The Critical National Need for Sustainable, High-Performance Computational and Data Science Resources in the Extreme-Scale Era

A Report from the Community Brainstorming HPCD Workshops Version 1.0

Contributors and Attendees

Stuart Anderson, LIGO/Caltech; Richard Arthur, GE Global Research; Dinshaw Balsara, University of Notre Dame; Bill Barth, TACC; Klaus Bartschat, Drake University; Greg Bauer, NCSA/University of Illinois; Cristina Beldica, NCSA; Jerzy Bernholc, North Carolina State University at Raleigh; Martin Berzins, SCI Institute University of Utah; Wes Bethel, LBNL; Tom Cheatham, University of Utah; David Dixon, University of Alabama; Thom Dunning, University of Washington; Said Elghobashi, University of California, Irvine; Jim Fonseca, Purdue University; Steven Gottlieb, Indiana University; Rama Govindaraju, Google; Bruce Harmon, Iowa State University; Anna Hasenfratz, University of Colorado Boulder; Thomas Hauser, University of Colorado Boulder; Victor Hazlewood, University of Tennessee – JICS; Curtis Hillegas, Princeton; Peter Kasson, University of Virginia; Fatemeh Khalili-Araghi, University of Illinois at Chicago; Kevin Kissel, Google; William Kramer, NCSA/University of Illinois; John Levesque, Cray Inc; Mike Levine, PSC; David Lifka, Cornell University; Honggao Liu, Louisiana State University; Paul Mackenzie, Fermilab; Pieter Maris, Iowa State University; Michael Norman, University of California San Diego; Steve Oberlin, NVIDIA; Nikolai Pogorelov, University of Alabama in Huntsville; Thomas Quinn, University of Washington; Jagnnathan Ramanujam, Center for Computation and Technology – LSU; Ralph Roskies, Pittsburgh Supercomputing Center; Mark Scheel, California Institute of Technology; Barry Schneider, NIST; Ed Seidel, NCSA/University of Illinois; Todd Simons, Rolls-Rovce Corporation; Dan Stanzione, University of Texas at Austin; John Towns, NCSA/University of Illinois; Frank Tsung, UCLA; Jorge Vinals, Minnesota Supercomputing Institute; Shaowen Wang, NCSA/University of Illinois; Ligiang Wang, University of Wyoming; Nancy Wilkins-Diehr, San Diego Supercomputer Center; Paul Woodward, University of Minnesota; Steve Wolff, Internet2; Donald Wuebbles, University of Illinois; Pui-kuen (P.K) Yeung, Georgia Institute of Technology; Shiwei Zhang, College of William & Mary

Workshop Organizers

William Kramer, Ed Seidel, Cristina Beldica, John Towns and Gregory Bauer

Workshop Funders

National Science Foundation, National Center for Supercomputing Applications

Editor and Contact Author

William Kramer National Center for Supercomputing Applications University of Illinois 1205 W. Clark Street Urbana, Illinois 61821 217 333-6260 wtkramer@illinois.edu

Ba	ackground	5
E	xecutive Summary	6
Fi	indings	9
	Shared Workshop 1 and 2 Findings	9
	Workshop 1 Specific Findings	10
	Workshop 2 Specific Findings	10
Sı	upporting Data	12
I.	Appendix A: Attendees	14
	Workshop 1	14
	Workshop 2	14
A]	ppendix B: Submitted White Papers available at the BrainstormHPCD website	15
A]	ppendix C: Attending Organizations	17
B	reakout Working reports from Workshop 1	17
	Breakout 1 – High Spectrum Computing and Analysis Needs and Priorities	17
	Characteristics of HPCD systems for integrated discovery	17
	Evaluate match of existing systems and needs (compute-intensive and data-intensive)	
	Impacts – if needs are unmet	18
	Prioritization guidelines	19
	What is needed (with the priorities)	19
	Breakout 2 – Access to Consistent Resources	20
	Summary	20
	Breakout 3 – Risks, Opportunities and NSF's Role Fostering High-Spectrum Sengineering	cience and 22
	Suggested NSF actions to improve sustained application performance and researcher pro	oductivity in 23
	Technical challenges for the next decade	23
	Ongoing Requirements Collection	24

Motivating Science Team Readiness	
Risks and Opportunities to U.S. Scientific Leadership	25
Breakout 4 – Computation and Data Analysis	25
Summary	25
Science Case studies	
Analysis Issues	
Breakout 5 – Storage & Data Analysis	40
Science Case Studies	
Persistence of Data	
Breakout 6 - Workflow & Methods	44
Current Assessment	
Challenges	
Requirements	
Science Cases	
Conclusions	
Breakout Working reports from Workshop 2	
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models	
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations	
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models	47 47 47 47 48
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods	
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives	47 47 47 47 48 48 49
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1	47 47 47 47 48 48 49 49
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1 Modality #2	47 47 47 47 48 48 49 49 50
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1 Modality #4	47 47 47 47 48 48 49 49 49 50 50
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1 Modality #2 Modality #4 Modality #5	47 47 47 47 48 48 49 49 49 50 50 51
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations. Alternative Design Models Breakout 2 - Alternative Provisioning Methods. Alternatives Modality #1. Modality #2. Modality #4. Modality #5. Additional Comments.	47 47 47 47 48 49 49 49 50 50 51 51 51
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1 Modality #2 Modality #2 Modality #4 Modality #5 Additional Comments Breakout 3 - Alternative Analysis Criteria	47 47 47 47 48 49 49 50 50 51 51 51 52 52
Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations Alternative Design Models Breakout 2 - Alternative Provisioning Methods Alternatives Modality #1 Modality #2 Modality #4 Modality #5 Additional Comments Criteria	47 47 47 48 49 49 50 50 51 51 51 52 52 52
 Breakout Working reports from Workshop 2 Breakout 1 - Alternative Design Models Summary of Recommendations. Alternative Design Models Breakout 2 - Alternative Provisioning Methods. Alternatives Modality #1. Modality #2. Modality #4. Modality #5. Additional Comments. Breakout 3 - Alternative Analysis Criteria Criteria Cost model components. 	47 47 47 48 49 49 49 50 51 51 51 52 52 52 52 52

Breakout 4 – NSF Administrative Implementation Alternatives	56
Criteria	56
Alternative Assessment	56
Breakout 6 – Alternative Evaluation and Assessment – 2	57
Support models vs. modalities	57
Support models vs. Criteria	
Support models vs. cost models	
Results of Workshop 1 Realtime Evening Brainstorming Session 1	59
"Elevator" Statements by Attendees	59
Results of Workshop 1 Realtime Evening Brainstorming Session 2	60
Revision Control	62

Abstract: Computation and data science are intrinsically intertwined in every aspect of science, engineering and research from theory to experiment. In Spring 2015, two community *BrainstormHPCD* workshops were held to identify the ongoing requirements among open science, engineering and research communities for future high-performance computational and data science (HPCD) resources and services. These workshops identified a persistent, broad and deep need for high-performance resources and services to enable leading-edge science, engineering and research investigations. The workshops provided insight into the changing requirements of pioneering investigations and explored alternative methods to provide these resources to the frontier research communities. This report summarizes results and conclusions of the two *BrainstormHPCD* workshops.

BACKGROUND

Several activities at the federal level are defining the science and research computational and data analysis efforts and activities for the U.S. research cyberinfrastructure. In 2014 the National Science Foundation (NSF) initiated a National Research Council report entitled *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020^a to "anticipate priorities and associated tradeoffs for advanced computing." Specifically, the committee was asked to report back on a) the contribution of high-end computing to U.S. leadership and competiveness in basic science and engineering and the role that NSF should play in sustaining this leadership; b) the expected future national-scale computing needs; c) complementarities and trade-offs that arise among investments in supporting advanced computing ecosystems; d) the range of operational models for delivering computational infrastructure; and e) the expected technical challenges to affordably delivering the capabilities needed for world-leading scientific and engineering research.*

Simultaneously, the Office of Science and Technology Policy was developing a plan that would define a government-wide strategy for high-performance computing, which was eventually announced in July 2015 as the National Strategic Computing Initiative.

The *BrainstormHPCD* workshops were designed to gather information and opinion on the direction NSF should take to support and further the use of high-performance computing and data in the quest for improved scientific discovery and insights and for national competiveness in science, engineering, research and commerce.

The first workshop—the Requirements and Needs Workshop—concentrated on gathering comprehensive information about the needs of the "NSF sponsored science and engineering research community" for advanced capacity through capability computing, including data-intensive and compute-intensive. The workshop focused on the research community needs at the high end of the spectrum and attempted to articulate the value, contributions and impacts of at-scale computing for the advancement of science, engineering and other research."

The second workshop—Implementation Alternatives Workshop—concentrated on "alternative operational models for providing computing capability, including the role of private industry, NSF centers, support of science workflows, and data management. This covered both how resources are provided to scientists and how they are managed."

^a See http://www.nap.edu/catalog/21886/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-2017-2020

In all, over 57 nationally recognized scientists, engineers and industry experts from 40 different institutions attended the workshops. Forty-three white papers were submitted to help start the discussions and support the conclusions. The appendix of this report lists the attendees, their institutions and the white papers. This information can also be found at http://www.brainstormHPCD.org as a permanent repository and should be considered as part of this report.

EXECUTIVE SUMMARY

The joint view of the workshop attendees is captured at a high level in the following recommendations and statements, with the first one being pre-eminent.

- 1. Because NSF has the broad national mission of advancing science and engineering and ensuring the scientific and economic competitiveness of the nation, it is vitally important that NSF provide a series of HPCD tier-1 and tier-2 systems over time to meet the increasing demands of high-spectrum digital science.
 - Scientific needs increasingly require high-performance capability for both modeling and simulation for data analysis and science.
 - $\circ~$ A variety of platforms are needed to be most efficient in addressing the breadth and depth of the national open-science research endeavors.
 - While diversity is important there is a point where dropping below a minimum size/capability of deployed national resources is counterproductive and inefficient to both the national science communities and the stakeholders. The national cyberinfrastructure **must** include systems of the scale and capability of the national labs' best systems. If NSF overfocuses on deploying smaller and/or specialized resources, there is a major risk that investigations will become incremental rather than best of breed.
 - The experiences of the 1980s, when large system funding was replaced with smaller local systems, should not be repeated. While there was a benefit for a while from these individual systems for small university research groups, this funding mode eventually led to what is commonly viewed as a fragmentary and less powerful computational landscape, which in turn led to incremental rather than "frontier" computational research. Fortunately, NSF recognized this problem and created the supercomputer centers in the second half of the 1980s.
- 2. The past sustained funding for HPCD has led to tremendous advances in many fields, and NSF cannot let that progress and momentum stop. It would be a serious blunder for NSF not to invest in high-end computing and data analysis infrastructure to support "frontier science."
 - 3. There are emerging curated petabyte-scale data repositories with thousands of people accessing them for both observational data, but also for simulated data. Examples of the former are the large experiment data repositories such as the Large Hadron Collider (LHC), Large Synoptic Survey Telescope (LSST), Laser Interferometer Gravitational-Wave Observatory (LIGO), and Square Kilometer Array (SKA). There are also computational data repositories such as Gauge Connection^b, hosted at the National Energy Research Scientific Computing Center (NERSC) at Berkeley Laboratory for QCD lattices, and the Program for Climate Model Diagnosis and Intercomparison (PCMDI)^c repositories at Lawrence Livermore National Laboratory. The need for sharing and

b http://qcd.nersc.gov/

^c http://www-pcmdi.llnl.gov

reanalysis of data will increase, and NSF should invest in resources to support community-based, curated, large data repositories with analysis capability to be used by hundreds to thousands of scientists.

- 4. It is important to enable complex workflows that are converging to be simultaneously computationintensive and data-intensive.
- 5. There is a compelling need to invest in software development and improvement programs and in the development of an effective workforce that can do this work.
 - In particular, NSF should invest in university programs related to the development and improvement of HPCD application software and tools. This education should include computer science classes and research and has to include discipline-specific computational science curricula (rather than solely computational modeling with tools like *MATLAB* and *Mathematica* in lieu of parallel programming).
 - While the exact configuration/architecture of future HPCD systems is uncertain, it will almost certainly be different than today. Hence, adapting these applications to effectively leverage future systems is critical for science discovery.
 - Development of software is crucial for proper exploitation of current resources and will become even more crucial for next-generation resources. Hence, funding for HPCD software development is crucial.
 - Algorithm development is critical and can at times be revolutionary in creating entirely new, more efficient methods for problem solving (e.g. adaptive mesh refinement is now used in many fields of science).
 - Workforce development is needed to train the current and next generation in the use of these advanced resources.
 - Just buying hardware is insufficient for productive science investigations, since hardware systems often are not used as effectively and efficiently as they could be with appropriate software optimization and training. NSF should continue to invest in intelligent algorithm methods and software that enable others to use supercomputers efficiently, such as asynchronous time integration, adaptive mesh refinement, load balancing, etc.
 - Developing and validating efficient codes is a long, labor-intensive process, particularly in the academic settings. Hence, there should be processes and funding that enable effective codes to continue to evolve and be used by wider communities.
- 6. Industry use of open HPCD resources to perform science and engineering work at scale, and therefore economic and/or advocacy support for such resources, will migrate to where the expertise exists in the facilities and centers. In order to help competitiveness, NSF-funded resources and facilities should encourage industrial engagement and facilitate use of these resources by industry. NSF facilities need to build this expertise to attract and keep industrial partners. Specifically, expertise is needed in:
 - Software tool chain, libraries, and algorithms are necessary enablers of high-end computing and support.
 - Validation of codes and reference data sets for validation are required for industry to transition to increased use of HPCD technologies.
 - NSF should pioneer allocation pathways for "non-discovery" use of high-end systems. This is critical for enhanced economic impact. An example of computationally intensive but "non-discovery" work that is important to industry is the meshing to create real-world complex

geometries and the validation of new codes and methods to establish the return on investment before being deployed in actual product creation.

- 7. The communities that rely on large-scale computing and data analysis need a comprehensive plan that will provide the delivery of multiple systems for the needs of scientists and engineers. The demand for higher performance resources is growing and will continue to grow over the next decade. This is due to increased use of modeling and simulation in all ranges of science and engineering, the rapid convergence of computation and data analysis, the increased workforce that can perform sophisticated programming, and the increasing effectiveness of software applications that can be used by more and more discipline scientists.
 - There are unique things that can be done only at the "top of the pyramid" and by best-ofbreed investigations. These activities have value. Examples of the size and scale are in the Science Case studies section below.
 - NSF needs to understand the diversity and breadth of the HPCD use cases to advance understanding in all areas of scientific investigation. Different tasks/types of work may best be done with different types of resources. Scientists and researchers need access to a range of systems at different scales and for data sharing. While focused on high-performance tier 1, NSF should not forget about funding for tier-2 and tier-3 systems. However, funding for tier-2 and tier-3 resources, or programs that do not supply computational/analysis resources, should not come at the expense of sustaining tier-1 resources. NSF should implement a mechanism focused on national institutional computing and data needs at larger scale than can be done with the Major Research Implementation (MRI) program.
 - NSF should support a pathway to expand the community capable of doing HPCD science. This requires campus and track-2 level steps for some, but there must be at least one leading-edge Tier-1 resource at the end of the path to make the journey worthwhile.
- 8. The Department of Energy (DOE) cannot support **all** HPCD science and engineering.
 - Teams asking for DOE HPCD resources have to argue the work is related to the DOE programs, which eliminates many worthwhile investigations.
 - DOE resources are designed, funded and allocated to serve the DOE programs and laboratory requirements. Hence, additional investment is required to support the rest of the science and research programs that are part of NSF's mission.
 - NSF has the mission for the nation to husband all fundamental scientific research and should provide the resources necessary to carry out that mission.
- 9. NSF needs to have a continuing plan. Not just the next step for 2017-2020.
- 10. It is not science if it is not reproducible and verifiable. Consequently there is a need for multiple teams/approaches to work at the same scale of problem to verify the science.
- 11. There are clear and negative economic impacts for the United States of not being at the global scientific leader. Intellectual propriety creation, whether open or proprietary, is enhanced by HPCD investment and supports the U.S. leadership in science and engineering.
- 12. NSF should understand the best options for open-source licensing.
- 13. To do at-scale science people have to be able to do software development at scale as well.
- 14. Return on investment is significant for HPCD for billion dollar experimental science. Computing to support both instrumentation design and experimental data understanding is necessary.
- 15. NSF should consider different methods for funding the infrastructure needed by all its divisions' programs so the entire foundation is doing coordinated investing in HPCD. NSF should investigate

funding mechanisms for HPCD that go beyond computing's traditional home in the Division of Advanced Cyberinfrastructure (ACI)/Directorate for Computer & Information Science and Engineering (CISE). Researchers across all NSF directorates can and do use HPCD resources, and funding should reflect the essential nature of these resources to discovery across the entire foundation.

- 16. It is essential for NSF to assure the community of the overall high-performance enterprise and critical for NSF to define both medium-range plans and a long-term commitment to computing at capability and capacity scales as well as data stewardship.
- 17. As science becomes more computationally and data intensive, a major expansion of computational and data resources becomes critical for advancing scientific discovery.
- 18. NSF should be less prescriptive in its solicitations for resources. NSF should be receptive to creative solutions for improving the national cyberinfrastructure ecosystem.
- 19. Coupling across scales makes HPC essential to science and enables remarkable steps forward.
- 20. There should be tighter coupling between the granting of science funding and the granting of computational and data resources for science teams.
 - When requesting funding for projects that require computation, analysis and/or storage, the proposals should include estimates of the resources needed in order to carry out the work, the likely providers of those resources (departmental or campus systems, national facilities, etc.). This information should be collected in a way that NSF can summarize and analyze the information to inform future investments.
 - The value of computing time and storage should be clear to the researchers receiving the allocations on shared resources that do not require reimbursement.
 - Innovative approaches should be used to influence teams to make timely and efficient use of resources. For example, the workshop discussed the idea of certifications of the team's skills to assist in determining awards of computational/data resources.
- 21. NSF should change the operational funding model toward an ongoing service and support model rather than the discrete project model that has been used for the past decade or more. Funding could be to an ongoing service center that science teams have "credits" to use at their choice. The "cloud service model" (not necessarily the cloud technology) could inform provisioning future resources and services.
- 22. Transferring new methods and algorithms in software to new systems/centers is non-trivial; research teams need help with these critical tasks. NSF should not be prescriptive as to the areas in which software development is done, whether libraries, community codes, individual science areas/teams, etc.

FINDINGS

Shared Workshop 1 and 2 Findings

- There is compelling evidence that computational and data-intensive computing are now implicitly intertwined with all areas of discovery.
- High-end computational and data-intensive computing resources are critical enablers for discovery in nearly all fields of science and engineering as well as for technological advancement.
- NSF has made great progress in expanding the use of high-end computing and data analysis throughout the research communities.
- The NSF is the primary agency responsible for basic science in the U.S. across all fields.
- NSF provides advanced computational and data resources for NSF grantees and researchers supported by other agencies such as the DoE, NIH, NASA, etc.

- At last count \sim 50% of all of the NSF-supported XSEDE compute cycles go to non-NSF-supported research.
- NSF will have decommissioned all its major HPCD resources by mid 2019. See Figure 2 NSF-Supported Computational Investments Reflect Increasing National Diversity (presented February 2015). There are no public roadmaps or time-tables for replacing these resources.
- NSF funding for HPCD resources^d, as reported in the OSTP/NITRD Federal Budget Supplements, has been close to flat for almost 15 years. See Figure 1 - NSF HEC Inflation Adjusted (CPI) Funding from NITRD Reports. At the same time, other federal agencies are investing more as shown in Figure 3 -NSF and DOE Office of Science HEC Funding Relative to Total Funding.

Workshop 1 Specific Findings

- Scientific needs increasingly require high-performance capabilities.
- The sustained funding for HPCD has led to tremendous advances in many fields and NSF cannot let that progress and momentum stop.
- Because NSF has the broad national mission of advancing science and engineering and ensuring the scientific and economic competitiveness of the nation, it is vitally important for NSF to provide a series of high-end computational and data analysis tier-1 and tier-2 systems to meet the increasing demands of high-spectrum digital science.
 - The NSF cyberinfrastructure must include systems of scale and capability.
 - There are benefits to deploying a variety of platforms, but more than two to four systems will create inefficiencies.
 - Minimum size/capability is required.
- "Frontier" science and engineering will be seriously inhibited if NSF does not to invest in high-end computing and data infrastructure.

Workshop 2 Specific Findings

- Computational and data-intensive computing are critical enablers for discovery in all fields of science and engineering and for technological advancement.
- The success of NSF's HPCD investments has created an environment where such HPCD resources are now assumed by large numbers of outstanding research teams.
- Investments in high-spectrum resources and services are necessary to maintain momentum in advanced research investigations.
- There is a need for an NSF-wide assessment of needs and risks with respect to provisioning high-spectrum resources and services to enable computational data science and engineering.
- Investments in HPCD resources/facilities should be viewed as equivalent to investments in experimental facilities (e.g. LIGO, Ocean Observatories Initiative, NEON, etc.) and should be implemented with a long-term investment plan.
- There needs to be a significant increase in coherence for NSF's investments in cyberinfrastructure.
 - NSF must evaluate and develop funding methods/programs for advanced CI that recognize the need for coherence and continuity for suppliers, providers and consumers of these capabilities in order to enhance the productivity of the scientific and engineering research endeavors.
- At least a real doubling of investment is needed in high-spectrum resources over the next five years.
- Geographical consolidation (2-4 major facilities with high-performance network access) is being driven by the need to co-locate compute and data resources.

^d Note the NITRD reports call HPCD High End Computing or HEC

- NSF must have a frontier science-driven digital ecosystem to support the broad open-science community. The duty to create this ecosystem for the open-science community cannot be ceded or delegated to other agencies because NSF would have no control of prioritization and management of the services.
- There are issues of access to resources at mission agencies, such as DOE. Worthy projects that fall outside these agencies' specific missions cannot access these resources.
- The existence of high-end CI resources enabled international competitiveness of U.S. researchers, which included expanding the expertise of the U.S. workforce.
- NSF must provide the resources necessary for researchers to keep pace with their collaborators.
- Technology paths forward for advanced computation are better defined than those for dataintensive science; therefore, the community needs more technology development efforts for resources supporting data-intensive science.

SUPPORTING DATA



Figure 1 - NSF HEC Inflation Adjusted (CPI) Funding from NITRD Reports



Figure 2 - NSF-Supported Computational Investments Reflect Increasing National Diversity (presented February 2015)



Figure 3 - NSF and DOE Office of Science HEC Funding Relative to Total Funding

I. APPENDIX A: ATTENDEES

Workshop 1

- 1. Stuart Anderson, LIGO/Caltech
- 2. Richard Arthur, GE Global Research
- 3. Klaus Bartschat, Drake University
- 4. Greg Bauer, NCSA/University of Illinois
- 5. Martin Berzins, University of Utah
- 6. Wes Bethel, LBNL
- 7. Tom Cheatham, University of Utah
- 8. Said Elghobashi, University of California, Irvine
- 9. Jim Fonseca, Purdue University
- 10. Steven Gottlieb, Indiana University
- 11. Bruce Harmon, Iowa State University
- 12. Anna Hasenfratz, University of Colorado Boulder
- 13. Peter Kasson, University of Virginia
- 14. Fatemeh Khalili-Araghi, University of Illinois at Chicago
- 15. William Kramer, NCSA/University of Illinois
- 16. John Levesque, Cray Inc.
- 17. David Lifka, Cornell University
- 18. Paul Mackenzie, Fermilab
- 19. Pieter Maris, Iowa State University
- 20. Michael Norman, University of California San Diego
- 21. Steve Oberlin, NVIDIA
- 22. Nikolai Pogorelov, University of Alabama in Huntsville
- 23. Thomas Quinn, University of Washington
- 24. Ralph Roskies, Pittsburgh Supercomputing Center
- 25. Mark Scheel, California Institute of Technology
- 26. Ed Seidel, NCSA/University of Illinois
- 27. Todd Simons, Rolls-Royce Corporation
- 28. Dan Stanzione, University of Texas at Austin
- 29. John Towns, NCSA/University of Illinois
- 30. Shaowen Wang, NCSA/University of Illinois
- 31. Paul Woodward, University of Minnesota
- 32. Donald Wuebbles, University of Illinois
- 33. Pui-kuen (P.K) Yeung, Georgia Institute of Technology

34. Shiwei Zhang, College of William & Mary

Workshop 2

- 1. Stuart Anderson, LIGO
- 2. Dinshaw Balsara, University of Notre Dame
- 3. Bill Barth, Texas Advanced Computing Center
- 4. Greg Bauer, NCSA/University of Illinois
- 5. Cristina Beldica, NCSA/University of Illinois
- 6. Jerzy Bernholc, North Carolina State University at Raleigh
- 7. David Dixon, University of Alabama
- 8. Thom Dunning, University of Washington
- 9. Rama Govindaraju, Google
- 10. Robert Harrison, NRC Observer/Stonybrook
- 11. Thomas Hauser, University of Colorado Boulder
- 12. Victor Hazlewood, University of Tennessee – JICS
- 13. Curtis Hillegas, Princeton
- 14. Fatemeh Khalili-Araghi, University of Illinois at Chicago
- 15. Kevin Kissel, Google
- 16. William Kramer, NCSA/University of Illinois
- 17. Mike Levine, Pittsburgh Supercomputing Center
- 18. Honggao Liu, Louisiana State University
- 19. Jagnnathan Ramanujam, Center for Computation and Technology – LSU
- 20. Barry Schneider, NIST
- 21. Ed Seidel, NCSA/University of Illinois
- 22. John Towns, NCSA/University of Illinois
- 23. Frank Tsung, UCLA
- 24. Jorge Vinals, Minnesota Supercomputing Institute
- 25. Liqiang Wang, University of Wyoming
- 26. Nancy Wilkins-Diehr, San Diego Supercomputer Center
- 27. Steve Wolff, Internet2
- 28. Paul Woodward, Minnesota Stellar Explosions

APPENDIX B: SUBMITTED WHITE PAPERS AVAILABLE AT THE BRAINSTORMHPCD WEBSITE

- 1. <u>Searches for Gravitational Waves with the Laser Interferometer Gravitational-Wave</u> <u>Observatory (LIGO) — Stuart Anderson, California Institute of Technology</u>
- 2. Position Paper to National Research Council Richard Arthur, GE Global Research
- 3. <u>PGAS Languages: An Easy-Entry Paradigm for Peta/ExaScale Computing</u> Dinshaw S. Balsara, <u>University of Notre Dame</u>
- 4. <u>Comments on Topics 4 and 6 on the NRC List Klaus Bartschat, Drake University</u>
- 5. <u>Quantum Simulations in Materials Design, Nano Science and Technology</u> Jerry Bernholc, <u>Center for High Performance Simulation and Department of Physics, NC State University</u>
- 6. <u>The Blurry Line Between Exascale and Big Data</u> E. Wes Bethel, Lawrence Berkeley National <u>Laboratory</u>
- 7. <u>Supplying Cycles for High Performance Computing with a Focus on Chemistry: Community</u> Needs to be Considered — David A. Dixon, Department of Chemistry, The University of Alabama
- 8. <u>Comments on Interim Report: "Future Directions for NSF Advanced Computing</u> <u>Infrastructure to Support U.S. Science and Engineering in 2017-2020" — Thom Dunning</u>, Pacific Northwest National Laboratory & University of Washington
- 9. <u>No Results Left Behind: The Case for Testing in Computational Science and a Test Case in</u> <u>Nanoelectronics Modeling — Jim Fonseca, Purdue University</u>
- 10. <u>Lattice Field Theory for High Energy Physics</u> Steven Gottlieb, Indiana University; Anna Hasenfratz, University of Colorado; Paul Mackenzie, Fermilab; Robert Sugar, UCSB
- 11. <u>Lattice Field Theory Computations</u> Steven Gottlieb, Indiana University; Anna Hasenfratz, <u>University of Colorado; Paul Mackenzie, Fermilab; Robert Sugar, UCSB</u>
- 12. NSF 2nd HPCD Workshop 2015 Rama Govindaraju, Google
- 13. <u>Change or Perish Bruce Harmon, Iowa State University</u>
- 14. <u>It's going to take more than hardware to advance the state of the art in scientific</u> <u>computation — John M. Levesque, Cray Inc.</u>
- 15. <u>Tighter development cycles between scientific applications and advanced computational</u> <u>infrastructure — Peter Kasson, University of Virginia</u>
- 16. <u>Molecular Dynamics simulations of Biological Systems</u> Fatemeh Khalili-Araghi, University of <u>Illinois at Chicago</u>
- 17. <u>Accelerating Scientific Discovery and Engineering Practice through Advanced, High</u> <u>Spectrum Computing and Data Analysis (tables)</u> — William Kramer, NCSA/University of Illinois; <u>et al.</u>
- 18. Thoughts on the NSF Future Directions Interim Report Glenn K. Lockwood
- 19. Position Paper for HPCD Brainstorm Workshop Pieter Maris, Iowa State University
- 20. <u>Answers from a user of advanced computing infrastructure</u> Pieter Maris, Iowa State <u>University</u>
- 21. <u>Needs and Opportunities In Cosmology In The Era of Ultra-Deep Surveys</u> Michael L. <u>Norman, UCSD</u>
- 22. The Other 90 Percent Steve Oberlin, NVIDIA
- 23. <u>Thoughts on The Future of NSF-supported Advanced Computing from a Numerical Relativity</u> <u>and Computational Astrophysics Perspective — Christian Ott and Mark Scheel, Caltech</u>
- 24. <u>High-Performance Computing Challenges in Space Physics Simulations</u> Nikolai V. <u>Pogorelov, University of Alabama in Huntsville</u>
- 25. <u>Computational Challenges for Galaxy Formation Thomas Quinn, University of Washington</u>

- 26. Position Paper for 2nd Brainstorming HPCD Workshop Barry I. Schneider, NIST
- 27. <u>High Performance Computing Policy Paper Todd Simons, Rolls-Royce</u>
- 28. <u>Reflections on an Observation of the Interim Report of the Committee on Future Directions</u> <u>for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017-20 — Dan</u> <u>Stanzione, Texas Advanced Computing Center</u>
- 29. <u>**Towards HPC in the Cloud**</u> Liqiang Wang, Department of Computer Science, University of Wyoming
- 30. <u>Scalable CyberGIS Analytics for Solving Complex Environmental, Geospatial, and Social</u> <u>Scientific Problems — Shaowen Wang, University of Illinois</u>
- 31. <u>Observations and Recommendations Regarding NSF's Support for High-Performance</u> <u>Computing — Nancy Wilkins-Diehr, Wayne Pfeiffer, and Richard L. Moore, SDSC</u>
- 32. <u>High-Resolution Earth System Global and Regional Modeling for Climate Assessment and</u> <u>Policymaking Require Advanced Computing Infrastructure</u> — Donald Wuebbles, University of <u>Illinois at Urbana-Champaign; Warren Washington, Gerald Meehl, and Tom Bettge, National Center</u> <u>for Atmospheric Research</u>
- **33.** Fluid Dynamics and Turbulence: the Case for NSF-HPC in 2017-2020 P.K. Yeung, Georgia <u>Tech</u>
- 34. <u>An Integrative, Cross-Foundation Cyberinfrastructure for Science & Engineering Research</u> <u>multiple authors</u>
- 35. <u>Simulating the First Galaxies and Quasars: The Bluetides Cosmological Simulation</u> Blue Waters PI: Tiziana Di Matteo, Carnegie Mellon University
- 36. Solving Prediction Problems in Earthquake System Science Blue Waters PI: Thomas H. Jordan, SCEC
- **37**. <u>Enabling Breakthrough Kinetic Simulations of the Magnetosphere</u> Blue Waters PI: <u>Homayoun Karimabadi, UCSD</u>
- 38. <u>Accelerating Nanoscale Transistor Innovation with NEM05</u> Blue Waters PI: Gerhard <u>Klimeck, Purdue University</u>
- 39. <u>Design and Management of Satellite Assets to Advance Space-Based Earth Science</u> Blue Waters PI: Patrick Reed, Cornell University
- 40. <u>Simulation of Turbulent Stellar Hydrodynamics</u> Blue Waters PI: Paul Woodward, University <u>of Minnesota</u>
- 41. <u>High-Resolution Climate Simulations</u> Blue Waters PI: Donald Wuebbles, University of Illinois <u>at Urbana-Champaign</u>

APPENDIX C: ATTENDING ORGANIZATIONS

- 1. California Institute of Technology
- 2. College of William & Mary
- 3. Cornell University
- 4. Cray Inc.
- 5. Drake University
- 6. Fermilab
- 7. GE Global Research
- 8. Georgia Institute of Technology
- 9. Indiana University
- 10. Iowa State University
- 11. LBNL
- 12. LIGO/Caltech
- 13. National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign
- 14. NVIDIA
- 15. Pittsburgh Supercomputing Center
- 16. Purdue University
- 17. Rolls-Royce Corporation
- 18. University of Alabama in Huntsville
- 19. University of California San Diego
- 20. University of California, Irvine
- 21. University of Colorado Boulder
- 22. University of Illinois at Chicago
- 23. University of Illinois at Urbana-Champaign
- 24. University of Minnesota
- 25. University of Texas at Austin
- 26. University of Utah
- 27. University of Virginia
- 28. University of Washington

BREAKOUT WORKING REPORTS FROM WORKSHOP 1

Breakout 1 – High Spectrum Computing and Analysis Needs and Priorities

Charge: Evaluate the match between resources and demand for the high spectrum of systems, for both compute- and data-intensive applications, and the impacts on the research community if NSF can no longer provide state-of-the-art computing and data analysis for its research community. Identify the characteristics for advanced high-spectrum computing and data infrastructure that enables integrated discovery involving experiments, observations, analysis, theory, and simulation.

Characteristics of HPCD systems for integrated discovery

- Data-intensive platforms are those on which large data sets are always close to the processors (on file system—not archived), similar to Google (running a Hadoop cluster). Example: 100TB disk on each compute node.
- A requirement for data-intensive computing is a low response time to data query.
- The U.S. does not have facilities for collaboration on very large data sets (Europe does, for cosmology—virtual observatories).
- The lattice QCD community has configurations that they share.

- This allows other researchers to use (and re-use) very large data sets that have been created at great expense.
- Data infrastructure is critical for:
 - Gravitational wave physics requires predominately computing throughput for processing analysis and large-scale simulations for modeling.
 - Climate modeling needs 10X computing resources from last assessment for higher resolution to understand local impact. Data sets will be 5-10 PB and will have a 5- to 8-year useful life with thousands of people exploring the data. This capability is there for CMAT5 (now 0.5 PB).
 - Environmental/geospatial science is more data-driven than compute-driven. There have been many innovations in data analytics. Now the research community is starting to look at computing-intensive applications (emergency response to storms, for example). Fast visualization is critical for discovery.
 - Nuclear physics needs larger computing and larger models.
 - Better uncertainty quantification (UQ) across multiple disciplines requires substantial data and simulation.
 - Atomic molecular optical physics needs computing and memory. Current data systems are sufficient.
- Systems need to be managed in a way that calculations can be run in a timely fashion.
- NSF has a mission to expand access to computing. Cutting back on computing would hurt a wide swatch of the scientific community. DOE's mission supports a narrower field of science.
- Industry needs support in scaling codes, both commercial third-party and in-house proprietary.
- Industry also needs support in developing and validating the modeling and simulation from low technology readiness.
- Turbulence research is both computationally intensive and data intensive.
- Longevity of large data sets generated on NSF machines is at risk for integrated discovery.
- NSF needs to be supporting fundamental research in fluid dynamics.

Evaluate match of existing systems and needs (compute-intensive and data-intensive)

• Blue Waters support has been very good for fluid dynamics research (PK Yeung)

Impacts – if needs are unmet

- Failing to expand computing would have disastrous impact on many fields of science and on many research teams.
- Scientific leadership would suffer in turbulence modeling and the ability to predict weather.
- Industrial research will migrate within global companies to Europe where governments are investing in industrial research. It will be increasingly difficult to reverse this trend. If companies are not global, they will be at an increasing disadvantage.
- The education of STEM workforce will suffer without continued investment in HPCD.
- Theoretical (computational) support for experimental research will also suffer without increased computational investment.
- Theoretical (computational) work is needed to support traditional experimental science.
- Without high-spectrum HPCD, society will miss out on the value of USGS topological LIDAR surveys. Agricultural benefits are significant (~\$10 billion). Development of a skilled workforce depends on increased investment.
- The speed of science will <u>decelerate</u> in climate science.
- Gravitational wave physics is investing \$1 billion in experiments. Not being able to analyze this data would be a loss to U.S. taxpayers. Other parts of the world would analyze this data. A proportionate investment in computing is needed.

- Biomolecular science is at a turning point where federally funded researchers will be unable to realize a transformation in how science is conducted.
- Cosmology is at a point where simulations are required to advance our understanding by proving hypotheses. Computations provide understanding where observations are insufficient.
- We have spent billions in collecting cosmological data but do not invest sufficiently in getting good value out of this data by analysis.
- Without computing allocations researchers will not be able to fulfill science proposals.

Prioritization guidelines

- Acquiring data modeling and simulation
- Acquiring data data analysis
- Software
- User support
- Data management

What is needed (with the priorities)

- 1. More powerful platforms
 - Compute intensive
 - Data intensive
 - Visualization improved
- 2. Funding for software development
 - Science applications
 - System software
 - Languages and tools
- 3. Support services at centers
 - Advanced training

What is needed (priorities - Lower is		Breakout Member 1	Breakout Member 2	Breakout Member 3	Breakout Member 4	Breakout Member 5	Breakout Member 6	Breakout Member 7	Average
1. More powerful									
- '	Compute – intensive	1	1	1	1	1	1	1	1.0
	Data – intensive	3	1	3	2	3	2	3	2.4
	Visualization	3	3	3	3	3	2	3	2.9
2. Funding for software development									
	Science applications	1	2	1	1	1	1	1	1.1
	System software - including workflow	2	2	2	2	2	3	2	2.1
	Languages and tools	3	3	3	3	3	3	3	3.0
3. Support at center									
	Advanced training	2	3	2	3	2	3	2	2.4

Other

- User support is very important to help researchers take advantage of computing resources. Computer science support allows codes to run more efficiently.
- User support is a critical part of the computing ecosystem.

Breakout 2 – Access to Consistent Resources

Charge: There are challenges facing researchers in obtaining access to advanced computational and data analysis resources. These include allocation processes, oversubscription of resources, and the need to frequently migrate from one resource to another. Some have said the easiest part is actually getting the at-scale application to work. In other cases, the communities may be self-limiting their requests to known resources rather than requesting what is truly needed. In reviewing the white papers and other sources, estimate the true computational requirements that would enable all areas of science and engineering to make timely progress in both best-of-breed problems and common-practice problems. Where possible, provide quantitative data on computing needs.

Attendees: not recorded

Summary

The need for high-performance research computing resources, at all levels, but particularly at the Track-1 level, is increasing at a rapid pace, with almost all domains of science being limited by access to these resources. Whether calling the resources cyberinfrastructure, centers, HPC facilities, and/or major research instruments does not change the fact that the NSF research communities rely on these advanced production systems to carry out their investigations.

Therefore, NSF should develop a coherent, transparent and coordinated strategic vision and implementation plan based on a full investigation of science needs, requirements and priorities. NSF and resource providers should constantly evaluate the portfolio of available resources from campuses up to Track-1 resources. Furthermore, NSF should recognize that science is requiring more and more research computing resources and that the growth in demand has (and will continue to) to outpace the growth in resources. Improvements in hardware and software lead to greater demand as new research teams can better include at-scale modeling and analysis.

Given the lack of community awareness of what will be available, it is impossible for research teams to plan feasible research goals or method improvements. The community has become reactive rather than proactive in developing long-term plans and goals. Researchers are being less ambitious due to perceived or actual limits in resources and therefore scientific goals are becoming more incremental (e.g. just enough to get the next paper accepted) rather than high impact and breakthrough. Some problems cannot be approached, to sufficient accuracy, without access to significantly more computing resources than is currently available or known to be part of NSF plans for future resources. The over-subscription of currently available resources is under-estimated since the process turns away users (science problems may not be able to be addressed, progress to success is too difficult). Researchers self-limit their proposals to computing and data allocations they think are available/awardable/expected. Over-subscription can inhibit scientific progress as well, making the system less efficient and introducing longer delays in time to solution.

Increased computer resources enable many people to do research. Due to difficulties in workflow/analyses, needed code development, etc., not everyone can suddenly use track-1 resources immediately that are 10x, 100x, 1,000x more powerful than individual systems. However, many, possibly even the majority, teams can or have already made the necessary improvements in their applications and envision problems that

need such levels of sustained performance. The expectation that large resource allocations are possible will encourage other teams to "think big" and pose new and exciting research. It has been shown in multiple environments that the lack of allocated time is at least as, if not more, inhibiting to teams wanting to do leadership-scale research as the readiness of codes or algorithms. So, lack of sufficient resources is stifling innovation and inhibiting exploration of creative new research directions.

NSF needs to address the balance of funding for moderate to large-scale (Track-1/2) computational resources against moderate to large-scale instruments (MREFC, MRI, etc.) and the computational and analysis resources required to meet the theoretical, experimental and simulation/modeling research needs. At the moment, the perception is that this balance is lacking and does not favor large scale computational and data analysis resources. If NSF adopts a long-term and sustainable approach to providing the necessary computational and data analysis resources, this will enable some projects with dedicated computing/analysis needs (e.g. LSST, LIGO) to move to shared infrastructure rather than creating their own dedicated infrastructure. It has been shown that even for experiments that need some amount of guaranteed or dedicated processing (such as real-time event notification) most processing needs can be met in a shared infrastructure that leverages cost-effectiveness, thereby reducing the overall costs to NSF. This approach would increase the cost-effectiveness of MREFCs and other long-term projects in NSF.

NSF should encourage and oversee proper, balanced, and efficient use of resources and discourage "improper" usage by providing the right resource for the right need or usage modality. An example is avoiding many single-node projects on Blue Waters or Stampede, unless there is a compelling need based on leading-edge usage of other system resources, such as the high-speed interconnect, memory or I/O subsystems.

NSF should consider an alternative for entry-level computational and data resources rather than the standard MRI. MRIs for computation seems to be at a disadvantage since there is a high degree of internal (campus-limited submissions) and external competition in the need for advanced resources and instrumentation (i.e. competing with cryoEM, sequencers, E- microscope, nanofab, etc.). Also, NSF should facilitate and actively guide the use of emerging architectures and software efforts to make useful impacts or end the investigation if it is not progressing in useful ways.

NSF should provide a mechanism to, at least, correlate research funding (dollars) with the required largescale computational and data analysis resource allocation. At present, there is a mismatch since the funded science may not (probably will not) get the computational and data analysis resources needed to maximize the impact and return on investment of the research funding.

NSF could consider a simplification of the allocation processes. Questions that should considered include:

- Are there benefits to a single unified process for requesting all NSF-wide cyber-infrastructure?
- If there is one unified submission process for all resources (XSEDE, Yellowstone, Blue Waters, Open Science Grid, MREFC computing, computer science and other research/prototype/testbed systems, etc.) should there be a tiered review process?
- If there are different allocation types (Track-1, domain-specific, industry, scalable applications, capacity applications, education and training, etc.) should there be different review processes?
- It is generally acknowledged that applying for allocations of computer time appears easier and/or more navigable than applying for funding for scientific research. However, there also is an impression that the additional expectation of a "heavy" review process for Blue Waters inhibits submitters and may impact the number of requests. Should there be precursor requirements or qualifications for projects and users of the most advanced systems, such as using significant XSEDE or Yellowstone resources?

- Estimating the resource requirement for the amount of computational/analysis time is very difficult for teams, and many are not trained in performance evaluation. It is difficult to predict what workflows will be useful for multi-year efforts. Often the use of resources is contingent on having the human resources available to enable the computations/analysis, but people are not always available when the computing resources are. Hence, many teams do not use their full allocations and often significantly underuse their allocations. Unlike funding, the computing time cannot be "banked." How should the allocation review process take these factors into account?
- Large groups join forces to get access to computing resources (e.g. USQCD) and are highly successful in obtaining awards, yet these groups are loosely coupled rather than focused teams undertaking coordinated efforts to achieve specific science outcomes. Individual investigators may perceive barriers to entry, leading a huge potential user base to opt out of competing for large-scale resources.
- NSF wants to encourage more industrial impact and competitiveness. Current panels and review criteria often are significantly biased toward academic and pure research projects and the criteria are often detrimental to industrial research. If NSF wants to have higher impact for industrial and commercial research it needs to develop new, equivalent criteria to review industrial proposals and ensure the review panel has a proper balance of people with academic and industrial research credentials.

NSF needs to be more clear and specific in how data related to NSF projects should be stored and preserved. NSF has no clear requirements on saving workflows, reproducibility of results or community curation of data repositories. While NSF has required data management plans for projects for several years, there is no guidance that infrastructure will be deployed to enable PIs to manage their data beyond the end of individual grants. There is no sustaining data infrastructure and little support for dissemination, sharing, and archiving of data. NSF should commission a study of the data management plans it has received, and the data storage needs of its research projects, and should develop an efficient and effective approach to curating and preserving important and required data.

NSF supports many science projects that use data not initiated or created by NSF-funded research (e.g. satellite data, NIH data, etc.). It is not NSF's role to manage or store data from other organizations, but it is an important responsibility to ensure NSF-funded results can be reproduced and extended.

Breakout 3 – Risks, Opportunities and NSF's Role Fostering High-Spectrum Science and Engineering

Charge: There are multiple technical challenges to building future, more capable advanced computing and data systems for the next decade. Technology limitations will make some approaches more difficult for applications to use in a productive manner. How should NSF best respond to the challenges for sustained application performance and researcher productivity? What are the risks in current U.S. plans for extreme-scale computing, and what can NSF do to address and reduce the risks?

Attendees: Richard Arthur (GE Global Research), Greg Bauer (NCSA/University of Illinois), Wes Bethel (LBNL), Said Elghobashi (University of California, Irvine), Jim Fonseca (Purdue University), Steven Gottlieb (Indiana University), Bruce Harmon (Iowa State University), John Levesque (Cray Inc.), David Lifka (Cornell University), Steve Oberlin (NVIDIA), Ralph Roskies (Pittsburgh Supercomputing Center), Dan Stanzione (University of Texas at Austin)

Suggested NSF actions to improve sustained application performance and researcher productivity in the future

- **Application performance**: training, support teams, research in computer science to support new modes in parallel programming.
- **Researcher productivity:** Training is a complex problem, but necessary if new code and algorithms are to be effectively implemented in a timely manner, particularly on machines with new architecture. Online training is one possible solution; but how to solicit and reward 'students'? Who writes 'the course'—vendors?
- **Sustained application performance:** If this means how to keep the code optimized in view of hardware improvements, then close interaction with vendors with training should be part of the acquisition priorities.

Technical challenges for the next decade

- In the late 1990s the disruption was the change from shared-memory multiprocessors to distributed systems. The current disruption is the end of Moore's Law, trends to extreme parallelization, need for vectorization, and multi-thread computing.
- Multiple levels of parallelism are required now.
- Deepening memory hierarchy: increasingly complex, no standards for accessing it, code performance and energy consumption increasingly dependent upon data movement
- Better tools to analyze vectorization, etc.

Will a research team's investment over time pay off? Researchers do not know what systems are coming, but they need to decide if they should invest time to support a particular programming scheme. Want performance portability across generations of technology, e.g., work done to build optimized code for a given platform would survive into future generations.

HPC system providers could spend time with large users to do deep analysis of code.

The desire is performance portability to allow for changing to a new architecture without significant recoding.

The vast majority of NSF users are not *coders*. How do we address those users?

Pair programming: Non-domain expert budgeted to help with making effective use of high-spectrum resources, including parallelization.

NSF announces new platforms a year in advance. But often this is insufficient lead time for teams to port their codes to the new platform. Also, there needs to be explicit funding for such efforts, to work on getting codes ready for next-generation platforms. Therefore, NSF should modify funding vehicles to allow the deployment of small early science systems before the new platforms are in service and a path for migration from existing systems to new ones.

In-depth support for both experienced and inexperienced science teams and investigators it beneficial in creating efficient applications and productive teams. Support models such as the XSEDE Extended Collaborative Support Services (ECCS) and/or the Blue Waters' Science and Engineering Applications Support (SEAS) programs are needed in the future.

Training materials exist. How to get people to use them? Is it lack of awareness or disinterest, something else?

Is it high-end computing if resources are not used efficiently? It may be a high-end scientific issue.

NSF is charged with advancing **all** science in the U.S., while other agencies (NIH, DOE) have more specific goals. NSF **not** addressing HPC means some area of science or some aspect of that broad mission is left unserved.

Productivity:

Bridge the gap between technology development and solving scientific problems. What other questions can be asked to make the science teams more productive? Includes workflow, human-computer interfaces, stream lining processes, etc.

We all want to reduce the time/effort required to perform complex computational activities, like setting up parameter studies and producing ensemble output, setting up analysis of ensemble results, uncertainty quantification, etc.

Industrial problems: pre- and post-processing take significant time. Input generation, mesh generation and validation are significant time-to-solution challenges. Best practices and workflows for *in situ*, i.e. steerable, technologies. There are different paradigms in industry and academia. Can lessons from industry be used to help shape computational modeling in academia?

Scaling codes from laptop to big machines is desired by some. How can that migration be simplified for researchers? Researchers may care more about time to answer than efficiency.

Some researchers, especially biologists, may go to the cloud, but many cannot.

Even though many researchers may not focus on performance, experts in the academic centers need to exist, otherwise everything will get farmed out to commercial cloud services.

NSF could formalize ways of having long-tail projects leverage commercial clouds. Clouds now offer continuity as well as abundant and (possibly) affordable compute and storage. Also, they provide ondemand access, whereas NSF and other academic research machines are oversubscribed. On-demand licenses for software are also available. Cost is shifted to end-user, so it comes out of their pocket whereas researchers see NSF resources as free.

NSF could give money to researchers rather than building centers. Fund an ongoing service with continuity in the model. Cloud business could be applied to some (smaller) applications, but it is not clear the economics are in favor of clouds.

Many software tools are available for a particular problem. Are they robust? Efficient? Will they be around in X years? It would be helpful if NSF focused on a key set of design and use patterns and focused software cultivation around those patterns. For example, in computational fluid dynamics, 12 large packages are available, but validation of codes is sometimes unclear.

Gap in ecosystem-funding for software quality. It is relatively easy to get money to start a code development project, but it is harder to sustain the code. Students may be interested in improving code at times, but may be dissuaded by advisors because it does not further their research/education goals.

There is a range of specialized resources needed to support high-end science: memory, I/O, etc.

Ongoing Requirements Collection

Currently, requirements are collected via NSF task forces and contact with resource users (surveys, calls, focus groups, audits, etc.). Although requirements are collected, around half of all computing applications receive allocations, and allocations are often much less than what is sought/required. It is notable that very few individuals from XSEDE or other service providers were included on the ACCI task force.

Motivating Science Team Readiness

Offering certification of computational performance literacy could encourage technical personnel and researchers to increase their skills and competency, particularly if allocation committees gave preference to certified researchers or if certified researchers received early access to new systems. Those with certification also would be expected to contribute to their research community's toolset.

Risks and Opportunities to U.S. Scientific Leadership

When only half the funded science projects can gain access to computational resources, science is being lost: Science projects do not have the ability to collect, store, analyze, share data, publish results, etc.

We suggest assessing the balance of the overall program so that science projects that are funded will have a reasonable likelihood of being able to access the resources (compute, data) they need to conduct their research.

An analogy is made to the way computing resources are decoupled from funding: scholarships for 100 doctors, but there is only lab space for only 10. That means 90 doctors will not be trained. The solution: increase lab space and decrease number of doctors to achieve better balance in the program.

Breakout 4 – Computation and Data Analysis

Charge: In reviewing the submitted white papers and other information, analyze the requirements for computation and data analysis for open science and engineering that will need to be met to enable high-spectrum science and engineering in the next decade. Provide a breakdown of the system and architectural requirements that are anticipated to enable high-spectrum science and engineering.

Attendees: not recorded

Summary

The group discussed the possibility of selecting two or three representative applications and deriving from them the requirements. However, the group concluded that the spectrum of applications and codes was too broad for this approach to properly capture the extent of the need. This situation reflects the diversity of scientific investigations being supported by high-performance computing today. It also reflects the very rapid rate of change in scientific computation that is being driven by the changes in computing hardware and the exciting new opportunities that those changes are opening up. Codes are in different states of change due to the simultaneous needs to produce today's research discoveries while also adapting to tomorrow's opportunities. What nevertheless was clear from the discussions was that the demand for high-end computing resources is strong across a broad range of science and engineering disciplines, and that the rapid developments in computing hardware at the high end will enable substantial advances.

There was universal agreement in the working group that the NSF should continue its policy of providing for the scientific community access to computing systems at the highest level of capability (NSF Track-1) as well as to a more diverse set of resources at the Track-2 level. The group felt that it could best illustrate the nature and breadth of this need by providing brief science cases. In addition to this set of examples, we include an extensive table that was compiled by the Blue Waters team from responses to a questionnaire by their science teams.

Before listing those examples, the group gives a few general observations that arose in our discussions.

First, it was clear to us that "one size will not fit all." We recommend that NSF continue its strategy of providing both a Track-1 system to enable researchers to carry out computations at the highest level of

capability and multiple Track-2 systems to accommodate a diversity of application types at a lower level of capability.

We agreed that at the present time an important trend in computing hardware, which has gained substantial momentum over recent years, is the development of many-core devices. These are single-chip devices containing large numbers of CPU or GPU cores running at clock speeds that do not generate excessive heat or consume excessive power yet nevertheless can deliver stunning performance to applications that are modified to exploit them. This is a trend that computational scientists cannot alter, because it is being driven by fundamental physical limitations as the feature sizes on silicon chips decrease to extreme levels. Although computing systems that do not incorporate these new generations of high-performance devices would place fewer demands on science teams for code modifications, the group agreed that ignoring these trends would expose our community to the risk of very substantial and important missed opportunities. Our group therefore agreed that adapting scientific research codes to exploit these new devices is well worth its cost, and that future Track-1 systems in particular should definitely incorporate such devices. This led us to conclude that it will not be enough to simply put in place a new Track-1 system containing these devices, but that NSF must also invest in the very necessary research code modifications that such a system will require if the full promise of its hardware is to be realized.

Science Case studies

Global climate studies (Donald Wuebbles, University of Illinois)

The goals in this work are projections for future climate with increased accuracy and confidence. Present work on Blue Waters uses 25 km grids (1/4 degree), which allow representation of hurricanes and local impacts. However, the accuracy of these simulations would dramatically improve if clouds could be resolved rather than entering the simulation only through model parameterizations. This requires grids with 4 km to 10 km resolution. The simulations need to be carried out covering 150 years of the past for validation purposes against historical data and then for 100 years into the future. The ocean component of these simulations needs to be treated at 0.1 degree resolution to capture and accurately describe important eddy scales that contribute to phenomena such as the Gulf Stream. Codes presently being run on Blue Waters incorporate a new spectral element dynamical core that allows them to scale well on the machine. The needed increase from 25 km to 10 km or 4 km resolution enters the computational cost in the third power, so that an increase of a factor of 15.6 to 216 results. In principle, the number of CPU cores that can be used can also scale up by this same factor from the size used in present runs, namely 64,000 cores.

Atomic, molecular, and optical physics (Klaus Bartschat, Drake University)

The goal of this work is *ab initio* solution of the quantum-mechanical few-body problem, including chargedparticle and (intense) laser interactions with atoms and molecules. This is basic research looking at quantum mechanical correlations. Such calculations support, for example, very expensive experiments on coherent control, i.e., attempts to affect the outcome of chemical reactions by stirring the electrons (to the extent allowed by quantum mechanics) as well as plasma physics (atomic and molecular data for modelling of things like ITER). There are also many applications in materials science and even medical physics (e.g., radiation damage by light or charged particles). One needs to solve the time-dependent or timeindependent Schröedinger equation, depending on whether there is an explicit time dependence or a steady-state situation. This involves diagonalization of large matrices, solving large systems of linear equations, or calculating the effect of exp{-i H \Delta t} on a known vector \Psi(t) to move it to \Psi (t + \Delta t). Problem sizes vary, depending on how much physics is included. Often, the accuracy of the results is limited by the space-time grid that can be handled, the number of coupled channels that can be included, or approximations made (e.g., fixed-nuclei approximation, even though we know that they move). Ever larger calculations are needed to increase the accuracy and to allow for reasonable uncertainty estimates of the theoretical predictions. This group is using Stampede and Gordon; others in the community run on NERSC and ORNL machines. A lot of activity also occurs outside of the U.S., using supercomputers in Europe, Asia, and Australia. This group has not applied to Blue Waters (yet), but is thinking about it. The currently largest calculations for electron collisions with complex atoms and ions are based on the R-matrix method to solve the so-called close-coupling equations, with the rate-limiting step usually being the construction and subsequent diagonalization of a dense, real-symmetric matrix, for which *all* eigenvalues are needed. Such calculations currently require about 2 million SU on Stampede. Scaling these codes to achieve higher accuracy and/or to handle more complex systems is challenging, but necessary for the theory to continue to guide experimental studies in this area.

Industrial generation of clean coal power (Martin Berzins, University of Utah)

Design of the next generation of industrial clean coal boilers with a capacity of, say, 5,000 cubic meters, requires the solution of models with 10**13 cells and 10**14 variables. Such models involve turbulent combustion, thermal radiation, and treatment of coal particles. Such a boiler would be a more efficient means of burning coal in a way that makes carbon sequestration possible. Such a model is about a factor of 1,000 larger than the models that can be solved today. It is challenging to compute the solution for even a fraction of a second. The global connectivity through both radiation and the solution of linear systems poses formidable challenges for machines with peak performance of hundreds of petaflops. In addition, the need to quantify the accuracy of the design requires the use of UQ (uncertainty quantification) techniques that in turn multiply the computing costs by two orders of magnitude beyond those of a single simulation.

Materials science

The goal of this work is accelerating targeted materials discovery. The computational requirements are driven by the need for fidelity in prediction of new materials properties (composition, structure, thermal stability). The simulated time duration depends upon the system complexity and collaboration with experiment. The number of simulations depends upon the number of atoms per cell and temperatures. A real space genetic algorithm is used that involves density function theory (DFT). This has been run on the ORNL machine with up to 100 atoms per cell at various high temperatures searching for materials with large magnetic anisotropy. The motivation is to identify and produce strong magnets whose properties rival or exceed the present rare earth magnetic applications. China presently controls the rare earth magnet economy. The algorithm is general and can be used to search for superior materials/properties. The present proof of principle results can be scaled to larger numbers of atoms and a wide range of strategic materials studied, where additions or substitutions of other elements can improve properties or lower costs.

Cosmology (Thomas Quinn, University of Washington)

The goal of this work is to determine the impact of baryon physics on the distribution of dark matter in the universe. These impacts are felt through processes such as star formation and supernova feedback on developing structures, such as galaxies and galaxy clusters, in the universe. To represent these processes in a simulation, resolution of about 10 parsecs is required. However, the volume of the simulation must be tens or hundreds of Mega-parsecs, thus at least a million times larger. The constraint on the simulation volume comes from the need to properly treat the 10 Mega-parsec minimum gravitational influence range in the problem. A spatial volume of 100 Mega-parsec across gives good statistics in this sort of simulation. We can therefore think of a single such simulation as producing a statistical ensemble of individual situations of galaxy formation, evolution, merger, etc., with the feedback from star formation and supernova explosions included. Such simulations will help us to determine what the dark matter in the universe is. We observe the effects of its gravity, but because we do not observe it in other ways, we cannot

yet determine its other properties. In the simulations, we can compute in detail the interactions between the dark and luminous matter using different potential dark matter models. By comparing the results to observations of the actual universe, we hope to be able to rule out some of these models and thus to increase our understanding of the nature of the dark matter. Simulations are now being run on Blue Waters with 100 parsec resolution within volumes of 25 Mega-parsec across. These simulations are run to a tenth of the age of the universe, which each requires about 3 million node-hours on Blue Waters (~100 million core hour equivalents.) To go to 10 parsec resolution will require a 10,000x increase in computing power. A new implementation of the most costly computations on GPUs is already known to be able to deliver a factor of 10 in the code performance. Use of the Charm++ runtime system for load balancing and overlapping of communication and computation in the simulations is expected to deliver another performance improvement. Thus the goal of this work is expected to be within reach of a next-generation Track-1 computing system.

Vaporization of liquid droplets in a turbulent flow (Said Elghobashi, University of California, Irvine)

The goal of this work is to accurately compute through direct numerical simulation (DNS) of the Navier-Stokes equations the interactions between thousands of vaporizing droplets and a turbulent flow. By fully resolving the droplets, it is possible to obtain first-principles computation of the effects of turbulence on the rate of vaporization and, in turn, the effects of vaporization on turbulence. This fundamental information is not known at present. The results of simulations of this type will greatly improve the understanding of phenomena such as cloud formation and liquid fuel combustion in engines of automobiles, aircraft, and rockets. It is expected that the improved understanding of the physics will lead to better control of pollutant formation and enhanced engine efficiency. The simulations are being carried out at present on meshes of 2,048³ points with 1,000 droplets at Reynolds numbers R_{λ} , based upon the Taylor length scale, of 250. A run takes about 20 wall hours on 100,000 cores. The goal for this work is to increase the mesh resolution to 4,096³ points and the number of droplets to 6,000. These finer resolution runs would reach Reynolds numbers, R_{λ} , of 800. Such a simulation with higher Reynolds number and larger number of droplets would approach situations in real engines. The higher Reynolds number allows the resolution of the small length- and time-scales of turbulence (adjacent to the liquid droplet surface) that are responsible for transport and mixing of the droplet vapor with the surrounding gas. Increasing the number of droplets per unit volume is essential for achieving a realistic fuel/air ratio observed in combustion engines.

Stellar hydrodynamics and nucleosynthesis (Paul Woodward, University of Minnesota)

The goal of this work is to simulate the complex processes in stellar interiors that determine the abundances of heavy elements that are produced and expelled into the surrounding medium. These processes determine the composition of later generations of stars and their planets, hence the chemical evolution of galaxies. Of particular interest are events in which it is possible that convection zones driven by the nuclear burning of different elements can merge as a result of material mixing at convection zone boundaries. Such events could have a profound influence on nucleosynthesis and on the state of a massive star just before it explodes, thus influencing the explosion as well. These simulations require that multiple burning shells be included in the simulations, and that a large number of reacting species be carefully tracked with sufficient resolution and accuracy. The simulations must also extend over larger times than those we do today, if the merger of burning shells is to have enough time to take place. Simulations today on Blue Waters run on over 440,000 cores at 0.42 Pflops/s sustained performance. Such a simulation, describing the ingestion of fresh fuel into a single convection zone, has been run for about four days on Blue Waters on a grid of 3.6 billion cells. To include, for example, both the oxygen and carbon burning shells in a massive star in a single simulation that resolves the inner oxygen burning shell on 3.6 billion cells, even using AMR (adaptive mesh refinement) we will need to at least triple the grid size. Running the simulation

for sufficient time to observe a shell merger, even with special techniques now under development, will involve at least a tripling of the cost of the simulation in computing power applied. Overall, we envision simulations of this type requiring roughly 10 times Blue Waters' present capability. This simulation code now runs well on Blue Waters' AMD processor nodes, and it has been tested on Stampede running on the Intel Xeon Phi and Sandy Bridge processors combined at about twice the per-node performance of the AMD nodes in Blue Waters. This code is now being implemented to run on Blue Waters' GPU nodes, on which it is hoped that the code performance will increase, on a per-node basis, by a substantial factor, although of course the number of such nodes available will then be smaller by a factor of five. This simulation code is poised to take advantage of the anticipated improvements in accelerator technologies from Intel and NVIDIA that give our best present idea of the hardware that might be incorporated in a future NSF Track-1 system.

Lattice QCD (Paul Mackenzie, Fermilab)

One of the ways that particle physicists search for evidence of physics beyond the standard model is to search for deviations in the relations between the parameters of the standard model (such as those governing the mixings of the quarks) from what the standard model predicts. The precision of some of these searches has been blurred by the fact that lattice QCD calculations are not yet precise enough to analyze the data properly. For example, the experimentally observed mixing of neutral B mesons with their anti-particles has been measured to a tenth of a percent accuracy. The lattice QCD calculations needed to derive quark mixing parameter relations from the experiments have only achieved order 1 percent accuracy so far, although they have used hundreds of millions of core-hours and have required up to 16 racks of the Blue Gene Q Mira. Over the next five to 10 years, with improved computing resources, lattice calculations will be improved to the required precision, improving the reach of these expensive searches for beyond-the-standard-model physics by an order of magnitude.

Materials physics (Shiwei Zhang, College of William and Mary)

The goal of this work is to determine the quantum mechanical behavior of a large number of interacting electrons, which is one of the grand challenges of modern science. The solution of such "many-electron problems" is important because electrons determine the physical properties of materials and molecules (e.g., hard or soft, reactive or inert, conducting or insulating, superconducting or magnetic, etc.). This challenge spans multiple disciplines, including condensed matter physics, chemistry, atomic and molecular physics, materials science, and biology. The rapid recent advances in theory and computational methods mean that a significant increase in computing resources would enable breakthrough progress toward the solution of these problems. An example would be a fully *ab initio* calculation of the simplest high-temperature superconductor to allow detailed comparisons with experiment, which would be a landmark for the materials genome initiative (materials design via predictive computing). Because of the breadth and diversity of these problems and the fundamental and technological importance, NSF must play a leading role in their support.

Solidification (Richard Arthur, GE Global Research)

This work simulates water molecules forming ice in the presence of a surface. A water droplet of about 1 cubic millimeter containing about 30 billion atoms with 300 local interactions per atom (assuming a 10 angstrom cutoff) is followed for a billion time-steps to the formation of ice. This requires 9,000 exaflops^e per case. A real water droplet contains roughly one mole of atoms, and is therefore far beyond the scale of

^e This is over 8 days running with perfect 100% efficiency on the entire Blue Waters system, but with more standard efficiency of 10% it is 80 days or 160 days at 5% of peak

the simulation just described. A realistic environmental minimum timeframe is 1 ms. Freezing can be coerced in 1 microsecond under pressures and magnetic fields (for experimental validation of the models). Time-steps in this simulation are 1 femtosecond, and hence the need for a billion time-steps. The coarse/fine time-step correction methods for time parallelization do not work in these kinds of problems because the nucleation events are stochastic and rare. We do employ a technique from LANL called Parallel Replicas to launch multiple universes and opportunistically harvest data from scenarios where the nucleation occurs earlier (with strong statistical basis to then map that to a generalized regime).

Full-wheel aircraft engine (Richard Arthur, GE, and Todd Simons, Rolls Royce)

The design of aircraft engines could be improved if we could accurately simulate the fluid flow and combustion processes inside them in complete detail. A present design goal is creating cooling laminar flows in order to keep the operating temperatures from melting the metal alloys. A full-wheel engine simulation that would address this particular design goal would need to include about 1,000 blades, with 5 million mesh points per blade, including the cavities in the blades that allow cooling air intake. Each mesh point would need to carry nine variables and be advanced for about 500 time-steps. The simulation would involve fluid dynamics, combustion, and some treatment of the fluid-structure interaction.

Provide a Breakdown of the System and Architectural Requirements that are Anticipated to Enable High-spectrum Science and Engineering

While it is possible to speculate about the form of future machines post 2017, the architectural uncertainty makes such speculation somewhat error-prone. A major challenge for the future is to provide platform portability across different "swim lanes." At this point, there appear to be two broad architectural paths, or "swim lanes," for future machines. One is based on the idea of a hybrid system consisting of multi-core CPUs and many-core GPUs. The other is an entirely multi-core/many-core CPU system. In both cases, systems can be thought of as consisting of processors having O(100-1000) cores/chip. Other architectural changes are also anticipated, including a deepening memory hierarchy, consisting of features like so-called High-Bandwidth Memory, which is high-bandwidth, (relatively) low-capacity on-chip memory; use of solid-state memory (e.g., burst buffers), and so forth.

Many code teams are, at present, expending significant effort to enhance and modify their codes so they perform well on present-day systems, which include both types of systems. Meanwhile, R&D in the computer science and computational science community focuses on methods for enabling applications to achieve both a high degree of performance and portability across different architectures. The idea is to present capabilities in the form of "processing motifs," often implemented as "domain-specific languages" (DSL), where an application expresses its algorithm using DSL constructs, which are then reduced to device-specific code and managed by advanced runtime systems, which take care of scheduling, resource mapping, and so forth. Such approaches are coupled with ideas such as runtime systems to enable the code generated by the DSL to execute across large machines. Such areas are being very actively explored in the DOE community as a way of dealing with architectural uncertainty. Many in the community believe this approach to be viable, and preferable to the type of highly labor-intensive efforts needed to make a given code work on a given platform, especially when considering that such effort would be multiplied by some number of platforms and across code projects.

Although the DSL approach lowers the programming burden in adapting to new high-end computing platforms, it cannot lower that burden completely. In the not-so-distant past, codes required substantial rethinking and restructuring to incorporate MPI messaging in order to adapt to distributed memory (MPP) systems, and later they required still more rethinking and restructuring to incorporate OpenMP directives (a form of DSL) in order to adapt to large shared-memory multiprocessor nodes. There is no question that

investments in the new generation of tools and DSLs is required, but experience indicates that a corresponding investment directly in the code teams to restructure their codes is also necessary.

Analysis Issues

- Changing balance between traditional *post hoc* and *in situ* analysis. With a widening gap between rates of computation and I/O, it will be increasingly likely that more code teams will, by necessity, need to perform an increasing amount of data analysis/processing *in situ*, rather than writing data to persistent storage for subsequent *post hoc* analysis. *In situ* and *post hoc* analysis is not necessarily mutually exclusive: some forms of *in situ* processing can produce reduced-size representations of data that are then used in *post hoc* analysis. Examples include feature detection, statistical summaries, data subset selection, compression, and so forth.
- Analysis pipelines grow increasingly complex in response to increasingly complex scientific data, increasingly complex lines of scientific inquiry, and the increasingly complicated computational environment. It is likely there will be an increasing reliance on new, emerging "workflow systems" to automate such processes, which include capture of provenance to facilitate reproduction, robustness, encapsulate complexity.

Tables from the Blue Waters White Paper

Science Category	Goal	Fidelity	Simulated Time Duration	Number of Simulation Problems	Data Integratio n and sharing	Algorithm & Workflow Re- engineering
Astronomy, Astrophysics and Space Science Best of Breed :	First 3D multi-physics general-relativistic magnetohydrodynamics simulations of core- collapse supernovae. Detailed microphysics: nuclear equation of state, neutrino heating/cooling.	~400 3D grid functions on 1 billion mesh zones, finely resolving magnetorotationally -driven outflows out to 2000 km. Simulation volume: 10^4 km	0.5 seconds 1 day Several hours	MHD box 5200'2880' 2800 AU cubed	Yes	In progress for GPU-CPU hybrid operation.
	Simulating a full star (radius ~10 ³ km); initial ignition point (r ~ 2 km); effective resolution 36,864 ³ (135 m/zone) ^f ; Full 3D general relativity; GR magneto-	1,536 ³ cells for 5.7 million time steps ^g . Simulating a full star (radius ~10 ³ km); initial ignition point (r ~ 2 km); effective resolution				Adaptive Mesh Refinement 5- 10 levels of AMR

Table 1. Example Science Area Drivers and	l Requirements over Time
---	--------------------------

^f Chris Malone's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^g Paul Woodward's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

	hydrodynamics with detailed nuclear equation of state and neutrino heating/cooling Several hours of simulated heliosphere time with full kinetics	36,864 ³ (135 m/zone) ^h ; Full 3D general relativity; GR magneto- hydrodynamics with detailed nuclear equation of state and neutrino heating/cooling Helio Science – fully kinetics – 10 ¹⁰ cells, 4x10 ¹² particles; Large Scale Hybrid Kinetic – 1.7x10 ¹⁰ cells, 2x10 ¹² particles ⁱ . Coupled continuum and kinetic equations – 10 ¹⁰ particles; local effective resolution on the MHD grid 5,000 ³ for a very large computational				Hybrid parallelization, optimization of the source term calculations in Monte Carlo simulations, improved load balancing for Monte Carlo algorithms
Biophysics and Biology Best of Breed :	All-atom molecular dynamics – e.g virus interaction, cell organelles (ribosomes, chromataphore) Ultra Coarse Grain Simulations – e.g. Virus interactions Multiple Copy Simulations ⁱ	box. 100M Atoms ^k 1M Atoms	Micro seconds	2 to 4 1	Yes	Higher Order PME interpolation Ultra Coarse Grain
Chemistry Best of Breed:		Systematic, predictive, but extremely expensive ("ab initio") methods	Interested in static properties of materials			The computational cost of density functional theory (DFT) is routine for

^h Chris Malone's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

ⁱ Vadim Roytershteyn's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^j Phillips, Roux and Schulten

^k Klaus Schulten's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

				-		
						many types of solids and is so enormous it is necessary to re-engineer the algorithms.
Climate and Atmospheric Sciences Best of Breed :		1/4° Atm, 1° ocean ¹ Radiative heating rates for real (i.e., 3D) clouds and their impact on dynamics. ^m	~250 years	1-4 for climate, 10 for others	Yes	
Geophysics/ Seismic Best of Breed :	Dynamic Rupture Simulation to understand physics of Earthquake Rupture Process	5-Hz simulation in small, fault- oriented, volume. 150kmx100kmx40k m 82K time steps	100 sec	5-10 runs; 37- billion mesh elements	Yes	Use Dynamic rupture simulation to create inputs for high frequency wave propagations
	High-frequency Scenario Earthquake Simulation	Wave propagation run: 2-Hz, 810x405x85km, 40m spacing, Min. S- wave Velocity 400m/s, 160K time steps	250 sec	2-4 runs 435 billion mesh points		Heterogeneous computing on Blue Waters
	Regional Probabilistic Seismic Hazard Analysis	1144 seismic hazard sites at 0.5 Hz UCERF 2 700,000 rupture variations used to simulate 500 million 2-component seismograms ⁿ		1.2 billion points/mesh per wave propagation simulation		Ground Motion Probability Models replaced with physics-based ground motion models; Seismic Reciprocity. 100 million jobs

¹ Don Weubbles' *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^m Larry Di Girolamo *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

ⁿ Tom Jordan's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

						12 TB data products
Fluids and Turbulence Best of Breed :	Turbulent Flows Ice nucleation in turbulent domains	8,192 ³ cells ^o (over 0.5 trillion grid points) 1M Atoms droplet	5-7 large- eddy turnover times	1 (over 3 years)	Yes - Sharing reference simulation data sets	Reducing Communicatio n costs and implementing many-core methods
Materials Science Best of Breed:	Transistor Materials Condensed Matter/ Large Scale Quantum Simulations ^p	100 Million atoms; 1, 2, 3D; any crystal; Spin and Classical Multi-Physics; Transport; Strain MVFF ^r , problem specific assumptions 10,000 atoms				DFT Monte Carlo
Particle Physics Best of Breed:	Quantum Monte Carloq	1,000 atoms HISQ gauge – 96 ³ ×256 Clover quark propagators - 32 ³ ×256 ^s		600 Configuration s 485 configuration s	Yes	
Social Science, GIS, Economics Best of Breed :		260,000 agents – US Scale ^t	1 to 2 Months	1 scenario	Yes	Agent based methods

^o P.K. Yeung's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^p Jerry Bernholc

^q e.g David Ceperly's work

r Gerhard Klimeck's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^s Steve Gottleib's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

Table 2: 2017-2018 Best of Breed and Co	ommunity Normal Science Drivers
---	---------------------------------

Science Categories	Goal	Increased Fidelity	Increased Simulated Time	Number of Simulation Problems	Data Integration and sharing	Algorithm and Workflow re- engineering
Astronomy, Astrophysics and Space Science					Yes	
Best of Breed:	Star Ignition	Simulations of local small-scale processes (magnetic reconnection, plasma instabilities, turbulence, etc.) in the framework of large-scale global MHD simulations.	10x increase – 1 day of simulated time	10x simulations		Optimized discrete event simulations; code-code coupling; semi-implicit 3D PIC with high order particles; development of accurate subgrid models
	Helio Science – Turbulence and global modeling	Currently, small- and micro-scale phenomena are addressed only locally, which makes difficult to predict their effect on the global solution.	Requires 104 more computing time than in the best contemporaneous simulations.	Total volume is not increasing, but the grid is 10 times finer in all directions locally.		AMR and load balancing improvements; I/O issues for very large data files (40-50 TB).
Community Standard:						
Observational Support:		Voyager and Interplanetary Boundary Explorer missions.				
Biophysics and Biology Best of Breed :	All-Atom Molecular Dynamics Ultra Coarse Grain Simulations – e.g. Virus interactions	Multiple Cellular Structures- 200- 300M all atom Polarizable force fields and O(N) long range dispersion forces	10 microseconds			O(N) electrostatis
Community Standard:	Multiple Copy Simulations ^u	10M Atoms	Microseconds	10s of problems		Streaming Data Flow Analysis

^t Keith Bisset's *Blue Waters 2014 Symposium* Presentation (http://bluewaters.ncsa.illinois.edu/symposium-2014-schedule-descriptions)

^u Roux and Tajkorshid

Observational Support:		Cellular Structures of 100M all atom Large all to all comparisons				
Climate and Atmospheric Sciences Best of Breed :		Climate 1/8° (12 km) Weather .1 km IPCC - Climate 1/8°	250-500 years 500 years	3 to 4	Yes	Atmospheric chemical tracers, indirect radiative balance, nested grids for regional impacts
		(12 km) Climate 1/4° (25 km	1,000 years	20-30		
Geophysics Best of Breed:	Physics of Earthquake Rupture Process	Dynamic rupture simulations: 100x50x50km, 2.5m spacing, 16 trillion mesh points, 200K time steps	150 secs	2-4 32 trillion points/mesh	Yes	Scalability and load rebalancing of discontinuous mesh algorithm. Wave propagation simulation is run longer than dynamic
	High Frequency Scenario Earthquake Simulation	Wave propagation simulations: 5-Hz, 810x405x85km, 10m spacing, min. S- wave Velocity 250m/s, 640K time steps	250 secs	80 billion points/mesh per wave propagation simulation		rupture simulation to allow waves to propagate across full simulation volume for wave propagation simulation
Community Standard:	Regional Probabilistic Seismic Hazard Analysis	3500 California statewide seismic hazard sites at 2.0 Hz using UCERF3 earthquake rupture forecast An estimated 23 million rupture variations used to simulate 50 billion 3-component seismograms Improved physics (frequency- dependent attenuation, fault roughness, near- fault plasticity, soil non-linearities, near-surface heterogeneities)	250 secs	3,500 hazard curve calculations		9.5 billion jobs 7.2 PB data products

Fluids and Turbulence Best of Breed: Community Standard:		8,192 ³ simulations cells 4,096 ³ simulations cells	Increased	4-5 10-20	Yes	Increased complexity, multiple parameters
Materials Science Best of Breed :	Systematic, predictive ab- initio methods					Methods for stochastic formulation and algorithms Fragment methods
Particle Physics Best of Breed :	Fully dynamical electromagnetic effects	HISQ gauge – 128 ³ ×256 Similar increase for Clover quark propagators and DFW actions	10x to match precision of latest experimental devices at Fermi Lab	1000 configurations		Inclusion of electromagnetic and isospin breaking effects
Community Standard:	Calculation of physical quantities of above configurations	(14*3) x increase in computational requirements				Improved inversion of sparse matrix
Social Science, GIS, Economics Best of Breed: Community Standard:		1,500,000,000 agents + increased agent sophistication – largest country	1 to 2 Months	1 to 2 scenarios	Yes	Dynamic Workload Balancing for scale
		260,000,000 agents - US scale	1 to 2 Months	100 scenarios	Yes	

Table 3 2020-2021 Best of Breed and Community Normal Science Drivers

Science Categories	Goal	Increased Fidelity	Increased Simulated Time	Number of Simulation Problems	Data Integration and sharing	Algorithm and Workflow Re- engineering
Astrophysics and Space Science Best of Breed :	Helio Science	Numerical simulation of the turbulent, multi-ion solar wind flow from the solar surface to the	The physical times to be covered do not increase, but the grids will be larger, to cover	The physical volume is the same. The problem will be in the division of the	Yes	Using new supercomputin g architecture will be essential to implement this

		Earth's location and its further interaction with the interstellar medium. This will be necessary for the Solar Probe Plus mission to be launched by NASA in 2018. Currently, only some pieces of this approach are implemented, sometimes only locally	different length scales involved in the simulation.	computational volume into sub- volumes to make the simulation feasible.		challenge because of the additional physics to be added.
Biophysics and Biology Best of Breed: Community Standard:	All Atom MD – e.g virus- cellular interactions	1-10 billion-atom virus-cellular interactions Multiple cellular structures – 200 M	1-10 Milliseconds Microseconds	2-4 10s to 100s	Yes	Full O(N) quantum mechanics methods for 1M atoms
Climate and Atmospheric Sciences Best of Breed :	Properly calculate radiative heating rates dynamically within weather and climate models	Atoms Climate – 1/16° (6 km) 1/8° (12 km) seasonal prediction of water resources, ultra fine-grain scale to tropical cyclones)	250 years 1,000 years	3-4 20-30	Yes	Nested grids and AMR
Community Standard:	Remote sensing problems for retrieving cloud and aerosol microphysical properties in the face of 3D radiative transfer Water Cycles, Ice/Cloud integration/Regional Severe Weather Prediction for Climate/{Carbon, Methane, Nitrogen} Cycles ^v	Weather Prediction	1 day 15 day	Many Regions		

^v http://extremecomputing.labworks.org/climate/reports/ClimateReport_6-24-09.pdf

	1					
	Severe Weather – 10m resolution					
Geophysics Best of Breed :	Physics of Earthquake Rupture Process	Dynamic rupture simulation: 100x50x50km, 1m spacing, 0.25 quadrillion mesh points, 1 million time steps	250 secs 250 secs 250 secs	2-4 runs 1.2 quadrillion points/mesh	Yes	Adopting high precision AMR code for new programming model, load balancing, <i>in-</i> <i>situ</i> analytics, fault tolerance
	High Frequency Scenario Earthquake Simulation	Wave propagation simulation: 10-Hz, 810x405x85km, 3m spacing, min. S- wave Velocity 150m/s, 2.56 million time step		2-4 runs 1.2 trillion points/mesh per wave propagation simulation		
Community Standard:	Regional Probabilistic Seismic Hazard Analysis	20,000 California statewide seismic hazard sites at 5.0 Hz using UCERF3 earthquake rupture forecast 23 million rupture variations used to simulate 285 billion 3- component seismograms Dynamic rupture models		20,000 hazard curve calculations		54 billion jobs 102 PB data products
Fluid Dynamics Best of Breed:	Ice nucleation in turbulent domains	16,384 ³ simulations cells				Long-term strategy to keep multi- Petabyte data collections
Community Standard:		Realistic droplet – 1 Trillion atoms				LES and DNS integration
Materials Science Best of Breed :	Multi-dimensions maps for Material Design					
Particle Physics Best of Breed:	Increase in 10-100x to keep pace with planned experimental fidelity	128 ^{3*} 256 HISQ gauge configurations with physical quark masses and lattice		1000 configurations		

	spacing				
Community Standard:	Similar increase for Clover quark propagators and DFW actions				
	(16*3)x increase in computational needs				
Social Science,				Yes	
Best of Breed:	10,000,000,000 agents + increased agent sophistication – world wide	1 to 2 Months	1 to 2 scenarios		Dynamic Workload Balancing for scale
Community Standard:	1,000,000,000 agents	1 to 2 Months	100 scenarios		

Breakout 5 – Storage & Data Analysis

Charge: In reviewing the submitted white papers and other information, analyze the requirements for storage and data movement for open science and engineering that will need to be met to enable high-spectrum science and engineering in the next decade. Data movement is expensive both in terms of energy, investment and human effort. Identify strategies to minimize data movement to make science and engineering teams as productive as possible. Discuss the current number of storage hierarchies (memory, rotating disk, tape) and project the types and levels of hierarchies that will be likely and effective over the next decades.

Attendees: not recorded

Science Case Studies

Cosmology

Typically a small number of groups produce "legacy" simulations, which are then analyzed by a much greater number of data consumers. Currently, the typical scale of this activity is each year one group produces a 1 petabyte data set which is then analyzed over five years by hundreds of users. A prototypical example is the Millennium simulation, run in 2005 and the basis of about 1,000 papers.

How is data analyzed: Post-processing does not require an HPC allocation; it can be done locally on workstations or high-throughput (HTC) resources.

Key issues:

- 1 BIG simulation/year: 80-100 million core-hours producing 1 PB of data of community interest
- 1 PB/year of raw data/year for community collaboration: allowing up to 1,000 users to analyze
- Biggest challenge is the sharing issue (not the producing and storage of the data)
- online close to processing
- shared nothing, MapReduce cluster
- analyze data in parallel

• Virtualization to allow researchers to deploy their own analysis & related tools

General Issues:

- Preservation: What happens to data (particularly BIG data) when your allocation is over?
- Where/how do large instruments store/share data and preserve it after funding ends? Example is the astronomy community (see below)
- Need to analyze BIG data on demand (not batch) data from simulations & instrumentation

Lattice Field Theory

Currently, the largest data-sets for lattice field theory are on the order of 250 TB. These consist of 15 to 20 ensembles of about 1,000 files (configurations) each. Individual file sizes range from 60 MB to 15 GB. These files are accessed by a number of groups for additional analysis in which particular physical observables are studied. While the creation of the configurations requires leadership-class computing, the additional analysis can be done on computers of moderate peak speed but considerable capacity. The total requirement for cycles for analysis jobs exceeds that of creation of the configurations. The analysis jobs may be done far from the computer on which the configurations were created. Archiving and fetching the configurations for further analysis requires considerable human effort, and it would be good to have better systems for automating this. For this task, Globus Online has been a great improvement. The analysis phase may last for up to five years after the configurations are generated, so it is important to have a mechanism for storing the configurations after they are generated. We expect that future needs will grow by not more than a factor of 16 over the next five years.

Experiment	2013	2020	2030+
Storage	1PB	6 PB	100-1500 PB
Cores	10^3	70x10^3	300+K
CPU Hr	3x10^6	2x10^8	10^9

Astrophysics and Cosmology "

Simulation	2013	2020	2030+
Storage	1-10 PB	10-100 PB	>100PB-1EB
Cores	0.1-1x10^6	10-100x10^6	>10^9
CPU Hr	200x10^6	>20x10^9	>10^11

Nuclear physics

Our calculations are typical large-scale calculations, producing nuclear wave functions that can be used for further post-processing by a small number of collaborators. Currently we produce a few hundred wave

^w Snowmass Computing Frontier: Computing for the Cosmic Frontier, Astrophysics, and Cosmology, *arXiv:1311.2841*

functions of the order of up to 100 GB each a year that we want to store, so we need tens (and in the future hundreds) of terabytes of archival storage. Post-processing generally happens on the same platform where the wave functions are generated. The useful lifetime for post-processing these wave functions is typically two to three years.

Genomic data

For genomic data, such as The Cancer Genome Atlas, the raw data is today less than 100 TB. The content can be captured in persistent databases that are constantly online next to adequate computing resources. Some centers can provide assurance that their policies are consistent with HIPAA requirements, although they may allow access only to certain users.

Materials genome

Through the Materials Project, the materials research community shares an extensive repository of data on known and predicted materials.

Astronomical Research

Astronomical research includes modeling and simulation, data analytics, and image processing. Many of the problems include automated feature search that compares past images to recent ones, looking for relevant new features. This area, that merges computation and data analysis, has grown in the last decade with projects such as the Sloan Digital Sky Survey, the Palomar Transient Factory and the Supernova Factory.

The efforts have some common characteristics, including quasi-real time image comparison, feature extraction, complex workflows and data movement, and longer term post analysis. In some cases, the image search is also combined with modeling and simulation in order to properly detect and classify transients. While there are many images that need to be moved and analyzed, each file is typically modest size and is more a constant data stream than bursts. Similarly, if the data can be kept online near an adequate analysis engine, the resulting data products are often very small and can be easily moved anywhere over the Internet.

Climate Science

Climate scientists are drinking from the proverbial fire hose—more data are being produced than can be adequately analyzed. Historical observational data (dominated by satellite measurements) already totals \sim 30 PB, and it is likely that \sim 1 PB per year will be added. These data are stored in a variety of locations around the world and are essential to help analyze the model results discussed below.

Every six to seven years, there is a set of major climate modeling analyses done as part of the preparations for the next Intergovernmental Panel on Climate Change (IPCC) assessment. These model intercomparisons, called CMIP, are taking ever-larger amounts of storage and are analyzed by several thousand atmospheric scientists, plus by many researchers from other disciplines (hydrology, agriculture, ecosystem, health, etc.). Note that there are other important climate model runs done outside of CMIP also being analyzed by many others

CMIP-5, completed in 2013, has \sim 0.4-0.5 PB. CMIP-6 is the next phase and will be completed around 2018-2019. The data output will likely be \sim 5-10 PB.

As a lead-in to CMIP-6, the research team led by Donald Wuebbles at the University of Illinois is doing runs with NCAR's Community Earth System Model on Blue Waters at ~25 km resolution (0.25 degree horizontal resolution around the planet) for 150 years in the past and 100 years in the future for multiple ensembles for one assumed (human-related) emissions scenario. This produces 1-2 PB of data that we and many others will want to analyze over the next few years. These runs are unique because no one has run at such a

long time with a global model at such a high resolution where one can capture severe storms like hurricanes with reasonable accuracy and also be able to do impacts studies at the local level across the planet. Most major climate model runs to this point have been done at 1 degree (100 km) or coarser resolution.

Data analysis and visualization are essential and require a large number of people analyzing the data because of the many types of complex queries and analyses being done. GIS and visualization tools are essential to the use of these analyses in a large variety of adaptation and mitigation policy decisions.

Biomolecular simulation

The rate-limiting steps in many biomolecular simulation experiments are no longer simply generating the data but simulation set-up and, more importantly, data analysis to derive information from the simulation results that helps to guide in understanding the models and implications of the results. Simulations are getting longer (representing longer time scales of sampling; μ s to ms), are getting bigger (10,000 to millions of atoms), and increasingly ensemble-based, leading to larger and larger datasets. Analyzing molecular dynamics trajectory data (i.e. the time course of atomic movements), and derived properties from this data, is an art and is approached in many varied and different ways. Also, data may have different uses for different people; whereas one group may be interested in modeling a biological process of interest, such as protein folding, another group may want to mine the data to guide the development of coarsegrained force fields, whereas another may want to probe structure correlations or other properties. Although many analysis tools and codes exist and are under active development, a priori, users must explore the data to "see" what the data reveals, and this exploration can be costly. With hundreds to thousands of ensemble instances, and time scales approaching μ s to ms, it can take hours to days to go through this data. Although the data can be reduced, how this is reduced and coalesced depends on what you want to learn from the data and what you are willing to omit, per analysis run. As the data is rich, and can have multiple uses, and often require significant computational resources to generate, it is worthwhile to consider making this available to a larger community to explore, mine and to attempt to derive new information. The scale at which this is done will critically depend on resource availability. Ideally:

- Interactive / *in-situ* / on-demand analysis being able to process the raw data and derived properties and to be able allow users to inspect the data in a low-latency fashion. For example, to correlate one property with another, visualize, then alter the lag or running averaging time or other parameters to see how data is altered, then visualize again without have to go through the raw data a second time.
- Parallel analysis of large data sets (likely with lightweight nodes)
- Means to pull derived properties from the data and to do further simulation work, for example, clustering conformations sampled and then instigating long time-scale molecular dynamics (MD) simulations on highly populated conformations to get kinetics information on large-scale GPU resources or using high-level QM methods in big memory nodes to analyze representative conformations from the clusters.
- Means to pull information from large sets of MD simulation data, for example to pull out all Ala-Gly-Pro peptide trimers sampled across a set of MD simulation to infer information (potential Hadoop/Spark). Requires development of annotation standards, means to identify the data, means to search the data, and means to extract/reduce/transform the data. A prototype is http://ibiomes.chpc.utah.edu.
- Ideally modes of data access include at data source, with computational capability to analyze, via Globus Transfer for large-scale transfer, and also via http: to enable easy access (likely for reduced data).

• Providing the means for users to choose how to generate derived and/or reduced data, move to local site, and do analysis.

Persistence of Data

The issue of long-term storage of data is a thorny one for XSEDE and service providers. There are really several issues.

The first is, when does a simulation data set become of broad interest and one that should be preserved? This is really a decision for the individual disciplines to make. We have to realize that there is a financial annual cost to keeping data, and the relevant community will have to decide whether this data merits that continued investment. The decision will rest not only on the intrinsic interest of the dataset, but also on whether the metadata is properly structured to permit convenient access to others.

The second issue is how to guarantee persistence of the availability of the data since a service provider (SP) may not be renewed beyond the expiration date of its current NSF award. This has precluded XSEDE and SPs from entering into reasonable arrangements with groups that have asked XSEDE/SPs to store their data. In most cases, the owners of the data do not insist that the data be held at the same physical location for the entire period for which they want to contract. NSF could alleviate the problem by assuring the community that storage resources would be available, even if a particular SP's award terminated. SP awards would include assurance that if their award was not renewed, funds would be provided to transfer the data they are holding to another site within the NSF program.

The third issue is the cost of storage. Storage costs are, by and large, not included in the operations budgets of the SPs. Owners of the data, or their communities, will have to find a way to fund the long-term storage of data they deem important. If a pricing system were established (\$/TB/year), SPs receiving data transferred from an SP being terminated would have to accept data from the latter SP provided that the owners paid the new SP the established price. Communities could then plan for the longer term storage of their data, independent of the long-term viability of a particular SP.

Breakout 6 - Workflow & Methods

Charge: In reviewing the submitted white papers and other information, analyze the requirements for advanced workflow and methods that will need to be met to enable high-spectrum science and engineering in the next decade. Time to insight in many science areas is not only related to the largest scale work steps, but may be dominated by other workflow steps. Today, many workflows have been developed on a project or team basis, with many different assumptions. Project, where possible, the commonalities of workflows and where their could be opportunities for synergies and optimizations.

Current Assessment

As computational science becomes more multidisciplinary and more data-intensive, automated workflow engines yield both quantitative and qualitative gains in scientific discovery. Every user has a workflow, even if they don't call it that. While manual workflows are sufficient for small-scale, "one-off" projects, interactive analyses and complex execution of tasks that can rapidly adapt to prior results require automation. There are several mature workflow managers available, including general-purpose packages such as Pegasus, Kepler and Swift, as well as domain-specific packages such as Copernicus. The adoption of such mature and well-developed workflow managers is hindered by several factors. Workflow managers are needed to automate repetitive tasks, reduce human error and increase the scope of the science being conducted through computations. Frequently research groups develop *ad-hoc* workflow managers to complete tasks, but because of limited time and funding such software is often neither robust, portable, or flexible for changing tasks. Lack of awareness is a significant hurdle for broader adoption of well-supported and robust software packages. Another challenge to adoption is the learning curve for making use of these products. In addition, specific areas of science require domain-specific specialization for these tools to be adopted by research groups.

Adoption of robust, efficient workflow managers will greatly enhance scientific discovery and increase both researcher productivity and efficient use of research facilities. Well-designed and developed workflow managers also offer the potential for better data management as research becomes more data-intensive.

Challenges

Awareness. Many workflows are developed by individual investigators and/or small groups who are not aware of mature workflow solutions. Home-grown workflows often have limited portability across machines and are fragile. Users should be made aware of existing workflow management software and its capabilities and should be encouraged to use this software where appropriate. In some cases, it may even be useful to shift work from job-level code to the workflow level.

Migration. The learning curve associated with workflow management systems may discourage their use, particularly by users who have been using HPC resources for years or decades, without the benefit of workflow management. Hence, there should be some mechanism for mitigating the learning curve and easing adaptation of workflow management systems by new users.

Discipline-specific versus general-purpose software. For some fields such as GIS, it can be valuable to incorporate discipline-specific knowledge into the workflow management system. In contrast, there is also a large demand for generic workflow management solutions, so both discipline-specific and generic workflow software should be considered.

Sharing and re-use. An effective workflow management system should encode information regarding reuse, archiving, and curation of intermediate datasets, should integrate datasets across sub-communities, should deal with nontrivial dependencies (such as recompilation of code), and must interact with job schedulers on individual compute environments in ways that do not overload the scheduler and/or cause suboptimal scheduling.

Scaling. At 100-1,000x the scale of current workflow jobs, the demands both on the workflow engine itself, on job schedulers and system resource managers will increase substantially. At best, these demands will limit efficient usage of large-scale compute platforms. At worst, they will effectively create distributed denial of service (DDOS) attack symptoms against the workflow platform. Bottlenecks and failure points associated with large-scale task execution scenarios must therefore be an integral part of planning for next-generation computing.

Requirements

- Workflow engines must singly or in combination accommodate both static workflows and adaptive workflows that alter the tasks to be executed in response to either external stimuli or computation results.
- Some workflows are by nature interactive and require low latency while others are latencytolerant.
- The ability to share workflows and/or results, either to collaborators or publicly, is essential.
- Workflow engines should be reusable across fields (question of generality versus organic development per field, but ability for reuse and dissemination is critical)
- There should be robust solutions for workflow communication with queuing systems and job schedulers.

- Resources should be efficiently allocated by any workload engine that matches tasks with execution environments and heterogeneous resource capabilities and takes into account data locality and I/O subsystem capabilities.
- A consistent workflow interface should be available across all platforms, from campus computing to the largest scale machines, to provide easy interoperability.

Science Cases

Workflows for different science realms have commonalities, and it may be advantageous to identify those to increase the amount of reusable work. Industry spends significant time to ensure efforts are not duplicated. Workflow management can also facilitate data sharing, as shown by the lattice QCD community. Workflow systems can facilitate cross-domain collaboration on "grand challenge" problems, such as the GIS example, below.

It has become increasingly important to incorporate domain science knowledge into future workflow management systems for optimally harnessing high-spectrum computing systems to tackle interdisciplinary and multidisciplinary grand scientific challenges. As an example, in emergency management, a domain that cuts across a number of scientific fields (e.g. environmental science and engineering, geography and social sciences, and geosciences), domain knowledge embedded in geospatial big data and related analytics is important to guide the construction and execution of scientific workflows. The integration of domain knowledge requires engagement of pertinent communities to understand the interfaces and synergies for coupling data and models with high-spectrum computing resources. CyberGIS workflow management systems (e.g. GISolve middleware) serve this integration purpose while capturing geospatial characteristics of computational intensity for optimal scheduling of heterogeneous high-spectrum computational and data resources.

Biomolecular simulation increasingly relies on multiple simulations, with job preparation and analysis consuming a large fraction of time to insight. Workflow automation of these tasks both increases reliability and frees researchers to focus on the science rather than data jockeying. In genomics, workflows have developed into community standards for best practices in large-scale data analysis. Materials research requires stochastic integration of what are relatively tractable individual solutions. Aggregate computation is the key to improved accuracy. Different job sizes lead to heterogeneous tasks, and efforts must be redirected based on intermediate discoveries. Across disciplines, workflow management is particularly important to uncertainty quantification and/or ensemble measurements. Time is of particular essence in some fields, such as emergency management preparedness. The cyber-GIS community workflow system facilitates decision making based on interactive qualitative results rather than quantitative results. In general, faster turnaround leading to interactive qualitative results rather than quantitative results can provide benefits for scientific productivity. In space science, there is record keeping on the fly with an eye towards future computation. Saving limited, compacted data enables efficient resumption, however more data-intensive computations are expected in the future.

Industrial applications of robust workflows include organizing the large numbers of medium calculations used in design of experiments (DOE), uncertainty quantification (UQ), and probabilistic analysis. Industry will benefit from robust workflow software solutions that are well documented and supported by a community of developers. This will provide more efficient development and support of workflows in both industry and academia. Support for workflow solutions is needed much in the same way support is needed for system software that schedule jobs.

The drawbacks to workflow design are that it is difficult to know beforehand what workflow components will be used a few times and which will be become a core part of the entire workflow. Additionally,

workflows are generally built incrementally as researchers focus on the weakest or most time-consuming link before moving on to the next. Many times these intermediate steps are accomplished with Python scripts and may be repeated by multiple researchers, even among collaborators. Some groups, like those involved in SciDAC, use more collaborative processes. How can those efforts be identified and shared?

As an example of bypassing the necessity of an overly complicated workflow pattern, space physicists can save data along all known spacecraft trajectories, as well as at Earth and other planets, in the process of data-driven numerical simulation. This makes it possible to save data storage and increase the output cadence for time-dependent results.

Workflow planners should provide enhanced support for data re-use, provenance, and data curation. These are import issues to support large-scale data-driven computations in order to increase efficient use of resources, e.g., reusing previously generated data products. In addition, as the scale and complexity of workflows are expected to increase over time it will be important for sufficient metadata of large workflows be captured to provide support for validation of computational results that consumed significant resources.

Conclusions

Easy-to-use, scalable workflow methods will enable more efficient science and better usage of diverse HPC resources. In several disciplines, workflow management has become essential to performing science at scale. Examples include GIS, where interactive response (and anticipatory execution of analyses to enable this) are critical; and LIGO, where the scale of data-driven workloads demands workflow execution. In many cases, adaptive variation of the workflow graph is either necessary or substantially increases efficiency, while in some others the workflows are predominantly static. An increasing number of scientific problems can be addressed **only** with the use of workflow or similar engines. For a broader class of problems, good workflow software can greatly enhance the science-for-the-money on high-end computing resources (both researcher productivity and platform utilization). Robust workflow engines (and ease of use, domain-appropriateness, and interfaces with compute resources) will thus be a critical and highly cost-effective part of next-generation infrastructure platforms.

BREAKOUT WORKING REPORTS FROM WORKSHOP 2

Breakout 1 - Alternative Design Models

Attendees: Michael Levine, Pittsburgh Supercomputing Center; Rama Govindaraju, Google; Bill Barth, The Texas Advanced Computing Center; Victor Hazlewood, University of Tennessee – JICS; Liqiang Wang, University of Wyoming; William Kramer, NCSA/University of Illinois; Steve Wolff, Internet2; Robert Harrison, Stonybrook.

Charge: Identify and describe the alternative design models for providing resources and services. Example models could be implementation by discipline orientation (common resources shared by many/all disciplines to resources subdivided for discipline-specific resources); system integration and scale (tightly integrated to very loosely integrated); resource distribution (localized into a few facilities to highly geographically distributed); etc. Identify the points of "diminishing returns" for each of the models and the potential strengths and weaknesses for each. Each alternative should be described in two to three sentences.

Summary of Recommendations

The conclusion was that NSF should invest in two to four high-end HPCD centers. Although this does not assume the "centers model" funding from the 1980s, the group did identify issues with the current systems

model; pros and cons of a variety of models are discussed in more detail below. The group concluded that NSF should track commercial cloud services and trends for hardware innovations and should learn and track computing resources requirements submitted in funding proposals. A community-based study could be used to track these three key areas.

Alternative Design Models

A. HPCD Center Model: NSF competitively awards a small number of university-led data centers/facilities to provide resources over a 4+4 or 5+5 year timeframe, with options to further extend based on quality of service and past performance. The suggested number is two to four data centers/facilities.

Strengths: long-term awards/commitments to a small set of centers, efficient investments in infrastructure, expert staff retention, reviews to allow for feedback on accountability, transparency and ending funding for non-performance

Weaknesses: perception of limited competition and/or cronyism, potential lack of innovation or perceived lack of innovation

Discussion: NSF-supported a center model from 1986 to 1996 where funding was given to a center to support resources for five years with option for additional five years, and the program was generally successful. The consensus of the discussion was that one data center/facility is too few and 10 or 100 is significantly too many due to dilution of the available funding to support HPCD at a scale larger than what universities provide. XSEDE is essentially funded this way.

We find no compelling argument for a specific number of centers. Supporting more centers entrains the creative and innovative talents of a larger number of people, allows the program to explore more technical approaches, and eliminates a possible 'single point of failure.' However, it reduces potential economies of scale and bounds the maximum size of an individual center for a given total aggregate budget. DOE studies have estimated the additional cost of going from one to two centers at ~5% of total budget, and from one to four centers at ~13% of the total budget. We feel the benefits of multiple centers listed above, along with the political benefits of distribution, argue for four centers, but not necessary of equal capability. We observe that the NSF Centers program started with five centers and reduced that number to four through performance reviews.

B. Individual resource model (current model): NSF funds individual resource awards to organizations based on competitive solicitations with specific HPCD criteria

Strengths: short award cycles

Weaknesses: loss of efficiency of investment, less stable expert staff retention, greater transition efforts, less certain roadmaps for science teams

C. Discipline-based model: NSF funds resources through competitive solicitations

Though there was discussion and some consensus that a discipline-based model has some promise, there was not consensus that a discipline-based model be a singular model for supporting the acquisition of HPCD resources. This model should be reserved for very focused and particular cases.

E. Cloud Technology Model: The NIST Cloud Computing definition from NIST Special Publication 800-145^x: "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a

[×] http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction..."

Strengths: Adaptability to support end user supplied (possibly "verified") software stack for their application, on-demand provisioning of resources, simplified access via thin client

Weaknesses: Current cloud resources and technologies do not map to all or even "a significant fraction of" NSF's comprehensive, diverse and growing range of scientific research. Large tightly coupled HPCD systems, with fast interconnects, remain essential for a variety of research challenges.

Discussion: It is expected that cloud technology will continue to be one of the resource choices for HPCD support of NSF-sponsored scientific research for the future, but there is <u>no consensus</u> from the workshop participants that cloud technology will fully or even substantially replace all HPCD support for the diverse NSF scientific research community by 2017 or 2025. Nor will clouds "provide a viable integrated solution for a significant fraction of (but not all) data- and compute-intensive and combined workloads" as described by the interim NRC report. There is promise to cloud technology, possibly for periods where "surging" resource use is necessary, but cloud technology is not a panacea for NSF's HPCD resource needs. NSF should continue to closely follow cloud technology for HPCD applicability and should make investments as appropriate based on the trends and capabilities of cloud technology. Furthermore, there are few reliable studies of the cost-effectiveness of clouds for large amounts of sustained computing, and the ones that exist indicate cloud computing is not less expensive and probably more expensive than equivalent capability in HPCD centers^y.

Breakout 2 - Alternative Provisioning Methods

Charge: Identify the primary alternatives that NSF could use to provide access to the necessary resources for high-spectrum science, engineering and research. Triage these alternatives to identify the three to five that would likely be mostly feasible and implementable. For each alternative, identify the role(s) that NSF, universities, private industry and other federal agencies might have in making the alternative succeed. Each alternative should be described in two to three sentences.

Alternatives

- University-run centers
- Other agency-run (DOE) centers
- o Commercial providers: Google, Amazon, etc., vendors

Modality #1

- very large scale parallel simulation
- o tightly coupled computation
- o requires access to closely coupled large data resource

Roles:

• *NSF*: provide sufficient funding for large-scale, university-led national facilities, including system operation and user support; catalyze the development of the required high-end software stack, including application software

^y An example analysis is the DOE ASCR Magellan report - /http://science.energy.gov/~/media/ascr/pdf/programdocuments/docs/Magellan_Final_Report.pdf

- o other federal agencies: leveraged resources, expertise, shared experience
- o *universities*: provide, staff, and operate the facilities
- o *states*: in-kind support
- *regional providers*: high-end networking
- *private industry*: equipment vendors, software vendors, collaboration, in some cases even facilities, data storage vendor

Responsibilities:

- *NSF*: provide funding for acquisition of an independently-managed resource (including power, cooling), resource operations, user support services
- *Universities*: provide facility (building, building maintenance/operation)
 - Propose solutions, not necessarily vendors
- *Other federal agencies*: collaborate with providers as appropriate, to optimize user productivity and science impact
- *Industry*: act as vendors to the providers

Modality #2

- very large-scale parallel simulation
- tightly coupled computation
- o modest data requirements

Roles:

- *NSF*: provide sufficient funding for large-scale university-led national facilities, including system operation and user support; catalyze the development of the required high-end software stack, including application software
- *Universities*: provide, staff, and operate the facilities
- o Other federal agencies: leveraged resources, expertise, shared experience
- States: in-kind support
- *Regional providers*: high-end networking
- *Industry*: equipment vendors, software vendors, collaboration, in some cases even facilities

Responsibilities:

- *NSF*: provide funding for acquisition of an independently-managed resource (including power, cooling), resource operations, user support services
- *Universities*: provide facility (building, building maintenance/operation)
 - Propose solutions, not necessarily vendors
- *Other federal agencies*: collaborate with providers as appropriate, to optimize user productivity and science impact
- *Industry*: act as vendors to the providers

Modality #3

- o large number of moderate-scale parallel ensemble simulations
- o tightly-coupled computation within each ensemble member
- o may or may not require access to large data resource

Roles:

• *NSF*: provide sufficient funding for large-scale university-led national facilities, including system operation and user support; catalyze the development of required high-end software stack, including application software

- Universities: provide, staff, and operate the facilities
- *Other federal agencies*: leveraged resources, expertise, shared experience
- *States*: in-kind support
- Regional providers: high-end networking
- *Industry*: equipment vendors, software vendors, collaboration, in some cases even facilities, in some cases a data storage vendor

Responsibilities:

- *NSF*: provide funding for acquisition of an independently-managed resource (including power, cooling), resource operations, user support services
- *University*: provide facility (building, building maintenance/operation)
 - Propose solutions, not necessarily vendors
- *Other federal agencies*: collaborate with providers as appropriate, to optimize user productivity and science impact
- *Industry*: act as vendors to the providers

Modality #4

- Large-scale, distributed simulation/analysis
- loosely coupled computation
- o requires access to closely-coupled, large data resource

Roles

- *NSF*: provide sufficient funding for large-scale university-led national facilities, including system operation and user support; catalyze the development of required high-end software stack, including application software
- Universities: provide, staff, and operate the facilities
- *Other federal agencies*: leveraged resources, expertise, shared experience
- *States*: in-kind support
- *Regional providers*: high-end networking
- *Industry*: equipment vendors, software vendors, collaboration, in some cases even facilities, data storage vendor

Responsibilities:

- *NSF*: provide funding for acquisition of an independently-managed resource (including power, cooling), resource operations, user support services
 - The requirements for the large, tightly-coupled data resource are different than those for the leading-edge HPC resource.
- *University*: provide facility (building, building maintenance/operation)
 - Propose solutions, not necessarily vendors
- *Other federal agencies*: collaborate with providers as appropriate, to optimize user productivity and science impact
- *Industry*: act as vendors to the providers

Modality #5

- large scale, distributed simulation/analysis
- loosely coupled computation
- no requirement for closely-coupled data

Roles

- *NSF*: software development, including application development and support.
- Universities: resources for science.
- Other federal agencies: expertise, shared experience
- *States*: in-kind support
- *Regional providers*: networking
- o *Industry*: equipment vendors, software vendors, collaboration

Recommendations

- *NSF*: support software development
- *University*: typical faculty and student support, computing resources
- Industry: vendor and cloud service provider

Additional Comments

- Different disciplines and modalities have different HPCD requirements and may require different system balances and/or policies, but many can co-exist on the same systems rather than having special systems/facilities.
- All providers funded by NSF should be coordinated in order to deliver optimal resources and services to support the community's science objectives.
- NSF must provide long-term commitments to centers/facilities or similar organizations, in order to support career paths for technical staff, continuity of support for science teams and long-term planning.
- Software development is a long-term (multi-decadal) effort, which must have stable, long-term support.
- Cross-fertilization of software development efforts must be supported.
- All resources provided, including hardware and software development, contribute to workforce training and development in STEM areas.
- NSF should ensure that a variety of resources remain available which enable key capabilities important to the continued advancement of the scientific community.
- The scientific continuum, from desktop to supercomputer, must be preserved.

Breakout 3 – Alternative Analysis Criteria

Charge: Develop the important evaluation criteria that should be used to assess and select alternative methods for NSF to deploy high-spectrum resources. These criteria should include perspectives from all stakeholders and potential uses, but need to be aggregated so that there are no more than 10 criteria for assessment. Additionally, identify the important cost model components and describe them so the alternatives can be assessed. Each criterion should be described in two to three sentences.

Criteria

1) Alignment with NSF's Mission

- a) Model should support the progress of science; to advance the national health, prosperity and welfare; and to secure the national defense.
- b) Model should support transformative, high impact research
- c) Model should "foster and support the development and use of computers ... primarily for research and education in the sciences"

2) Open Access

- a) Model should lead to projects that support NSF-funded researchers, regardless of nationality
- b) Model should lead to projects that support non-NSF-funded researchers (e.g., campus-funded, NIH-funded, NEH-funded)

c) Model should lead to projects that support the private sector (e.g., industry, health care)

3) Strategic Horizon

- *a)* Model should enable research on a strategic arc of a longer term than private industry, individual institutions, or regional entities can/will fund
- *b)* Model should enable work with researchers on developing computing technologies that advance simulation and data science

4) Appropriateness

- a) Model should provide researchers with the type of computing/data/networking resources needed to advance science & engineering
- b) Model should evaluate programs based on full costs—financial, social, environmental, etc.

5) Robustness and Sustainability

- a) Model should provide continuity in resources and services with some degree of certainty in the capabilities that will become available.
- b) Model should withstand changes in funding mechanisms and sources, or changing technologies and priorities.

6) Flexibility

- a) Model should allow programs to respond to innovations in computing (hardware, software, operating modes, application domains, algorithms, etc.).
- b) Model should enable programs to respond to new categories of scientific exploration
- c) Model should be able to adapt to changing science priorities across NSF

7) Diversity

- a) Model should support a broad range of disciplinary research and computing modalities
- b) Model should provide a range of technologies and platforms

8) Total Cost of Ownership

- a) Model should encourage provision of cost-effective resources and services.
- b) Model should include a range of stakeholders in supporting long-term operations.

9) Complementarity

a) Model should support activities that are important to NSF's mission that are not funded by other agencies or the private sector.

10) Workforce development

- a) Model should support the development of the next generation of scientists and engineers who will advance high-performance computing and/or simulation science/data science.
- b) Model should provide career paths for HPC-trained scientists.

11) Computing/Data/Networking Ecosystem

- a) Model should integrate across the full scale of computing, data and networking resources—national to campus
- b) Model should enable transfer of knowledge and expertise between different disciplines and research domains, and between academia and the private sector.

Cost model components

The current relative investments in hardware, software, personnel, etc., may not be optimal in the future. It was difficult for us to determine what these relative investments should be in the future, but we note that more effective use of existing computing systems can be as valuable as more powerful computers.

- 1. Hardware
 - a. Computer systems
 - b. Data storage systems
 - c. Data analytics and display systems

- d. Networks and networking systems
- 2. Software
 - a. System software
 - b. Applications software
- 3. Personnel
 - a. Operations
 - b. Scientific support
 - c. Educational support
- 4. R&D costs
 - a. Materials
 - b. Components
 - c. Algorithms
 - d. Software
- 5. Ecosystem Costs
 - a. Tying the compute/data/network campus-national systems together
- 6. Operating Costs
 - a. Facility (purchase or lease)
 - b. Compute
 - c. Data storage
 - d. Data transfer
 - e. Maintenance
 - f. Utilities
- 7. Opportunity Costs
 - a. Costs in scientific progress from changes in investments

Scorecard Template

Model Being Evaluated:		
Criterion		Satisfaction Level, Comments
0 Mission	a	
MISSION	b	
	с	
1 Openness	а	
openness	b	
	с	
2 Horizon	а	
	b	
3 Appropriateness	a	
Appropriateness		
4 Robustness	а	
	b	
5 Flexibility	а	
	b	
	с	
6 Diversity	a	
	b	
7 Total Cost of Ownership	a	
_	b	
8 Complementarity	а	
9 Workforce	a	
	b	
10 Fcosystem	a	
Leosystem	b	

Breakout 4 – NSF Administrative Implementation Alternatives

Attendees: Mike Levine, Stuart Anderson (F), Victor Hazlewood, Liqiang Wang, Thomas Hauser (S), Honggao Liu, Jerzy Bernholc, Bill Kramer, Ed Seidel

Charge: Identify the primary implementation alternatives that NSF might implement to provision the necessary resources for high-spectrum science, engineering and research. What processes are available and feasible for NSF to implement (current methods of site/equipment awards, contracts, intergovernmental agreements, MREFC processes, etc.). Triage these alternatives to identify the three to five that would likely be mostly feasible and implementable. For each alternative, identify the role(s) that NSF, universities, private industry and other federal agencies can have in making the alternative succeed.

Criteria

- Long-term stability with strong review process; to address retention of expert staff
- Flexibility to make technology decisions as innovations occur
- Support of highly integrated services including storage, networking, compute, technical and scientific expertise
- National scale, broad support of multiple communities across all disciplines supporting open science
- Enable synergies and national coordination between other CI projects
- Scientific innovation
- Integration of industry and university research
- Education

Alternative Assessment

Commercial service provider model - outsourcing to commercial cloud providers

Pros

- Can be included in the other funding approaches
- Flexible scale of management overhead depending on needs
- Commercial cloud providers could provide service, but commitment and pricing unknown
- Hardware refresh managed by cloud provider
- Management and mining of large-scale data

Cons

- Currently does not cover most tightly coupled modalities of high-spectrum research
- Current pricing more expensive
- Unknown long-term price structure trending down investigate future price structure
- Extreme-scale highly coupled simulation currently not possible
- Decouple operations of hardware from services provided to researchers

Academia/other agencies/industry centers similar to UICRC - not a separate model

Major Research Equipment Facility and Construction program (MREFC)

Pros

- Long-term funding model from a separate account at NSF/not a research account
- Research and Related Activities (R&RA) would only have to pay for operating expenses
- Other large projects (NEON, LIGO) can leverage the CI investments in a CI MREFC

• Adding additional funding and coordinate CI approach to create an effective leading-edge national infrastructure with a diverse portfolio of resources and services

Cons

- current process too heavy weight, needs to be more flexible for provisioning CI resources
- Lack of innovation possible
- Adding and removing resource providers

Compete for every machine (i.e. current Track-2 program)

Pros

• Refresh of ideas and approaches

Cons

- Lack of long-term commitment
 - Limited retention of expert staff because of short funding cycles
 - Loss of infrastructure investments
- Coupling award with vendor

Longer-term funding of an organization (e.g. DOE and DoD model)

Pros

- Long-term commitment
 - Extract better value of infrastructure investment
 - Better retention of staff
 - Decoupling vendor from services

Cons

0

• Requires comprehensive review to avoid stagnation/inefficiencies

Federally Funded R&D Center, e.g NCAR model

Cons

• This was not considered a feasible option since no new FFRDCs have been created since the 1990s

Breakout 6 – Alternative Evaluation and Assessment – 2

Attendees: Barry Schneider, Dinshaw Balsara (F), David Dixon, Curtis Hillegas (S), Jorge Vinals, Steve Wolff, Frank Tsung, Jagannathan Ramanujam, Cristina Beldica, Greg Bauer

Charge: Given the alternatives developed in Breakouts # 1 and 2, and the evaluation criteria developed in Breakout #3, evaluate the alternatives with strengths and weaknesses, risks and uncertainties and perform an initial qualitative cost assessment.

Support models vs. modalities

- 1) Center-based and individual resource
 - existing working model
 - a) can encompass #2
 - b) Modality #1, #2, #3, #4
 - c) Commercial provider run the center?

2) discipline specific

• when is this an acceptable model?

- a) More cost effective?
- b) Modality #1, #2, #3, #4
- 3) loosely coupled multiple sites
 - a) Modality #3, #5
- 4) cloud model (4a academic, 4b commercial)
 - Modality #1, #2, #3, #4, #5 possibly in the future if we consider that it will be a commercial "center"
 - a) Do the commercial cloud providers want to become a player in this area? Could they serve an area that fits their model?

	0	1	2	3	4	5	6	7	8	9	10
1 Center	Х	Х	Х	Х	X?	Х	Х	I	Х	XX	Х
2 Discipline	Х	Х	Х	Х	X?	Х		Ι	Х	XX	
3 Loose	Х	Х	Х	Х	X?	XX	XX	Ι	Х	XX	Х
4a EDU Cloud	Х	Х	Х	Х	?	XX	XX	Ι	х	XX	Х
4b Commercial Cloud	X?	Х?	?	Х	?	Х	XX	Ι	Х	Х	Х

Support models vs. Criteria

X-Meets; XX-Really meets; X?-May meet, but some caveats; ?-not sure; I-Point of information. 0-Alignment with mission; 1-Open Access; 2-Strategic Horizon; 3-Appropriateness; 4-Robustness & Sustainability; 5-Flexibility; 6-Diversity; 7-Total Cost of Ownership; 8-Complementarity; 9-Workforce Development; 10-Ecosystem

The difference between getting a cycle allocation vs. getting money to pay for cloud cycles is a huge cultural change. There is concern about how this funding for cloud use would be treated by institutions. For example, would universities apply overhead costs that would diminish the amount of funding to the research team? Another complexity is how to implement the acquisition/purchase of cloud services, which would probably require extra work by researchers.

Over the larger horizon, commercial providers may become financially feasible, but at that time consideration would need to be taken to determine if they can meet all the uses and needs of the community (including educational mission). Bring them in as a "Tier 2" system and see if they can function in this space?

Have to include software development and user support costs into the assessment. Typically cloud service providers do not provide the range of intellectual services (consulting, training, porting, optimization, etc.) that HPCD centers do. That would have to be provided in other ways to make research teams productive

NIH does not provide their own computing infrastructure. They use the NSF systems, but do not contribute or pay to use.

XSEDE is an example of how criterion 10 (Ecosystem) could be implemented.

	Hardware	Software	Personnel	R&D	Ecosystem	Operating	Opportunity
Model 1	Н	Н					
Model 2	L,M,H	Н					

Support models vs. cost models

Model 3	М	Н			
Model 4a	L	Н			
Model 4b	L	Н			

An idea suggested was to locate domain experts at only one center. For example, support for chemistry would come from TACC (regardless of what facility codes were run on.) Physics might be at NCSA, etc. This would possible improve synergy of support staff. But disassociating any given domain expert from an individual source or location presents coordination challenges for the expertise across all resources. Another challenge may be how to support new and emerging disciplines.

Do we need personnel onsite 24x7?

RESULTS OF WORKSHOP 1 REALTIME EVENING BRAINSTORMING SESSION 1

"Elevator" Statements by Attendees

What would you tell NSF about this workshop so far?

- The global nature of industry and science makes investing in computational research an imperative for American competitiveness and our future position in the world order.
- I enjoyed writing the position paper and the workshop is off to a good start—everyone is engaged and the breakout sessions were fruitful. Looking forward to the next 1.5 days.
- It is important that we emphasized that NSF should support method and software development in coordination with stronger support for hardware.
- There is a clear need for both a compute-intensive and a data-intensive high-end supercomputer, as well as support for code development for and training on these platforms.
- We generated a list of potential recommendations to improve allocation and efficiency of use of NSF computational resources and determined that the most valuable step would be to double NSF spending on cyberinfrastructure.
- Major outcomes today were: (1) science and technology require {major?} further advancements in computation hardware and software; (2) major resources are needed to handle the ever-increasing data being produced; and (3) similar advances are needed in visualization and presentation, communication tools especially for some HPC uses, such as climate change.
- NSF support for HPC to advance our scientific knowledge is indispensable.
- The whole science/engineering community is desperate for NSF to continue, even enhance, support for HPC. American scientific leadership in the world is at stake.
- Merged compute and data-intensive high-performance platforms are needed to maintain U.S. scientific leadership and derive full value from billion dollar experiments and observatories.
- It is important that NSF capture the computational needs of the full set of NSF research activities as early in the approval process as possible to match the aggregate need with the total available computing resources.
- Since computing is becoming an integral part of many areas of science, NSF should support computing for NSF-funded research, on all scales from campus computing (via MRI-style support dedicated to computing) to tera/peta/exa scale.
- Extremely satisfied that our community understands the importance of data storage and sharing. Very soon this will become essential for realistic high-performance simulations.

- NSF could commit to a "cloud service" business model without using commercial cloud providers.
- There are numerous potential interaction and collaboration opportunities between DOE and NSF computing/data programs.
- Data-driven complex workflows are important to many science communities and represent a high priority for the future high-spectrum systems to support.
- Demands for research computing are increasing across many science domains, while oversubscription is real.
- I learned that the HPCD community needs a range of resources, including the largest capability systems and also very large systems for medium-scale jobs.
- The meeting confirmed my impressions that the HPCD needs are very divers and not the same among different use groups. Consequently, it seems difficult to prioritize the needs. In any case, NSF must continue to invest in HPCD. It's hard to believe why so many dollars go into large experimental setups while there does not seem to be a clear mechanism in place to ensure theoretical/computational support.
- In the breakout session, it was discussed that in fact the number of applications for Blue Waters resources through the PRAC is well below what is expected. People simply don't apply because they don't know that their code might be scalable or their project suitable for Blue Waters, and that is why the machine is undersubscribed.
- Despite the wide diversity of scientific applications, there is a surprising degree of convergence of HPCD needs into a few different compute and analytic paradigms.
- Many different groups are concerned about porting their codes to new architectures, particularly to addressing the issue of performance portability.
- HPCD is a diverse area and there is used for different architectures to satisfy the need of the various groups and areas. Expecting that all groups can use the same architectures effectively is counterproductive.
- NSF support for HPC to advance scientific knowledge is indispensable.
- To drive home to the disciplinary program directors that HPC resources are limited, they should be informed of how much their grantees have been cut back from resources and allocations. And to scrutinize them further, perhaps applications for funding should indicate how many SUs they will request if the project is funded.

Results of Workshop 1 Realtime Evening Brainstorming Session 2

Task: Describe your discipline's scientific goals that require further investment in HPCD resources and services in one sentence that is understandable and motivating to the general scientifically supportive public.

- In 10 years, we will be able to predict the molecular mechanisms for protein conformational changes on the time and space scales observable via light microscopy.
- For moderately sized bio-molecular systems (riboswitch, enzyme) we will be able to fully elucidate the conformational ensemble and assess/calibrate experimental interpretations.
- Starting at the initial fluctuations from the Big Bang, we will be able to model the shapes and sizes of galaxies in all environments done the smallest scales observable.
- The synthesis of lattice QCD calculations and new experimental results from the Large Hadron Collider will shed light on physics that lies beyond the standard model of particle physics and the Higgs Boson.
- Emergency management will be able to take advantage of geospatial big data, cyberGIS analytics and high-spectrum computing for achieving scalable planning, preparedness and response.

- With exascale computers, scientists will be able to finally model rain formation reliably, taking into account small-scale turbulence, phase change, thermodynamics and atmospheric science. This will be a breakthrough highly relevant to predication of the effects of hurricanes.
- Without advanced computing, experimentalists and modelers in turbulence will not be able to validate their results, and design of practical engineering devices will suffer accordingly.
- In five years it will be possible to have adaptive, scalable software that enables computer engineering to be computed on computers with capabilities of running at 100-200 Petaflops.
- In 10 years, lattice QCD calculations will increase in precision from 1% to 0.1% to match the precision of particle physics experiments searching for evidence of physics beyond the standard model.
- Understanding the function of proteins by being able to simulate them with atomic resolution for realistic time scales (3-4 orders of magnitude longer than achievable today) leads to a fundamental understanding of their malfunction in various diseases.
- A major goal in computational nuclear physics is to understand from first principles the fusion of three alpha-particles to form carbon, a key process in the formation of elements (nucleosynthesis).
- Solve the quantum mechanical few body problem.
- I would like to see all science codes running at >25% of peak of the limiting resource (scalar/parallel Flops, memory bandwidth, interconnect, etc.)
- Possibly make a material or device design insight to keep Moore's Law alive.
- Do a totally *ab-initio* calculation of the simplest high-temperature superconductor and make comparisons with experiments, which would service as a landmark in materials genome (discovery by predictive computing)
- Supercomputing will enable climate scientists to make accurate enough predictions to improve agricultural practices to enable the earth to sustain all people on it.
- Advances in high-spectrum computing are needed to improve our understanding of everyday physical phenomena, such as turbulence, to improve our ability to apply science for the good of mankind.
- The primary science goal of gravitational wave physics in the next five years is the detection of gravitational waves and the beginning of observational gravitational wave astronomy with the advanced LIGO instrument.
- In the next five years, we aim to carry out enough long, accurate full general relativistic simulations of black hole/black hole and neutron star/black hole binaries to provide wave forms to test general relativity via LIGO observations to the full accuracy of the instrument. In the next 10 years, a key goal is to carry out general relativity + relativistic magnetohydrodynamics + neutrino transport simulations with enough accuracy and enough physics to solve the supernova mechanism problems, as well as nucleo-synthesis from neutron stare binary mergers.
- In the next five years, the space physics community should expect hybrid, multi-scale numerical models capable of combining an magnetohydrodynamics approach globally while solving the kinetic Boltzmann equations locally, which will need exascale supercomputers and petascale data files. Solving the Boltzmann equations directly will require six-dimensional parallelization strategies and innovative data-handling technologies.
- Over the next 5 years, the focus in climate science will be on getting very high-resolution models ready and running for the next major international climate model intercomparison (CMIP6), which will likely produce 5-10 PB of data. The analyses of past, present, and future projections of climate from CMIP6 will provide important input into the next IPCC and NCA assessments as well as contributing extensively to impacts and policy-related analyses. Our aim for CMIP6 is to have some global models run at as high a resolution as 1/4° (25 km resolution). Over the next 10 years, we will

be aiming for even higher resolution, perhaps more like 4-10 km models that are cloud-resolving. Although improvements in the representation of physical, chemical, and biological processes in the models are also important, high-resolution simulations are imperative because of the enhanced, detailed information they can give through better representation of orography and local processes, and resulting effects on synoptic scale weather structures across the planet.

REVISION CONTROL

Revision Number	Revision Date	Revision Description	Person Making the Revision
1.0	7/15/2016	First complete and final version of the public workshop report.	William Kramer