities					
High Performance Production Computing (NERSC)	86,000	92,145	80,000	92,000	94,000
Leadership Computing Facilities					
Leadership Computing Facility at ANL (ALCF)	77,000	80,000	100,000	100,000	100,000
Exascale	*(...)*	*(...)*	*(100,000)*	*(100,000)*	*(100,000)*
Leadership Computing Facility at ORNL (OLCF)	104,317	110,000	149,321	112,000	150,000
Exascale	*(...)*	*(...)*	*(50,000)*	*(...)*	*(50,000)*
Total, Leadership Computing Facilities	181,317	190,000	249,321	212,000	250,000
Research and Evaluation Prototypes	121,471	25,301	24,452	24,452	24,559
CSGF	*(10,000)*	*(10,000)*	*(10,000)*	*(10,000)*	*(10,000)*
High Performance Network Facilities and Testbeds (ESnet)	38,000	45,000	45,000	65,000	79,000
SBIR/STTR	15,036	13,162	14,728	14,526	14,049
Total, High Performance Computing and Network Facilities	441,824	365,608	413,501	407,978	461,608
Exascale Computing					
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	157,894	164,000	196,580	170,000	184,000
Total, Advanced Scientific Computing Research	**621,000**	**647,000**	**722,010**	**694,200**	**763,000**
Computational Sciences Workforce Programs, with WDTS (non-add)	*(10,000)*	*(10,000)*	*(10,000)*	*(10,000)*	*(10,000)*
Exascale Computing Crosscut (non-add)	*(157,894)*	*(164,000)*	*(346,580)*	*(282,000)*	*(334,000)*

Components of the Exascale Program

- **Exascale Computing Initiative (ECI)**
  - The ECI was initiated in FY 2016 to support research, development and computer system procurements to deliver an exascale (10^{18} ops/sec) computing capability by the early to mid-2020s.
  - It is a partnership between SC and NNSA, addressing science and national security missions.
  - **In the FY2018 President's Budget request, ECI includes the SC/ASCR and NNSA/ASC facility investments in site preparations and non-recurring engineering activities needed for delivery of early to mid-2020s exascale systems.**

- **Exascale Computing Project (ECP)**
  - Beginning in FY 2017, the ASCR ECI funding was transitioned to the DOE project (ECP), which is managed according to the principles of DOE Order 413.3B.
  - The ECP subprogram in ASCR (SC-ECP) includes only support for research and development activities in applications, and in partnership with NNSA, investments in software and hardware technology and co-design required for the design of capable exascale computers .
  - The NNSA/ASC Advanced Technology Development and Mitigation (ATDM) program supports the development of applications and, in collaboration with SC/ASCR, investments in software and hardware technology and co-design required for the design of exascale capable computers**.**

Exascale FY2017 → FY2018

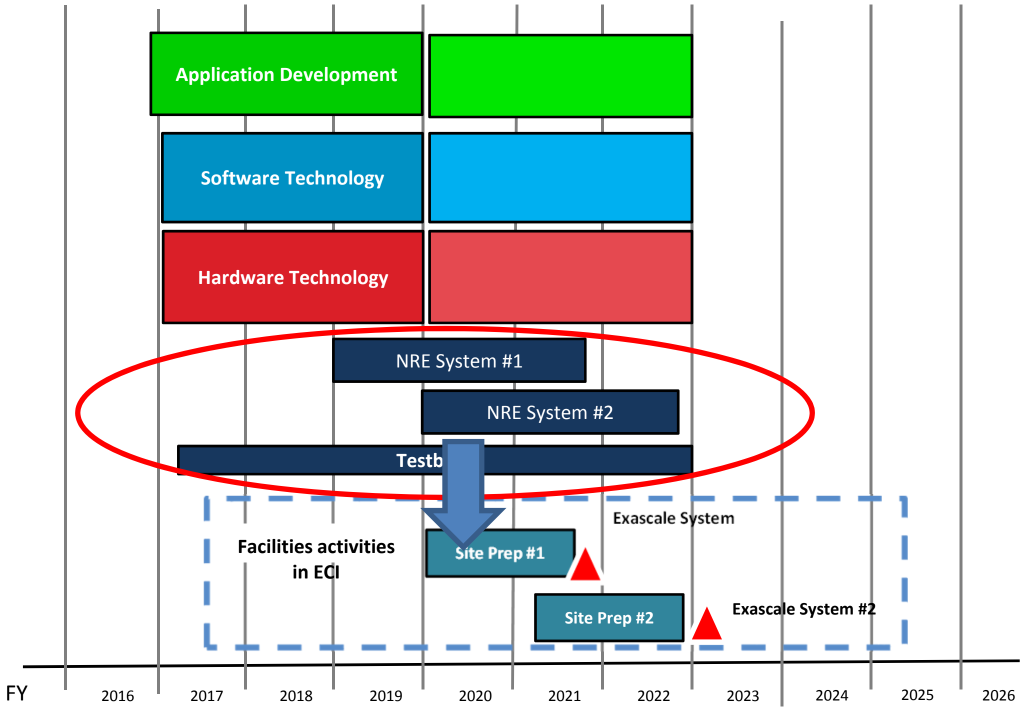

Moves to ECI through facility Investments

- **Engagement between facilities and ECP is critical to--**
  - Articulate common principles and shared values among the DOE Computing Facilities as well as the ECP Leadership to set context for specific laboratory-project agreements. (Includes NERSC and the LCFs)
  - Establish integrated teams between ECP and Facilities to align objectives and share milestones as necessary
  - Perform periodic gap analyses to identify areas of mutual interest, fostering collaboration, not duplication of efforts

ALCF 2021 EXASCALE SUPERCOMPUTER – A21

Intel/Cray Aurora supercomputer planned for 2018 shifted to 2021

- Scaled up from **180 PF to over 1000 PF**

Support for three “pillars”

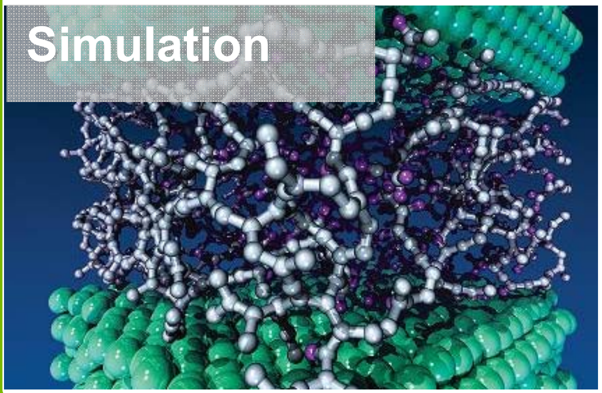

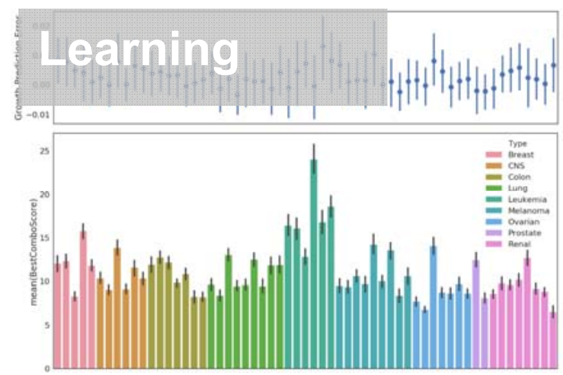

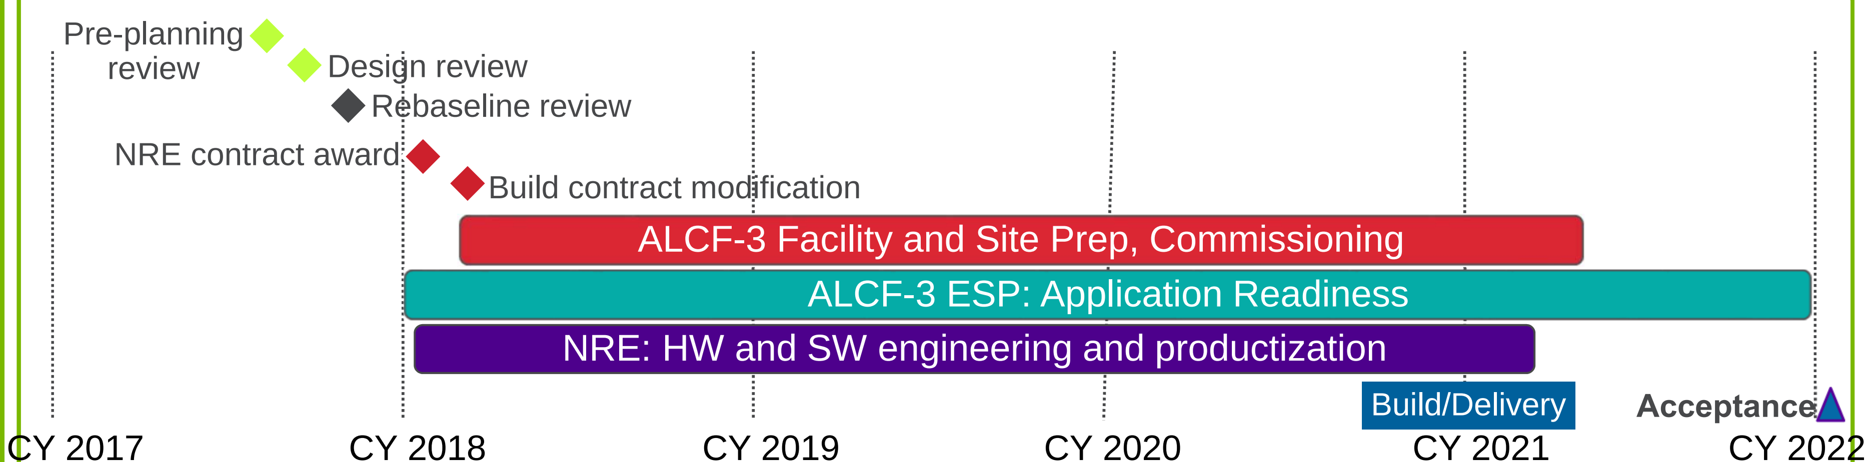

ALCF-3 Rebaseline Update

Design Review 9/20-21: Charge questions and results

1. Will this new system design, as described today, meet the requirements as stated and provide a productive system for science? **Yes, with 4 recs.**
2. Is the project on track to complete a draft Build SOW that clearly describes subcontractor and ALCF obligations for architecture, performance targets, user environment, system support, integration, and acceptance? **Yes, with 1 rec.**
3. Do the planned ALCF-3 facility enhancements clearly identify and support the power, space, and cooling requirements of the proposed system? **Yes, with 2 recs.**
4. Have the appropriate technical risks been defined? Are the risk mitigation strategies proposed appropriate? **Yes, with 3 recs.**

- **"The system as presented is exciting with many novel technology choices that can change the way computing is done. The committee supports the bold strategy and innovation, which is required to meet the targets of exascale computing. The committee sees a credible path to success."**
- **"The hardware choices/design within the node is extremely well thought through. Early projections suggest that the system will support a broad workload."**

Rebaseline Independent Project Review scheduled for 11/7-9

House Energy and Water Development Report HR 3266

- The Advanced Scientific Computing Research program develops and hosts some of the world's fastest computing and network capabilities to enable science and energy modeling, simulation, and research. **Exascale Computing Project.—The recommendation includes $170,000,000 for exascale activities.** High Performance Computing and Network Facilities.—In addition to the long-term exascale initiative, the Committee supports continued upgrade and operation of the Leadership Computing Facilities at Argonne and Oak Ridge national laboratories and of the High Performance Production Computing capabilities at Lawrence Berkeley National Laboratory. **The recommendation includes $100,000,000 for the Argonne Leadership Computing Facility, $112,000,000 for the Oak Ridge Leadership Computing Facility, and $92,000,000 for the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory. Within available funds, the recommendation includes $65,000,000 to support necessary infrastructure upgrades and operations for ESnet.** The Committee is concerned that the deployment plan for an exascale machine has undergone major changes without an appropriately defined cost and performance baseline. *The Department is directed to provide to the Committees on Appropriations of both Houses of Congress not later than 90 days after the enactment of this Act an update to the exascale plan that includes a detailed cost and performance baseline, taking into account flat and slightly increasing funding assumptions, for the technological challenges remaining to be solved to deliver an exascale machine.*

Senate Energy and Water Development Report (115-132)

- The Committee recommends $763,000,000 for Advanced Scientific Computing Research. The Committee is supportive of the plan to accelerate delivery of at least one exascale-capable system in 2021, reasserting U.S. leadership in this critical area. The Committee **recommends $184,000,000 for the Exascale Computing Project.** In addition, the Committee **recommends $150,000,000 for the Oak Ridge Leadership Computing Facility, $100,000,000 for the Argonne Leadership Computing Facility, $94,000,000 for the National Energy Research Scientific Computing Center, and $79,000,000 for ESnet.** Further, the Committee **recommends $10,000,000 for the Computational Sciences Graduate Program.** The Committee **recommends $24,559,000 for Research and Evaluation Prototypes.** The Committee recommends **not less than $117,392,000 for Mathematical, Computational, and Computer Sciences Research**.

Staffing Changes

THE OFFICE OF

ADVANCED SCIENTIFIC COMPUTING RESEARCH

Functional Organization Chart

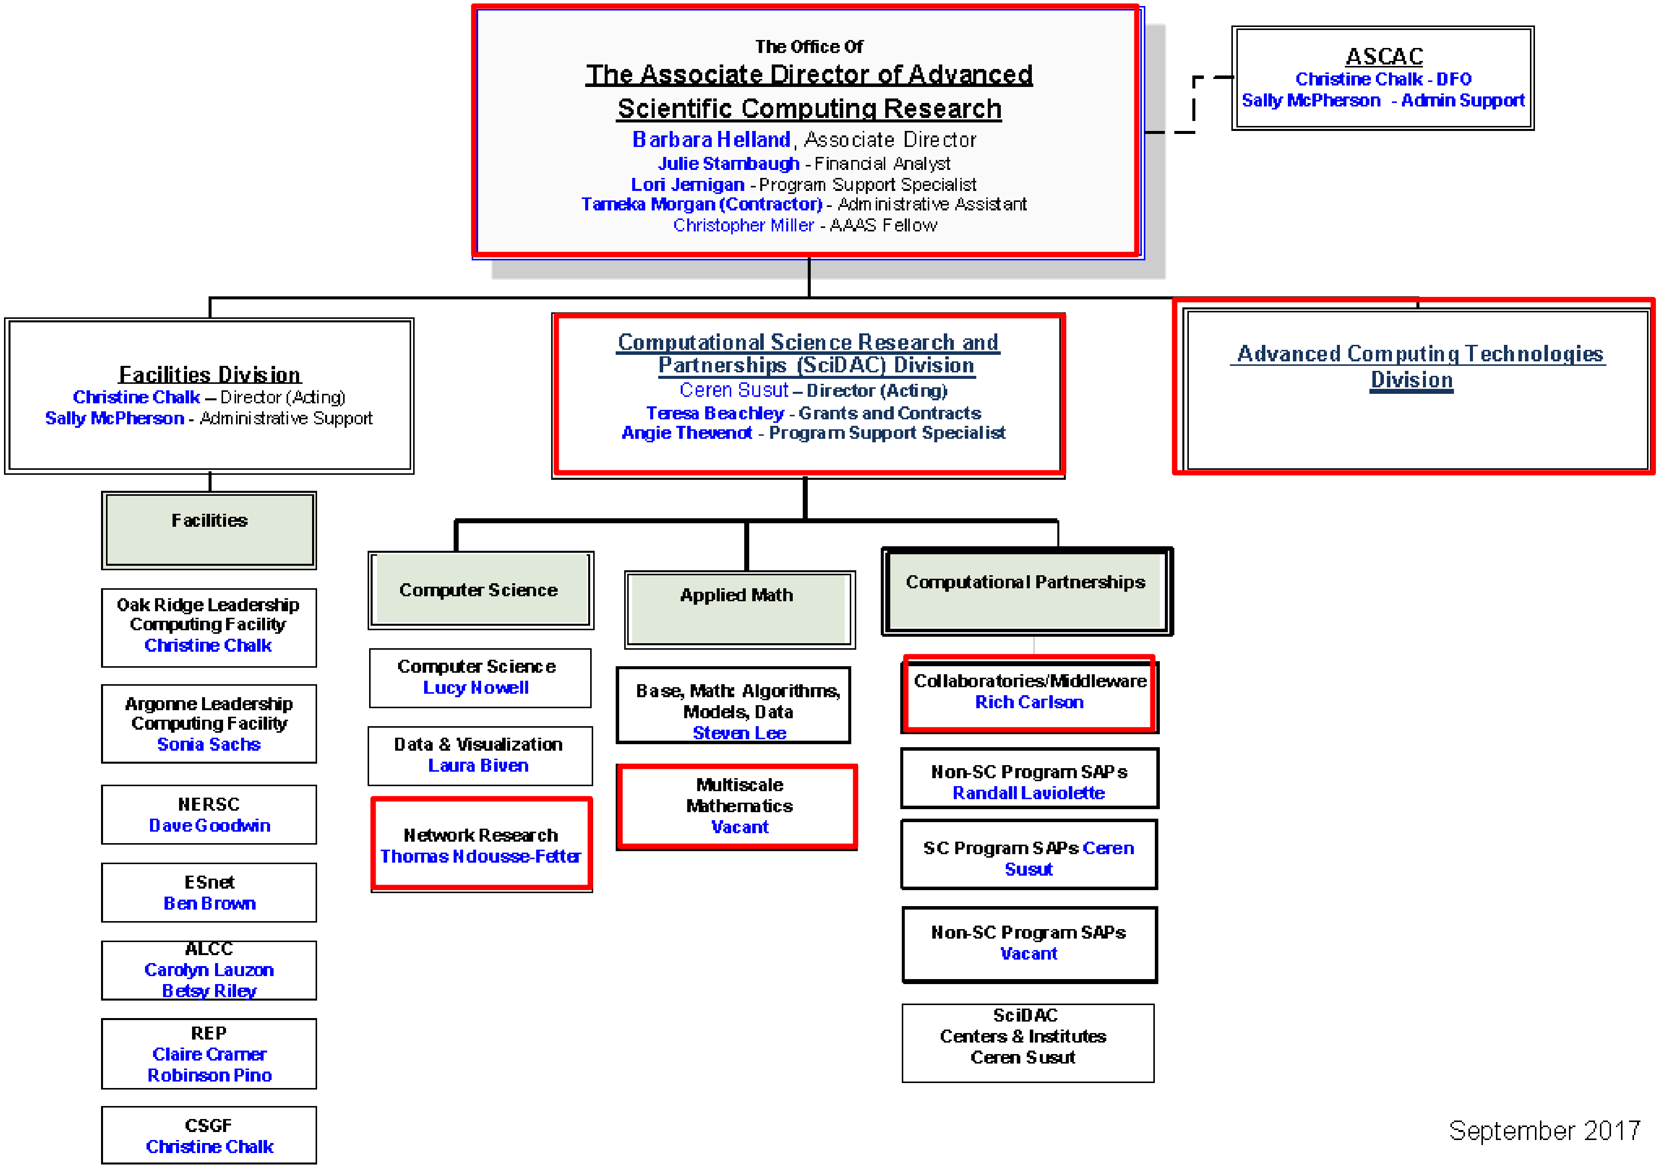

September 2017

FY2017: ASCR Year in Review

Awards – FY17 ASCR Early Career Research Program

Name	Institution	Title	Program	Topic Area
Lin Lin **PhD in 2011**	University of California, Berkeley	Green's function methods for multiphysics simulations	Applied Mathematics	Multiscale Mathematics
James Ostrowski **PhD in 2009**	University of Tennessee, Knoxville	Symmetric convex sets: Theory, algorithms, and application	Applied Mathematics	Algorithms, Solvers, and Optimization
Siqian Shen **PhD in 2011**	University of Michigan	Extreme-scale stochastic optimization and simulation via learning-enhanced decomposition & parallelization	Applied Mathematics	Algorithms, Solvers, and Optimization
Timothy Wildey **PhD in 2007**	Sandia National Laboratories	Enabling beyond forward simulation for predictive multiscale modeling	Applied Mathematics	Multiscale Mathematics
Tom Peterka **PhD in 2007**	Argonne National Laboratory	A continuous model of discrete scientific data	Computer Science	Resiliency, Data, and Visualization
Tiark Rompf **PhD in 2012**	Purdue University	Program generators for exascale and beyond	Computer Science	Extreme Scale - Performance portability
Miriam Kiran **PhD in 2010**	Lawrence Berkeley National Laboratory	Large-scale deep learning for intelligent networks	Computer Science	Advanced Network Architectures

Pavel Bochev (SNL) wins national award in Computational Fluid Dynamics

Dr. Pavel Bochev (Sandia National Laboratories) has been awarded the Thomas J.R. Hughes Medal by the U.S. Association for Computational Mechanics.

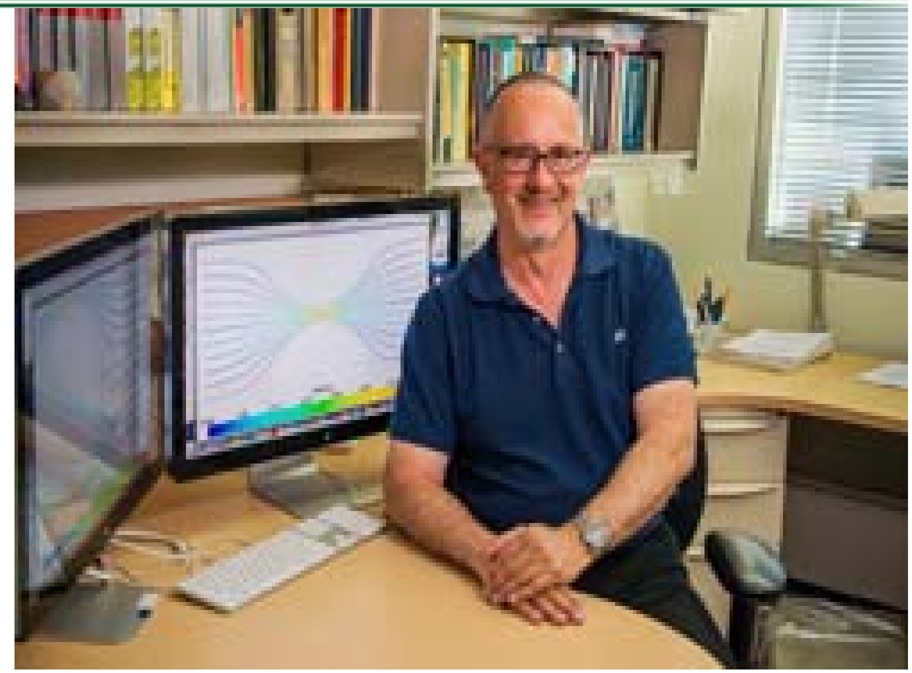

Thomas J.R. Hughes Medal

- Awarded every two years
- Contributions in the form of important research results that significantly advance the understanding of theories and methods impacting CFD.
- Industrial applications and engineering analyses that advance CFD shall also represent accomplishments worthy of recognition.

Citation: "Foundational contributions to numerical partial differential equations, especially advances in the development and analysis of new stabilized and compatible finite element methods, and software design for advanced discretizations."

ASCR Applied Mathematics Principal Investigators Meeting September 11-12, 2017 at Rockville Hilton

PI Meeting Website: http://www.orau.gov/ascr-appliedmath-pi2017

- **Keynote speakers: Bruce Hendrickson (LLNL), Steve Binkley (DOE/SC)**
- **~130 attendees: PIs & teams, DOE HQ, lab management, stakeholders**
- **102 posters from the PIs & their Applied Math research projects**
- **59 whitepapers on emerging applied math & scientific computing trends**
- **8 brainstorming discussions & report-outs (whitepaper topics)**
- **Lunchtime discussion: sharing best practices in research & training**

Whitepaper topics

1. **Toward predictive simulation, optimization, & design for complex systems**
2. **Convergence of data- & model-driven research**
3. **Sustaining applied mathematics workforce & products**
4. **Applied mathematics for future computing directions**

PI Meeting Feedback: Format, ideas
Next Annual PI Meeting: Summer 2018

Co-Chairs: Jeff Hittinger (LLNL) and Lois Curfman McInnes (ANL) **Committee:** Abani Patra (U Buffalo), Nathan Baker, Miranda Holmes-Cerfon, Barney Maccabe, Esmond Ng, Michael Parks, Pieter Swart, Karen Willcox. **ASCR Program Manager:** Steven Lee

Cori is in production for Office of Science

NERSC

Cray XC40 system with 9,688 Intel Xeon Phi "Knights Landing" manycore nodes

68 cores,16 GB high BW memory on chip

96 GB DDR memory / node

Supports the entire Office of Science research community

Transitioning the Office of Science workload to energy efficient architectures

CD-4 approval on Sep. 19, 2017

Entered production on July 1, 2017

Data Intensive Science Support

10 Haswell processor cabinets (Phase 1)

2,388 nodes, 128 GB DDR

NVRAM Burst Buffer 1.5 PB, 1.5 TB/sec

30 PB of disk, >700 GB/sec I/O bandwidth

Integrated with Cori Haswell nodes on Aries network for data / simulation / analysis on one system

Software Defined Networking for Enhanced external network connectivity

Deep Learning on Cori KNL

NERSC is actively exploring Deep Learning for Science

- Collaborating with leading vendors to optimize and deploy stack
- Collaborating with leading research institutions to develop methods
- Drive real science use cases

Deep Learning at 15 PF on NERSC Cori (Cray + Intel KNL)

- Trained in 10s of minutes on 10 terabyte datasets, millions of Images
- 9600 nodes, optimized on KNL with IntelCaffe and MKL (NERSC / Intel collaboration)
- Synch + Asynch parameter update strategy for multi-node scaling (NERSC / Stanford)

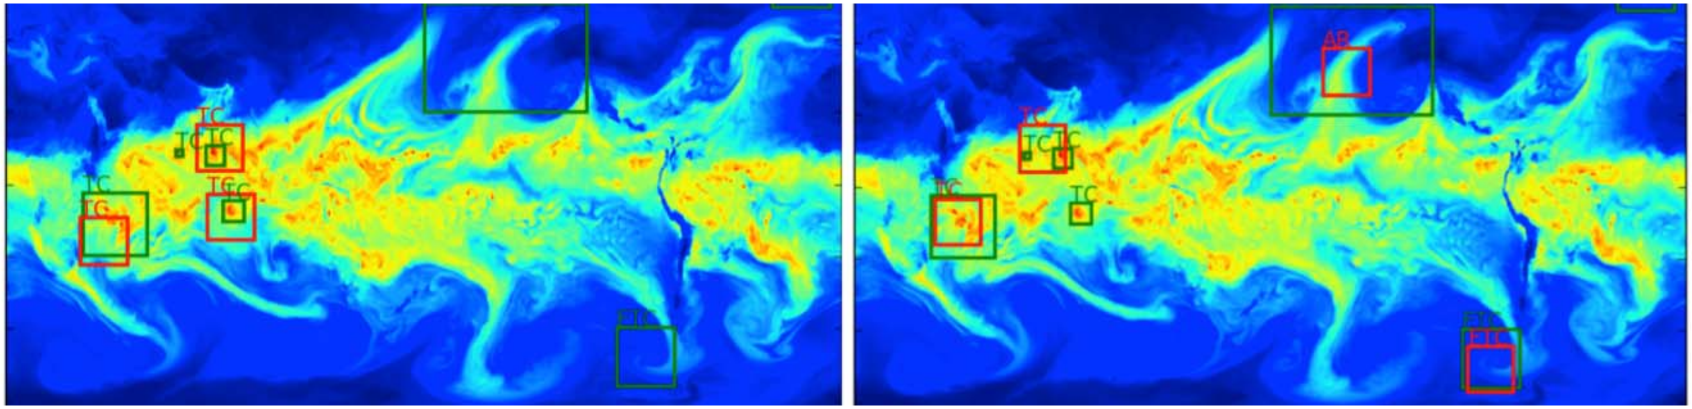

Identified extreme climate events using supervised (left) and semisupervised (right) deep learning. Green = ground truth, Red = predictions (confidence > 0.8). [NIPS 2017]

summit Begins at OLCF

Start of Installation

Installation Complete

Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Nodes	18,688	~4,600
Node performance	1.4 TF	> 40 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM
NV memory per Node	0	1600 GB
System Interconnect	Gemini (6.4 GB/s)	Dual Rail EDR-IB (23 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™

Machine Learning for Science at OLCF

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Neutrino Detection

Fermilab

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Analytics of Deep Learning Hyperparameter search running on 15000 Nodes of Titan.

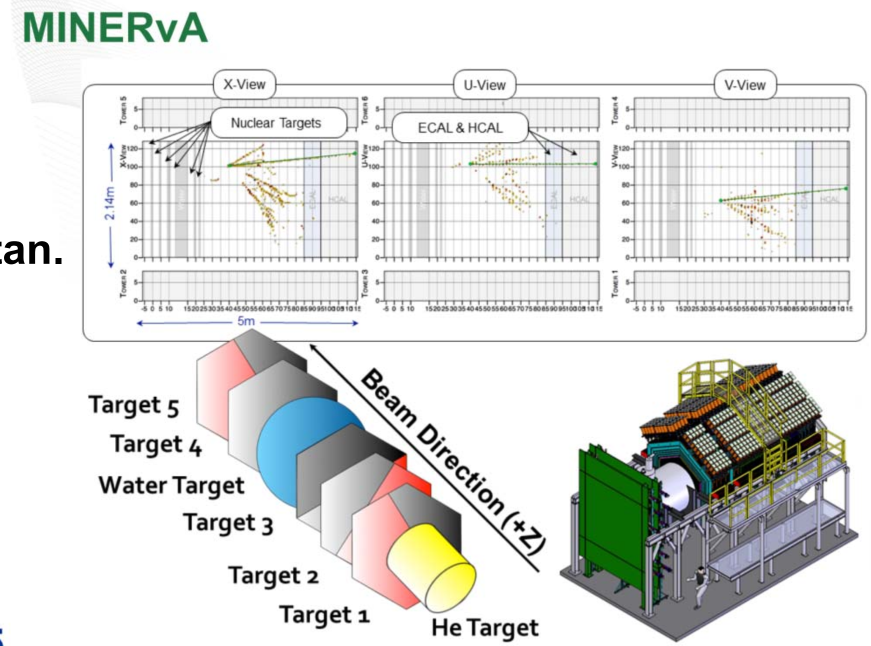

Fusion Experiment

Deep Learning executing on ~6000 GPUs with TensorFlow+MPI.

Tensorflow+MPI (using Singularity containers), CUDA7.5,

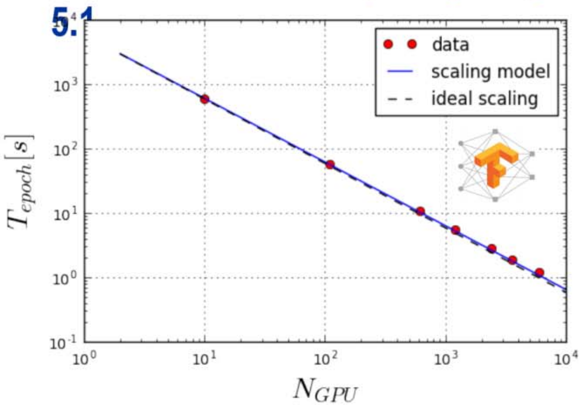

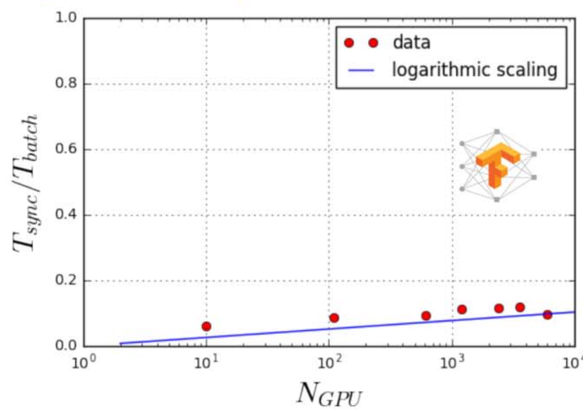

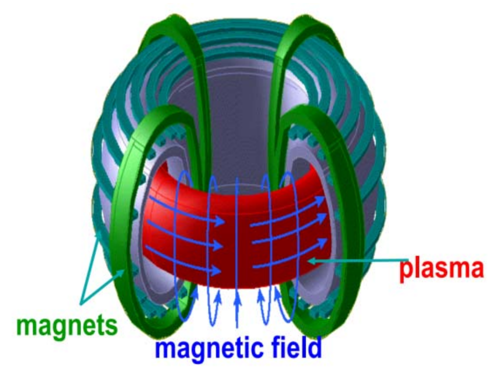

Theta – A New Architecture for Simulation and Data Science

Features Intel processors and interconnect technology, a new memory architecture, and a Lustre-based parallel filesystem – all integrated by Cray's HPC software stack

Already supporting:

- 10 ASCR Leadership Computing Challenge (ALCC) projects (419 million core-hours)
- 4 ALCF Data Science Program (ADSP) projects (100 million core-hours)
- 12 Theta Early Science Program (ESP) projects (436 million core hours)
- Several Director's Discretionary projects

New capabilities include:

- MCDRAM, the high bandwidth significantly alleviates the memory bandwidth bottleneck
- AVX512, 1st instance of the new instruction set that will be in a number of future Intel products
- 64 cores, demonstrates scalability to significantly increased core counts on a socket

Theta

Cray XC40
– 3,624 nodes
– 231,935 cores
– 56 TB MCDRAM
– 679 TB DDR4
– 453 TB SSD
– Peak flop rate: 9.65 PF

- **Iota – T&D**
- **Cray XC40**
  – 44 nodes
  – 2,816 cores
- **Storage**
- **– ~18 PB GPFS/Lustre filesystem**
- **– ~240 GB/s**

Expanding Leadership Computing Reach

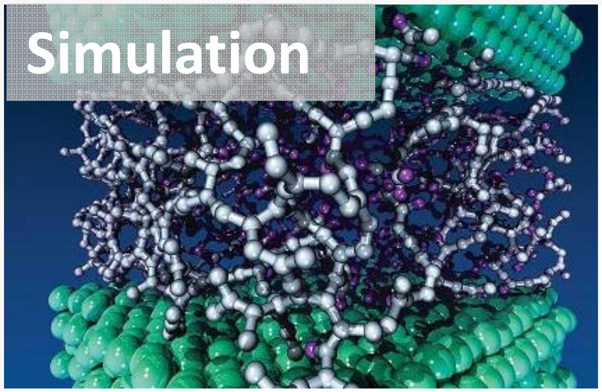

Reactive Mesoscale Simulations of Tribological Interfaces

PI: S. Sankaranarayanan, ANL
Insight to the complex processes that make oils, coatings, electrodes, and other electrochemical interfaces effective. Using Mira, this team discovered a self-healing, anti-wear coating that drastically reduces friction. Their findings are being used to virtually test other potential self-regenerating catalysts.

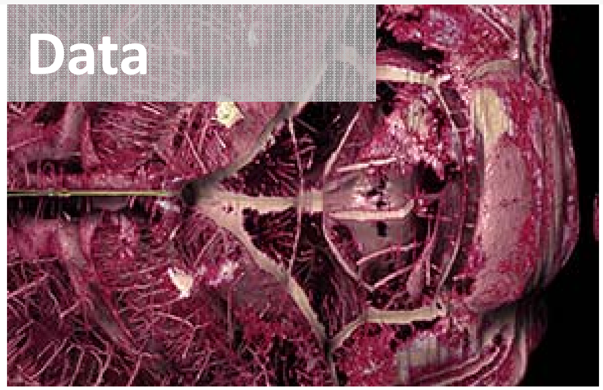

Large-Scale Computing on the Connectomes of the Brain

PI: D. Gursoy, ANL
3D reconstructions of high-resolution imaging will provide a clearer understanding of how even the smallest changes to the brain play a role in the onset and evolution of neurological diseases, such as Alzheimer's and autism, and perhaps lead to improved treatments or even a cure.

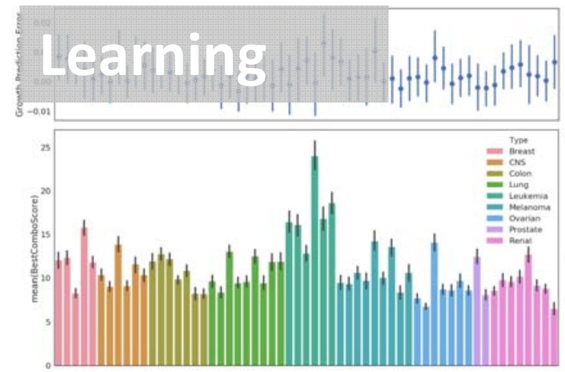

CANcer Distributed Learning Environment (CANDLE)

PI: R. Stevens, ANL
CANDLE is tackling the hardest deep learning problems in cancer research. Its first architecture release for large-scale model hyperparameter exploration uses representative problems--coded as deep learning problems--at the core of the predictive oncology challenge. Future data parallelism work will allow the training of a single model across several nodes.

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

ALCF DATA SCIENCE PROGRAM

ECP EXASCALE COMPUTING PROJECT

U.S. DEPARTMENT OF ENERGY | Office of Science

SciDAC-4 Overview

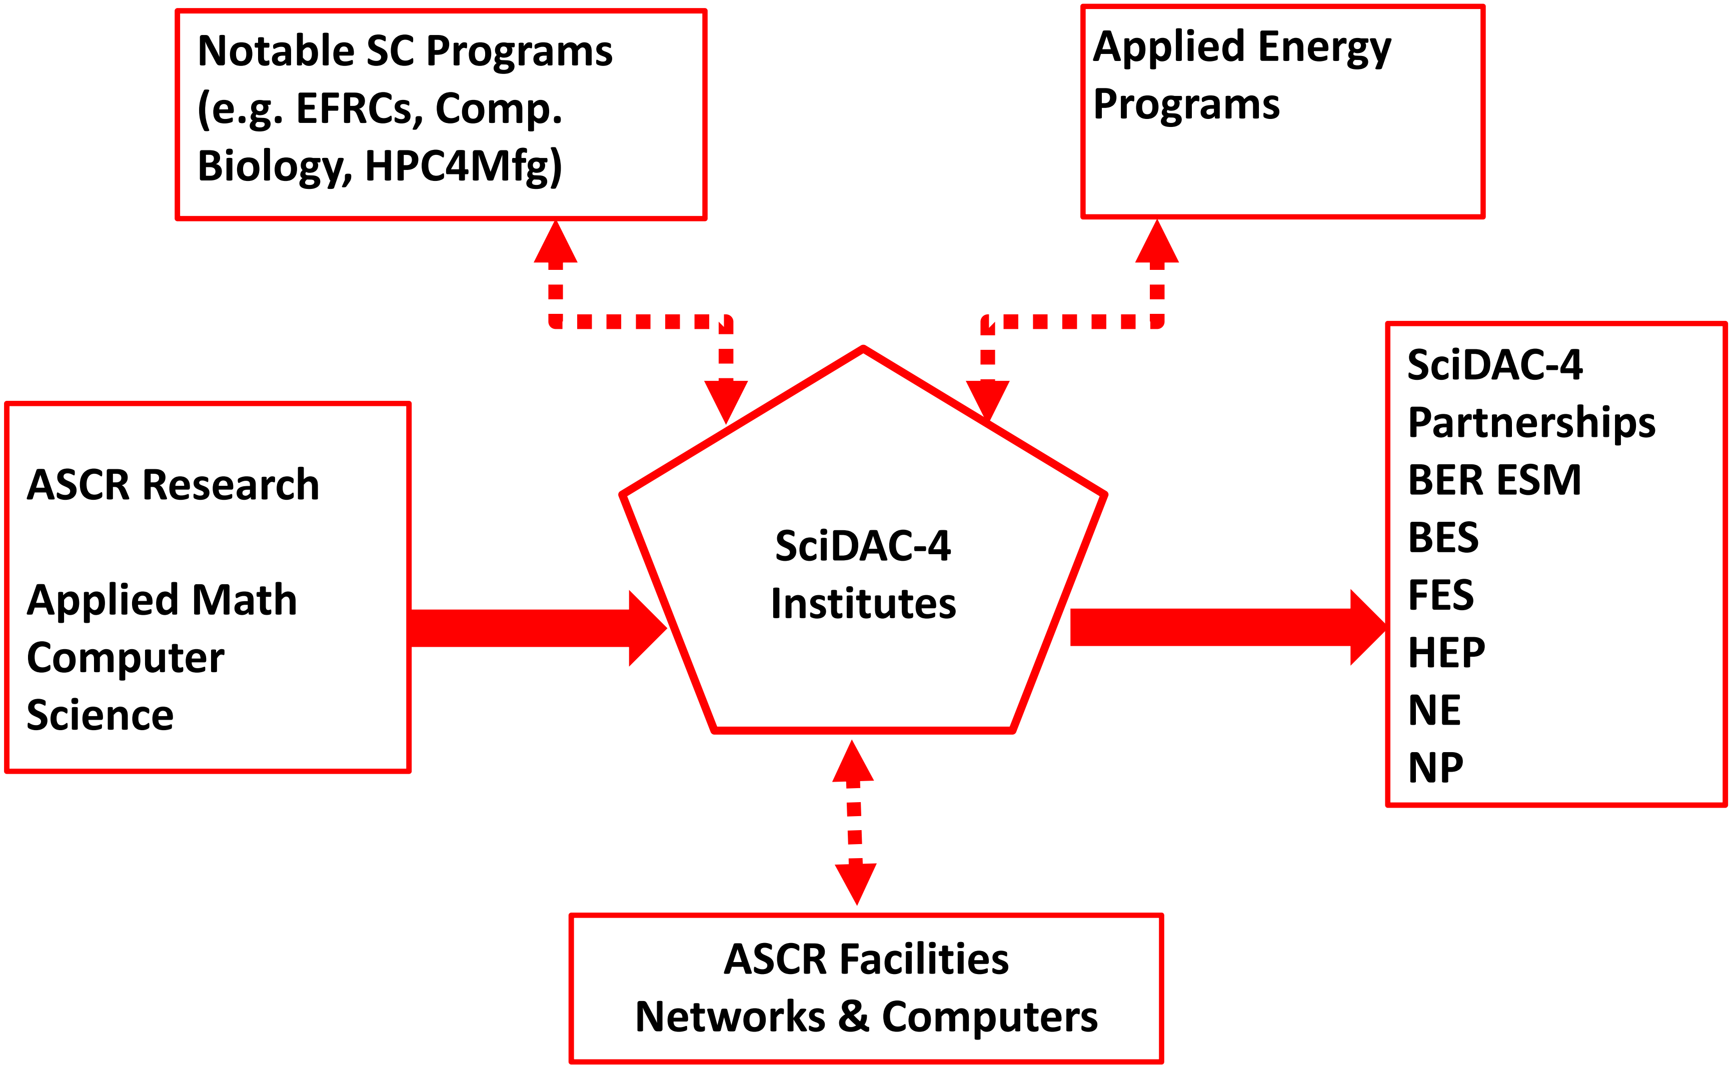

SciDAC-4 Partnerships – Review Process

Partner (PM)	Received	Closed (open)	Mail-in + Panel	Reviews (reviewers)	Recommended for co-funding
BES Davenport	19 @ 4-yr	21 Jun. (10 May)	20 July	103 (49)	3 @ 4-yr
BER Koch	6 @ 5-yr 28 @ 2½-yr	15 Mar. (4 Nov.)	3,4 May	32 (15) 134 (35)	2 @ 5-yr 4 @ 2½-yr
FES Mandrekas	17 @ 5-yr	21 Feb. (16 Nov.)	19-21 April	126 (47)	7 @ 5-yr
HEP Chatterjee	14 @ 5-yr	27 Feb. (4 Nov.)	7 April	66 (29)	3 @ 5-yr 2 @ 2-yr
NP Barnes	7 @ 5-yr	24 Feb. (10 Nov.)	3 April	63 (12)	3 @ 5-yr
NE Funk	5 @ 5-yr	5 Apr. (16 Dec.)	21 July (no panel)	36 (10)	1 @ 5-yr
TOTAL	96			560 (197)	25

Representative SciDAC-4 Partnership Topics

- **BES:** catalysis, quantum materials
- **BER:** atmosphere dynamics, component-coupling and performance; hi-res. sea ice; ice-land-sea interaction; sensor placement; atmospheric physics convergence
- **FES:** tokamak WDM, energetic particles in burning plasmas, core & edge dynamics, injection & RF coupling, plasma-materials interactions, disruptions
- **HEP:** accelerator physics, cosmology, event analysis & simulation
- **NP:** LQCD, supernovae (II) & mergers, nuclear structure of elements
- **NE:** radiation effects in materials & fuels

- **Equations**
  - CFD, solid mechanics, gyrokinetic, Newton's, Maxwell's, Schrödinger's, kinetics, QCD, ill-posed inversion
- **Applied Math**
  - numerical optimization, linear algebra, PDE, ODE, SDE (solvers & theory), MC, FFT/wavelets, UQ, statistics, quadrature
  - code coupling for multi-scale/-physics
- **Computer Science**
  - code coupling, memory management, in-situ data movement, auto-tuning, resilience
  - portability (between LCF swim lanes)

ASCR support for SciDAC-4 Partnerships is widely distributed

Partner	#	ANL	BNL	FNAL	INL	LBNL	LLNL	LANL	ORNL	PNNL	PPPL	SNL	SLAC	TJNAF
BES	**3**					☒			☒				☒	
BER	**6**	☒				☒	☒	☒	☒	☒		☒		
FES	**7**	☒				☒	☒	☒	☒	☒	☒	☒		
HEP	**5**	☒	☒	☒		☒		☒					☒	
NP	**3**	☒	☒			☒		☒	☒	☒				☒
NE	**1**	☒			☒				☒			☒		

- 22 collaborations are Laboratory-led
  - Some Laboratory budgets provide subcontracts for additional Universities and companies
- ASCR supports 62 PIs at 20 institutions (in 25 collaborations)
  - 53 PIs at 13 DOE (SC, NNSA, NE) Laboratories w/ support from ASCR
  - 9 PIs (at UMD, MIT, MSU, RPI, UTn-K, UTx-A & GA) receive direct support from ASCR grants
  - FY17: $19,352,000 for 76 awards processed

SciDAC-4 Institutes

Purpose: To provide intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to the SC and DOE

Emphasis:

Application-aware: Support SciDAC-4 Partnerships and equipped to interact with other communities (e.g. EFRC's, HPC4Mfg, Computational Biosciences);

Architecture-aware: Existing and emerging DOE HPC systems such as those existing and planned for at OLCF, ALCF and NERSC over the next 3 years. Application-portability and software engineering best practices;

Institutes-aware: Processes for effectively leveraging basic research advances from applied math and computer science & expertise and results from the other Institute.

Management structure: Built-in flexibility, no duplication within the Institute.

Timeline & Proposals: A DOE National Laboratory Announcement was issued in June 2017. 3 proposals were received in July 2017.

2 Institutes @ Total $12M/year:

RAPIDS-Resource and Application Productivity through Computation, Information, and Data Science. Director: Rob Ross, ANL, Deputy Director: Lenny Oliker, LBNL;

FASTMath-Frameworks, Algorithms and Scalable Technologies for Mathematics. Director: Lori Diachin, LLNL, Deputy Director: Esmond Ng, LBNL.

Awards to 7 Labs; support to 11 universities and 2 companies through subcontracts.

The RAPIDS Institute

- **Resource and Application Productivity through Computation, Information, and Data Science**

Objective: Solve computer science and data challenges for Office of Science application teams to achieve science breakthroughs on DOE platforms.

- **Application Engagement**
  - *Tiger Teams* engage experts in multiple technology areas to work with science teams and codes
  - Outreach activities connect with broader community
- **Technology Focus Areas**
  - **Data Understanding** – scalable methods, robust infrastructure, machine learning
  - **Platform Readiness** – hybrid programming, deep memory hierarchy, autotuning, correctness
  - **Scientific Data Management** – I/O libraries, coupling, knowledge management

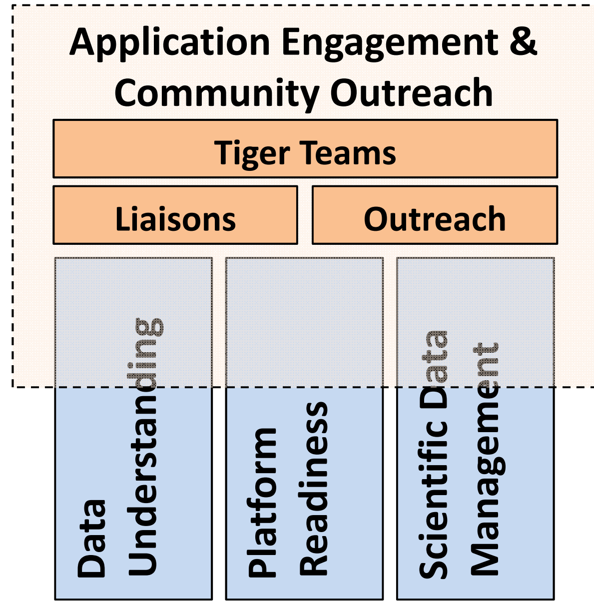

FASTMath brings leading edge computational mathematics technologies to the SciDAC Program

FASTMATH

Develop advanced numerical techniques for DOE applications

Deploy high-performance software on DOE supercomputers

FASTMath Objective: Reduce the barriers facing application scientists

Demonstrate basic research technologies from applied mathematics

Engage and support of the computational science community

8 Core Technology Areas

- Structured Mesh Spatial Discretization
- Unstructured Mesh Spatial Discretization
- Time Integrators
- Solution of Linear Systems
- Solution of Eigenvalue Problems
- Numerical Optimization
- Uncertainty Quantification
- Data Analytics

For more information contact: Lori Diachin, LLNL diachin2@llnl.gov

Our team comprises over 50 researchers from 5 national laboratories and 5 universities

BERKELEY LAB, Argonne National Laboratory, LLNL, MIT, Sandia National Laboratories, Oak Ridge National Laboratory, USC, SMU, CU Boulder, Rensselaer

New Understanding of Chemistry and Dynamics in Li-ion Batteries

ASCR-BES SciDAC-3 Highlight

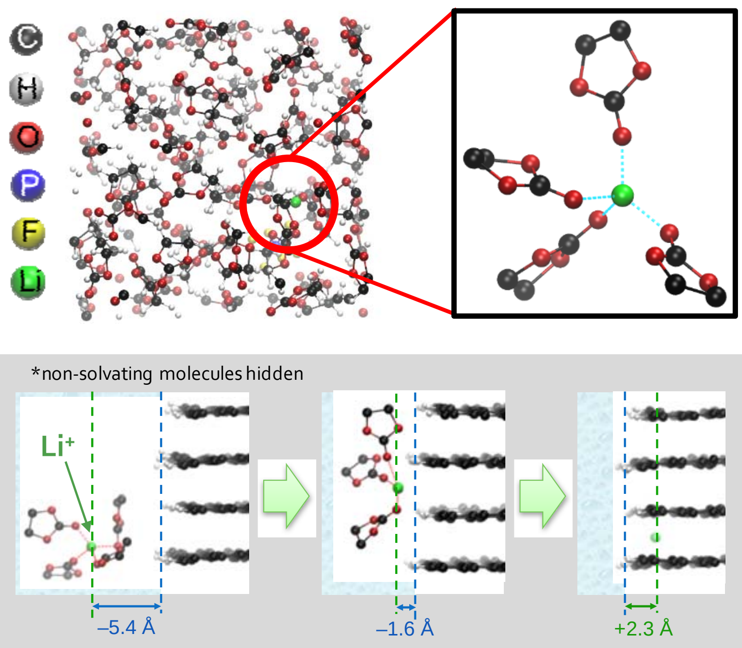

Top: Strong tetrahedral bonding of Li^+ by compact cyclic ethylene carbonate (EC) molecules reduces the ion mobility in EC compared to less-structured solvation in linear-chain ethyl-methyl carbonate (EMC) solvent. **Bottom:** During Li^+ migration from electrolyte to electrode, a series of desolvation steps occurs that cost energy.

Ong et al., J. Phys. Chem. B 119, 1535-1545, 2014.
Ong et al., J. Phys. Chem. C 121, 6589-6595, 2017.
Ong et al., LLNL-PROC-678868, 2015.
Pham et al., 2017, in preparation.

Scientific Achievement

We developed and employed new and existing methodologies to carry out quantum molecular dynamics (QMD) simulations of electrolytes and anode-electrolyte interfaces in Li-ion battery systems at unprecedented accuracy and scale, revealing both the chemistry and dynamics of solvation, diffusion, and intercalation

Significance and Impact

New insights into solvation, diffusion, and intercalation in Li-ion batteries suggest multiple avenues for fundamental advances in device performance, lifetime, and safety

Research Details

- QMD simulations included explicit models of all components: ions, counter-ions, solvent, and electrodes
- Exchange and correlation interactions beyond conventional local-density (LDA) and generalized-gradient (GGA) were taken into account to better model common anode materials
- Significant findings include:
  - Organic solvents forming weaker solvation shells yield increased ion mobility
  - Ethylene carbonate (EC) can stabilize dissociated ions, while ethyl methyl carbonate (EMC) prefers paired ions
  - Relative angle analysis reveals distinct ion dynamics at different time scales: from ballistic to caged to Brownian
  - Anode termination strongly affects intercalation barriers
  - Na^+ and K^+ ions show 2-3x higher mobility than Li^+
  - Incorporating applied voltage into simulations is necessary for quantitative predictions in device configurations

Up to 40% Performance Improvement from New Hybrid Load Balancing

PI: C.S. Chang, Fusion SciDAC Center for Edge Physics Simulation (EPSI)

ASCR−FES SciDAC-3 Highlight: Collaboration between Fusion-EPSI and SUPER

Objectives

- **Address performance degradation due to load imbalances in i) particle time-advance and ii) nonlinear collision calculation for XGC1 on DOE Leadership Computing Systems.**

Impact

- **Low overhead automatic adjustment of parallel decomposition improves computational performance robustly and with minimal user input.**

Accomplishment highlight

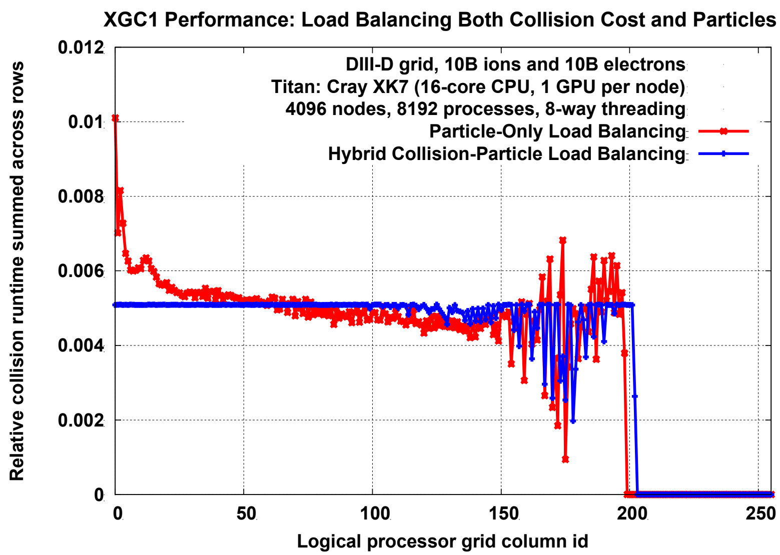

Example load imbalance in collision operator cost across columns of logical 2D processor grid, comparing load-balancing only particle distribution (red) with also load-balancing collision cost (blue). Cost is summed over rows of grid. Full model performance improvement is 30% for this example.

<u>Challenge</u>

- **Existing particle load balancing algorithm does not adequately equidistribute the collision cost in parallel decomposition.**
- **Both particle count and collision cost per grid cell distributions evolve with the simulation.**

<u>Solution</u>

- **Two level automated optimization strategy: (a) balance collision cost subject to constraint on particle load imbalance, (b) optimize XGC1 performance by varying constraint periodically, converging to the optimum if distributions are static and adapting to the changing distributions otherwise.**

<u>Result</u>

- **10%-40% improvement for production runs.**
- **Could be generalized to other similar codes.**

U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR Computational Partnerships
FY17 Renewal Projects

Scientific Workflow Analysis

4 Projects @ Total FY17 Funding $4.0M:

- **Integrated End-to-end Performance Prediction and Diagnosis for Extreme Scientific Workflows (IPPD).** Lead: PNNL (Nathan Tallent), Collaborators: BNL (Kerstin Kleese van Dam), UCSD (Ilkay Altintas), UCSC (Darrell Long)
- **Panorama 360: Performance Data Capture and Analysis for End-to-End Scientific Workflows.** Lead: USC (Ewa Deelman), Collaborators: ORNL (Jeff Vetter), LBNL (Mariam Kiran), RENCI (Anirban Mandal)
- **RAMSES: Robust Analytical Models for Science at Extreme Scales.** Lead: ANL (Ian Foster), Collaborator: ORNL (Rao Nageswara)
- **Nested Task-Parallel Workflows for Scientific Applications.** Lead: ANL (Tom Peterka)
- **X-Swap: Extreme-Scale Scientific Workflow Analysis and Prediction.** Lead: LBNL (Erich Strohmaier), Collaborator: SLAC (Amedeo Perazzo)

Mathematical Multifaceted Integrated Capability Centers LAB 17-1766

MMICCs Purpose: Basic research that address fundamental mathematical challenges within the DOE mission & from a perspective that requires new integrated efforts across multiple mathematical, statistical, and computational disciplines.

Emphasis: MMICC proposals must

- Advance multifaceted, integrated mathematics that appropriately spans novel formulations, discretizations, algorithm development, data analysis, uncertainty quantification, optimization, and other mathematical and statistical approaches
- Address mathematical problems with clear relevance and significant impact to DOE
- Actively engage in community building events to rapidly disseminate scientific advances and maintain clear channels of communication to the DOE user community

Timeline

May 5: DOE National Laboratory Announcement issued

June 5: Fourteen pre-proposals submitted

July 11: Fourteen full proposals submitted

July 31- August 1: Panel reviews conducted

Project awards: 1

MACSER: Multifaceted Mathematics for Rare, High Impact Events in Complex Energy and Environment Systems (Project Director: Mihai Anitescu, ANL)

ANL (lead), PNNL, LLNL, and 3 universities supported through Lab subcontracts.

ASCR Applied Mathematics FY17 Renewal Projects

10 Projects @ Total $6,600K/year:

Title	Lead PI	Institution
High-Order Methods for High-Performance Multiphysics Simulations	Paul Fischer	ANL
Efficient Error Estimation & Propagation in Complex ODE/DAE/PDE Simulations	Barry Smith	ANL
Simulation and Analysis of Reacting Flows	John Bell	LBNL
High-Resolution and Adaptive Numerical Algorithms for PDEs	Phil Colella	LBNL
Multilevel Iterative Temporal Integration and Time Parallelism for PDEs	Michael Minion	LBNL
Frontiers in Computation: New Methods for Fluids, Structures and Interfaces, Advanced Materials, and Stochastics	James Sethian	LBNL
High-Resolution Methods for Phase Space Problems in Complex Geometries	Milo Dorr & Phil Colella	LLNL LBNL
New Multigrid Advances for Highly Concurrent Architectures	Rob Falgout	LLNL
Formulation, Analysis and Computation of Heterogeneous Numerical Methods	Pavel Bochev	SNL
Enabling Multiphysics Plasma Simulations by the Development of Stable, Accurate, and Scalable Computational Formulations and Solution Methods	John Shadid & Luis Chacon	SNL LANL

Lab Base Math Portfolio sustains foundational research projects in PDEs, Multiphysics modeling & simulation, Time integration, Multigrid methods, & Computational mathematics.

High-performance multigrid-based algorithms provide foundation for adaptive parallel-in-time integration
Lead PI: Rob Falgout (LLNL)

Objectives

- Develop multigrid algorithms in space and time to address solver challenges anticipated in DOE science simulations on future architectures
  - O(N) algorithms are crucial
- Develop parallel implementations in the *hypre* and XBraid libraries
  - Readily available to DOE scientists

Impact

- Enable new science by dramatically speeding up linear system solve times in DOE codes
  - Fast linear solves are often essential
- Enable new science by creating parallelism in the time dimension
  - Up to 50x speedup on some problems (so far)

Moving Mesh

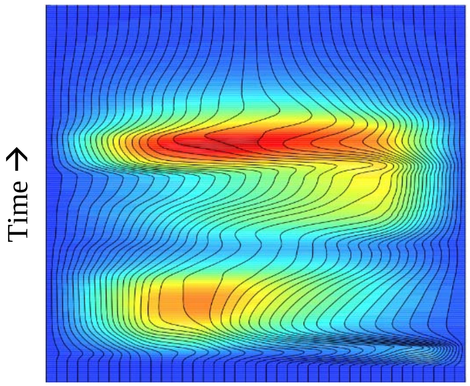

Navier-Stokes
7.5x speedup in Strand2D

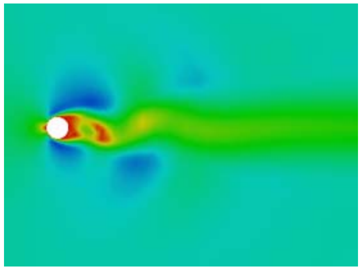

1D Inviscid Burgers

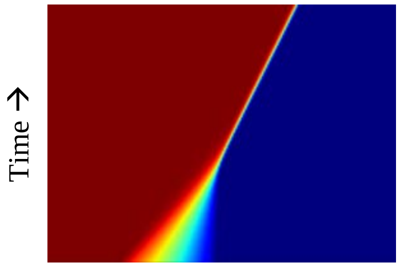

Highlights and Accomplishments

- Developed MGRIT non-intrusive parallel time integration algorithm and XBraid library
  - Builds on existing codes and technologies
  - Supports adaptivity in time, spatial coarsening, implicit/explicit/multi-step/multi-stage methods, moving meshes, low-storage options
- Demonstrated XBraid effectiveness and potential in a variety of codes and applications
  - Codes: Strand2D, Cart3D, LifeV, CHeart, GridDyn
  - Apps: Navier-Stokes, linear/nonlinear diffusion, power grid DAE sims, inviscid Burgers (shocks)
- Sharp predictive convergence theory for MGRIT
- New AMG algorithms and techniques decrease communication & improve parallel performance

ASCR Computer Science
FY17 Renewal Projects

Scientific Data Management, Analysis and Visualization

6 Projects @ Total $4,350K/year:

- **EOD-HDF5: Advancing HDF5 for Managing Experimental and Observational Data.** Lead: LBNL (Surendra Byna), Collaborator: The HDF Group (John Mainzer)
- **Sample-based, Perceptually- and Cognitively-driven Visual Analysis of Massive Scientific Data Using an Asynchronous Tasking Engine.** Lead: LANL (James Ahrens), Collaborators: University of Texas (Greg Abram), University of New Hampshire (Colin Ware)
- **Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery (SENSEI).** Lead: LBNL (Wes Bethel), Collaborators: ANL (Nicola Ferrier), ORNL (Matthew Wolf), Intelligent Light (Earl Duque), Kitware (Patrick O'Leary)
- **Nested Task-Parallel Workflows for Scientific Applications.** Lead: ANL (Tom Peterka)
- **Visual Analytics for Large Scale Scientific Ensemble Datasets.** Lead: LANL (Jonathan Woodring), Collaborators: ANL (Tom Peterka), Ohio State University (Han-Wei Shen)
- **Scalable Data-Computing Convergence and Scientific Knowledge Discovery.** Lead: LBNL (Wes Bethel)

Performance Evaluation of Large-scale I/O on Cori Burst Buffer
PI: S. Byna, Q. Koziol (LBNL)

BERKELEY LAB

Scientific Achievement

Tuning parallel I/O on burst buffers (BB) of upcoming supercomputer architectures is challenging because BB software is still evolving. Moreover, existing I/O software, such as MPI-IO and HDF5, have not been tuned for use on the BB. LBNL's ExaHDF5 project team participated in NERSC's BB Early User Program to study performance of large-scale parallel I/O in a plasma physics simulation code to identify bottlenecks and to optimize performance.

Significance and Impact

Identified that previously tuned plasma physics simulation code did not scale well using BB on Cori because an SSD-based BB performs differently than a disk-based Lustre file system. BB-specific optimizations avoid performance degradations and perform ~5X better than Lustre. Results from this study contributed to a paper that won the ***best* paper award** at the 2016 Cray Users Group meeting.

Research Details

– Optimized writing of data by a plasma physics simulation code using burst buffers on the NERSC Cori-Phase 1 system

– Devised a strategy for automatically selecting performant tuning parameters for I/O software libraries

Work was performed on Cori Phase I system at NERSC

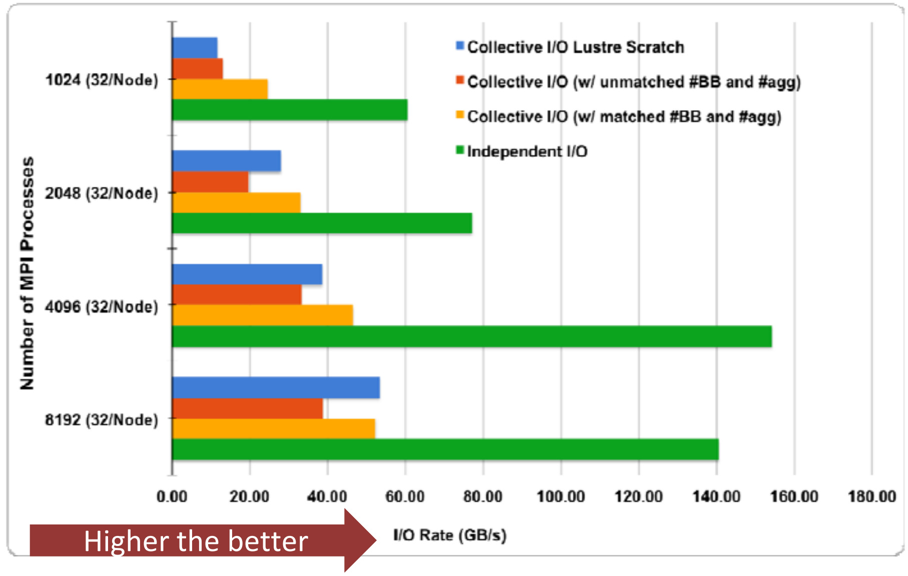

Performance improvements obtained by selecting appropriate tuning parameters for I/O libraries: This plot shows the first large-scale scientific benchmark to exercise parallel I/O on the Cori burst buffer. Our optimized I/O mini-app, extracted from the VPIC plasma physics space weather simulation, performs 2.5X to 5X better than Lustre on Cori. Our tuning also performs 4.5X better on burst buffers (compared to running the code with default parameters).

- W. Bhimji, D. Bard, et al., Accelerating Science with the NERSC Burst Buffer Early User Program, *CUG 2016*, ***Best Paper Winner,*** *https://sdm.lbl.gov/~sbyna/research/papers/201605_CUG_NERSC-BB-EUP.pdf*

Quantum Computing Applications for SC Grand Challenges

QIS Task Force identified SC-wide grand challenges that will potentially be transformed by quantum computing applications.

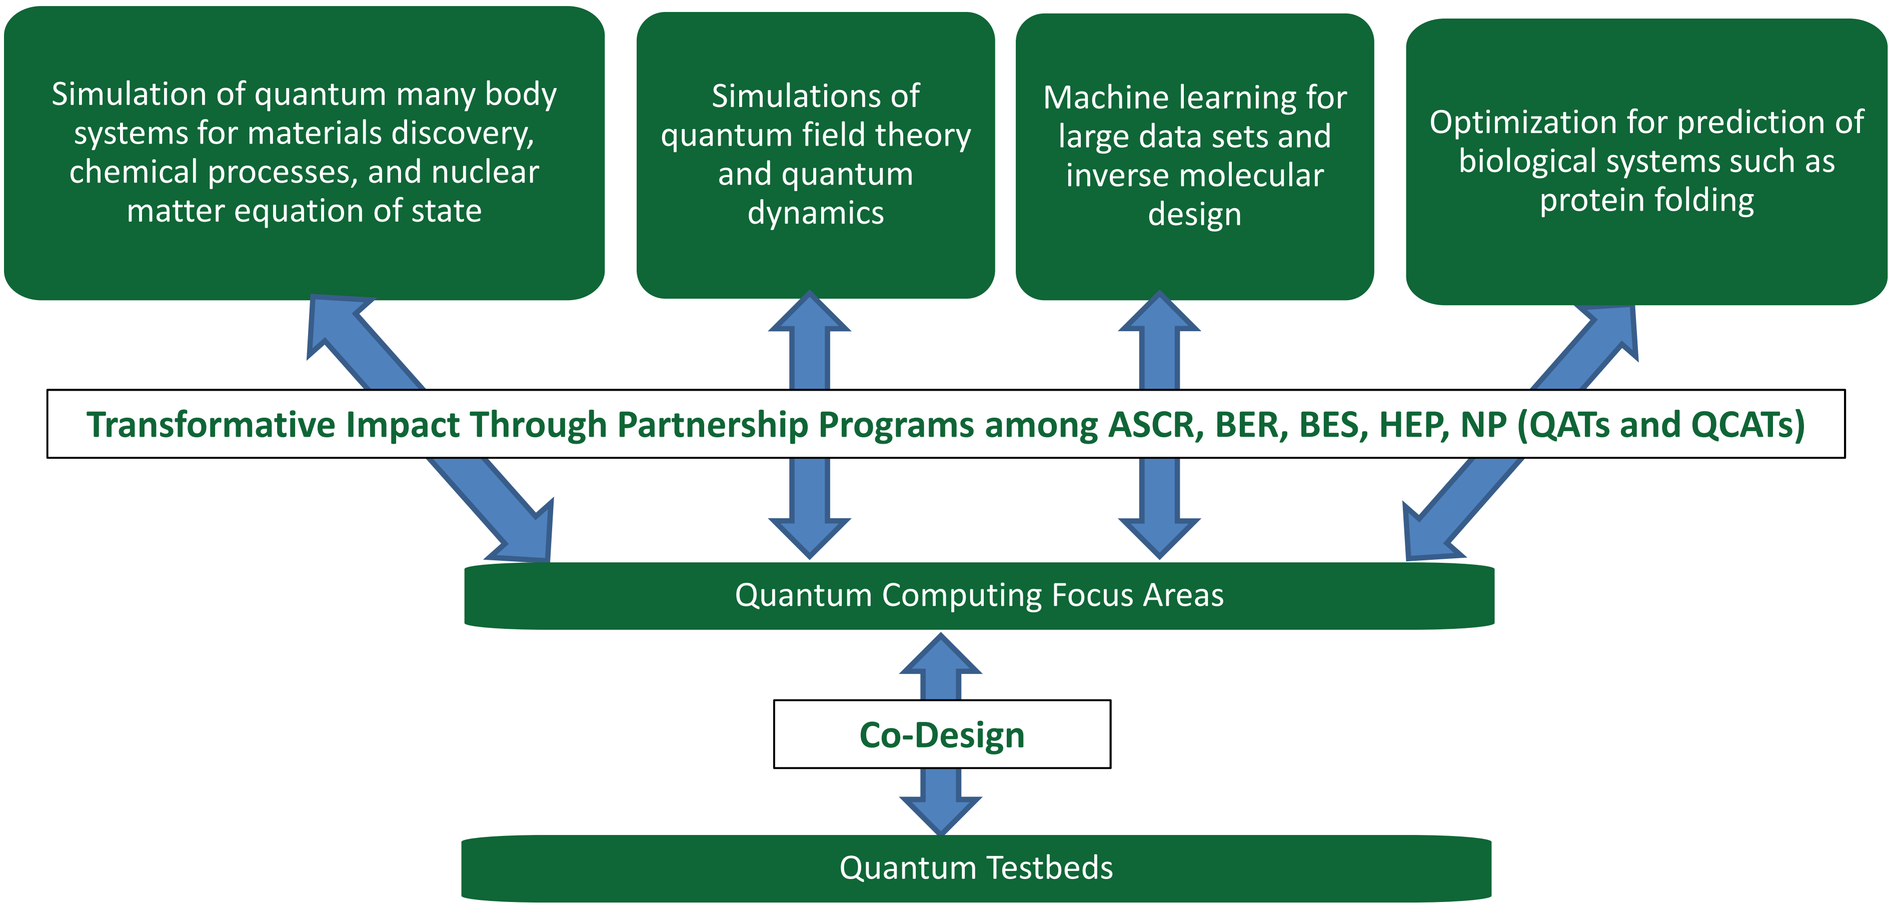

Quantum Algorithm Teams (QATs)

Purpose: To stimulate early investigations of quantum simulation and machine learning algorithms by focusing on key topics of research with relevance to problems of interest to SC

Emphasis: Interdisciplinary teams of quantum information science (QIS) experts, applied mathematicians and computer scientists that target specific application areas for quantum computing and analog quantum simulation

Timeline & Proposals: A DOE National Laboratory Announcement was issued in May 2017. 13 highly competitive proposals were received in July 2017.

3 Projects @ Total $4M/year:

Quantum Algorithms, Mathematics and Compilation Tools for Chemical Sciences. Lead: LBNL (Bert de Jong), Collaborators: ANL (Stefan Wild), Harvard University (Alán Aspuru-Guzik);

Heterogeneous Digital-Analog Quantum Dynamics Simulations. Lead: ORNL (Pavel Lougovski), Collaborator: University of Washington (Martin Savage);

Quantum Algorithms from the Interplay of Simulation, Optimization, and Machine Learning. Lead: SNL (Ojas Parekh), Collaborators: LANL (Rolando Somma), California Institute of Technology (John Preskill), University of Maryland (Andrew Childs), Virginia Commonwealth University (Sevag Gharibian)

FY 2017: Quantum Testbed Pathfinder

Purpose: To provide decision support for future investments in quantum computing (QC) hardware and increase both breadth and depth of expertise in QC hardware in the DOE community

Emphasis: Research in the relationship between device architecture and application performance, including development of meaningful metrics for evaluating device performance.

Timeline & Proposals: A DOE National Laboratory Announcement was issued in June 2017. 6 proposals were received in July 2017.

2 Awards:

<u>Advanced Quantum-Enabled Simulation</u> (LBNL, LLNL, UC Berkeley);

<u>Methods and Interfaces for Quantum Acceleration of Scientific Applications</u> (ORNL, IBM, IonQ, Georgia Tech);

Science Internet of Things (S-IoT)

Purpose: To exploit the next wave of Internet evolution (IoT) to develop advanced capabilities that will define a new paradigm of interactions between scientists and machines in hyper-connected smart environment for science

Technologies Trends
- IoT, SDN, Smart Sensors
- BigData analytics
- Autonomic computing
- Machine Learning

→

Proposed Capabilities
- Self-Management
- IoT Software Stack
- Embedded ML
- IoT security

→

Smart Science Things
- Smart Instruments
- Self-aware Networks
- Intelligent Workflows
- Smart edge computing

→

Smart Superfacility
(Science Internet of Things)

Benefits

- Improve access and utilization of connected scientific resources
- Reduce scientific resources management and operation complexities
- Simplify complex scientific workflow automation
- Enable machine –to – machine communication

Project Title: Architecture and Management for Autonomic Science Ecosystems (AMASE). Lead: ANL (Pete Beckman), Collaborator: LBNL (Alex Sim), Northwestern University (Alok Choudhary)

Laboratory Reverse Site Visits

Purpose: To inform ASCR's planning process to maintain and improve a robust research program in the near future and to provide each Laboratory an opportunity to communicate its plans to leverage and maintain its unique expertise and core capabilities in order to advance research in support of ASCR's mission

Structure: Each Laboratory presented its plans for the near future during a half-day visit to Germantown (2 weeks in early May 2017)

Emphasis: The template for the presentations highlighted the following:

- Mission Relevance: ASCR-mission relevant core research capabilities and proposed strategic shifts
- Self-Analysis of Strengths and Limitations: How the Laboratory's plans will leverage the unique capabilities that distinguish the Laboratory and address any current problems
- Realistic Budget Scenarios: Flat FY16 budget scenario and impact of a 17% reduction in ASCR Research funding

Outcome & Future: Very informative presentations and engaging discussions. Format and template will evolve based on discussions with the Labs.

What does the Future Hold: Strategic Vision for ASCR's Research Program

Emerging trends are pointing to a future that is increasingly
1. **Instrumented**: Sensors, satellites, drones, offline repositories
2. **Interconnected**: Internet of Things, composable infrastructure, heterogeneous resources
3. **Automated**: Complexity, real-time, machine learning
4. **Accelerated**: Faster &flexible research pathways for science & research insights

What is the role of ASCR's Research Program in transforming the way we carry out energy & science research?

1. **Post-Moore technologies**: Need basic research in new algorithms, software stacks, and programming tools for quantum and neuromorphic systems
2. **Extreme Heterogeneity**: Need new software stacks, programming models to support the heterogeneous systems of the future
3. **Adaptive Machine Learning, Modeling, & Simulation for Complex Systems**: Need algorithms and tools that support automated decision making from intelligent operating systems, in situ workflow management, improved resilience and better computational models.
4. **Uncertainty Quantification:** Need basic research in uncertainty quantification and artificial intelligence to enable statistically and mathematically rigorous foundations for advances in science domain-specific areas.
5. **Data Tsunami:** Need to develop the software and coordinated infrastructure to accelerate scientific discovery by addressing challenges and opportunities associated with research data management, analysis, and reuse.

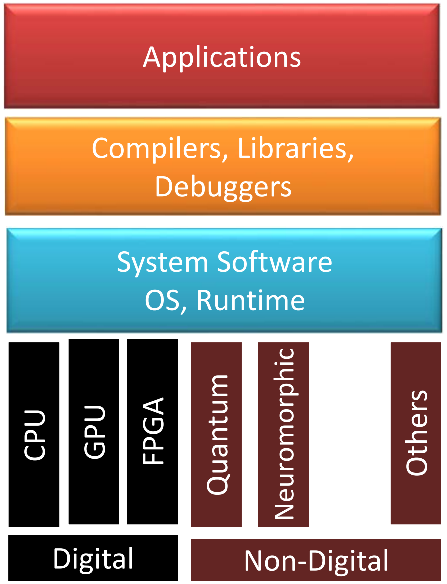

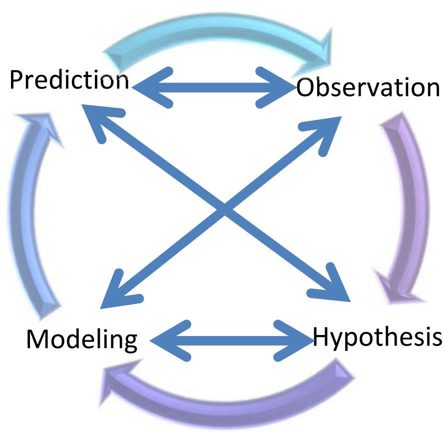

FY18 Plans

Extreme Heterogeneity Workshop

Tentatively planned for Jan. 23-25, 2018, in D.C. area.

- **POC: Lucy Nowell(Lucy.Nowell@science.doe.gov)**
- **Goal: Identify Priority Research Directions for Computer Science** needed to make future supercomputers *usable, useful and secure* for science applications in the 2025-2035 timeframe
- Primary focus on the software stack and programming models/environments
- 120 expected participants: DOE Labs, academia, & industry
- Observers from DOE and other federal agencies
- Planning: Factual Status Document (FSD) is under development, with outreach planned.
  - White papers to be solicited to contribute to the FSD, identify potential participants, and help refine the agenda
  - Report due early May 2018

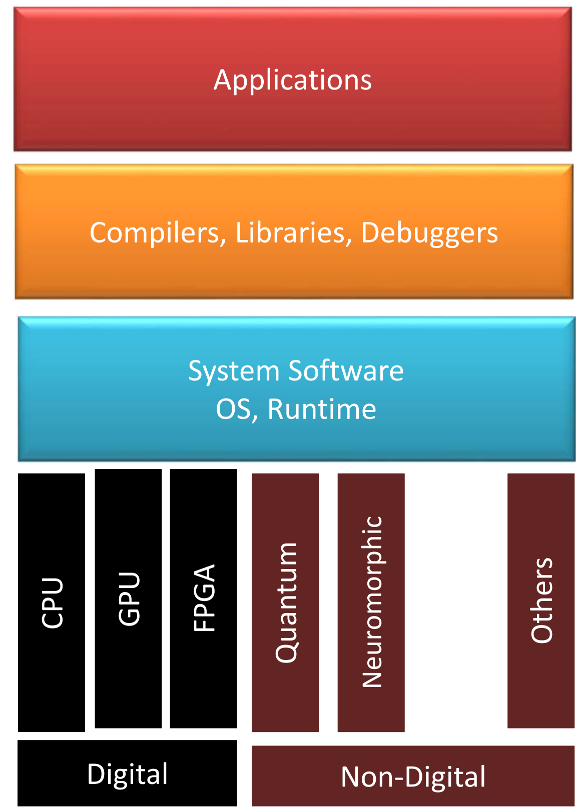

ASCR Machine Learning Workshop

ASCR Machine Learning workshop tentatively planned for 2-3 days in late January 2018 (DC area)

- **POC**: Steven Lee (Steven.Lee@science.doe.gov)
- **Purpose**: Identify Priority Research Directions and the role of applied mathematics in enabling greater machine learning capabilities for DOE-mission challenges.
- Factual Status Document – Status and recent trends in the underlying mathematical, statistical, and computational foundations of HPC machine learning techniques for scientific data analysis. Identify challenges and opportunities for high-impact through fundamental advances in mathematical modeling and algorithms.
- **Cross-cutting research themes:** optimization, linear algebra, UQ, discrete (tensors, graphs, networks), statistical approaches, ensembles, game theory, validation, scientific method
- **Number of participants:** About 80-100 researchers.
- White papers will be solicited to broaden community input.
- Observers from DOE and other federal agencies.
- Final Workshop Report due in March-April 2018

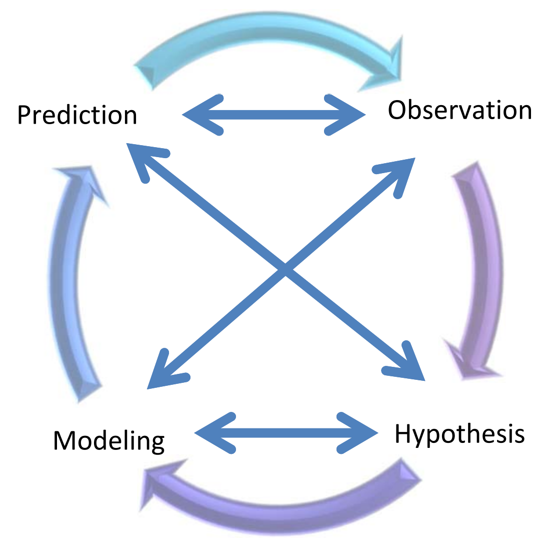

FY 2018 SBIR/STTR Recurring Topics

- **Advanced Digital Network Technologies and Middleware Services**
  - 3 subtopics focused on developing network performance monitoring and analysis tools and services
- **Increasing Adoption of HPC Modeling and Simulation in the Advanced Manufacturing and Engineering Industries**
  - 2 subtopics focused on increasing the effectiveness and productivity of manufacturing and engineering businesses through the use of HPC.

  New in FY18: HPC applications that address engineering challenges related to the design, integration, and fabrication of new devices for Beyond Moore's Law computing technologies including quantum computing.
- **HPC Cybersecurity**
  - Single topic focused on advanced methods and tools to protect HPC systems and centers
- **Smart Devices and Technologies for Science, Engineering, and Manufacturing**
  - 2 subtopics focused on accelerating the integration of intelligent sensors and devices into environments of interest to DOE

Phase I Release I Topics:
https://science.energy.gov/~/media/sbir/pdf/TechnicalTopics/FY2018_Phase_1_Release_1_Topics.pdf

Manny Oliver talk this afternoon!

FY 2018 SBIR/STTR Innovative Topics

- **Big Data Technologies for Science, Engineering, and Manufacturing**
  - Multi-program office topics focused on managing and analyzing large amounts of data
  - Topics by ASCR, BER, and BES

- **Collaborative Development Projects**
  - Multi-company collaborative topic focused on demonstrating a photonic replacement for the PCI Express (PCIe) electrical interconnect
  - University/Lab PIs may participate as subcontractors

Phase I Release I Topics:
https://science.energy.gov/~/media/sbir/pdf/TechnicalTopics/FY2018_Phase_1_Release_1_Topics.pdf

Manny Oliver talk this afternoon!

Some Agenda Details

- **UPDATE ON THE EXASCALE COMPUTING PROJECT –** *Paul Messina, ECP Director*
- **UPDATE ON CURRENT CHARGES**
  - *Committee of visitors – Susan Gregurick, ASCAC*
  - *Future Technologies – Vivek Sarkar, ASCAC*
- **ECP APPLICATION –** *Paul Kent, Oak Ridge National Laborato*ry
- **ESnet 6 Upgrade Planning –** *Inder Monga Lawrence Berkeley National Laboratory*
- **MACHINE LEARNING –** *Prabhat, Lawrence Berkley National Laboratory*
- ***DOE SMALL BUSINESS INNOVATIVE RESEARCH PROGRAM –*** *Manny Oliver, Office of Science*
- **40th Anniversary of DOE – Celebrating ASCR**
  - *CSGF – Bob Voigt*
  - *Partnership between NNSA and SC in HPC – Paul Messina, ANL*
  - *DOE Supported Computing Technologies that Made a Difference – Van Jacobson (TCP/IP)*, Rusty Lusk (MPI/MPICH), Phil Colella (AMR), Scott Klasky (ADIOS), Barry Smith (PETsc), John Wu (FastBit) and Buddy Bland (HPSS)*

* Invited

ECP Project Leadership Team as of October 1, 2017